

Tensors in computations

Lek-Heng Lim

*Computational and Applied Mathematics Initiative,
University of Chicago, Chicago, IL 60637, USA
E-mail: lekheng@uchicago.edu*

The notion of a tensor captures three great ideas: equivariance, multilinearity, separability. But trying to be three things at once makes the notion difficult to understand. We will explain tensors in an accessible and elementary way through the lens of linear algebra and numerical linear algebra, elucidated with examples from computational and applied mathematics.

CONTENTS

1	Introduction	555
2	Tensors via transformation rules	566
3	Tensors via multilinearity	595
4	Tensors via tensor products	640
5	Odds and ends	746
	References	750

1. Introduction

We have two goals in this article: the first is to explain in detail and in the simplest possible terms what a tensor is; the second is to discuss the main ways in which tensors play a role in computations. The two goals are interwoven: what defines a tensor is also what makes it useful in computations, so it is important to gain a genuine understanding of tensors. We will take the reader through the three common definitions of a tensor: as an object that satisfies certain transformation rules, as a multilinear map, and as an element of a tensor product of vector spaces. We will explain the motivations behind these definitions, how one definition leads to the next, and how they all fit together. All three definitions are useful in computational mathematics but in different ways; we will intersperse our discussions of each definition with considerations of how it is employed in computations, using the latter as impetus for the former.

Tensors arise in computations in one of three ways: equivariance under coordinate changes, multilinear maps and separation of variables, each corresponding to one of the aforementioned definitions.

Separability. The last of the three is the most prevalent, arising in various forms in problems, solutions and algorithms. One may exploit separable structures in ordinary differential equations, in integral equations, in Hamiltonians, in objective functions, *etc.* The separation-of-variables technique may be applied to solve partial differential and integral equations, or even finite difference and integro-differential equations. Separability plays an indispensable role in Greengard and Rokhlin's fast multipole method, Grover's quantum search algorithm, Hartree–Fock approximations of various stripes, Hochbaum and Shanthikumar's non-linear separable convex optimization, Smolyak's quadrature, and beyond. It also underlies tensor product construction of various objects – bases, frames, function spaces, kernels, multiresolution analyses, operators and quadratures among them.

Multilinearity. Many common operations, ranging from multiplying two complex numbers or matrices to the convolution or bilinear Hilbert transform of functions, are bilinear operators. Trilinear functionals arise in self-concordance and numerical stability, and in fast integer and fast matrix multiplications. The last leads us to the matrix multiplication exponent, which characterizes the computational complexity of not just matrix multiplication but also inversion, determinant, null basis, linear systems, LU/QR/eigenvalue/Hessenberg decompositions, and more. Higher-order multilinear maps arise in the form of multidimensional Fourier/Laplace/Z/discrete cosine transforms and as cryptographic multilinear maps. Even if a multivariate function is not multilinear, its derivatives always are, and thus Taylor and multipole series are expansions in multilinear maps.

Equivariance. Equivariance is an idea as old as tensors and has been implicitly used throughout algorithms in numerical linear algebra. Every time we transform the scalars, vectors or matrices in a problem from one form to another via a sequence of Givens or Jacobi rotations, Householder reflectors or Gauss transforms, or through a Krylov or Fourier or wavelet basis, we are employing equivariance in the form of various 0-, 1- or 2-tensor transformation rules. Beyond numerical linear algebra, the equivariance of tools from interior point methods to deep neural networks is an important reason for their success. The equivariance of Newton's method allows it to solve arbitrarily ill-conditioned optimization problems in theory and highly ill-conditioned ones in practice. AlphaFold 2 conquered the protein structure prediction problem with an equivariant neural network.

Incidentally, out of the so-called 'top ten algorithms of the twentieth century' (Dongarra and Sullivan 2000), fast multipole (Board and Schulten 2000), fast Fourier transform (Rockmore 2000), Krylov subspaces (van der Vorst 2000), Francis's QR (Parlett 2000) and matrix decompositions (Stewart 2000) all draw from these three aforementioned tensorial ideas in one way or another. Nevertheless,

by ‘computations’ we do not mean just algorithms, although these will constitute a large part of our article; the title of our article also includes the use of tensors as a tool in the analysis of algorithms (*e.g.* self-concordance in polynomial-time convergence), providing intractability guarantees (*e.g.* cryptographic multilinear maps), reducing complexity of models (*e.g.* equivariant neural networks), quantifying computational complexity (*e.g.* exponent of matrix multiplication) and in yet other ways. It would be prudent to add that while tensors are an essential ingredient in the aforementioned computational tools, they are rarely the *only* ingredient: it is usually in combination with other concepts and techniques – calculus of variations, Gauss quadrature, multiresolution analysis, the power method, reproducing kernel Hilbert spaces, singular value decomposition, *etc.* – that they become potent tools.

The article is written with accessibility and simplicity in mind. Our exposition assumes only standard undergraduate linear algebra: vector spaces, linear maps, dual spaces, change of basis, *etc.* Knowledge of numerical linear algebra is a plus since this article mainly targets computational mathematicians. As physicists have played an outsize role in the development of tensors, it is inevitable that motivations for certain aspects, explanations for certain definitions, *etc.*, are best understood from a physicist’s perspective and to this end we will include some discussions to provide context. When discussing applications or the physics origin of certain ideas, it is inevitable that we have to assume slightly more background, but we strive to be self-contained and limit ourselves to the most pertinent basic ideas. In particular, it is not our objective to provide a comprehensive survey of all things tensorial. We focus on the foundational and fundamental; results from current research make an appearance only if they illuminate these basic aspects. We have avoided a formal textbook-style treatment and opted for a more casual exposition that hopefully makes for easier reading (and in part to keep open the possibility of a future book project).

1.1. Overview

There are essentially three ways to define a tensor, reflecting the chronological evolution of the notion through the last 140 years or so:

- ① as a multi-indexed object that *satisfies certain transformation rules*,
- ② as a multilinear map,
- ③ as an element of a tensor product of vector spaces.

The key to the coordinate-dependent definition in ① is the emphasized part: the transformation rules are what define a tensor when one chooses to view it as a multi-indexed object, akin to the group laws in the definition of a group. The more modern coordinate-free definitions ② and ③ automatically encode these transformation rules. The multi-indexed object, which could be a hypermatrix, a polynomial, a differential form, *etc.*, is then a coordinate representation of either ② or ③ with respect to a choice of bases, and the corresponding transformation rule is a change-of-basis theorem.

Take the following case familiar to anyone who has studied linear algebra. Let \mathbb{V} and \mathbb{W} be finite-dimensional vector spaces over \mathbb{R} . A linear map $f: \mathbb{V} \rightarrow \mathbb{W}$ is an order-2 tensor of covariant order 1 and contravariant order 1; this is the description according to definition ②. We may construct a new vector space $\mathbb{V}^* \otimes \mathbb{W}$, where \mathbb{V}^* denotes the dual space of \mathbb{V} , and show that f corresponds to a unique element $T \in \mathbb{V}^* \otimes \mathbb{W}$; this is the description according to definition ③. With a choice of bases $\mathcal{B}_{\mathbb{V}}$ and $\mathcal{B}_{\mathbb{W}}$, f or T may be represented as a matrix $A \in \mathbb{R}^{m \times n}$, where $m = \dim \mathbb{V}$ and $n = \dim \mathbb{W}$. Note, however, that the matrix A is not unique and depends on our choice of bases, and different bases $\mathcal{B}'_{\mathbb{V}}$ and $\mathcal{B}'_{\mathbb{W}}$ would give a different matrix representation $A' \in \mathbb{R}^{m \times n}$. The change-of-basis theorem states that the two matrices are related via the transformation rule¹ $A' = X^{-1}AY$, where X and Y are the change-of-basis matrices on \mathbb{W} and \mathbb{V} respectively. Definition ① is essentially an attempt to define a linear map using its change-of-basis theorem – possible but awkward. The reason for such an awkward definition is one of historical necessity: definition ① had come before any of the modern notions we now take for granted in linear algebra – vector space, dual space, linear map, bases, *etc.* – were invented.

We stress that if one chooses to work with definition ①, then it is the transformation rule/change-of-basis theorem and not the multi-indexed object that defines a tensor. For example, depending on whether the transformation rule takes $A \in \mathbb{R}^{n \times n}$ to XAY^{-1} , XAY^{\top} , $X^{-1}AY^{-\top}$, XAX^{-1} , XAX^{\top} or $X^{-1}AX^{-\top}$, we obtain different tensors with entirely different properties. Also, while we did not elaborate on the change-of-basis matrices X and Y , they play an important role in the transformation rule. If \mathbb{V} and \mathbb{W} are vector spaces without any additional structures, then X and Y are just required to be invertible; but if \mathbb{V} and \mathbb{W} are, say, norm or inner product or symplectic vector spaces, then X and Y would need to preserve these structures too. More importantly, every notion we define, every property we study for any tensor – rank, determinant, norm, product, eigenvalue, eigenvector, singular value, singular vector, positive definiteness, linear system, least-squares problems, eigenvalue problems, *etc.* – must conform to the respective transformation rule. This is a point that is often lost; it is not uncommon to find mentions of ‘tensor such and such’ in recent literature that makes no sense for a tensor.

As we will see, the two preceding paragraphs extend in a straightforward way to order- d tensors (henceforth d -tensors) of contravariant order p and covariant order $d - p$ for any integers $d \geq p \geq 0$, of which a linear map corresponds to the case $p = 1, d = 2$.

While discussing tensors, we will also discuss their role in computations. The most salient applications are often variations of the familiar *separation-of-variables* technique that one encounters when solving ordinary and partial differential equations, integral equations or even integro-differential equations. Here the relevant perspective of a tensor is that in definition ③; we will see that

¹ Equivalently, $A = XA'Y^{-1}$. In numerical linear algebra, there is a tendency to view this as matrix decomposition as opposed to change of basis.

$L^2(X_1) \otimes L^2(X_2) \otimes \cdots \otimes L^2(X_d) = L^2(X_1 \times X_2 \times \cdots \times X_d)$, that is, multivariate functions are tensors in an appropriate sense. Take $d = 3$ for simplicity. Separation of variables, in its most basic manifestation, exploits functions $f: X \times Y \times Z \rightarrow \mathbb{R}$ that take a multiplicatively separable form

$$f(x, y, z) = \varphi(x)\psi(y)\theta(z), \quad (1.1)$$

or, equivalently,

$$f = \varphi \otimes \psi \otimes \theta. \quad (1.2)$$

For real-valued functions, this defines the tensor product \otimes , that is, (1.2) simply means (1.1). A tensor taking the form in (1.2) is called *decomposable* or *pure* or *rank-one* (the last only if it is non-zero). This deceptively simple version is already an essential ingredient in some of the most important algorithms, such as Greengard and Rokhlin's fast multipole method, Grover's quantum search algorithm, Hochbaum and Shanthikumar's non-linear separable convex optimization, and more.

A slightly more involved version assumes that f is a sum of separable functions,

$$f(x, y, z) = \sum_{i=1}^r \varphi_i(x)\psi_i(y)\theta_i(z),$$

or, equivalently,

$$f = \sum_{i=1}^r \varphi_i \otimes \psi_i \otimes \theta_i, \quad (1.3)$$

and this underlies various fast algorithms for evaluating bilinear operators, which is a 3-tensor in the sense of definition ②. If one allows for a slight extension of ② to modules (essentially a vector space where the field of scalars is a ring like \mathbb{Z}), then various fast integer multiplication algorithms such as those of Karatsuba, Toom and Cook, Schönhage and Strassen, and Fürer may also be viewed in the light of ②. But even with the standard definition of allowing only scalars from \mathbb{R} or \mathbb{C} , algorithms that exploit (1.3) include Gauss's algorithm for complex multiplication, Strassen's algorithm for fast matrix multiplication/inversion and various algorithms for structured matrix–vector products (*e.g.* Toeplitz, Hankel, circulant matrices). These algorithms show that a combination of both perspectives in definitions ② and ③ can be fruitful.


An even more involved version of the separation-of-variables technique allows more complicated structures on the indices such as

$$f(x, y, z) = \sum_{i,j,k=1}^{p,q,r} \varphi_{ij}(x)\psi_{jk}(y)\theta_{ki}(z),$$

or, equivalently,

$$f = \sum_{i,j,k=1}^{p,q,r} \varphi_{ij} \otimes \psi_{jk} \otimes \theta_{ki}. \quad (1.4)$$

This is the celebrated *matrix product state* in the tensor network literature. Techniques such as DMRG simplify computations of eigenfunctions of Schrödinger operators by imposing such structures on the desired eigenfunction. Note that (1.4), like (1.3), is also a decomposition into a sum of separable functions but the indices are captured by the following graph:



More generally, for any undirected graph G on d vertices, we may decompose a d -variate function into a sum of separable functions in a similar manner whereby the indices are summed over the edges of G and the different variables correspond to the different vertices. Such representations of a tensor are often called *tensor network states* and the graph G a *tensor network*. They often make an effective ansatz for analytical and numerical solutions of PDEs arising from quantum chemistry.

The oldest definition of a tensor, *i.e.* definition ①, encapsulates a fundamental notion, namely that of *equivariance* (and also *invariance*, as a special case) under coordinate transformations, which continues to be relevant in modern applications. For example, the affine invariance of Newton's method and self-concordant barrier functions plays an important role in interior point methods for convex optimization (Nesterov and Nemirovskii 1994); exploiting equivariance in deep convolutional neural networks has led to groundbreaking performance on real-world tasks such as classifying images in the CIFAR10 database (Cohen and Welling 2016) and, more recently, predicting the shape of protein molecules in the CASP14 protein-folding competition (Callaway 2020, Service 2020).² More generally, we will see that the notion of equivariance has guided the design of algorithms and methodologies through the ages, and is implicit in many classical algorithms in numerical linear algebra.

The three notions we raised – separability, multilinearity, equivariance – largely account for the usefulness of tensors in computations, but there are also various measures of the 'size' of a tensor, notably tensor ranks and tensor norms, that play an indispensable auxiliary role. Of these, two notions of rank and three notions of

² Research papers that describe AlphaFold 2 are still unavailable at the time of writing but the fact that it uses an equivariant neural network (Fuchs, Worrall, Fischer and Welling 2020) may be found in Jumper's slides at https://predictioncenter.org/casp14/doc/presentations/2020_12_01_TS_predictor_AlphaFold2.pdf.

norm stand out for their widespread applicability. Given the decomposition (1.5), a natural way to define rank is evidently

$$\text{rank}(f) := \min \left\{ r : f = \sum_{i=1}^r \varphi_i \otimes \psi_i \otimes \theta_i \right\}, \quad (1.5)$$

and a natural way to define a corresponding norm is

$$\|f\|_v := \inf \left\{ \sum_{i=1}^r \|\varphi_i\| \|\psi_i\| \|\theta_i\| : f = \sum_{i=1}^r \varphi_i \otimes \psi_i \otimes \theta_i \right\}. \quad (1.6)$$

If we have an inner product, then its dual norm is

$$\|f\|_\sigma := \sup \frac{|\langle f, \varphi \otimes \psi \otimes \theta \rangle|}{\|\varphi\| \|\psi\| \|\theta\|}. \quad (1.7)$$

Definitions (1.5), (1.6) and (1.7) are the *tensor rank*, *nuclear norm* and *spectral norm* of f respectively. They generalize to tensors of arbitrary order d in the obvious way, and for $d = 2$ reduce to rank, nuclear norm and spectral norm of linear operators and matrices. The norms in (1.6) and (1.7) are in fact the two standard ways to turn a tensor product of Banach spaces into a Banach space in functional analysis, and in this context they are called projective and injective norms respectively.

From a computational perspective, we will see in Sections 3 and 4 that in many instances there is a tensor, usually of order three, at the heart of a computational problem, and issues regarding complexity, stability, approximability, *etc.*, of that problem can often be translated to finding or bounding (1.5), (1.6) and (1.7). One issue with (1.5), (1.6) and (1.7) is that they are all NP-hard to compute when $d \geq 3$. As computationally tractable alternatives, we also consider the *multilinear rank*

$$\mu\text{rank}(f) := \min \left\{ (p, q, r) : \sum_{i,j,k=1}^{p,q,r} \varphi_i \otimes \psi_j \otimes \theta_k \right\},$$

and the *Frobenius norm* (also called the *Hilbert–Schmidt norm*)

$$\|f\|_F = \sqrt{\langle f, f \rangle}.$$

While there are many other notions in linear algebra such as eigenvalues and eigenvectors, singular values and singular vectors, determinants, characteristic polynomials, *etc.*, that do generalize to higher orders $d \geq 3$, their value in computation is somewhat limited relative to ranks and norms. As such we will only mention these other notions in passing.

1.2. Goals and scope

We view our article in part as an elementary alternative to the heavier treatments in Bürgisser, Clausen and Shokrollahi (1997) and Landsberg (2012, 2017, 2019), all

excellent treatises on tensors with substantial discussions of their role in computations. Nevertheless they require a heavy dose of algebra and algebraic geometry, or at least a willingness to learn these topics. They also tend to gloss over analytic aspects such as tensor norms. There are also excellent treatments such as those of Light and Cheney (1985), Defant and Floret (1993), Diestel, Fourie and Swart (2008) and Ryan (2002) that cover the analytic perspective, but they invariably focus on tensors in infinite-dimensional spaces, and with that comes the many complications that can be avoided in a finite-dimensional setting. None of the aforementioned references are easy reading, being aimed more at specialists than beginners, with little discussion of elementary questions.

Our article takes a middle road, giving equal footing to both algebraic and analytic aspects, that is, we discuss tensors over vector spaces (algebraic) as well as tensors over norm or inner product spaces (analytic), and we explain why they are different and how they are related. As we target readers whose main interests are in one way or another related to computations – numerical linear algebra, numerical PDEs, optimization, scientific computing, theoretical computer science, machine learning, *etc.* – and such a typical target reader would tend to be far more conversant with analysis than with algebra, the manner in which we approach algebraic and analytic topics is calibrated accordingly. We devote considerable length to explaining and motivating algebraic notions such as modules or commutative diagrams, but tend to gloss over analytic ones such as distributions or compact operators. We assume some passing familiarity with computational topics such as quadrature or von Neumann stability analysis but none with ‘purer’ ones such as formal power series or Littlewood–Richardson coefficients. We also draw almost all our examples from topics close to the heart of computational mathematicians: kernel SVMs, Krylov subspaces, Hamilton equations, multipole expansions, perturbation theory, quantum chemistry and physics, semidefinite programming, wavelets, *etc.* As a result almost all our examples tend to have an analytic bent, although there are also exceptions such as cryptographic multilinear maps or equivariant neural networks, which are more algebraic.

Our article is intended to be completely elementary. Modern studies of tensors largely fall into four main areas of mathematics: algebraic geometry, differential geometry, representation theory and functional analysis. We have avoided the first three almost entirely and have only touched upon the most rudimentary aspects of the last. The goal is to show that even without bringing in highbrow subjects, there is still plenty that could be said about tensors; all we really need is standard undergraduate linear algebra and some multivariate calculus – vector spaces, linear maps, change of basis, direct sums and products, dual spaces, derivatives as linear maps, *etc.* Among this minimal list of prerequisites, one would find that dual spaces tend to be a somewhat common blind spot of our target readership, primarily because they tend to work over inner product spaces where duality is usually a non-issue. Unfortunately, duality cannot be avoided if one hopes to gain any reasonably complete understanding of tensors, as tensors defined over inner product spaces

represent a relatively small corner of the subject. A working familiarity with dual spaces is a must: how a dual basis behaves under a change of basis, how to define the transpose and adjoint of abstract linear maps, how taking the dual changes direct sums into direct products, *etc.* Fortunately these are all straightforward, and we will remind readers of the relevant facts as and when they are needed. Although we will discuss tensors over infinite-dimensional spaces, modules, bundles, operator spaces, *etc.*, we only do so when they are relevant to computational applications. When discussing the various properties and notions pertaining to tensors, we put them against a backdrop of familiar analogues in linear algebra and numerical linear algebra, by either drawing a parallel or a contrast between the multilinear and the linear.

Our approach towards explaining the roles of tensors in computations is to present them as examples alongside our discussions of the three standard definitions of tensors on page 557. The hope is that our reader would gain both an appreciation of these three definitions, each capturing a different feature of a tensor, as well as how each is useful in computations. While there are a very small number of inevitable overlaps with Bürgisser *et al.* (1997) and Landsberg (2012, 2017, 2019) such as fast matrix multiplication, we try to offer a different take on these topics. Also, most of the examples in our article cannot be found in them.

Modern textbook treatments of tensors in algebra and geometry tend to sweep definition ① under the rug of the coordinate-free approaches in definitions ② and ③, preferring not to belabour the coordinate-dependent point of view in definition ①. For example, there is no trace of it in the widely used textbook on algebra by Lang (2002) and that on differential geometry by Lee (2013). Definition ① tends to be found only in much older treatments (Borisenko and Tarapov 1979, Brand 1947, Hay 1954, Lovelock and Rund 1975, McConnell 1957, Michal 1947, Schouten 1951, Spain 1960, Synge and Schild 1978, Wrede 1963) that are often more confusing than illuminating to a modern reader. This has resulted in a certain amount of misinformation. As a public service, we will devote significant space to definition ① so as to make it completely clear. For us, belabouring the point is the point. By and large we will discuss definitions ①, ② and ③ in chronological order, as we think there is some value in seeing how the concept of tensors evolved through time. Nevertheless we will occasionally get ahead of ourselves to point out how a later definition would clarify an earlier one or revisit an earlier definition with the hindsight of a later one.

There is no lack of reasonably elementary, concrete treatments of tensors in physics texts (Abraham, Marsden and Ratiu 1988, Martin 1991, Misner, Thorne and Wheeler 1973, Wald 1984) but the ultimate goal in these books is invariably *tensor fields* and not tensors. The pervasive use of Einstein's summation convention,³ the use of both upper and lower indices (sometimes also left and right, *i.e.* indices

³ No doubt a very convenient shorthand, and its avoidance of the summation symbol \sum a big typesetting advantage in the pre- \TeX era.

on all four corners) in these books can also be confusing, especially when one needs to switch between Einstein's conventions and the usual conventions in linear algebra. Do the raised indices refer to the inverse of the matrix or are they there simply because they needed to be summed over? How does one tell the inverse apart from the inverse transpose of a matrix? The answers depend on the source, somewhat defeating the purpose of a convention. In this article, tensor fields are mentioned only in passing and we do not use Einstein's summation convention. Instead, we frame our discussions of tensors in notation standard in linear algebra and numerical linear algebra or slight extensions of these and, by so doing, hope that our readers would find it easier to see where tensors fit in.

Our article is not about algorithms for solving various problems related to higher-order tensors. Nevertheless, we take the opportunity to make a few points in Section 5.2 about the prospects of such algorithms. The bottom line is that we do not think such problems, even in low-order ($d = 3$ or 4) low-dimensional ($n < 10$) instances, have efficient or even just provably correct algorithms. This is unfortunate as there are various 'tensor models' for problems in areas from quantum chemistry to data analysis that are based on the optimistic assumption that there exist such algorithms. It is fairly common to see claims of extraordinary things that follow from such 'tensor models', but upon closer examination one would invariably find an NP-hard problem such as

$$\min_{\varphi, \psi, \theta} \|f - \varphi \otimes \psi \otimes \theta\|$$

embedded in the model as a special case. If one could efficiently solve this problem, then the range of earth-shattering things that would follow is limitless.⁴ Therefore, these claims of extraordinariness are hardly surprising given such an enormous caveat; they are simply consequences of the fact that if one could efficiently solve *any* NP-hard problems, one could efficiently solve all NP-complete problems.

Nearly all the materials in this article are classical. It could have been written twenty years ago. We would not have been able to mention the resolution of the salmon conjecture or discuss applications such as AlphaFold, the matrix multiplication exponent and complexity of integer multiplication would be slightly higher, and some of the books cited would be in their earlier editions. But 95%, if not more, of the content would have remained intact. We limit our discussion in this article to results that are by-and-large rigorous, usually in the mathematical sense of having a proof but occasionally in the sense of having extensive predictions consistent with results of scientific experiments (like the effectiveness of DMRG). While this article contains no original research, we would like to think that it offers abundant new

⁴ In the words of a leading authority (Fortnow 2013, pp. ix and 11), 'society as we know it would change dramatically, with immediate, amazing advances in medicine, science, and entertainment and the automation of nearly every human task,' and 'the world will change in ways that will make the Internet seem like a footnote in history.'

insights on existing topics: our treatment of multipole expansions in Example 4.45, separation of variables in Examples 4.32–4.35, stress tensors in Example 4.8, our interpretation of the universal factorization property in Section 4.3, discussions of the various forms of higher-order derivatives in Examples 3.2, 3.4, 4.29, the way we have presented and motivated the tensor transformation rules in Section 2, *etc.*, have never before appeared elsewhere to the best of our knowledge. In fact, we have made it a point to not reproduce anything verbatim from existing literature; even standard materials are given a fresh take.

1.3. Notation and convention

Almost all results, unless otherwise noted, will hold for tensors over both \mathbb{R} and \mathbb{C} alike. In order to not have to state ‘ \mathbb{R} or \mathbb{C} ’ at every turn, we will denote our field as \mathbb{R} with the implicit understanding that all discussions remain true over \mathbb{C} unless otherwise stated. We will adopt numerical linear algebra convention and regard vectors in \mathbb{R}^n as column vectors, *i.e.* $\mathbb{R}^n \equiv \mathbb{R}^{n \times 1}$. Row vectors in $\mathbb{R}^{1 \times n}$ will always be denoted as v^\top with $v \in \mathbb{R}^n$. To save space, we also adopt the convention that an n -tuple delimited with parentheses always denotes a column vector, that is,

$$(a_1, \dots, a_n) = \begin{bmatrix} a_1 \\ \vdots \\ a_n \end{bmatrix} \in \mathbb{R}^n.$$

All bases in this article will be ordered bases. We write $\text{GL}(n)$ for the general linear group, that is, the set of $n \times n$ invertible matrices and $\text{O}(n)$ for the orthogonal group of $n \times n$ orthogonal matrices. Other matrix groups will be introduced later. For any $X \in \text{GL}(n)$, we write

$$X^{-\top} := (X^{-1})^\top = (X^\top)^{-1}.$$

The vector space of polynomials in variables x_1, \dots, x_n with real coefficients will be denoted by $\mathbb{R}[x_1, \dots, x_n]$. For Banach spaces \mathbb{V} and \mathbb{W} , we write $\mathcal{B}(\mathbb{V}; \mathbb{W})$ for the set of all bounded linear operators $\Phi: \mathbb{V} \rightarrow \mathbb{W}$. We abbreviate $\mathcal{B}(\mathbb{V}; \mathbb{V})$ as $\mathcal{B}(\mathbb{V})$. A word of caution may be in order. We use the term ‘separable’ in two completely different and unrelated senses with little risk of confusion: a Banach or Hilbert space is separable if it has a countable dense subset whereas a function is separable if it takes the form in (1.1).

We restrict some of our discussions to 3-tensors for a few reasons: (i) the sacrifice in generality brings about an enormous reduction in notational clutter, (ii) the generalization to d -tensors for $d > 3$ is straightforward once we present the $d = 3$ case, (iii) $d = 3$ is the first unfamiliar case for most readers given that $d = 0, 1, 2$ are already well treated in linear algebra, and (iv) $d = 3$ is the most important case for many of our applications to computational mathematics.

2. Tensors via transformation rules

We will begin with definition ①, which first appeared in a book on crystallography by the physicist Woldemar Voigt (1898):

An abstract entity represented by an array of components that are functions of coordinates such that, under a transformation of coordinates, the new components are related to the transformation and to the original components in a definite way.

While the idea of a tensor had appeared before in works of Cauchy and Riemann (Conrad 2018, Yau 2020) and also around the same time in Ricci and Levi-Civita (1900), this is believed to be the earliest appearance of the word ‘tensor’. Although we will refer to the quote above as ‘Voigt’s definition’, it is not a direct translation of a specific sentence from Voigt (1898) but a paraphrase attributed to Voigt (1898) in the *Oxford English Dictionary*. Voigt’s definition is essentially the one adopted in all early textbooks on tensors such as Brand (1947), Hay (1954), Lovelock and Rund (1975), McConnell (1957), Michal (1947), Schouten (1951), Spain (1960), Synge and Schild (1978) and Wrede (1963).⁵ This is not an easy definition to work with and likely contributed to the reputation of tensors being a tough subject to master. Famously, Einstein struggled with tensors (Earman and Glymour 1978) and the definition he had to dabble with would invariably have been definition ①. Nevertheless, we should be reminded that linear algebra as we know it today was an obscure art in its infancy in the 1900s when Einstein was learning about tensors. Those trying to learn about tensors in modern times enjoy the benefit of a hundred years of pedagogical progress. By building upon concepts such as vector spaces, linear transformations, change of basis, *etc.*, that we take for granted today but were not readily accessible a century ago, the task of explaining tensors is significantly simplified.

In retrospect, the main issue with definition ① is that it is a ‘physicist’s definition’, that is, it attempts to define a quantity by describing the change-of-coordinates rules that the quantity must satisfy without specifying the quantity itself. This approach may be entirely natural in physics where one is interested in questions such as ‘Is stress a tensor?’ or ‘Is electromagnetic field strength a tensor?’, that is, the definition is always applied in a way where some physical quantity such as stress takes the place of the unspecified quantity, but it makes for an awkward definition in mathematics. The modern definitions ② and ③ remedy this by stating unequivocally what this unspecified quantity is.

2.1. Transformation rules illustrated with linear algebra

The defining property of a tensor in definition ① is the ‘transformation rules’ alluded to in the second and third lines of Voigt’s definition, that is, how the coordinates

⁵ These were originally published in the 1940s–1960s but all have been reprinted by Dover Publications and several remain in print.

transform under change of basis. In modern parlance, these transformation rules express the notion of *equivariance* under a linear action of a matrix group, an idea that has proved to be important even in trendy applications such as deep convolutional and attention networks. There is no need to go to higher order tensors to appreciate them; we will see that all along we have been implicitly working with these transformation rules in linear algebra and numerical linear algebra. Instead of stating them in their most general forms right away, we will allow readers to discover these transformation rules for themselves via a few simple familiar examples; the meaning of the terminology appearing in italics may be surmised from the examples and will be properly defined in due course.

Eigenvalues and eigenvectors. An eigenvalue $\lambda \in \mathbb{C}$ and a corresponding eigenvector $0 \neq v \in \mathbb{C}^n$ of a matrix $A \in \mathbb{C}^{n \times n}$ satisfy the eigenvalue/eigenvector equation $Av = \lambda v$. For any invertible $X \in \mathbb{C}^{n \times n}$,

$$(XAX^{-1})Xv = \lambda Xv. \quad (2.1)$$

The transformation rules here are $A \mapsto XAX^{-1}$, $v \mapsto Xv$ and $\lambda \mapsto \lambda$. An eigenvalue is an *invariant* 0-tensor and an eigenvector is a *contravariant* 1-tensor of a *mixed* 2-tensor, that is, eigenvalues of A and XAX^{-1} are always the same whereas an eigenvector v of A transforms into an eigenvector Xv of XAX^{-1} . This says that eigenvalues and eigenvectors are defined for mixed 2-tensors.

Matrix product. For matrices $A \in \mathbb{C}^{m \times n}$ and $B \in \mathbb{C}^{n \times p}$, we have

$$(XAY^{-1})(YBZ^{-1}) = X(AB)Z^{-1} \quad (2.2)$$

for any invertible $X \in \mathbb{C}^{m \times m}$, $Y \in \mathbb{C}^{n \times n}$, $Z \in \mathbb{C}^{p \times p}$. The transformation rule is again of the form $A \mapsto XAY^{-1}$. This says that the standard matrix–matrix product is defined on mixed 2-tensors.

Positive definiteness. A matrix $A \in \mathbb{R}^{n \times n}$ is positive definite if and only if

$$XAX^T \quad \text{or} \quad X^{-T}AX^{-1}$$

is positive definite for any invertible $X \in \mathbb{R}^{n \times n}$. This says that positive definiteness is a property of covariant 2-tensors that have transformation rule $X \mapsto XAX^T$ or of *contravariant* 2-tensors that have transformation rule $X \mapsto X^{-T}AX^{-1}$.

Singular values and singular vectors. A singular value $\sigma \in \mathbb{R}$ and its corresponding left singular vector $0 \neq u \in \mathbb{R}^m$ and right singular vector $0 \neq v \in \mathbb{R}^n$ of a matrix $A \in \mathbb{R}^{m \times n}$ satisfy the singular value/singular vector equation

$$\begin{cases} Av = \sigma u, \\ A^T u = \sigma v. \end{cases}$$

For any *orthogonal* matrices $X \in \mathbb{R}^{m \times m}$, $Y \in \mathbb{R}^{n \times n}$, we have

$$\begin{cases} (XAY^T)Yv = \sigma Xu, \\ (XAY^T)^T Xu = \sigma Yv. \end{cases}$$

Singular values are invariant and singular vectors are contravariant as they transform as $\sigma \mapsto \sigma$, $u \mapsto Xu$, $v \mapsto Yv$ for a matrix that transforms as $A \mapsto XAY^T$. This tells us that singular values are *Cartesian* 0-tensors, left and right singular vectors are contravariant Cartesian 1-tensors, defined on contravariant Cartesian 2-tensors. ‘Cartesian’ means that these transformations use orthogonal matrices instead of invertible ones.

Linear equations. Let $A \in \mathbb{R}^{m \times n}$ and $b \in \mathbb{R}^m$. Clearly $Av = b$ has a solution if and only if

$$(XAY^{-1})Yv = Xb$$

has a solution for any invertible $X \in \mathbb{R}^{m \times m}$ and $Y \in \mathbb{R}^{n \times n}$. If $m = n$, then we may instead want to consider

$$(XAX^{-1})Xv = Xb,$$

and if furthermore A is symmetric, then

$$(XAX^T)X^{-T}v = Xb,$$

where $X \in \mathbb{R}^{n \times n}$ is either invertible or orthogonal (for the latter $X^{-T} = X$). We will examine this ambiguity in transformation rules later. The solution in the last case is a covariant 1-tensor, that is, it transforms as $v \mapsto X^{-T}v$ with an invertible X .

Ordinary and total least-squares. Let $A \in \mathbb{R}^{m \times n}$ and $b \in \mathbb{R}^m$. The ordinary least-squares problem is given by

$$\min_{v \in \mathbb{R}^n} \|Av - b\|^2 = \min_{v \in \mathbb{R}^n} \|(XAY^{-1})Yv - Xb\|^2 \quad (2.3)$$

for any orthogonal $X \in \mathbb{R}^{m \times m}$ and invertible $Y \in \mathbb{R}^{n \times n}$. Unsurprisingly, the normal equation $A^T Av = A^T b$ has the same property:

$$(XAY^{-1})^T(XAY^{-1})Yv = (XAY^{-1})^T Xb.$$

The total least-squares problem is defined by

$$\begin{aligned} & \min\{\|E\|^2 + \|r\|^2 : (A + E)v = b + r\} \\ & = \min\{\|XEY^T\|^2 + \|Xr\|^2 : (XAY^T + XEY^T)Yv = Xb + Xr\} \end{aligned}$$

for any orthogonal $X \in \mathbb{R}^{m \times m}$ and orthogonal $Y \in \mathbb{R}^{n \times n}$. Here the minimization is over $E \in \mathbb{R}^{m \times n}$, $r \in \mathbb{R}^m$, and the constraint is interpreted as the linear system being consistent. Both ordinary and total least-squares are defined on tensors – minimum values transform as invariant 0-tensors, v, b, r as contravariant 1-tensors and A, E as mixed 2-tensors – but the 2-tensors involved are different as Y is not required to be orthogonal in (2.3).

Rank, determinant, norm. Let $A \in \mathbb{R}^{m \times n}$. Then

$$\text{rank}(XAY^{-1}) = \text{rank}(A), \quad \det(XAY^{-1}) = \det(A), \quad \|XAY^{-1}\| = \|A\|,$$

where X and Y are, respectively, invertible, special linear or orthogonal matrices.

Here $\|\cdot\|$ may denote either the spectral, nuclear or Frobenius norm and the determinant is regarded as identically zero whenever $m \neq n$. Rank, determinant and norm are defined on mixed 2-tensors, special linear mixed 2-tensors and Cartesian mixed 2-tensors respectively.

The point that we hope to make with these familiar examples is the following. The most fundamental and important concepts, equations, properties and problems in linear algebra and numerical linear algebra – notions that extend far beyond linear algebra into other areas of mathematics, science and engineering – all conform to tensor transformation rules. These transformation rules are not merely accessories to the definition of a tensor: they are the very crux of it and are what make tensors useful.

On the one hand, it is intuitively clear what the transformation rules are in the examples on pages 567–568: They are transformations that preserve either the *form* of an equation (e.g. if we write $A' = XAX^{-1}$, $v' = Xv$, $\lambda' = \lambda$, then (2.1) becomes $A'v' = \lambda'v'$), or the *value* of a quantity such as rank/determinant/norm, or a *property* such as positive definiteness. In these examples, the ‘multi-indexed object’ in definition ① can be a matrix like A or E , a vector like u, v, b, r or a scalar like λ ; the coordinates of the matrix or vector, i.e. a_{ij} and v_i , are sometimes called the ‘components of the tensor’. The matrices X, Y play a different role in these transformation rules and should be distinguished from the multi-indexed objects; for reasons that will be explained in Section 3.1, we call them change-of-basis matrices.

On the other hand, these examples also show why the transformation rules can be confusing: they are ambiguous in multiple ways.

- (a) The transformation rules can take several different forms. For example, 1-tensors may transform as

$$v' = Xv \quad \text{or} \quad v' = X^{-T}v,$$

2-tensors may transform as

$$A' = XAY^{-1}, \quad A' = XAY^T, \quad A' = XAX^{-1}, \quad A' = XAX^T,$$

or yet other possibilities we have not discussed.

- (b) The change-of-basis matrices in these transformation rules may also take several different forms, most commonly invertible or orthogonal. In the examples, the change-of-basis matrices may be a single matrix⁶

$$X \in \text{GL}(n), \text{SL}(n), \text{O}(n),$$

and a pair of them

$$(X, Y) \in \text{GL}(m) \times \text{GL}(n), \text{SL}(m) \times \text{SL}(n), \text{O}(m) \times \text{O}(n),$$

⁶ For those unfamiliar with this matrix group notation, it will be defined in (2.14).

or, as we saw in the case of ordinary least-squares (2.3), $(X, Y) \in O(m) \times GL(n)$. There are yet other possibilities for change-of-basis matrices we have not discussed, such as Lorentz, symplectic, unitary, *etc.*

- (c) Yet another ambiguity on top of (a) is that the roles of v and v' or A and A' are sometimes swapped and the transformation rules written as

$$v' = X^{-1}v, \quad v' = X^T v$$

or

$$A' = X^{-1}AY, \quad A' = X^{-1}AY^{-T}, \quad A' = X^{-1}AX, \quad A' = X^{-1}AX^{-T}$$

respectively. Note that the transformation rules here and those in (a) are all but identical: the only difference is in whether we label the multi-indexed object on the left or that on the right with a prime.

We may partly resolve the ambiguity in (a) by introducing covariance and contravariance type: tensors with transformation rules of the form $A' = XAX^T$ are covariant 2-tensors or $(0, 2)$ -tensors; those of the form $A' = XAX^{-1}$ are mixed 2-tensors or $(1, 1)$ -tensors; those of the form $A' = X^{-T}AX^{-1}$ are contravariant 2-tensors or $(2, 0)$ -tensors. Invariance, covariance and contravariance are all special cases of *equivariance* that we will discuss later. Nevertheless, we are still unable to distinguish between $A' = XAX^T$ and $A' = XAY^T$: both are legitimate transformation rules for covariant 2-tensors.

These ambiguities (a), (b), (c) are the result of a defect in definition ①. One ought to be asking: What is the tensor in definition ①? The answer is that it is actually missing from the definition. The multi-indexed object A *represents* the tensor whereas the transformation rules on A *defines* the tensor but the tensor itself has been left unspecified. This is a key reason why definition ① has been so confusing to early learners. Fortunately, it is easily remedied by definition ② or ③. We should, however, bear in mind that definition ① predated modern notions of vector spaces and linear maps, which are necessary for definitions ② and ③.

When we introduce definitions ② or ③, we will see that the ambiguity in (a) is just due to different transformation rules for different tensors. Getting ahead of ourselves, for vector spaces \mathbb{V} and \mathbb{W} , the transformation rule $A' = XAX^{-1}$ applies to tensors in $\mathbb{V} \otimes \mathbb{V}^*$ whereas $A' = XAY^{-1}$ applies to those in $\mathbb{V} \otimes \mathbb{W}^*$; $A' = XAX^T$ and $A' = X^{-T}AX^{-1}$ apply to tensors in $\mathbb{V} \otimes \mathbb{V}$ and $\mathbb{V}^* \otimes \mathbb{V}^*$ respectively. The matrices A and A' are representations of a tensor with respect to two different bases, and the ambiguity in (c) is just a result of which basis is regarded as the 'old' basis and which as the 'new' one.

The ambiguity in (b) is also easily resolved with definitions ② or ③. The matrices X and Y are change-of-basis matrices on \mathbb{V} and \mathbb{W} and are always invertible, but they also have to preserve additional structures on \mathbb{V} and \mathbb{W} such as inner products or norms. For example, singular values and singular vectors are only defined for inner

product spaces and so we require the matrices X and Y to preserve inner products; for the Euclidean inner product, this simply means that X and Y are orthogonal matrices. Tensors with transformation rules involving orthogonal matrices are sometimes called ‘Cartesian tensors’.

These examples on pages 567–568 illustrate why the transformation rules in definition ① are as crucial in mathematics as they are in physics. Unlike in physics, we do not use the transformation rules to check whether a physical quantity such as stress or strain is a tensor; instead we use them to ascertain whether an equation (*e.g.* an eigenvalue/eigenvector equation), a property (*e.g.* positive definiteness), an operation (*e.g.* a matrix product) or a mathematical quantity (*e.g.* rank) is defined on a tensor. In other words, when working with tensors, it is not a matter of simply writing down an expression involving multi-indexed quantities, because if the expression does not satisfy the transformation rules, it is undefined on a tensor. For example, multiplying two matrices via the Hadamard (also known as Schur) product

$$\begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \circ \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix} = \begin{bmatrix} a_{11}b_{11} & a_{12}b_{12} \\ a_{21}b_{21} & a_{22}b_{22} \end{bmatrix} \quad (2.4)$$

seems a lot more obvious than the standard matrix product

$$\begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix} = \begin{bmatrix} a_{11}b_{11} + a_{12}b_{21} & a_{11}b_{12} + a_{12}b_{22} \\ a_{21}b_{11} + a_{22}b_{21} & a_{21}b_{12} + a_{22}b_{22} \end{bmatrix}. \quad (2.5)$$

However, (2.5) defines a product on tensors whereas (2.4) does not, the reason being that the standard matrix product satisfies the transformation rules in (2.2) whereas the Hadamard product does not. While there may be occasional scenarios where the Hadamard product could be useful, the standard matrix product is far more prevalent in all areas of mathematics, science and engineering. This is not limited to matrix products; as we can see from the list of examples, the most useful and important notions we encounter in linear algebra are invariably tensorial notions that satisfy various transformation rules for 0-, 1- and 2-tensors. We will have more to say about these issues in the next few sections.

2.2. Transformation rules for d -tensors

We now state the transformation rules in their most general form, laying out the full details of definition ①. Let $A = [a_{j_1 \dots j_d}] \in \mathbb{R}^{n_1 \times \dots \times n_d}$ be a d -dimensional matrix or *hypermatrix*. Readers who require a rigorous definition may take it as a real-valued function $A: \{1, \dots, n_1\} \times \dots \times \{1, \dots, n_d\} \rightarrow \mathbb{R}$, a perspective we will discuss in Example 4.5, but those who are fine with a ‘ d -dimensional matrix’ may picture it as such. For matrices, *i.e.* the usual two-dimensional matrices, $X = [x_{i_1 j_1}] \in \mathbb{R}^{m_1 \times n_1}$, $Y = [y_{i_2 j_2}] \in \mathbb{R}^{m_2 \times n_2}$, \dots , $Z = [z_{i_d j_d}] \in \mathbb{R}^{m_d \times n_d}$, we define

$$(X, Y, \dots, Z) \cdot A = B, \quad (2.6)$$

where $B = [b_{i_1 \dots i_d}] \in \mathbb{R}^{m_1 \times \dots \times m_d}$ is given by

$$b_{i_1 \dots i_d} = \sum_{j_1=1}^{n_1} \sum_{j_2=1}^{n_2} \dots \sum_{j_d=1}^{n_d} x_{i_1 j_1} y_{i_2 j_2} \dots z_{i_d j_d} a_{j_1 \dots j_d}$$

for $i_1 = 1, \dots, n_1, i_2 = 1, \dots, n_2, \dots, i_d = 1, \dots, n_d$. We call the operation (2.6) multilinear matrix multiplication; as we will discuss after stating the tensor transformation rules, the notation (2.6) is chosen to be consistent with standard notation for group action. For $d = 1$, it reduces to $Xa = b$ for $a \in \mathbb{R}^n, b \in \mathbb{R}^m$, and for $d = 2$, it reduces to

$$(X, Y) \cdot A = XAY^T.$$

For now, we will let the hypermatrix $A \in \mathbb{R}^{n_1 \times \dots \times n_d}$ be the multi-indexed object in definition ① to keep things simple; we will later see that this multi-indexed object does not need to be a hypermatrix.

Let $X_1 \in GL(n_1), X_2 \in GL(n_2), \dots, X_d \in GL(n_d)$. The covariant d -tensor transformation rule is

$$A' = (X_1^T, X_2^T, \dots, X_d^T) \cdot A. \tag{2.7}$$

The contravariant d -tensor transformation rule is

$$A' = (X_1^{-1}, X_2^{-1}, \dots, X_d^{-1}) \cdot A. \tag{2.8}$$

The mixed d -tensor transformation rule is⁷

$$A' = (X_1^{-1}, \dots, X_p^{-1}, X_{p+1}^T, \dots, X_d^T) \cdot A. \tag{2.9}$$

We say that the transformation rule in (2.9) is of *type* or *valence* $(p, d - p)$ or, more verbosely, of contravariant order p and covariant order $d - p$. As such, covariance is synonymous with type $(0, d)$ and contravariance with type $(d, 0)$.

For $d = 1$, the hypermatrix is just $a \in \mathbb{R}^n$. If it transforms as $a' = X^{-1}a$, then it is the coordinate representation of a contravariant 1-tensor or *contravariant vector*; if it transforms as $a' = X^T a$, then it is the coordinate representation of a covariant 1-tensor or *covariant vector*. For $d = 2$, the hypermatrix is just a matrix $A \in \mathbb{R}^{m \times n}$; writing $X_1 = X, X_2 = Y$, the transformations in (2.7), (2.8), (2.9) become $A' = X^T A Y$, $A' = X^{-1} A Y^{-T}$, $A' = X^{-1} A Y$, which are the transformation rules for a covariant, contravariant, mixed 2-tensor respectively. These look different from the transformation rules in the examples on pages 567–568 as a result of the ambiguity (c) on page 570, which we elaborate below.

⁷ We state it in this form for simplicity. For example, we do not distinguish between the 3-tensor transformation rules $A' = (X^{-1}, Y^T, Z^T) \cdot A$, $A' = (X^T, Y^{-1}, Z^T) \cdot A$ and $A' = (X^T, Y^T, Z^{-1}) \cdot A$. All three are of type $(1, 2)$.

Clearly, the equalities in the middle and right columns below only differ in which side we label with a prime but are otherwise identical:

$$\begin{array}{lll}
 \text{contravariant 1-tensor} & a' = X^{-1}a, & a' = Xa, \\
 \text{covariant 1-tensor} & a' = X^T a, & a' = X^{-T}a, \\
 \text{contravariant 2-tensor} & A' = X^{-1}AY^{-T}, & A' = XAY^T, \\
 \text{covariant 2-tensor} & A' = X^T AY, & A' = X^{-T}AY^{-1}, \\
 \text{mixed 2-tensor} & A' = X^{-1}AY, & A' = XAY^{-1}.
 \end{array}$$

We will see in Section 3.1 after introducing definition ② that these transformation rules come from the change-of-basis theorems for vectors, linear functionals, dyads, bilinear functionals and linear operators, respectively, with X and Y the change-of-basis matrices. The versions in the middle and right column differ in terms of how we write our change-of-basis theorem. For example, take the standard change-of-basis theorem in a vector space. Do we write $a_{\text{new}} = X^{-1}a_{\text{old}}$ or $a_{\text{old}} = Xa_{\text{new}}$?

Observe, however, that there is no repetition in either column: the transformation rule, whether in the middle or right column, uniquely identifies the type of tensor. It is not possible to confuse, say, a contravariant 2-tensor with a covariant 2-tensor just because we use the transformation rule in the middle column to describe one and the transformation rule in the right column to describe the other.

The version in the middle column is consistent with the names ‘covariance’ and ‘contravariance’, which are based on whether the hypermatrix ‘co-varies’, *i.e.* transforms with the same X , or ‘contra-varies’, *i.e.* transforms with its inverse X^{-1} . This is why we have stated our transformation rules in (2.7), (2.8), (2.9) to be consistent with those in the middle column. But there are also occasions, as in the examples on pages 567–568, when it is more natural to express the transformation rules as those in the right column.

When $n_1 = n_2 = \dots = n_d = n$, the transformation rules in (2.7), (2.8), (2.9) may take on a different form with a single change-of-basis matrix $X \in \text{GL}(n)$ as opposed to d of them. In this case the hypermatrix $A \in \mathbb{R}^{n \times \dots \times n}$ is hypercubical, *i.e.* the higher-order equivalent of a square matrix, and the covariant tensor transformation rule is

$$A' = (X^T, X^T, \dots, X^T) \cdot A, \quad (2.10)$$

the contravariant tensor transformation rule is

$$A' = (X^{-1}, X^{-1}, \dots, X^{-1}) \cdot A, \quad (2.11)$$

and the mixed tensor transformation rule is

$$A' = (X^{-1}, \dots, X^{-1}, X^T, \dots, X^T) \cdot A. \quad (2.12)$$

Again, we will see in Section 3.1 after introducing definition ② that the difference between, say, (2.7) and (2.10) is that the former expresses the change-of-basis

theorem for a multilinear map $f: \mathbb{V}_1 \times \cdots \times \mathbb{V}_d \rightarrow \mathbb{R}$ whereas the latter is for a multilinear map $f: \mathbb{V} \times \cdots \times \mathbb{V} \rightarrow \mathbb{R}$.

For $d = 2$, we have $A \in \mathbb{R}^{n \times n}$, and each transformation rule again takes two different forms that are identical in substance:

$$\begin{aligned} \text{covariant 2-tensor} & \quad A' = X^T A X, & \quad A' = X^{-T} A X^{-1}, \\ \text{contravariant 2-tensor} & \quad A' = X^{-1} A X^{-T}, & \quad A' = X A X^T, \\ \text{mixed 2-tensor} & \quad A' = X^{-1} A X, & \quad A' = X A X^{-1}. \end{aligned}$$

Note that either form uniquely identifies the tensor type.

Example 2.1 (covariance versus contravariance). The three transformation rules for 2-tensors are $A' = X A X^T$, $A' = X^{-T} A X^{-1}$, $A' = X A X^{-1}$. We know from linear algebra that there is a vast difference between the first two (congruence) and the last one (similarity). For example, eigenvalues and eigenvectors are defined for mixed 2-tensors by virtue of (2.1) but undefined for covariant or contravariant 2-tensors since the eigenvalue/eigenvector equation $A v = \lambda v$ is incompatible with the first two transformation rules. While both the contravariant and covariant transformation rules describe congruence of matrices, the difference between them is best seen to be the difference between the quadratic form and the second-order partial differential operator:

$$\sum_{i=1}^n \sum_{j=1}^n a_{ij} v_i v_j \quad \text{and} \quad \sum_{i=1}^n \sum_{j=1}^n a_{ij} \frac{\partial^2}{\partial v_i \partial v_j}.$$

$A' = X^{-T} A X^{-1}$ and $A' = X A X^T$, respectively, describe how their coefficient matrices transform under a change of coordinates $v' = X v$.

As we have mentioned in Section 2.1, definition ① leaves the tensor unspecified. In physics, this is to serve as a placeholder for a physical quantity under discussion such as deformation or rotation or strain. In mathematics, if we insist on not bringing in definitions ② or ③, then the way to use definition ① is best illustrated with an example.

Example 2.2 (transformation rule determines tensor type). We say the eigenvalue/eigenvector equation $A v = \lambda v$ conforms to the transformation rules since $(X A X^{-1}) X v = \lambda X v$, and thus A transforms like a mixed 2-tensor $A' = X A X^{-1}$, the eigenvector v transforms like a covariant 1-tensor $v' = X v$ and the eigenvalue λ transforms like an invariant 0-tensor. It is only with context (here we have eigenvalues/eigenvectors) that it makes sense to speak of A , v , λ as tensors. Is $A = \begin{bmatrix} 1 & 2 & 3 \\ 2 & 3 & 4 \\ 3 & 4 & 5 \end{bmatrix}$ a tensor? It is a tensor if it represents, say, stress (in which case A transforms as a contravariant 2-tensor) or, say, we are interested in its eigenvalues and eigenvectors (in which case A transforms as a mixed 2-tensor). Stripped of such contexts, the question as to whether a matrix or hypermatrix is a tensor is not meaningful.

We have stressed that the transformation rules define the tensor but one might think that this only applies to the type of the tensor, *i.e.* whether it is covariant or contravariant or mixed. This is not the case: it defines *everything* about the tensor.

Example 2.3 (transformation rule determines tensor order). Does $A = \begin{bmatrix} 1 & 2 & 3 \\ 2 & 3 & 4 \\ 3 & 4 & 5 \end{bmatrix}$ represent a 1-tensor or a 2-tensor? One might think that since a matrix is a doubly indexed object, it must represent a 2-tensor. Again this is a fallacy. If the transformation rules we apply to a matrix $A \in \mathbb{R}^{m \times n}$ are of the forms

$$A' = XA = [Xa_1, \dots, Xa_n], \quad A' = X^{-T}A = [X^{-T}a_1, \dots, X^{-T}a_n], \quad (2.13)$$

where $a_1, \dots, a_n \in \mathbb{R}^m$ are the columns of A and $X \in \text{GL}(m)$, then A represents a contravariant or covariant 1-tensor respectively, that is, the *order* of a tensor is also determined by the transformation rules. In numerical linear algebra, algorithms such as Gaussian elimination and Householder or Givens QR are all based on transformations of a matrix as a 1-tensor as in (2.13); we will have more to say about this in Section 2.4. In equivariant neural networks, it does not matter that filters are represented as five-dimensional arrays (Cohen and Welling 2016, Section 7.1) or that inputs are a matrix of vectors of ‘irreducible fragments’ (Kondor, Lin and Trivedi 2018, Definition 1): the transformation rules involved are ultimately just those of contravariant 1-tensors and mixed 2-tensors, as we will see in Example 2.16.

The change-of-basis matrices X_1, \dots, X_d in (2.7), (2.8), (2.9) and X in (2.10), (2.11), (2.12) do not need to be just invertible matrices, *i.e.* elements of $\text{GL}(n)$: they could instead come from any classical groups, most commonly the orthogonal group $\text{O}(n)$ or unitary group $\text{U}(n)$, possibly with unit-determinant constraints, *i.e.* $\text{SL}(n)$, $\text{SO}(n)$, $\text{SU}(n)$. The reason is that as the name implies, these are change-of-basis matrices for the vector spaces in definitions ② or ③, and if those vector spaces are equipped with additional structures, we expect the change-of-basis matrices to preserve those structures.

Example 2.4 (Cartesian and Lorentzian tensors). For the tensor transformation rule over \mathbb{R}^4 equipped with the Euclidean inner product

$$\langle x, y \rangle = x_0y_0 + x_1y_1 + x_2y_2 + x_3y_3,$$

we want the change-of-basis matrices to be from the orthogonal group $\text{O}(4)$ or special orthogonal group $\text{SO}(4)$, but if \mathbb{R}^4 is equipped instead with the Lorentzian scalar product in relativity,

$$\langle x, y \rangle = x_0y_0 - x_1y_1 - x_2y_2 - x_3y_3,$$

then the Lorentz group $\text{O}(1, 3)$ or proper Lorentz group $\text{SO}(1, 3)$ or restricted Lorentz group $\text{SO}^+(1, 3)$ is more appropriate. Names like Cartesian tensors or Lorentzian tensors are sometimes used to indicate whether the transformation rules involve orthogonal or Lorentz groups but sometimes physicists would assume that these are self-evident from the context and are left unspecified, adding to the

confusion for the uninitiated. The well-known *electromagnetic field tensor* or Faraday tensor, usually represented in matrix form by

$$F = \begin{bmatrix} 0 & -E_x/c & -E_y/c & -E_z/c \\ E_x/c & 0 & -B_z & B_y \\ E_y/c & B_z & 0 & -B_x \\ E_z/c & -B_y & B_x & 0 \end{bmatrix},$$

where we have written $(x_0, x_1, x_2, x_3) = (ct, x, y, z)$, is a contravariant Lorentz 2-tensor as it satisfies a transformation rule of the form $F' = XF X^T$ with a change-of-basis matrix $X \in O(1, 3)$.

For easy reference, here is a list of the most common matrix Lie groups that the change-of-basis matrices X_1, \dots, X_d or X may belong to:

$$\begin{aligned} \text{GL}(n) &= \{X \in \mathbb{R}^{n \times n} : \det(X) \neq 0\}, \\ \text{SL}(n) &= \{X \in \mathbb{R}^{n \times n} : \det(X) = 1\}, \\ \text{O}(n) &= \{X \in \mathbb{R}^{n \times n} : X^T X = I\}, \\ \text{SO}(n) &= \{X \in \mathbb{R}^{n \times n} : X^T X = I, \det(X) = 1\}, \\ \text{U}(n) &= \{X \in \mathbb{C}^{n \times n} : X^* X = I\}, \\ \text{SU}(n) &= \{X \in \mathbb{C}^{n \times n} : X^* X = I, \det(X) = 1\}, \\ \text{O}(p, q) &= \{X \in \mathbb{R}^{n \times n} : X^T I_{p,q} X = I_{p,q}\}, \\ \text{SO}(p, q) &= \{X \in \mathbb{R}^{n \times n} : X^T I_{p,q} X = I_{p,q}, \det(X) = 1\}, \\ \text{Sp}(2n, \mathbb{R}) &= \{X \in \mathbb{R}^{2n \times 2n} : X^T J X = J\}, \\ \text{Sp}(2n) &= \{X \in \mathbb{C}^{2n \times 2n} : X^T J X = J, X^* X = I\}. \end{aligned} \tag{2.14}$$

In the above, \mathbb{R} may be replaced by \mathbb{C} to obtain corresponding groups of complex matrices, p and q are positive integers with $p + q = n$, and $I = I_n$ denotes the $n \times n$ identity matrix. The matrices $I_{p,q} \in \mathbb{C}^{n \times n}$ and $J \in \mathbb{C}^{2n \times 2n}$, respectively, are

$$I_{p,q} := \begin{bmatrix} I_p & 0 \\ 0 & -I_q \end{bmatrix}, \quad J := \begin{bmatrix} 0 & I \\ I & 0 \end{bmatrix}.$$

There are some further possibilities that deserve a mention here, if only for cultural reasons. While not as common as the classical Lie groups above, depending on the application at hand, the groups of invertible lower/upper or unit lower/upper triangular matrices, Heisenberg groups, *etc.*, may all serve as the group of change-of-basis matrices. Another three particularly important examples are the general

affine, Euclidean and special Euclidean groups

$$\begin{aligned} \text{GA}(n) &= \left\{ \begin{bmatrix} X & y \\ 0 & 1 \end{bmatrix} \in \mathbb{R}^{(n+1) \times (n+1)} : X \in \text{GL}(n), y \in \mathbb{R}^n \right\}, \\ \text{E}(n) &= \left\{ \begin{bmatrix} X & y \\ 0 & 1 \end{bmatrix} \in \mathbb{R}^{(n+1) \times (n+1)} : X \in \text{O}(n), y \in \mathbb{R}^n \right\}, \\ \text{SE}(n) &= \left\{ \begin{bmatrix} X & y \\ 0 & 1 \end{bmatrix} \in \mathbb{R}^{(n+1) \times (n+1)} : X \in \text{SO}(n), y \in \mathbb{R}^n \right\}, \end{aligned} \quad (2.15)$$

which encode a linear transformation by X and translation by y in \mathbb{R}^n . Note that even though the matrices involved are $(n+1) \times (n+1)$, they act on vectors $v \in \mathbb{R}^n$ embedded as $(v, 0) \in \mathbb{R}^{n+1}$. As we will see in Example 2.16, the last two groups play a special role in equivariant neural networks for image classification ($n = 2$) and protein structure predictions ($n = 3$).

The above discussions extend to an infinite-dimensional setting in two different ways. In general relativity, the tensors are replaced by tensor fields and the change-of-basis matrices are replaced by change-of-coordinates matrices whose entries are functions of coordinates; these are called *diffeomorphisms* and they form infinite-dimensional Lie groups. In quantum mechanics, the tensors over finite-dimensional vector spaces are replaced by tensors over infinite-dimensional Hilbert spaces and the change-of-basis matrices are replaced by *unitary operators*, which may be regarded as infinite-dimensional matrices (see Example 4.5). In this case the tensor transformation rules are particularly important as they describe the evolution of the quantum system. We will have more to say about these in Examples 3.13 and 4.49.

We are now ready to state a slightly modernized version of definition ①.

Definition 2.5 (tensors via transformation rules). A hypermatrix $A \in \mathbb{R}^{n_1 \times \dots \times n_d}$ represents a d -tensor if it satisfies one of the transformation rules in (2.7)–(2.12) for some matrix groups in (2.14).

Notice that this only defines what it means for a hypermatrix to represent a tensor: it does not fix the main drawback of definition ①, *i.e.* the lack of an object to serve as the tensor. Definition ① cannot be easily turned into a standalone definition and is best understood in conjunction with either definition ② or ③; these transformation rules then become change-of-basis theorems for a multilinear map or an element in a tensor product of vector spaces. Nevertheless, definition ① has its value: the transformation rules encode the important notion of equivariance that we will discuss next over Sections 2.3 and 2.4.

There are some variants of the tensor transformation rules that come up in various areas of physics. We state them below for completeness and as examples of hypermatrices that do not represent tensors.

Example 2.6 (pseudotensors and tensor densities). For a change-of-basis matrix $X \in \text{GL}(n)$ and a hypermatrix $A \in \mathbb{R}^{n \times \dots \times n}$, the pseudotensor transformation

rule (Borisenko and Tarapov 1979, Section 3.7.2), the tensor density transformation rule (Plebański and Krasieński 2006, Section 3.8) and the pseudotensor density transformation rule are, respectively,

$$\begin{aligned} A' &= \operatorname{sgn}(\det X) (X^{-1}, \dots, X^{-1}, X^\top, \dots, X^\top) \cdot A, \\ A' &= (\det X)^p (X^{-1}, \dots, X^{-1}, X^\top, \dots, X^\top) \cdot A, \\ A' &= \operatorname{sgn}(\det X)(\det X)^p (X^{-1}, \dots, X^{-1}, X^\top, \dots, X^\top) \cdot A. \end{aligned} \quad (2.16)$$

Evidently they differ from the tensor transformation rule (2.12) by a scalar factor. The power $p \in \mathbb{Z}$ is called the weight with $p = 0$ giving us (2.12). An alternative name for a pseudotensor is an ‘axial tensor’ and in this case a true tensor is called a ‘polar tensor’ (Hartmann 1984, Section 4). Note that the transformation rules in (2.16) are variations of (2.12). While it is conceivable to define similar variations of (2.9) with $\det(X_1 \cdots X_d) = \det X_1 \cdots \det X_d$ in place of $\det X$, we are unable to find a reference for this and would therefore refrain from stating it formally.

A word about notation. Up to this point we have introduced only one notion that is not from linear algebra, namely (2.6) for multilinear matrix multiplication; this notion dates back to at least Hitchcock (1927) and the notation is standard too. We denote it so that (2.7), (2.8), (2.9) reduce to the standard notation $g \cdot x$ for a group element $g \in G$ acting on an element x of a G -set. Here our group G is a product of other groups $G = G_1 \times \cdots \times G_d$, and so an element takes the form $g = (g_1, \dots, g_d)$. We made a conscious decision to cast everything in this article in terms of *left* action so as not to introduce another potential source of confusion. We could have introduced a *right* multilinear matrix multiplication of $A \in \mathbb{R}^{n_1 \times \cdots \times n_d}$ by matrices $X \in \mathbb{R}^{n_1 \times m_1}, Y \in \mathbb{R}^{n_2 \times m_2}, \dots, Z \in \mathbb{R}^{n_d \times m_d}$; note that the dimensions of these matrices are now the transposes of those in (2.6),

$$A \cdot (X, Y, \dots, Z) = B,$$

with $B \in \mathbb{R}^{m_1 \times \cdots \times m_d}$ given by

$$b_{i_1 \dots i_d} = \sum_{j_1=1}^{n_1} \sum_{j_2=1}^{n_2} \cdots \sum_{j_d=1}^{n_d} a_{j_1 \dots j_d} x_{i_1 j_1} y_{i_2 j_2} \cdots z_{i_d j_d}.$$

This would have allowed us to denote (2.7) and (2.9) without transposes, with the latter in a two-sided product. Nevertheless, we do not think it is worth the trouble.

2.3. Importance of the transformation rules

When we say that a group is a set with an operation satisfying certain axioms, these axioms play a crucial role and cannot be disregarded. It is the same with tensors. The importance of the transformation rules in definition ① should be abundantly clear at this point. In general, a formula in terms of the entries of the hypermatrix is undefined for the tensor since that formula is unlikely to conform to the stringent transformation rules; this mistake is unfortunately commonplace. Here we will

discuss some erroneous ideas that result from treating a tensor as no more than the hypermatrix that represents it, oblivious of the transformation rules it must satisfy.

Example 2.7 ('multiplying higher-order tensors'). There have been many recent attempts at finding a formula for 'multiplying higher-order tensors' that purports to extend the standard matrix–matrix product (2.5). These proposed formulas are justified with a demonstration that they are associative, distributive, have a multiplicative identity, *etc.*, just like matrix multiplication. This is misguided. We need go no further than 2-tensors to see the problem: both (2.4) and (2.5) are associative, distributive and have multiplicative identities, but only the formula in (2.5) defines a product of 2-tensors because it conforms to the transformation rules (2.2). In fact no other formula aside from the standard matrix–matrix product defines a product for 2-tensors, a consequence of the first fundamental theorem of invariant theory; see De Concini and Procesi (2017, Chapter 9) or Goodman and Wallach (2009, Section 5.3). This theorem also implies that there is no formula that will yield a product for any odd-ordered tensors. In particular, one may stop looking for a formula to multiply two $n \times n \times n$ matrices to get a third $n \times n \times n$ hypermatrix that defines a product for 3-tensors, because such a formula does not exist. Incidentally, the transformation rules also apply to adding tensors: we may not add two matrices representing two tensors of different types because they satisfy different transformation rules; for example, if $A \in \mathbb{R}^{n \times n}$ represents a mixed 2-tensor and $B \in \mathbb{R}^{n \times n}$ a covariant 2-tensor, then $A + B$ is rarely if ever meaningful.

Example 2.8 ('identity tensors'). The identity matrix $I \in \mathbb{R}^{3 \times 3}$ is of course

$$I = \sum_{i=1}^3 e_i \otimes e_i \in \mathbb{R}^{3 \times 3}, \quad (2.17)$$

where $e_1, e_2, e_3 \in \mathbb{R}^3$ are the standard basis vectors. What should be the extension of this notion to d -tensors? Take $d = 3$ for illustration. It would appear that

$$A = \sum_{i=1}^3 e_i \otimes e_i \otimes e_i \in \mathbb{R}^{3 \times 3 \times 3} \quad (2.18)$$

is the obvious generalization but, as with the Hadamard product, obviousness does not necessarily imply correctness when it comes to matters tensorial. One needs to check the transformation rules. Note that (2.17) is independent of the choice of orthonormal basis: we obtain the same matrix with any orthonormal basis $q_1, q_2, q_3 \in \mathbb{R}^3$, a consequence of

$$(Q, Q) \cdot I = QIQ^T = I$$

for any $Q \in O(3)$, that is, the identity matrix is well-defined as a Cartesian tensor. On the other hand the hypermatrix A in (2.18) does not have this property. One may show that up to scalar multiples, $M = I$ is the unique matrix satisfying

$$(Q, Q) \cdot M = M \quad (2.19)$$

for any $Q \in O(3)$, a property known as *isotropic*. An isotropic 3-tensor would then be one that satisfies

$$(Q, Q, Q) \cdot T = T. \tag{2.20}$$

Up to scalar multiples, (2.20) has a unique solution given by the hypermatrix

$$J = \sum_{i=1}^3 \sum_{j=1}^3 \sum_{k=1}^3 \varepsilon_{ijk} e_i \otimes e_j \otimes e_k \in \mathbb{R}^{3 \times 3 \times 3},$$

where ε_{ijk} is the Levi-Civita symbol

$$\varepsilon_{ijk} = \begin{cases} +1 & \text{if } (i, j, k) = (1, 2, 3), (2, 3, 1), (3, 1, 2), \\ -1 & \text{if } (i, j, k) = (1, 3, 2), (2, 1, 3), (3, 2, 1), \\ 0 & \text{if } i = j, j = k, k = i, \end{cases}$$

which is evidently quite different from (2.18). More generally, n -dimensional isotropic d -tensors, *i.e.* tensors with the same hypermatrix representation regardless of choice of orthonormal bases in \mathbb{R}^n , have been explicitly determined for $n = 3$ and $d \leq 8$ in Kearsley and Fong (1975), and studied for arbitrary values of n and d in Weyl (1997). From a tensorial perspective, isotropic tensors, not a ‘hypermatrix with ones on the diagonal’ like (2.18), extend the notion of an identity matrix.

Example 2.9 (hyperdeterminants). The simplest order-3 analogue for the determinant, called the Cayley hyperdeterminant for $A = [a_{ijk}] \in \mathbb{R}^{2 \times 2 \times 2}$, is given by

$$\begin{aligned} \text{Det}(A) = & a_{000}^2 a_{111}^2 + a_{001}^2 a_{110}^2 + a_{010}^2 a_{101}^2 + a_{011}^2 a_{100}^2 \\ & - 2(a_{000} a_{001} a_{110} a_{111} + a_{000} a_{010} a_{101} a_{111} + a_{000} a_{011} a_{100} a_{111} \\ & + a_{001} a_{010} a_{101} a_{110} + a_{001} a_{011} a_{110} a_{100} + a_{010} a_{011} a_{101} a_{100}) \\ & + 4(a_{000} a_{011} a_{101} a_{110} + a_{001} a_{010} a_{100} a_{111}), \end{aligned}$$

which looks nothing like the usual expression for a matrix determinant. Just as the matrix determinant is preserved under a transformation $A' = (X, Y) \cdot A = XAY^T$ for $X, Y \in \text{SL}(n)$, this is preserved under a transformation of the form $A' = (X, Y, Z) \cdot A$ for any $X, Y, Z \in \text{SL}(2)$. This notion of hyperdeterminant has been extended to any $A \in \mathbb{R}^{n_1 \times \dots \times n_d}$ with

$$n_i - 1 \leq \sum_{j \neq i} (n_j - 1), \quad i = 1, \dots, d,$$

in Gel’fand, Kapranov and Zelevinsky (1992).

As these three examples reveal, extending any common linear algebraic notion to higher order will require awareness of the transformation rules. Determining the higher-order tensorial analogues of notions such as linear equations or eigenvalues

or determinant is not a matter of simply adding more indices to

$$\sum_{j=1}^n a_{ij}x_j = b_i, \quad \sum_{j=1}^n a_{ij}x_j = \lambda x_i, \quad \sum_{\sigma \in \mathfrak{S}_n} \text{sgn}(\sigma) \prod_{i=1}^n a_{i\sigma(i)}.$$

Of course, one might argue that anyone is free to concoct any formula for ‘tensor multiplication’ or call any hypermatrix an ‘identity tensor’ even if they are undefined on tensors: there is no right or wrong. But there is. They would be wrong in the same way adding fractions as $a/b + c/d = (a + b)/(c + d)$ is wrong, or at least far less useful than the standard way of adding fractions. Fractions are real, and adding half a cake to a third of a cake does not give us two-fifths of a cake. Likewise tensors are real: we discovered the transformation rules for tensors just as we discovered the arithmetic rules for adding fractions, that is, we did not invent them arbitrarily.

While we have limited our examples to mathematical ones, we will end with a note about the physics perspective. In physics, these transformation rules are just as if not more important. One does not even have to go to higher-order tensors to see this: a careful treatment of *vectors* in physics already requires such an approach.

Example 2.10 (tensor transformation rules in physics). As we saw at the beginning of Section 2, the tensor transformation rules originated from physics as change-of-coordinates rules for certain physical quantities. Not all physical quantities satisfy tensor transformation rules; those that do are called ‘tensorial quantities’ or simply ‘tensors’ (Voigt 1898, p. 20). We will take the simplest physical quantity, *displacement*, *i.e.* the location q of a point particle in space, to illustrate. Once we select our x -, y - and z -axes, space becomes \mathbb{R}^3 and q gets coordinates $a = (a_1, a_2, a_3) \in \mathbb{R}^3$ with respect to these axes. If we change our axes to a new set of x' -, y' - and z' -axes – note that nothing physical has changed, the point q is still where it was, the only difference is in the directions we decided to pick as axes – the coordinates must therefore change accordingly to some $a' = (a'_1, a'_2, a'_3) \in \mathbb{R}^3$ to compensate for this change in the choice of axes (see Figure 2.1). If the matrix $X \in \text{GL}(n)$ transforms the x -, y - and z -axes to the x' -, y' - and z' -axes, then the coordinates must transform in the opposite way $a' = X^{-1}a$ to counteract this change so that on the whole everything stays the same. This is the transformation rule for a contravariant 1-tensor on page 573, and we are of course just paraphrasing the change-of-basis theorem for a vector in a vector space. This same transformation rule applies to many other physical quantities: any time derivatives of displacement such as velocity \dot{q} , acceleration \ddot{q} , momentum $m\dot{q}$, force $m\ddot{q}$. So the most familiar physical quantities we encountered in introductory mechanics are contravariant 1-tensors.

Indeed, contravariant 1-tensors are how many physicists would regard vectors (Weinreich 1998): a vector is an object represented by $a \in \mathbb{R}^n$ that satisfies the transformation rule $a' = X^{-1}a$, possibly with the additional requirement that




Figure 2.1. Linear transformation of coordinate axes.

the change-of-basis matrix X be in $O(n)$ (Feynman, Leighton and Sands 1963, Chapter 11) or $X \in O(p, q)$ (Rindler 2006, Chapter 4).

The transformation rules perspective has proved to be very useful in physics. For example, special relativity is essentially the observation that the laws of physics are invariant under Lorentz transformations $X \in O(1, 3)$ (Einstein 2002). In fact, a study of the contravariant 1-tensor transformation rule under the $O(1, 3)$ -analogue of Givens rotations,

$$\begin{bmatrix} \cosh \theta & -\sinh \theta & 0 & 0 \\ -\sinh \theta & \cosh \theta & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \begin{bmatrix} \cosh \theta & 0 & -\sinh \theta & 0 \\ 0 & 1 & 0 & 0 \\ -\sinh \theta & 0 & \cosh \theta & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \begin{bmatrix} \cosh \theta & 0 & 0 & -\sinh \theta \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ -\sinh \theta & 0 & 0 & \cosh \theta \end{bmatrix},$$

is enough to derive most standard results of special relativity; see Friedberg, Insel and Spence (2003, Section 6.9) and Woodhouse (2003, Chapters 4 and 5).

Covariant 1-tensors are also commonplace in physics. If the coordinates of q transform as $a' = X^{-1}a$, then the coordinates of its derivative $\partial/\partial q$ or ∇_q would transform as $a' = X^T a$. So physical quantities that satisfy the covariant 1-tensor transformation rule tend to have a contravariant 1-tensor ‘in the denominator’ such as conjugate momentum $p = \partial L/\partial \dot{q}$ or electric field $E = -\nabla\phi - \partial A/\partial t$. In the former, L is the Lagrangian, and we have the velocity \dot{q} ‘in the denominator’; in the latter, ϕ and A , respectively, are the scalar and vector potentials, and the gradient ∇ is a derivative with respect to spatial variables, and thus have displacement q ‘in the denominator’.

More generally, what we said above applies to more complex physical quantities: when the axes – also called the reference frame in physics – are changed, coordinates

of physical quantities must change in a way that preserves the laws of physics; for tensorial quantities, this would be (2.7)–(2.12), the ‘definite way’ in Voigt’s definition on page 566. This is the well-known maxim that the laws of physics should not depend on coordinates, a version of which the reader will find on page 595. The reader will also find higher-order examples in Examples 2.4, 3.12 and 4.8.

2.4. Tensors in computations I: equivariance

These transformation rules are no less important in computations. The most fundamental algorithms in numerical linear algebra are invariably based implicitly on one of these transformation rules combined with special properties of the change-of-basis matrices. Further adaptations for numerical stability and efficiency can sometimes obscure what is going on in these algorithms, but the basic principle is always to involve one of these transformation rules, or, in modern lingo, to take advantage of *equivariance*.

Example 2.11 (equivalence, similarity, congruence). The transformation rules for 2-tensors have better-known names in linear algebra and numerical linear algebra.

- Equivalence of matrices: $A' = XAY^{-1}$.
- Similarity of matrices: $A' = XAX^{-1}$.
- Congruence of matrices: $A' = XAX^T$.

The most common choices for X and Y are either orthogonal matrices or invertible ones (similarity and congruence are identical if X is orthogonal). These transformation rules have *canonical forms*, and this is an important feature in linear algebra and numerical linear algebra alike, notwithstanding the celebrated fact (Golub and Wilkinson 1976) that some of these cannot be computed in finite precision: Smith normal form for equivalence, singular value decomposition for orthogonal equivalence, Jordan and Weyr canonical forms for similarity (O’Meara, Clark and Vinsonhaler 2011), Turnbull–Aitken and Hodge–Pedoe canonical forms for congruence (De Terán 2016). One point to emphasize is that the transformation rules determine the canonical form for the tensor: it makes no sense to speak of Jordan form for a contravariant 2-tensor or a Turnbull–Aitken form for a mixed 2-tensor, even if both tensors are represented by exactly the same matrix $A \in \mathbb{R}^{n \times n}$.

At this juncture it is appropriate to highlight a key difference between 2-tensors and higher-order ones: while 2-tensors have canonical forms, higher-order tensors in general do not (Landsberg 2012, Chapter 10). This is one of several reasons why we should not expect an extension of linear algebra or numerical linear algebra to d -dimensional hypermatrices in a manner that resembles the $d = 2$ versions.

Before we continue, we will highlight two simple properties.

- (i) The change-of-basis matrices may be multiplied and inverted: if X and Y are orthogonal or invertible, then so is XY and so is X^{-1} , that is, the set of all change-of-basis matrices $O(n)$ or $GL(n)$ forms a group.
- (ii) The transformation rules may be composed: if we have $a' = X^{-T}a$ and $a'' = Y^{-T}a'$, then $a'' = (YX)^{-T}a$; if we have $A' = XAX^{-1}$ and $A'' = YA'Y^{-1}$, then $A'' = (YX)A(YX)^{-1}$, that is, the transformation rule defines a group action.

These innocuous observations say that to get a matrix A into a desired form B , we may just work on a ‘small part’ of the matrix A , e.g. a 2×2 -submatrix or a fragment of a column, by applying a transformation that affects that ‘small part’. We then repeat it on other parts to obtain a sequence of transformations:

$$\begin{aligned}
 A &\rightarrow X_1 A \rightarrow X_2 (X_1 A) \rightarrow \cdots \rightarrow B, \\
 A &\rightarrow X_1^{-T} A \rightarrow X_2^{-T} (X_1^{-T} A) \rightarrow \cdots \rightarrow B, \\
 A &\rightarrow X_1 A X_1^T \rightarrow X_2 (X_1 A X_1^T) X_2^T \rightarrow \cdots \rightarrow B, \\
 A &\rightarrow X_1 A X_1^{-1} \rightarrow X_2 (X_1 A X_1^{-1}) X_2^{-1} \rightarrow \cdots \rightarrow B, \\
 A &\rightarrow X_1 A Y_1^{-1} \rightarrow X_2 (X_1 A Y_1^{-1}) Y_2^{-1} \rightarrow \cdots \rightarrow B,
 \end{aligned} \tag{2.21}$$

and piece all change-of-basis matrices together to get the required X as either $X_m X_{m-1} \dots X_1$ or its limit as $m \rightarrow \infty$ (likewise for Y). Algorithms for computing standard matrix decompositions such as LU, QR, EVD, SVD, Cholesky, Schur, *etc.*, all involve applying a sequence of such transformation rules (Golub and Van Loan 2013). In numerical linear algebra, if m is finite, *i.e.* X may be obtained in finitely many steps (in exact arithmetic), then the algorithm is called *direct*, whereas if it requires $m \rightarrow \infty$, *i.e.* X may only be approximated with a limiting process, then it is called *iterative*. Furthermore, in numerical linear algebra, one tends to see such transformations as giving a *matrix decomposition*, which may then be used to solve other problems involving A . This is sometimes called ‘the decompositional approach to matrix computation’ (Stewart 2000).

Designers of algorithms for matrix computations, even if they were not explicitly aware of these transformation rules and properties, were certainly observing them implicitly. For instance, it is rare to find algorithms that mix different transformation rules for different types of tensors, since what is incompatible for tensors tends to lead to meaningless results. Also, since eigenvalues are defined for mixed 2-tensors but not for contravariant or covariant 2-tensors, the transformation $A' = XAX^{-1}$ is pervasive in algorithms for eigenvalue decomposition but we rarely if ever find $A' = XAX^T$ (unless of course X is orthogonal).

In numerical linear algebra, the use of the transformation rules in (2.21) goes hand in hand with a salient property of the group of change-of-basis matrices.

Example 2.12 (Givens rotations, Householder reflectors, Gauss transforms).

Recall that these are defined by

$$G = \begin{bmatrix} 1 & \cdots & 0 & \cdots & 0 & \cdots & 0 \\ \vdots & \ddots & \vdots & & \vdots & & \vdots \\ 0 & \cdots & \cos \theta & \cdots & -\sin \theta & \cdots & 0 \\ \vdots & & \vdots & \ddots & \vdots & & \vdots \\ 0 & \cdots & \sin \theta & \cdots & \cos \theta & \cdots & 0 \\ \vdots & & \vdots & & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & \cdots & 0 & \cdots & 1 \end{bmatrix} \in \text{SO}(n),$$

$$H = I - \frac{2vv^T}{v^T v} \in \text{O}(n), \quad M = I - ve_i^T \in \text{GL}(n),$$

with the property that the Givens rotation $a' = Ga$ is a rotation of a in the (i, j) -plane by an angle θ , the Householder reflection $a' = Ha$ is the reflection of a in the hyperplane with normal $v/\|v\|$ and, for a judiciously chosen v , the Gauss transform $a' = Ma$ is in $\text{span}\{e_{i+1}, \dots, e_n\}$, that is, it has $(i + 1)$ th to n th coordinates zero.⁸ Standard algorithms in numerical linear algebra such as Gaussian elimination for LU decomposition and Cholesky decomposition, Givens and Householder QR for QR decomposition, Francis's QR or Rutishauser's LR algorithms for EVD, Golub–Kahan bidiagonalization for SVD, *etc.*, all rely in part on applying a sequence of transformation rules as in (2.21) with one of these matrices playing the role of the change-of-basis matrices. The reason this is possible is that:

- any $X \in \text{SO}(n)$ is a product of Givens rotations,
- any $X \in \text{O}(n)$ is a product of Householder reflectors,
- any $X \in \text{GL}(n)$ is a product of elementary matrices,
- any unit lower triangular $X \in \text{GL}(n)$ is a product of Gauss transforms.

In group-theoretic lingo, these matrices are *generators* of the respective matrix Lie groups; in the last case, the set of all unit lower triangular matrices, *i.e.* ones on the diagonal, is also a subgroup of $\text{GL}(n)$.

Whether one seeks to solve a system of linear equations or find a least-squares solution or compute eigenvalue or singular value decompositions, the basic underlying principle in numerical linear algebra is to transform the problem in such a way that the solution of the transformed problem is related to the original solution in a definite way; note that this is practically a paraphrase of Voigt's definition of a tensor on page 566. Any attempt to give a comprehensive list of examples will simply result in our reproducing a large fraction of [Golub and Van Loan](#)

⁸ If $m = \alpha e_j$, then MA adds an α multiple of the j th row to the i th row of A ; so the Gauss transform includes the elementary matrices that perform this operation.

(2013), so we will just give a familiar example viewed through the lens of tensor transformation rules.

Example 2.13 (full-rank least-squares). As we saw in Section 2.1, the least-squares problem (2.3) satisfies the transformation rule of a mixed 2-tensor $A' = XAY^{-1}$ with change-of-basis matrices $(X, Y) \in O(m) \times GL(n)$. Suppose $\text{rank}(A) = n$. Then, applying a sequence of covariant 1-tensor transformation rules

$$A \rightarrow Q_1^T A \rightarrow Q_2^T (Q_1^T A) \rightarrow \cdots \rightarrow Q^T A = \begin{bmatrix} R \\ 0 \end{bmatrix}$$

given by the *Householder QR algorithm*, we get

$$A = Q \begin{bmatrix} R \\ 0 \end{bmatrix}.$$

As the minimum value is an invariant Cartesian 0-tensor,

$$\begin{aligned} \min \|Av - b\|^2 &= \min \|Q^T(Av - b)\|^2 = \min \left\| \begin{bmatrix} R \\ 0 \end{bmatrix} v - Q^T b \right\|^2 \\ &= \min \left\| \begin{bmatrix} R \\ 0 \end{bmatrix} v - \begin{bmatrix} c \\ d \end{bmatrix} \right\|^2 = \min \|Rv - c\|^2 + \|d\|^2 = \|d\|^2, \end{aligned}$$

where we have written

$$Q^T b = \begin{bmatrix} c \\ d \end{bmatrix}.$$

In this case the solution of the transformed problem $Rv = c$ is in fact equal to that of the least-squares problem, and may be obtained through *back-substitution*, that is, another sequence of contravariant 1-tensor transformation rules

$$c \rightarrow Y_1^{-1} c \rightarrow Y_2^{-1} (Y_1^{-1} c) \rightarrow \cdots \rightarrow R^{-1} c = v,$$

where the Y_i are Gauss transforms. As noted above, the solution method reflects Voigt's definition: we transform the problem $\min \|Av - b\|^2$ into a form where the solution of the transformed problem $Rv = c$ is related to the original solution in a definite way. Here we obtained the change-of-basis matrices $X = Q \in O(m)$ and $Y = R^{-1} \in GL(n)$ via Householder QR and back-substitution respectively.

As we just saw, there are distinguished choices for the change-of-basis matrices that aid the solution of a numerical linear algebra problem. We will mention one of the most useful choices below.

Example 2.14 (Krylov subspaces). Suppose we have $A \in \mathbb{R}^{n \times n}$, with all eigenvalues distinct and non-zero for simplicity. Take an arbitrary $b \in \mathbb{R}^n$. Then the matrix K whose columns are *Krylov basis* vectors of the form

$$b, Ab, A^2 b, \dots, A^{n-1} b$$

is invertible, *i.e.* $K \in \text{GL}(n)$, and using it as the change-of-basis matrix gives us

$$A = K \begin{bmatrix} 0 & 0 & \cdots & 0 & -c_0 \\ 1 & 0 & \cdots & 0 & -c_1 \\ 0 & 1 & \cdots & 0 & -c_2 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & -c_{n-1} \end{bmatrix} K^{-1}. \quad (2.22)$$

This is a special case of the rational canonical form, a canonical form under similarity. The seemingly trivial observation (2.22), when combined with other techniques, becomes a powerful iterative method for a wide variety of computational tasks such as solving linear systems, least-squares, eigenvalue problems or evaluating various matrix functions (van der Vorst 2000). Readers unfamiliar with numerical linear algebra may find it odd that we do not use another obvious canonical form, one that makes the aforementioned problems trivial to solve, namely, the eigenvalue decomposition

$$A = X \begin{bmatrix} \lambda_1 & 0 & 0 & \cdots & 0 \\ 0 & \lambda_2 & 0 & \cdots & 0 \\ 0 & 0 & \lambda_3 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & \lambda_n \end{bmatrix} X^{-1}, \quad (2.23)$$

where the change-of-basis matrix $X \in \text{GL}(n)$ has columns given by the eigenvectors of A . The issue is that this is more difficult to compute than (2.22). In fact, to compute it, one way is to implicitly exploit (2.22) and the relation between the two canonical forms:

$$\begin{bmatrix} \lambda_1 & 0 & 0 & \cdots & 0 \\ 0 & \lambda_2 & 0 & \cdots & 0 \\ 0 & 0 & \lambda_3 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & \lambda_n \end{bmatrix} = V \begin{bmatrix} 0 & 0 & \cdots & 0 & -c_0 \\ 1 & 0 & \cdots & 0 & -c_1 \\ 0 & 1 & \cdots & 0 & -c_2 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & -c_{n-1} \end{bmatrix} V^{-1},$$

where

$$V = \begin{bmatrix} 1 & \lambda_1 & \lambda_1^2 & \cdots & \lambda_1^{n-1} \\ 1 & \lambda_2 & \lambda_2^2 & \cdots & \lambda_2^{n-1} \\ 1 & \lambda_3 & \lambda_3^2 & \cdots & \lambda_3^{n-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \lambda_n & \lambda_n^2 & \cdots & \lambda_n^{n-1} \end{bmatrix}. \quad (2.24)$$

The pertinent point here is that the approach of solving a problem by finding an appropriate computational basis is also an instance of the 2-tensor transformation rule. Aside from the Krylov basis, there are simpler examples such as diagonalizing

a circulant matrix

$$C = \begin{bmatrix} c_0 & c_{n-1} & \cdots & c_2 & c_1 \\ c_1 & c_0 & c_{n-1} & & c_2 \\ \vdots & c_1 & c_0 & \ddots & \vdots \\ c_{n-2} & & \ddots & \ddots & c_{n-1} \\ c_{n-1} & c_{n-2} & \cdots & c_1 & c_0 \end{bmatrix}$$

by expressing it in the Fourier basis, *i.e.* with $\lambda_j^k = e^{2(j-1)k\pi i/n}$ in (2.24), and there are far more complex examples such as the wavelet approach to the fast multipole method of [Beylkin, Coifman and Rokhlin \(1991\)](#). This computes, in $O(n)$ operations and to arbitrary accuracy, a matrix–vector product $v \mapsto Av$ with a dense $A \in \mathbb{R}^{n \times n}$ that is a finite-dimensional approximation of certain special integral transforms. The algorithm iteratively computes a special wavelet basis $X \in GL(n)$ so that ultimately $X^{-1}AX = B$ gives a banded matrix B where $v \mapsto Bv$ can be computed in time $O(n \log(1/\varepsilon))$ to ε -accuracy, and where $v \mapsto Xv, v \mapsto X^{-1}v$ are both computable in time $O(n)$. One may build upon this algorithm to obtain an $O(n \log^2 n \log \kappa(A))$ algorithm for the pseudoinverse ([Beylkin 1993](#), Section 6) and potentially also to compute other matrix functions such as square root, exponential, sine and cosine, *etc.* ([Beylkin, Coifman and Rokhlin 1992](#), Section X). We will discuss some aspects of this basis in Example 4.47, which is constructed as the tensor product of multiresolution analyses.

Our intention is to highlight the role of the tensor transformation rules in numerical linear algebra but we do not wish to overstate it. These rules are an important component of various algorithms but almost never the only one. Furthermore, far more work is required to guarantee correctness in finite-precision arithmetic ([Higham 2002](#)), although the tensor transformation rules can sometimes help with that too, as we will see next.

Example 2.15 (affine invariance of Newton’s method). The discussion below is essentially the standard treatment in [Boyd and Vandenberghe \(2004\)](#), supplemented by commentary pointing out the tensor transformation rules. Consider the equality-constrained convex optimization problem

$$\begin{aligned} & \text{minimize} && f(v) \\ & \text{subject to} && Av = b \end{aligned} \tag{2.25}$$

for a strongly convex $f \in C^2(\Omega)$ with $\beta I \leq \nabla^2 f(v) \leq \gamma I$. The Newton step $\Delta v \in \mathbb{R}^n$ is defined as the solution to

$$\begin{bmatrix} \nabla^2 f(v) & A^\top \\ A & 0 \end{bmatrix} \begin{bmatrix} \Delta v \\ \Delta \lambda \end{bmatrix} = \begin{bmatrix} -\nabla f(v) \\ 0 \end{bmatrix}$$

and the Newton decrement $\lambda(v) \in \mathbb{R}$ is defined as

$$\lambda(v)^2 := \nabla(v)^\top \nabla f(v)^{-1} \nabla f(v).$$

Let $X \in GL(n)$ and suppose we perform a linear change of coordinates $Xv' = v$. Then the Newton step in these new coordinates is given by

$$\begin{bmatrix} X^T \nabla^2 f(Xv) X & X^T A^T \\ AX & 0 \end{bmatrix} \begin{bmatrix} \Delta v' \\ \Delta \lambda' \end{bmatrix} = \begin{bmatrix} -X^T \nabla f(Xv) \\ 0 \end{bmatrix}.$$

We may check that $X\Delta v' = \Delta v$ and thus the iterates are related by $Xv'_k = v_k$ for all $k \in \mathbb{N}$ as long as we initialize with $Xv'_0 = v_0$ (Boyd and Vandenberghe 2004, Section 10.2.1). Note that steepest descent satisfies no such property no matter which 1-tensor transformation rule we use: $v' = Xv$, $v' = X^{-1}v$, $v' = X^T v$, or $v' = X^{-T}v$. We also have that $\lambda(Xv)^2 = \lambda(v)^2$, which is used in the stopping condition of Newton’s method; thus the iterations stop at the same point $Xv'_k = v_k$ when $\lambda(v'_k)^2 = \lambda(v_k)^2 \leq 2\varepsilon$ for a given $\varepsilon > 0$. In summary, if we write $g(v') = f(Xv)$, then we have the following relations:

coordinates	contravariant 1-tensor	$v' = X^{-1}v,$
gradient	covariant 1-tensor	$\nabla g(v') = X^T \nabla f(Xv),$
Hessian	covariant 2-tensor	$\nabla^2 g(v') = X^T \nabla^2 f(Xv) X,$
Newton step	contravariant 1-tensor	$\Delta v' = X^{-1} \Delta v,$
Newton iterate	contravariant 1-tensor	$v'_k = X^{-1} v_k,$
Newton decrement	invariant 0-tensor	$\lambda(v'_k) = \lambda(v_k).$

Strictly speaking, the gradient and Hessian are tensor fields, and we will explain the difference in Example 3.12. We may extend the above discussion with the general affine group $GA(n)$ in (2.15) in place of $GL(n)$ to incorporate translation by a vector.

To see the implications on computations, it suffices to take, for simplicity, an unconstrained problem with $A = 0$ and $b = 0$ in (2.25). Observe that the condition number of $X^T \nabla^2 f(Xv) X$ can be scaled to any desired value in $[1, \infty)$ with an appropriate $X \in GL(n)$. In other words, any Newton step is independent of the condition number of $\nabla^2 f(v)$ in exact arithmetic. In finite precision arithmetic, this manifests itself as insensitivity to condition number. Newton’s method gives solutions of high accuracy for condition number as high as $\kappa \approx 10^{10}$ when steepest descent already stops working at $\kappa \approx 20$ (Boyd and Vandenberghe 2004, Section 9.5.4). Strictly speaking, κ refers to the condition number of sublevel sets $\{v \in \Omega: f(v) \leq \alpha\}$. For any convex set $C \subseteq \mathbb{R}^n$, this is given by

$$\kappa(C) := \frac{\sup_{\|v\|=1} w(C, v)^2}{\inf_{\|v\|=1} w(C, v)^2}, \quad w(C, v) := \sup_{u \in C} u^T v - \inf_{u \in C} u^T v.$$

However, if $\beta I \leq \nabla^2 f(v) \leq \gamma I$, then $\kappa(\{v \in \Omega: f(v) \leq \alpha\}) \leq \gamma/\beta$ (Boyd and Vandenberghe 2004, Section 9.1.2), and so it is ultimately controlled by $\kappa(\nabla^2 f(v))$.

Our last example shows that the tensor transformation rules are as important in the information sciences as they are in the physical sciences.

Example 2.16 (equivariant neural networks). A feed-forward neural network is usually regarded as a function $f: \mathbb{R}^n \rightarrow \mathbb{R}^n$ obtained by alternately composing affine maps $\alpha_i: \mathbb{R}^n \rightarrow \mathbb{R}^n, i = 1, \dots, k$, with a non-linear function $\sigma: \mathbb{R}^n \rightarrow \mathbb{R}^n$:

$$f = \alpha_k \circ \sigma \circ \alpha_{k-1} \circ \dots \circ \sigma \circ \alpha_2 \circ \sigma \circ \alpha_1$$

The *depth*, also known as the number of *layers*, is k and the *width*, also known as the number of *neurons*, is n . We assume that our neural network has constant width throughout all layers. The non-linear function σ is called *activation* and we may assume that it is given by the ReLU function $\sigma(t) := \max(t, 0)$ for $t \in \mathbb{R}$ and, by convention,

$$\sigma(v) = (\sigma(v_1), \dots, \sigma(v_n)), \quad v = (v_1, \dots, v_n) \in \mathbb{R}^n. \tag{2.26}$$

The affine function is defined by $\alpha_i(v) = A_i v + b_i$ for some $A_i \in \mathbb{R}^{n \times n}$ called the *weight* matrix and some $b_i \in \mathbb{R}^n$ called the *bias* vector. We assume that $b_k = 0$ in the last layer.

Although convenient, it is somewhat misguided to be lumping the bias and weight together in an affine function. The biases b_i are intended to serve as *thresholds* for the activation function σ (Bishop 2006, Section 4.1.7) and should be part of it, detached from the weights A_i that transform the input. If one would like to incorporate translations, one could do so with weights from a matrix group such as $SE(n)$ in (2.15). Hence a better but mathematically equivalent description of f would be as

$$f = A_k \sigma_{b_{k-1}} A_{k-1} \dots \sigma_{b_2} A_2 \sigma_{b_1} A_1$$

where we identify $A_i \in \mathbb{R}^{n \times n}$ with the linear operator $\mathbb{R}^n \rightarrow \mathbb{R}^n, v \mapsto A_i v$, and for any $b \in \mathbb{R}^n$ we define $\sigma_b: \mathbb{R}^n \rightarrow \mathbb{R}^n$ by $\sigma_b(v) = \sigma(v + b) \in \mathbb{R}^n$. We have dropped the composition symbol \circ to avoid clutter and potential confusion with the Hadamard product (2.4), and will continue to do so below. For a fixed $\theta \in \mathbb{R}$,

$$\sigma_\theta(t) = \begin{cases} t - \theta & t \geq \theta, \\ 0 & t < \theta, \end{cases}$$

plays the role of a threshold for activation as was intended by Rosenblatt (1958, p. 392) and McCulloch and Pitts (1943, p. 120).

A major computational issue with neural networks is the large number of unknown parameters, namely the $kn^2 + k(n - 1)$ entries of the weights and biases, that have to be fitted with data, especially for deep neural networks where k is large.

Thus successful applications of neural networks require that we identify, based on the problem at hand, an appropriate low-dimensional subset of $\mathbb{R}^{n \times n}$ from which we will find our weights A_1, \dots, A_k . For instance, the very successful convolutional neural networks for image recognition (Krizhevsky, Sutskever and Hinton 2012) relies on restricting A_1, \dots, A_k to some block-Toeplitz–Toeplitz-block or BTTB matrices (Ye and Lim 2018a, Section 13) determined by a very small number of parameters. It turns out that convolutional neural networks are a quintessential example of equivariant neural networks (Cohen and Welling 2016), and in fact every equivariant neural network may be regarded as a generalized convolutional neural network in an appropriate sense (Kondor and Trivedi 2018). We will describe a simplified version that captures its essence and illustrates the tensor transformation rules.

Let $G \subseteq \mathbb{R}^{n \times n}$ be a matrix group. A function $f: \mathbb{R}^n \rightarrow \mathbb{R}^n$ is said to be equivariant if it satisfies the condition that

$$f(Xv) = Xf(v) \quad \text{for all } v \in \mathbb{R}^n, X \in G. \quad (2.27)$$

An equivariant neural network is simply a feed-forward neural network $f: \mathbb{R}^n \rightarrow \mathbb{R}^n$ that satisfies (2.27). The key to constructing an equivariant neural network is the trivial observation that

$$\begin{aligned} f(Xv) &= A_k \sigma_{b_{k-1}} A_{k-1} \cdots \sigma_{b_2} A_2 \sigma_{b_1} A_1 Xv \\ &= X(X^{-1} A_k X)(X^{-1} \sigma_{b_{k-1}} X)(X^{-1} A_{k-1} X) \cdots \\ &\quad \cdots (X^{-1} \sigma_{b_2} X)(X^{-1} A_2 X)(X^{-1} \sigma_{b_1} X)(X^{-1} A_1 X)v \\ &= X A'_k \sigma'_{b_{k-1}} A'_{k-1} \cdots \sigma'_{b_2} A'_2 \sigma'_{b_1} A'_1 v \end{aligned}$$

and the last expression equals $Xf(v)$ as long as we have

$$A'_i = X^{-1} A_i X = A_i, \quad \sigma'_{b_i} = X^{-1} \sigma_{b_i} X = \sigma_{b_i} \quad (2.28)$$

for all $i = 1, \dots, k$, and for all $X \in G$. The condition on the right is satisfied by any activation as long as it is applied coordinate-wise (Cohen and Welling 2016, Section 6.2) as we did in (2.26). We will illustrate this below. The condition on the left, which says that A_i is invariant under the mixed 2-tensor transformation rule on page 574, is what is important. The condition limits the possible weights for f to a (generally) much smaller subspace of matrices that commute with all elements of G . However, finding this subspace (in fact a subalgebra) of *intertwiners*,

$$\{A \in \mathbb{R}^{n \times n} : AX = XA \text{ for all } X \in G\}, \quad (2.29)$$

is usually where the challenge lies, although it is also a classical problem with numerous powerful tools from group representation theory (Bröcker and tom Dieck 1995, Fulton and Harris 1991, Mackey 1978).

There are many natural candidates for the group G from the list in (2.14) and their subgroups, depending on the problem at hand. But we caution the reader that G will generally be a very low-dimensional subset of $\mathbb{R}^{n \times n}$. It will be pointless to

pick, say, $G = \text{SO}(n)$ as the set in (2.29) will then be just $\{\lambda I \in \mathbb{R}^{n \times n} : \lambda \in \mathbb{R}\}$, clearly too small to serve as meaningful weights for any neural network. Indeed, G is usually chosen to be an image of a much lower-dimensional group H ,

$$G = \rho(H) = \{\rho(h) \in \text{GL}(n) : h \in H\}$$

for some group homomorphism $\rho : H \rightarrow \text{GL}(n)$; here ρ is called a *representation* of the group H . In image recognition applications, possibilities for H are often discrete subgroups of $\text{SE}(2)$ or $\text{E}(2)$ such as the group of translations

$$H = \left\{ \begin{bmatrix} 1 & 0 & m_1 \\ 0 & 1 & m_2 \\ 0 & 0 & 1 \end{bmatrix} \in \mathbb{R}^{3 \times 3} : m_1, m_2 \in \mathbb{Z} \right\}$$

for convolutional neural networks; the $p4$ group that augments translations with right-angle rotations,

$$H = \left\{ \begin{bmatrix} \cos(k\pi/2) & -\sin(k\pi/2) & m_1 \\ \sin(k\pi/2) & \cos(k\pi/2) & m_2 \\ 0 & 0 & 1 \end{bmatrix} \in \mathbb{R}^{3 \times 3} : k = 0, 1, 2, 3, m_1, m_2 \in \mathbb{Z} \right\}$$

in Cohen and Welling (2016, Section 4.2); the $p4m$ group that further augments $p4$ with reflections,

$$H = \left\{ \begin{bmatrix} (-1)^j \cos(k\pi/2) & (-1)^{j+1} \sin(k\pi/2) & m_1 \\ \sin(k\pi/2) & \cos(k\pi/2) & m_2 \\ 0 & 0 & 1 \end{bmatrix} \in \mathbb{R}^{3 \times 3} : \begin{array}{l} k = 0, 1, 2, 3, \\ j = 0, 1, m_1, m_2 \in \mathbb{Z} \end{array} \right\}$$

in Cohen and Welling (2016, Section 4.3). Other possibilities for H include the rotation group $\text{SO}(3)$ for 3D shape recognition (Kondor *et al.* 2018), the rigid motion group $\text{SE}(3)$ for chemical property (Fuchs *et al.* 2020) and protein structure (see page 560) predictions and the Lorentz group $\text{SO}(1, 3)$ for identifying top quarks in high-energy physics experiments (Bogatskiy *et al.* 2020), *etc.*


When people speak of $\text{SO}(3)$ - or $\text{SE}(3)$ - or Lorentz-equivariant neural networks, they are referring to the group H and not $G = \rho(H)$. A key step of these works is the construction of an appropriate representation ρ for the problem at hand, or, equivalently, constructing a linear action of H on \mathbb{R}^n . In these applications \mathbb{R}^n should be regarded as the set of real-valued functions on a set S of cardinality n , a perspective that we will introduce in Example 4.5. For concreteness, take an image recognition problem on a 60 000 collection of 28×28 -pixel images of handwritten digits in greyscale levels $0, 1, \dots, 255$ (Deng 2012). Then $S \subseteq \mathbb{Z}^2$ is the set of $n = 28^2 = 784$ pixel indices and an image is encoded as $v \in \mathbb{R}^{784}$ whose coordinates take values from 0 (pitch black) to 255 (pure white). Note that this is why the first three H above are discrete: the elements $h \in H$ act on the pixel indices $S \subseteq \mathbb{Z}^2$ instead of \mathbb{R}^2 . These 60 000 images are then used to fit the neural network f , *i.e.* to find the parameters $A_1, \dots, A_k \in \mathbb{R}^{784 \times 784}$ and $b_1, \dots, b_{k-1} \in \mathbb{R}^{784}$.

We end this example with a note on why a non-linear (and not even multilinear) function like ReLU can satisfy the covariant 2-tensor transformation rule, which sounds incredible but is actually obvious once explained. Take a greyscale image $v \in \mathbb{R}^n$ drawn with a black outline (greyscale value 0) but filled with varying shades

of grey (greyscale values $1, \dots, 255$) and consider the activation

$$\sigma(t) = \begin{cases} 255 & t > 0, \\ 0 & t \leq 0, \end{cases} \quad t \in \mathbb{R},$$

so that applying σ to the image v produces an image $\sigma(v) \in \mathbb{R}^n$ with all shadings removed, leaving just the black outline. Now take a 45° rotation matrix $R \in \text{SO}(2)$ and let $X = \rho(R) \in \text{GL}(n)$ be the corresponding matrix that rotates any image v by 45° to Xv :



The bottom line is that R acts on the indices of v whereas σ acts on the values of v and the two actions are always independent, which is why $\sigma' = X^{-1}\sigma X = \sigma$.

Our definition of equivariance in (2.27) is actually a simplification but suffices for our purposes. We state a formal definition here for completeness. A function $f: \mathbb{V} \rightarrow \mathbb{W}$ is said to be H -equivariant if there are two homomorphisms $\rho_1: H \rightarrow \text{GL}(\mathbb{V})$ and $\rho_2: H \rightarrow \text{GL}(\mathbb{W})$ so that

$$f(\rho_1(h)v) = \rho_2(h)f(v) \quad \text{for all } v \in \mathbb{V}, h \in H.$$

In (2.27), we chose $\mathbb{V} = \mathbb{W} = \mathbb{R}^n$ and $\rho_1 = \rho_2$. The condition $f(Xv) = Xf(v)$ could have been replaced by other tensor or pseudotensor transformation rules, say, $f(Xv) = \det(X)X^{-\top}f(v)$, as long as they are given by homomorphisms. In fact, with slight modifications we may even allow for *antihomomorphisms*, i.e. $\rho(h_1h_2) = \rho(h_2)\rho(h_1)$.

There are many other areas in computations where the tensor transformation rules make an appearance. For example, both the primal and dual forms of a cone programming problem over a symmetric cone $\mathbb{K} \subseteq \mathbb{V}$ – which include as special cases linear programming (LP), convex quadratic programming (QP), second-order cone programming (SOCP) and semidefinite programming (SDP) – conform to the transformation rules for Cartesian 0-, 1- and 2-tensors. However, the change-of-basis matrices would have to be replaced by a linear map from the *orthogonal group of the cone* (Hauser and Lim 2002):

$$\text{O}(\mathbb{K}) := \{\varphi: \mathbb{V} \rightarrow \mathbb{V}: \varphi \text{ linear, invertible and } \varphi^* = \varphi^{-1}\}.$$

Here the vector space \mathbb{V} involved may not be \mathbb{R}^n – for SDP, $\mathbb{K} = \mathbb{S}_{++}^n$ and $\mathbb{V} = \mathbb{S}^n$, the space of $n \times n$ symmetric matrices – making the linear algebra notation we have been using to describe definition ① awkward and unnatural. Instead we should make provision to work with tensors over arbitrary vector spaces, for example the space of Toeplitz or Hankel or Toeplitz-plus-Hankel matrices, the space of polynomials

or differential forms or differential operators, or, in the case of equivariant neural networks, the space of L^2 -functions on homogeneous spaces (Cohen and Welling 2016, Kondor and Trivedi 2018). This serves as another motivation for definitions ② and ③.

2.5. Fallacy of ‘tensor = hypermatrix’

We conclude our discussion of definition ① with a few words about the fallacy of identifying a tensor with the hypermatrix that represents it. While we have alluded to this from time to time over the last few sections, there are three points that we have largely deferred until now.

Firstly, the ‘multi-indexed object’ in definition ① is not necessarily a hypermatrix; they could be polynomials or exterior forms or non-commutative polynomials. For example, it is certainly more fruitful to represent symmetric tensors as polynomials and alternating tensors as exterior forms. For one, we may take derivatives and integrals or evaluate them at a point – operations that become awkward if we simply regard them as hypermatrices with symmetric or skew-symmetric entries.

Secondly, if one subscribes to this fallacy, then one would tend to miss important tensors hiding in plain sight. The object of essence in each of the following examples is a 3-tensor, but one sees no triply indexed quantities anywhere.

(i) Multiplication of complex numbers:

$$(a + ib)(c + id) = (ac - bd) + i(bc + ad).$$

(ii) Matrix–matrix products:

$$\begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix} = \begin{bmatrix} a_{11}b_{11} + a_{12}b_{21} & a_{11}b_{12} + a_{12}b_{22} \\ a_{21}b_{11} + a_{22}b_{21} & a_{21}b_{12} + a_{22}b_{22} \end{bmatrix}.$$

(iii) Grothendieck’s inequality:

$$\max_{\|x_i\|=\|y_j\|=1} \left| \sum_{i=1}^m \sum_{j=1}^n a_{ij} \langle x_i, y_j \rangle \right| \leq K_G \max_{|\varepsilon_i|=|\delta_j|=1} \left| \sum_{i=1}^m \sum_{j=1}^n a_{ij} \varepsilon_i \delta_j \right|.$$

(iv) Separable representations:

$$\begin{aligned} \sin(x + y + z) &= \sin(x) \cos(y) \cos(z) + \cos(x) \cos(y) \sin(z) \\ &\quad + \cos(x) \sin(y) \cos(z) - \sin(x) \sin(y) \sin(z). \end{aligned}$$

Thirdly, the value of representing a d -tensor as a d -dimensional hypermatrix is often overrated.

- (a) *It may not be meaningful.* For tensors in, say, $L^2(\mathbb{R}) \otimes L^2(\mathbb{R}) \otimes L^2(\mathbb{R})$, where the ‘indices’ are continuous, the hypermatrix picture is not meaningful.
- (b) *It may not be computable.* Given a bilinear operator, writing down its entries as a three-dimensional hypermatrix with respect to bases is in general a #P-hard problem, as we will see in Example 3.5.

- (c) *It may not be useful.* In applications such as tensor networks, d is likely large. While there may be some value in picturing a three-dimensional hypermatrix, it is no longer the case when $d \gg 3$.
- (d) *It may not be possible.* There are applications such as Examples 3.8 and 3.14 that involve tensors defined over modules (vector spaces whose field of scalars is replaced by a ring). Unlike vector spaces, modules may not have bases and such tensors cannot be represented as hypermatrices.

By far the biggest issue with thinking that a tensor is just a hypermatrix is that such an overly simplistic view disregards the transformation rules, which is the key to definition ① and whose importance we have already discussed at length. As we saw, $\begin{bmatrix} 1 & 2 & 3 \\ 2 & 3 & 4 \\ 3 & 4 & 5 \end{bmatrix}$ may represent a 1-tensor or a 2-tensor, it may represent a covariant or contravariant or mixed tensor, or it may not represent a tensor at all. What matters in definition ① is the transformation rules, not the multi-indexed object. If one thinks that a fraction is just a pair of numbers, adding them as $a/b + c/d = (a + c)/(b + d)$ is going to seem perfectly fine. It is the same with tensors.

3. Tensors via multilinearity

The main deficiency of definition ①, namely that it leaves the tensor unspecified, is immediately remedied by definition ②, which states unequivocally that the tensor is a multilinear map. The multi-indexed object in definition ① is then a coordinate representation of the multilinear map with respect to a choice of bases and the transformation rules become a change-of-basis theorem. By the 1980s, definition ② had by and large supplanted definition ① as the standard definition of a tensor in textbooks on mathematical methods for physicists (Abraham *et al.* 1988, Choquet-Bruhat, DeWitt-Morette and Dillard-Bleick 1982, Hassani 1999, Martin 1991), geometry (Boothby 1986, Helgason 1978), general relativity (Misner *et al.* 1973, Wald 1984) and gauge theory (Bleecker 1981), although it dates back to at least Temple (1960).

Definition ② appeals to physicists for two reasons. The first is the well-known maxim that the laws of physics do not depend on coordinates, and thus the tensors used to express these laws should not depend on coordinates either. This has been articulated in various forms in many places, but we quote a modern version that appears as the very first sentence in Kip Thorne's lecture notes (now a book).

Geometric principle. The laws of physics must all be expressible as geometric (coordinate-independent and reference-frame-independent) relationships between geometric objects (scalars, vectors, tensors, ...) that represent physical entities (Thorne and Blandford 2017).

Most well-known equations in physics express a relation between tensors. For example, Newton's second law $F = ma$ is a relation between two contravariant 1-tensors – force F and acceleration a – and Einstein's field equation $G + \Lambda g = \kappa T$

is a relation between three covariant 2-tensors – the Einstein tensor G , the metric tensor g and the energy–momentum tensor T , with constants Λ and κ . Since such laws of physics should not depend on coordinates, a tensor ought to be a coordinate-free object, and definition ② meets this requirement: whether or not a map $f: \mathbb{V}_1 \times \cdots \times \mathbb{V}_d \rightarrow \mathbb{W}$ is multilinear does not depend on a choice of bases on $\mathbb{V}_1, \dots, \mathbb{V}_d, \mathbb{W}$.

For the same reason, we expect laws of physics expressed in terms of the Hadamard product to be rare if not non-existent, precisely because the Hadamard product is coordinate-dependent, as we saw at the end of Section 2.3. Take two linear maps $f, g: \mathbb{V} \rightarrow \mathbb{V}$: it will not be possible to define the Hadamard product of f and g without choosing some basis on \mathbb{V} and the value of the product will depend on that chosen basis. So one advantage of using a coordinate-free definition of tensors such as definitions ② or ③ is that it automatically avoids pitfalls such as performing coordinate-dependent operations that are undefined on tensors.

The second reason for favouring definition ② in physics is that the notion of multilinearity is central to the subject. We highlight another two maxims.

Linearity principle. Almost any natural process is linear in small amounts almost everywhere (Kostrikin and Manin 1997, p. vii).

Multilinearity principle. If we keep all but one factor constant, the varying factor obeys the linearity principle.

It is probably fair to say that the world around us is understandable largely because of a combination of these two principles. The universal importance of linearity needs no elaboration; in this section we will discuss the importance of multilinearity, particularly in computations.

In mathematics, definition ② appeals because it is the simplest among the three definitions: a multilinear map is trivial to define for anyone who knows about linear maps. More importantly, definition ② allows us to define tensors over arbitrary vector spaces.⁹ It is a misconception that in applications there is no need to discuss more general vector spaces because one can always get by with just two of them, \mathbb{R}^n and \mathbb{C}^n . The fact of the matter is that other vector spaces often carry structures that are destroyed when one artificially identifies them with \mathbb{R}^n or \mathbb{C}^n . Just drawing on the examples we have mentioned in Section 2, semidefinite programming and equivariant neural networks already indicate why it is not a good idea to identify the vector space of symmetric $n \times n$ matrices with $\mathbb{R}^{n(n+1)/2}$ or an m -dimensional subspace of real-valued functions $f: \mathbb{Z}^2 \rightarrow \mathbb{R}$ with \mathbb{R}^m . We will elaborate on these later in the context of definition ②. Nevertheless, we will also see that definition ② has one deficiency that will serve as an impetus for definition ③.

⁹ Note that we do not mean *abstract* vector spaces. There are many *concrete* vector spaces aside from \mathbb{R}^n and \mathbb{C}^n that are useful in applications and computations. Definitions ② and ③ allow us to define tensors over these.

3.1. Multilinear maps

Again we will begin from tensors of orders zero, one and two, drawing on familiar concepts from linear algebra – vector spaces, linear maps, dual vector spaces – before moving into less familiar territory with order-three tensors followed by the most general order- d version. In the following we let $\mathbb{U}, \mathbb{V}, \mathbb{W}$ be vector spaces over a field of scalars, assumed to be \mathbb{R} for simplicity but it may also be replaced by \mathbb{C} or indeed any ring such as \mathbb{Z} or $\mathbb{Z}/n\mathbb{Z}$, and we will see that such cases are also important in computations and applications.

Definition ② assumes that we are given at least one vector space \mathbb{V} . A tensor of order zero is simply defined to be a *scalar*, i.e. an element of the field \mathbb{R} . A tensor of order one may either be a *vector*, i.e. an element of the vector space $v \in \mathbb{V}$, or a *covector*, i.e. an element of a dual vector space $\varphi \in \mathbb{V}^*$. A covector, also called a dual vector or linear functional, is a map $\varphi: \mathbb{V} \rightarrow \mathbb{R}$ that satisfies

$$\varphi(\lambda v + \lambda' v') = \lambda \varphi(v) + \lambda' \varphi(v') \tag{3.1}$$

for all $v, v' \in \mathbb{V}, \lambda, \lambda' \in \mathbb{R}$. To distinguish between these two types of tensors, a vector is called a tensor of *contravariant* order one and a covector a tensor of *covariant* order one. To see how these relate to definition ①, we pick a basis $\mathcal{B} = \{v_1, \dots, v_m\}$ of \mathbb{V} . This gives us a dual basis $\mathcal{B}^* = \{v_1^*, \dots, v_m^*\}$ on \mathbb{V}^* , that is, $v_i^*: \mathbb{V} \rightarrow \mathbb{R}$ is a linear functional satisfying

$$v_i^*(v_j) = \begin{cases} 1 & i = j, \\ 0 & i \neq j, \end{cases} \quad j = 1, \dots, m.$$

Any vector v in \mathbb{V} may be uniquely represented by a vector in \mathbb{R}^m ,

$$\mathbb{V} \ni v = a_1 v_1 + \dots + a_m v_m \quad \longleftrightarrow \quad [v]_{\mathcal{B}} := \begin{bmatrix} a_1 \\ \vdots \\ a_m \end{bmatrix} \in \mathbb{R}^m, \tag{3.2}$$

and any covector φ in \mathbb{V}^* may also be uniquely represented by a vector in \mathbb{R}^m ,

$$\mathbb{V}^* \ni \varphi = b_1 v_1^* + \dots + b_m v_m^* \quad \longleftrightarrow \quad [\varphi]_{\mathcal{B}^*} := \begin{bmatrix} b_1 \\ \vdots \\ b_m \end{bmatrix} \in \mathbb{R}^m.$$

The vector $[v]_{\mathcal{B}} \in \mathbb{R}^m$ is the *coordinate representation* of v with respect to the basis \mathcal{B} ; likewise $[\varphi]_{\mathcal{B}^*} \in \mathbb{R}^m$ is the coordinate representation of φ with respect to \mathcal{B}^* . Here $[v]_{\mathcal{B}}$ and $[\varphi]_{\mathcal{B}^*}$ are the multi-indexed (in this case single-indexed) object in definition ① but, as we have emphasized repeatedly, the multi-indexed object is not the crux of the definition. For example, take $m = 3$ and $(3, -1, 2) \in \mathbb{R}^3$: which vector does it represent? The answer is that it can represent any vector except the zero vector in \mathbb{V} or any covector except the zero covector in \mathbb{V}^* , that is, given any non-zero $v \in \mathbb{V}$, there exists a basis $\{v_1, v_2, v_3\}$ such that $v = 3v_1 - v_2 + 2v_3$; ditto

for non-zero covectors. Knowing that the multi-indexed object is $(3, -1, 2)$ tells us absolutely nothing except that it is non-zero. One needs to know (i) the basis and (ii) the transformation rules, which are given by the change-of-basis theorem, stated below for easy reference.

Recall that if $\mathcal{C} = \{v'_1, \dots, v'_m\}$ is another basis of \mathbb{V} and $X \in \mathbb{R}^{m \times m}$ is such that

$$v'_j = \sum_{i=1}^m x_{ij}v_i, \quad j = 1, \dots, m, \tag{3.3}$$

then we must have $X \in \text{GL}(m)$ and

$$[v]_{\mathcal{C}} = X^{-1}[v]_{\mathcal{B}}, \quad [\varphi]_{\mathcal{C}^*} = X^T[\varphi]_{\mathcal{B}^*} \tag{3.4}$$

for any $v \in \mathbb{V}$ and any $\varphi \in \mathbb{V}^*$. More precisely, if

$$\begin{aligned} a_1v_1 + \dots + a_mv_m &= a'_1v'_1 + \dots + a'_mv'_m, \\ b_1v_1^* + \dots + b_mv_m^* &= b'_1v_1'^* + \dots + b'_mv_m'^* \end{aligned}$$

then the vectors of coefficients $a, a', b, b' \in \mathbb{R}^m$ must satisfy

$$a' = X^{-1}a, \quad b' = X^Tb.$$

We recover the transformation rules for contravariant and covariant 1-tensors, respectively, in the middle column of the table on page 573, which we know are simply the change-of-basis theorems for vectors and covectors in linear algebra (Berberian 2014, Friedberg *et al.* 2003). It is now clear why we have been calling the matrix X a *change-of-basis matrix* when discussing the transformation rules in Section 2.

The most familiar 2-tensor is a *linear operator*,¹⁰ *i.e.* a map $\Phi: \mathbb{U} \rightarrow \mathbb{V}$ that satisfies (3.1). Let $\dim \mathbb{U} = n$ and $\dim \mathbb{V} = m$. Then, for any bases $\mathcal{A} = \{u_1, \dots, u_n\}$ on \mathbb{U} and $\mathcal{B} = \{v_1, \dots, v_m\}$ on \mathbb{V} , the linear operator Φ has a matrix representation with respect to these bases:

$$[\Phi]_{\mathcal{A}, \mathcal{B}} = A \in \mathbb{R}^{m \times n}, \tag{3.5}$$

where the entries a_{ij} are the coefficients in

$$\Phi(u_j) = \sum_{i=1}^m a_{ij}v_i, \quad j = 1, \dots, n.$$

Let \mathcal{A}' be another basis on \mathbb{U} and \mathcal{B}' another basis on \mathbb{V} . Let

$$[\Phi]_{\mathcal{A}', \mathcal{B}'} = A' \in \mathbb{R}^{m \times n}.$$

¹⁰ Following convention, linear and multilinear maps are called *functionals* if they are scalar-valued, *i.e.* the codomain is \mathbb{R} , and *operator* if they are vector-valued, *i.e.* the codomain is a vector space of arbitrary dimension.

The change-of-basis theorem for linear operators tells us that if $X \in \text{GL}(m)$ is a change-of-basis matrix on \mathbb{V} , *i.e.* defined according to (3.3), and likewise $Y \in \text{GL}(n)$ is a change-of-basis matrix on \mathbb{U} , then

$$A' = X^{-1}AY.$$

For the special case $\mathbb{U} = \mathbb{V}$ with the same bases $\mathcal{A} = \mathcal{B}$ and $\mathcal{A}' = \mathcal{B}'$, we obtain

$$A' = X^{-1}AX.$$

Thus we have recovered the transformation rules for mixed 2-tensors in the middle column of the tables on pages 573 and 574.

The next most common 2-tensor is a *bilinear functional*, that is, a map $\beta: \mathbb{U} \times \mathbb{V} \rightarrow \mathbb{R}$ that satisfies

$$\begin{aligned} \beta(\lambda u + \lambda' u', v) &= \lambda\beta(u, v) + \lambda'\beta(u', v), \\ \beta(u, \lambda v + \lambda' v') &= \lambda\beta(u, v) + \lambda'\beta(u, v') \end{aligned} \quad (3.6)$$

for all $u, u' \in \mathbb{U}$, $v, v' \in \mathbb{V}$, $\lambda, \lambda' \in \mathbb{R}$. In this case the matrix representation of β with respect to bases \mathcal{A} and \mathcal{B} is particularly easy to describe:

$$[\beta]_{\mathcal{A}, \mathcal{B}} = A \in \mathbb{R}^{m \times n}$$

is simply given by

$$a_{ij} = \beta(u_i, v_j), \quad i = 1, \dots, m, \quad j = 1, \dots, n.$$

We have the corresponding change-of-basis theorem for bilinear forms: if

$$[\beta]_{\mathcal{A}', \mathcal{B}'} = A' \in \mathbb{R}^{m \times n},$$

then

$$A' = X^T AY,$$

and when $\mathbb{U} = \mathbb{V}$, $\mathcal{A} = \mathcal{B}$, $\mathcal{A}' = \mathcal{B}'$, it reduces to

$$A' = X^T AX.$$

We have recovered the transformation rules for covariant 2-tensors in the middle column of the tables on pages 573 and 574.

Everything up to this point is more or less standard in linear algebra (Berberian 2014, Friedberg *et al.* 2003), summarized here for easy reference and to fix notation. The goal is to show that, at least in these familiar cases, the transformation rules in definition ① are the change-of-basis theorems of the multilinear maps in definition ②.

What about contravariant 2-tensors? A straightforward way is simply to say that they are bilinear functionals on dual spaces $\beta: \mathbb{U}^* \times \mathbb{V}^* \rightarrow \mathbb{R}$ and we will recover the transformation rules

$$A' = X^{-1}AY^{-T} \quad \text{and} \quad A' = X^{-1}AX^{-T}$$

via similar change-of-basis arguments. We would like to revive an older alternative,

useful in engineering (Chou and Pagano 1992, Tai 1997) and physics (Morse and Feshbach 1953, Part I, Section 1.6), which has become somewhat obscure. Given vector spaces \mathbb{U} and \mathbb{V} , we define the *dyadic product* by

$$(a_1u_1 + \dots + a_nu_n)(b_1v_1 + \dots + b_mv_m) := \sum_{i=1}^n \sum_{j=1}^m a_i b_j u_i v_j.$$

Here the juxtaposed $u_i v_j$ is not given any further interpretation, and is simply taken to be an element (called a *dyad*) in a new vector space of dimension mn called the dyadic product of \mathbb{U} and \mathbb{V} . While this appears to be a basis-dependent notion, it is actually not, and dyadics satisfy a contravariant 2-tensor transformation rule. This is a precursor to definition ③; we will see that once we insert a \otimes between the vectors to get $u_i \otimes v_j$, the dyadic product of \mathbb{U} and \mathbb{V} is just the tensor product $\mathbb{U} \otimes \mathbb{V}$ that we will discuss in Section 4.

Before moving to higher-order tensors, we highlight another impetus for definition ③. Note that there are many other maps that are also 2-tensors according to definition ②. Aside from $\Phi: \mathbb{U} \rightarrow \mathbb{V}$, we have other linear operators

$$\Phi_1: \mathbb{U}^* \rightarrow \mathbb{V}, \quad \Phi_2: \mathbb{U} \rightarrow \mathbb{V}^*, \quad \Phi_3: \mathbb{U}^* \rightarrow \mathbb{V}^*, \tag{3.7}$$

and, aside from $\beta: \mathbb{U} \times \mathbb{V} \rightarrow \mathbb{R}$, other bilinear functionals

$$\beta_1: \mathbb{U}^* \times \mathbb{V} \rightarrow \mathbb{R}, \quad \beta_2: \mathbb{U} \times \mathbb{V}^* \rightarrow \mathbb{R}, \quad \beta_3: \mathbb{U}^* \times \mathbb{V}^* \rightarrow \mathbb{R}, \tag{3.8}$$

among other more esoteric maps we would not go into here. We begin to see an inkling of why definition ②, while certainly the simplest definition of a tensor, may not be the best one. After all, according to definition ①, there ought to be only three types of 2-tensors: covariant, contravariant and mixed. If we choose bases on \mathbb{U} and \mathbb{V} , we will find that a matrix representing the bilinear functional $\beta: \mathbb{U} \times \mathbb{V} \rightarrow \mathbb{R}$ and a matrix representing the linear operator $\Phi_2: \mathbb{U} \rightarrow \mathbb{V}^*$ satisfy exactly the same transformation rules, namely that of a covariant 2-tensor $A' = X^T A Y$. One benefit of definition ① is that it classifies tensors according to their transformation rules, but this is lost in definition ②: there are many more types of multilinear maps than there are types of tensors. For example, the linear and bilinear maps above that are constructed out of two vector spaces \mathbb{U} and \mathbb{V} may be neatly categorized into three different types of 2-tensors:

covariant 2-tensor	$\Phi_2: \mathbb{U} \rightarrow \mathbb{V}^*,$	$\beta: \mathbb{U} \times \mathbb{V} \rightarrow \mathbb{R},$
contravariant 2-tensor	$\Phi_1: \mathbb{U}^* \rightarrow \mathbb{V},$	$\beta_3: \mathbb{U}^* \times \mathbb{V}^* \rightarrow \mathbb{R},$
mixed 2-tensor	$\Phi: \mathbb{U} \rightarrow \mathbb{V},$	$\beta_2: \mathbb{U} \times \mathbb{V}^* \rightarrow \mathbb{R},$
	$\Phi_3: \mathbb{U}^* \rightarrow \mathbb{V}^*,$	$\beta_1: \mathbb{U}^* \times \mathbb{V} \rightarrow \mathbb{R}.$

Any two maps in the same category satisfy the same change-of-basis theorem, that is, the matrices that represent them satisfy the same tensor transformation rule.

The number of possibilities grows as the order of the tensor increases. A 3-tensor is most commonly a *bilinear operator*, i.e. a map $\mathbb{B}: \mathbb{U} \times \mathbb{V} \rightarrow \mathbb{W}$ that satisfies the

rules in (3.6) but is vector-valued. The next most common 3-tensor is a trilinear functional, *i.e.* a map $\tau: \mathbb{U} \times \mathbb{V} \times \mathbb{W} \rightarrow \mathbb{R}$ that satisfies

$$\begin{aligned}\tau(\lambda u + \lambda' u', v, w) &= \lambda \tau(u, v, w) + \lambda' \tau(u', v, w), \\ \tau(u, \lambda v + \lambda' v', w) &= \lambda \tau(u, v, w) + \lambda' \tau(u, v', w), \\ \tau(u, v, \lambda w + \lambda' w') &= \lambda \tau(u, v, w) + \lambda' \tau(u, v, w')\end{aligned}\quad (3.9)$$

for all $u, u' \in \mathbb{U}$, $v, v' \in \mathbb{V}$, $w, w' \in \mathbb{W}$, $\lambda, \lambda' \in \mathbb{R}$. However, there are also many other possibilities, including bilinear operators

$$\mathbf{B}: \mathbb{U}^* \times \mathbb{V} \rightarrow \mathbb{W}, \quad \mathbf{B}: \mathbb{U} \times \mathbb{V}^* \rightarrow \mathbb{W}, \dots, \mathbf{B}: \mathbb{U}^* \times \mathbb{V}^* \rightarrow \mathbb{W}^* \quad (3.10)$$

and trilinear functionals

$$\tau: \mathbb{U}^* \times \mathbb{V} \times \mathbb{W} \rightarrow \mathbb{R}, \quad \tau: \mathbb{U} \times \mathbb{V}^* \times \mathbb{W} \rightarrow \mathbb{R}, \dots, \tau: \mathbb{U}^* \times \mathbb{V}^* \times \mathbb{W}^* \rightarrow \mathbb{R}, \quad (3.11)$$

among yet more complicated maps. For example, if $\mathcal{L}(\mathbb{U}; \mathbb{V})$ denotes the vector space of all linear maps from \mathbb{U} to \mathbb{V} , then operator-valued linear operators or linear operators on operator spaces such as

$$\Phi_1: \mathbb{U} \rightarrow \mathcal{L}(\mathbb{V}; \mathbb{W}), \quad \Phi_2: \mathcal{L}(\mathbb{U}; \mathbb{V}) \rightarrow \mathbb{W} \quad (3.12)$$

or bilinear functionals such as

$$\beta_1: \mathbb{U} \times \mathcal{L}(\mathbb{V}; \mathbb{W}) \rightarrow \mathbb{R}, \quad \beta_2: \mathcal{L}(\mathbb{U}; \mathbb{V}) \times \mathbb{W} \rightarrow \mathbb{R} \quad (3.13)$$

are also 3-tensors. We may substitute any of \mathbb{U} , \mathbb{V} , \mathbb{W} with \mathbb{U}^* , \mathbb{V}^* , \mathbb{W}^* in (3.12) and (3.13) to obtain yet other 3-tensors.

Organizing this myriad of multilinear maps, which increases exponentially with order, is the main reason why we would ultimately want to adopt definition ③, so that we may classify the many different types of multilinear maps into a smaller number of tensors of different types. For the treatments that rely on definition ② – a formal version of which will appear later as Definition 3.1 – such as Abraham *et al.* (1988), Bleecker (1981), Boothby (1986), Choquet-Bruhat *et al.* (1982), Hassani (1999), Helgason (1978), Martin (1991), Misner *et al.* (1973) and Wald (1984), this myriad of possibilities for a d -tensor is avoided by requiring that the codomain of the multilinear maps be \mathbb{R} or \mathbb{C} , that is, only multilinear *functionals* are allowed; we will see why this approach makes for an awkward definition of tensors.

The change-of-basis theorems for order-3 tensors are no longer standard in linear algebra, but it is easy to extrapolate from the standard change-of-basis theorems for vectors, linear functionals, linear operators and bilinear functionals we have discussed above. Let \mathbb{U} , \mathbb{V} , \mathbb{W} be vector spaces over \mathbb{R} of finite dimensions m , n , p , and let

$$\mathcal{A} = \{u_1, \dots, u_m\}, \quad \mathcal{B} = \{v_1, \dots, v_n\}, \quad \mathcal{C} = \{w_1, \dots, w_p\} \quad (3.14)$$

be arbitrary bases. We limit ourselves to bilinear operators $\mathbf{B}: \mathbb{U} \times \mathbb{V} \rightarrow \mathbb{W}$ and trilinear functionals $\tau: \mathbb{U} \times \mathbb{V} \times \mathbb{W} \rightarrow \mathbb{R}$, as other cases are easy to deduce from

these. Take any vectors $u \in \mathbb{U}$, $v \in \mathbb{V}$, $w \in \mathbb{W}$ and express them in terms of linear combinations of basis vectors

$$u = \sum_{i=1}^m a_i u_i, \quad v = \sum_{j=1}^n b_j v_j, \quad w = \sum_{k=1}^p c_k w_k.$$

Since τ is trilinear, we get

$$\tau(u, v, w) = \sum_{i=1}^m \sum_{j=1}^n \sum_{k=1}^p a_i b_j c_k \tau(u_i, v_j, w_k),$$

that is, τ is completely determined by its values on $\mathcal{A} \times \mathcal{B} \times \mathcal{C}$ in (3.14). So if we let

$$a_{ijk} := \tau(u_i, v_j, w_k), \quad i = 1, \dots, m, \quad j = 1, \dots, n, \quad k = 1, \dots, p,$$

then the hypermatrix representation of τ with respect to bases \mathcal{A} , \mathcal{B} and \mathcal{C} is

$$[\tau]_{\mathcal{A}, \mathcal{B}, \mathcal{C}} = A \in \mathbb{R}^{m \times n \times p}. \tag{3.15}$$

Note that this practically mirrors the discussion for the bilinear functional case on page 599 and can be easily extended to any multilinear functional $\varphi: \mathbb{V}_1 \times \dots \times \mathbb{V}_d \rightarrow \mathbb{R}$ to get a hypermatrix representation $A \in \mathbb{R}^{n_1 \times \dots \times n_d}$ with respect to any bases \mathcal{B}_i on $\mathbb{V}_i, i = 1, \dots, d$.

The argument for the bilinear operator B requires an additional step. By bilinearity, we get

$$B(u, v) = \sum_{i=1}^m \sum_{j=1}^n \sum_{k=1}^p a_i b_j c_k B(u_i, v_j),$$

and since $B(u_i, v_j) \in \mathbb{W}$, we may express it as a linear combination

$$B(u_i, v_j) = \sum_{k=1}^p a_{ijk} w_k, \quad i = 1, \dots, m, \quad j = 1, \dots, n, \quad k = 1, \dots, p. \tag{3.16}$$

By the fact that \mathcal{C} is a basis, the above equation uniquely defines the values a_{ijk} and the hypermatrix representation of B with respect to bases \mathcal{A} , \mathcal{B} and \mathcal{C} is

$$[B]_{\mathcal{A}, \mathcal{B}, \mathcal{C}} = A \in \mathbb{R}^{m \times n \times p}. \tag{3.17}$$

The hypermatrices in (3.15) and (3.17), even if they are identical, represent different types of 3-tensors. Let $A' = [a'_{ijk}] \in \mathbb{R}^{m \times n \times p}$ be a hypermatrix representation of τ with respect to bases $\mathcal{A}', \mathcal{B}', \mathcal{C}'$ on $\mathbb{U}, \mathbb{V}, \mathbb{W}$, i.e. $a'_{ijk} = \tau(u'_i, v'_j, w'_k)$. Then it is straightforward to deduce that

$$A' = (X^T, Y^T, Z^T) \cdot A$$

with change-of-basis matrices $X \in GL(m)$, $Y \in GL(n)$, $Z \in GL(p)$ similarly defined as in (3.3). Hence, by (2.7), trilinear functionals are covariant 3-tensors.

On the other hand, had $A' \in \mathbb{R}^{m \times n \times p}$ been a hypermatrix representation of B , then

$$A' = (X^T, Y^T, Z^{-1}) \cdot A.$$

Hence, by (2.9), bilinear operators are mixed 3-tensors of covariant order 2 and contravariant order 1. The extension to other types of 3-tensors in (3.10)–(3.13) and to arbitrary d -tensors may be carried out similarly.

For completeness, we state a formal definition of multilinear maps, if only to serve as a glossary of terms and notation and as a pretext for interesting examples.

Definition 3.1 (tensors via multilinearity). Let $\mathbb{V}_1, \dots, \mathbb{V}_d$ and \mathbb{W} be real vector spaces. A *multilinear map*, or more precisely a d -linear map, is a map $\Phi: \mathbb{V}_1 \times \dots \times \mathbb{V}_d \rightarrow \mathbb{W}$ that satisfies

$$\Phi(v, \dots, \lambda v_k + \lambda' v'_k, \dots, v_d) = \lambda \Phi(v_1, \dots, v_k, \dots, v_d) + \lambda' \Phi(v_1, \dots, v'_k, \dots, v_d) \quad (3.18)$$

for all $v_1 \in \mathbb{V}_1, \dots, v_k, v'_k \in \mathbb{V}_k, \dots, v_d \in \mathbb{V}_d, \lambda, \lambda' \in \mathbb{R}$, and all $k = 1, \dots, d$. The set of all such maps will be denoted $\mathcal{M}^d(\mathbb{V}_1, \dots, \mathbb{V}_d; \mathbb{W})$.

Note that $\mathcal{M}^d(\mathbb{V}_1, \dots, \mathbb{V}_d; \mathbb{W})$ is itself a vector space: a linear combination of two k -linear maps is again a k -linear map; in particular $\mathcal{M}^1(\mathbb{V}; \mathbb{W}) = \mathcal{L}(\mathbb{V}; \mathbb{W})$. If the vector spaces $\mathbb{V}_1, \dots, \mathbb{V}_d$ and \mathbb{W} are endowed with norms, then

$$\|\Phi\|_\sigma := \sup_{v_1, \dots, v_d \neq 0} \frac{\|\Phi(v_1, \dots, v_d)\|}{\|v_1\| \cdots \|v_d\|} = \sup_{\|v_1\| = \dots = \|v_d\| = 1} \|\Phi(v_1, \dots, v_d)\| \quad (3.19)$$

defines a norm on $\mathcal{M}^d(\mathbb{V}_1, \dots, \mathbb{V}_d; \mathbb{W})$ that we will call a *spectral norm*; we have slightly abused notation by using the same $\|\cdot\|$ to denote norms on different spaces.

In Section 3.3, we will see various examples of why (higher-order) multilinear maps are useful in computations. In part to provide necessary background for those examples, we will review their role as higher derivatives in multivariate calculus.

Example 3.2 (higher-order derivatives). Let \mathbb{V} and \mathbb{W} be norm spaces and let $\Omega \subseteq \mathbb{V}$ be an open subset. For any function $F: \Omega \rightarrow \mathbb{W}$, recall that the (total) *derivative* at $v \in \Omega$ is a linear operator $DF(v): \mathbb{V} \rightarrow \mathbb{W}$ that satisfies

$$\lim_{h \rightarrow 0} \frac{\|F(v+h) - F(v) - [DF(v)](h)\|}{\|h\|} = 0, \quad (3.20)$$

or, if there is no such linear operator, then F is not differentiable at v . The definition may be recursively applied to obtain derivatives of arbitrary order, assuming that they exist on Ω : since $DF(v) \in \mathcal{L}(\mathbb{V}; \mathbb{W})$, we may apply the same definition to $DF: \Omega \rightarrow \mathcal{L}(\mathbb{V}; \mathbb{W})$ and get $D^2F(v): \mathbb{V} \rightarrow \mathcal{L}(\mathbb{V}; \mathbb{W})$ as $D(DF)$, that is,

$$\lim_{h \rightarrow 0} \frac{\|DF(v+h) - DF(v) - [D^2F(v)](h)\|}{\|h\|} = 0. \quad (3.21)$$

Doing this recursively, we obtain

$$DF(v) \in \mathcal{L}(\mathbb{V}; \mathbb{W}), \quad D^2F(v) \in \mathcal{L}(\mathbb{V}; \mathcal{L}(\mathbb{V}; \mathbb{W})), \quad D^3F(v) \in \mathcal{L}(\mathbb{V}; \mathcal{L}(\mathbb{V}; \mathcal{L}(\mathbb{V}; \mathbb{W}))),$$

and so on. Avoiding such nested spaces of linear maps is a good reason to introduce multilinear maps. Note that we have

$$\mathcal{L}(\mathbb{V}; \mathcal{M}^{d-1}(\mathbb{V}, \dots, \mathbb{V}; \mathbb{W})) = \mathcal{M}^d(\mathbb{V}, \dots, \mathbb{V}; \mathbb{W}) \tag{3.22}$$

since, for a linear map Φ on \mathbb{V} taking values in $\mathcal{M}^{d-1}(\mathbb{V}, \dots, \mathbb{V}; \mathbb{W})$,

$$[\Phi(h)](h_1, \dots, h_{d-1})$$

must be linear in h for any fixed h_1, \dots, h_{d-1} and $(d - 1)$ -linear in h_1, \dots, h_{d-1} for any fixed h , that is, it is d -linear in the arguments h, h_1, \dots, h_{d-1} . Thus we obtain the d th-order derivative of $F: \Omega \rightarrow \mathbb{W}$ at a point $v \in \Omega$ as a d -linear map,

$$D^d F(v): \mathbb{V} \times \dots \times \mathbb{V} \rightarrow \mathbb{W}. \tag{3.23}$$

With some additional arguments, we may show that as long as $D^d F(v)$ exists in an open neighbourhood of v , then it must be a *symmetric* multilinear map in the sense that

$$[D^d F(v)](h_1, h_2, \dots, h_d) = [D^d F(v)](h_{\sigma(1)}, h_{\sigma(2)}, \dots, h_{\sigma(d)}) \tag{3.24}$$

for any permutation $\sigma \in \mathfrak{S}_d$, the permutation group on d objects; and in addition we have Taylor’s theorem,

$$F(v + h) = F(v) + [DF(v)](h) + \frac{1}{2}[D^2 F(v)](h, h) + \dots + \frac{1}{d!}[D^d F(v)](h, \dots, h) + R(h),$$

where the remainder term $\|R(h)\|/\|h\|^d \rightarrow 0$ as $h \rightarrow 0$ (Lang 2002, Chapter XIII, Section 6). Strictly speaking, the ‘=’ in (3.22) should be ‘ \cong ’, but elements in both spaces satisfy the same tensor transformation rules, so as tensors, the ‘=’ is perfectly justified; making this statement without reference to the tensor transformation rules is a reason why we would ultimately want to bring in definition ③.

Following the convention in Abraham *et al.* (1988), Bleecker (1981), Boothby (1986), Choquet-Bruhat *et al.* (1982), Hassani (1999), Helgason (1978), Martin (1991), Misner *et al.* (1973) and Wald (1984), a d -tensor is then defined as a d -linear functional, *i.e.* by setting $\mathbb{W} = \mathbb{R}$. The following is a formal version of definition ②.

Definition 3.3 (tensors as multilinear functionals). Let $p \leq d$ be non-negative integers and let $\mathbb{V}_1, \dots, \mathbb{V}_d$ be vector spaces over \mathbb{R} . A *tensor of contravariant order p and covariant order $d - p$* is a multilinear functional

$$\varphi: \mathbb{V}_1^* \times \dots \times \mathbb{V}_p^* \times \mathbb{V}_{p+1} \times \dots \times \mathbb{V}_d \rightarrow \mathbb{R}. \tag{3.25}$$

We say that φ is a tensor of *type $(p, d - p)$* and of *order d* .

This is a passable but awkward definition. For $d = 2$ and 3, the bilinear and trilinear functionals in (3.8) and (3.11) are tensors, but the linear and bilinear

operators in (3.7) and (3.10) strictly speaking are not; sometimes this means one needs to convert operators to functionals, for example by identifying a linear operator $\Phi: \mathbb{V} \rightarrow \mathbb{W}$ as the bilinear functional defined by $\mathbb{V} \times \mathbb{W}^* \rightarrow \mathbb{R}$, $(v, \varphi) \mapsto \varphi(\Phi(v))$, before one may apply a result. Definition 3.3 is also peculiar considering that by far the most common 1-, 2- and 3-tensors are vectors, linear operators and bilinear operators respectively, but the definition excludes them at the outset. So instead of simply speaking of $v \in \mathbb{V}$, one would need to regard it as a linear functional on the space of linear functionals, *i.e.* $v^{**}: \mathbb{V}^* \rightarrow \mathbb{R}$, in order to regard it as a 1-tensor in the sense of Definition 3.3. In fact, it is often more useful to do the reverse, that is, given a d -linear functional, we prefer to convert it into a $(d - 1)$ -linear operator, and we will give an example.

Example 3.4 (higher-order gradients). We consider the higher-order derivatives of a *real-valued* function $f: \Omega \rightarrow \mathbb{R}$, where $\Omega \subseteq \mathbb{V}$ is an open set and where \mathbb{V} is equipped with an inner product $\langle \cdot, \cdot \rangle$. This is a special case of Example 3.2 with $\mathbb{W} = \mathbb{R}$ and \mathbb{V} an inner product space. Since

$$D^d f(v): \underbrace{\mathbb{V} \times \cdots \times \mathbb{V}}_{d \text{ copies}} \rightarrow \mathbb{R}$$

is a multilinear *functional*, for any fixed first $d - 1$ arguments,

$$[D^d f(v)](h_1, \dots, h_{d-1}, \cdot): \mathbb{V} \rightarrow \mathbb{R}$$

is a linear functional and, by the Riesz representation theorem, there must be a vector, denoted $[\nabla^d f(v)](h_1, \dots, h_{d-1}) \in \mathbb{V}$ as it depends on h_1, \dots, h_{d-1} , such that

$$[D^d f(v)](h_1, \dots, h_{d-1}, h_d) = \langle [\nabla^d f(v)](h_1, \dots, h_{d-1}), h_d \rangle. \quad (3.26)$$

Since $D^d f(v)$ is a d -linear functional,

$$\nabla^d f(v): \underbrace{\mathbb{V} \times \cdots \times \mathbb{V}}_{d-1 \text{ copies}} \rightarrow \mathbb{V}$$

is a $(d - 1)$ -linear operator. The symmetry in (3.24) shows that our argument does not depend on our having fixed the first $d - 1$ arguments: we could have fixed any $d - 1$ arguments and still obtained the same $\nabla^d f(v)$, which must itself be symmetric in its $d - 1$ arguments. We will call $\nabla^d f(v)$ the d th gradient of f ; note that it depends on the choice of inner product. For $d = 1$, the argument above is essentially how one would define a gradient $\nabla f: \Omega \rightarrow \mathbb{V}$ in Riemannian geometry. For $d = 2$, we see that $\nabla^2 f = D(\nabla f)$, which is how a Hessian is defined in optimization (Boyd and Vandenberghe 2004, Renegar 2001). In fact, we may show more generally that, for $d \geq 2$,

$$\nabla^d f = D^{d-1}(\nabla f), \quad (3.27)$$

which may be taken as an alternative definition of higher gradients. The utility of a multilinear map approach is familiar to anyone who has attempted to find the

gradient and Hessian of functions such as

$$f(X) = \text{tr}(X^{-1}) \quad \text{or} \quad g(x, Y) = x^T Y^{-1} x,$$

where $x \in \mathbb{R}^n$ and $X, Y \in \mathbb{S}_{++}^n$ (Boyd and Vandenberghe 2004, Examples 3.4 and 3.46). The definitions in terms of partial derivatives,

$$\begin{aligned} \nabla f &= \begin{bmatrix} \frac{\partial f}{\partial x_1} \\ \vdots \\ \frac{\partial f}{\partial x_n} \end{bmatrix}, & \nabla^2 f &= \begin{bmatrix} \frac{\partial^2 f}{\partial x_1^2} & \cdots & \frac{\partial^2 f}{\partial x_1 \partial x_n} \\ \vdots & \ddots & \vdots \\ \frac{\partial^2 f}{\partial x_n \partial x_1} & \cdots & \frac{\partial^2 f}{\partial x_n^2} \end{bmatrix}, \\ \nabla^d f &= \left[\frac{\partial^d f}{\partial x_i \partial x_j \cdots \partial x_k} \right]_{i,j,\dots,k=1}^n, \end{aligned} \quad (3.28)$$

provide little insight when functions are defined on vector spaces other than \mathbb{R}^n . On the other hand, using $f(X) = \text{tr}(X^{-1})$ for illustration, with (3.20) and (3.21), we get

$$Df(X): \mathbb{S}^n \rightarrow \mathbb{R}, \quad H \mapsto -\text{tr}(X^{-1}HX^{-1}),$$

$$D^2f(X): \mathbb{S}^n \times \mathbb{S}^n \rightarrow \mathbb{R}, \quad (H_1, H_2) \mapsto \text{tr}(X^{-1}H_1X^{-1}H_2X^{-1} + X^{-1}H_2X^{-1}H_1X^{-1}),$$

and more generally

$$D^d f(X): \mathbb{S}^n \times \cdots \times \mathbb{S}^n \rightarrow \mathbb{R},$$

$$(H_1, \dots, H_d) \mapsto (-1)^d \text{tr} \left[\sum_{\sigma \in \mathfrak{S}_d} X^{-1} H_{\sigma(1)} X^{-1} H_{\sigma(2)} X^{-1} \cdots X^{-1} H_{\sigma(d)} X^{-1} \right].$$

By the cyclic invariance of trace,

$$\begin{aligned} [Df(X)](H) &= \text{tr}[-X^{-2}H], \\ [D^2f(X)](H_1, H_2) &= \text{tr}[(X^{-1}H_1X^{-2} + X^{-2}H_1X^{-1})H_2], \\ [D^3f(X)](H_1, H_2, H_3) &= \text{tr}[(-X^{-1}H_1X^{-1}H_2X^{-2} - X^{-1}H_2X^{-2}H_1X^{-1} \\ &\quad - X^{-2}H_1X^{-1}H_2X^{-1} - X^{-1}H_2X^{-1}H_1X^{-2} \\ &\quad - X^{-1}H_1X^{-2}H_2X^{-1} - X^{-2}H_2X^{-1}H_1X^{-1})H_3], \end{aligned}$$

and so by (3.26),

$$\begin{aligned} \nabla f(X) &= -X^{-2}, \\ [\nabla^2 f(X)](H) &= X^{-1}HX^{-2} + X^{-2}HX^{-1}, \\ [\nabla^3 f(X)](H_1, H_2) &= -X^{-1}H_1X^{-1}H_2X^{-2} - X^{-1}H_2X^{-2}H_1X^{-1} \\ &\quad - X^{-2}H_1X^{-1}H_2X^{-1} - X^{-1}H_2X^{-1}H_1X^{-2} \\ &\quad - X^{-1}H_1X^{-2}H_2X^{-1} - X^{-2}H_2X^{-1}H_1X^{-1}, \end{aligned}$$

all without having to calculate a single partial derivative. Example 3.6 will provide a similar illustration of how the use of multilinear maps makes such calculations routine.

Returning to Definition 3.3, suppose we choose any bases $\mathcal{B}_1 = \{u_1, \dots, u_{n_1}\}$, $\mathcal{B}_2 = \{u_2, \dots, u_{n_2}\}$, ..., $\mathcal{B}_d = \{w_1, \dots, w_{n_d}\}$ for the vector spaces $\mathbb{V}_1, \mathbb{V}_2, \dots, \mathbb{V}_d$ and let $A \in \mathbb{R}^{n_1 \times n_2 \times \dots \times n_d}$ be given by

$$a_{ij\dots k} = \varphi(u_i, v_j, \dots, w_k).$$

Then a consideration of change of basis similar to the various cases we discussed for $d = 1, 2, 3$ recovers the covariant d -tensor transformation rules (2.7) for

$$\varphi: \mathbb{V}_1 \times \dots \times \mathbb{V}_d \rightarrow \mathbb{R},$$

the contravariant d -tensor transformation rules (2.8) for

$$\varphi: \mathbb{V}_1^* \times \dots \times \mathbb{V}_d^* \rightarrow \mathbb{R},$$

and the mixed d -tensor transformation rules (2.9) for the general case in (3.25). Furthermore,

$$[\varphi]_{\mathcal{B}_1, \dots, \mathcal{B}_d} = A \in \mathbb{R}^{n_1 \times \dots \times n_d}$$

is exactly the hypermatrix representation of φ in Definition 2.5. The important special case where $\mathbb{V}_1 = \dots = \mathbb{V}_d = \mathbb{V}$ and where we pick only a single basis \mathcal{B} gives us the transformation rules in (2.10), (2.11) and (2.12).

At this juncture it is appropriate to highlight two previously mentioned points (page 594) regarding the feasibility and usefulness of representing a multilinear map as a hypermatrix.

Example 3.5 (writing down a hypermatrix is #P-hard). As we saw in (3.17), given bases \mathcal{A} , \mathcal{B} and \mathcal{C} , a bilinear operator B may be represented as a hypermatrix A . Writing down the entries a_{ijk} as in (3.16) appears to be a straightforward process, but this is an illusion: the task is #P-hard in general. Let $0 \leq d_1 \leq d_2 \leq \dots \leq d_n$ be integers. Define the generalized Vandermonde matrix

$$V_{(d_1, \dots, d_n)}(x) := \begin{bmatrix} x_1^{d_1} & x_2^{d_1} & \dots & x_n^{d_1} \\ x_1^{d_2} & x_2^{d_2} & \dots & x_n^{d_2} \\ \vdots & \vdots & \ddots & \vdots \\ x_1^{d_{n-1}} & x_2^{d_{n-1}} & \dots & x_n^{d_{n-1}} \\ x_1^{d_n} & x_2^{d_n} & \dots & x_n^{d_n} \end{bmatrix},$$

observing in particular that

$$V_{(0,1,\dots,n-1)}(x) = \begin{bmatrix} 1 & 1 & \dots & 1 \\ x_1 & x_2 & \dots & x_n \\ \vdots & \vdots & \ddots & \vdots \\ x_1^{n-2} & x_2^{n-2} & \dots & x_n^{n-2} \\ x_1^{n-1} & x_2^{n-1} & \dots & x_n^{n-1} \end{bmatrix}$$

is the usual Vandermonde matrix. Suppose $d_i \geq i$ for each $i = 1, \dots, n$; then it is not hard to show using the well-known formula $\det V_{(0,1,\dots,n-1)}(x) = \prod_{i < j} (x_i - x_j)$ that $\det V_{(d_1,d_2,\dots,d_n)}(x)$ is divisible by $\det V_{(0,1,\dots,n-1)}(x)$. So, for any integers $0 \leq p_1 \leq p_2 \leq \dots \leq p_n$,

$$s_{(p_1,p_2,\dots,p_n)}(x) := \frac{\det V_{(p_1,p_2+1,\dots,p_n+n-1)}(x)}{\det V_{(0,1,\dots,n-1)}(x)}$$

is a multivariate polynomial in the variables x_1, \dots, x_n . These are symmetric polynomials, *i.e.* homogeneous polynomials s with

$$s(x_1, x_2, \dots, x_n) = s(x_{\sigma(1)}, x_{\sigma(2)}, \dots, x_{\sigma(n)})$$

for any $\sigma \in \mathfrak{S}_n$, the permutation group on n objects. Let $\mathbb{U}, \mathbb{V}, \mathbb{W}$ be the vector spaces of symmetric polynomials of degrees d, d' and $d + d'$ respectively, and let $B: \mathbb{U} \times \mathbb{V} \rightarrow \mathbb{W}$ be the bilinear operator given by polynomial multiplication, *i.e.* $B(s(x), t(x)) = s(x)t(x)$ for any symmetric polynomials $s(x)$ of degree d and $t(x)$ of degree d' . A well-known basis of the vector space of degree- d symmetric polynomials is the *Schur basis* given by

$$\{s_{(p_1,p_2,\dots,p_n)}(x) \in \mathbb{U} : p_1 \leq p_2 \leq \dots \leq p_n \text{ is an integer partition of } d\}.$$

Let $\mathcal{A}, \mathcal{B}, \mathcal{C}$ be the respective Schur bases for $\mathbb{U}, \mathbb{V}, \mathbb{W}$. In this case the coefficients in (3.16) are called *Littlewood–Richardson coefficients*, and determining their values is a #P-complete problem¹¹ (Narayanan 2006). In other words, determining the hypermatrix representation A of the bilinear operator B is #P-hard. Littlewood–Richardson coefficients are not as esoteric as one might think but have significance in linear algebra and numerical linear algebra alike. Among other things they play a central role in the resolution of Horn’s conjecture about the eigenvalues of a sum of Hermitian matrices (Klyachko 1998, Knutson and Tao 1999).

As we mentioned earlier, a side benefit of definition ② is that it allows us to work with arbitrary vector spaces. In the previous example, $\mathbb{U}, \mathbb{V}, \mathbb{W}$ are vector spaces of degree- d symmetric polynomials for various values of d ; in the next one we will have $\mathbb{U} = \mathbb{V} = \mathbb{W} = \mathbb{S}^n$, the vector space of $n \times n$ symmetric matrices.

¹¹ This means it is as intractable as evaluating the permanent of a matrix whose entries are zeros and ones (Valiant 1979). #P-complete problems include all NP-complete problems, for example, deciding whether a graph is 3-colourable is NP-complete, but computing its chromatic number is #P-complete.

Aside from showing that hypermatrix representations may be neither feasible nor useful, Examples 3.5 and 3.6 also show that there are usually good reasons to work intrinsically with whatever vector spaces one is given.

Example 3.6 (higher gradients of log determinant). The key to all interior point methods is a *barrier function* that traps iterates within the feasible region of a convex program. In semidefinite programming, the optimal barrier function with respect to iteration complexity is the log barrier function for the cone of positive definite matrices \mathbb{S}_{++}^n ,

$$f: \mathbb{S}_{++}^n \rightarrow \mathbb{R}, \quad f(X) = -\log \det X.$$

Using the characterization in Example 3.4 with the inner product $\mathbb{S}^n \times \mathbb{S}^n \rightarrow \mathbb{R}$, $(H_1, H_2) \mapsto \text{tr}(H_1^T H_2)$, we may show that its gradient is given by

$$\nabla f: \mathbb{S}_{++}^n \rightarrow \mathbb{S}^n, \quad \nabla f(X) = -X^{-1}, \quad (3.29)$$

and its Hessian at any $X \in \mathbb{S}_{++}^n$ is the linear map

$$\nabla^2 f(X): \mathbb{S}^n \rightarrow \mathbb{S}^n, \quad H \mapsto X^{-1} H X^{-1}, \quad (3.30)$$

expressions we may also find in Boyd and Vandenberghe (2004) and Renegar (2001). While we may choose bases on \mathbb{S}^n and artificially write the gradient and Hessian of f in the forms (3.28), interested readers may check that they are a horrid mess that obliterates all insights and advantages proffered by (3.29) and (3.30). Among other things, (3.29) and (3.30) allow one to exploit specialized algorithms for matrix product and inversion.

The third-order gradient $\nabla^3 f$ also plays an important role as we need it to ascertain *self-concordance* in Example 3.16. By our discussion in Example 3.4, for any $X \in \mathbb{S}_{++}$, this is a bilinear operator

$$\nabla^3 f(X): \mathbb{S}^n \times \mathbb{S}^n \rightarrow \mathbb{S}^n,$$

and by (3.27), we may differentiate (3.30) to get

$$[\nabla^3 f(X)](H_1, H_2) = -X^{-1} H_1 X^{-1} H_2 X^{-1} - X^{-1} H_2 X^{-1} H_1 X^{-1}. \quad (3.31)$$

Repeatedly applying (3.27) gives the $(d - 1)$ -linear operator

$$\begin{aligned} \nabla^d f(X): \mathbb{S}^n \times \cdots \times \mathbb{S}^n &\rightarrow \mathbb{S}^n, \\ (H_1, \dots, H_{d-1}) &\mapsto (-1)^d \sum_{\sigma \in \mathfrak{S}_{d-1}} X^{-1} H_{\sigma(1)} X^{-1} H_{\sigma(2)} X^{-1} \cdots X^{-1} H_{\sigma(d-1)} X^{-1}, \end{aligned}$$

as the d th gradient of f . As interested readers may again check for themselves, expressing this as a d -dimensional hypermatrix is even less illuminating than expressing (3.29) and (3.30) as one- and two-dimensional hypermatrices. Multilinear maps are essential for discussing higher derivatives and gradients of multivariate functions.

In Examples 3.4 and 3.6, the matrices in \mathbb{S}^n are 1-tensors even though they are doubly indexed objects. This should come as no surprise after Example 2.3: the order of a tensor is not determined by the number of indices. We will see this in Example 3.7 again, where we will encounter a hypermatrix with $3n + 3$ indices which nonetheless represents a 3-tensor.

We conclude this section with some infinite-dimensional examples. Much like linear operators, there is not much that one could say about multilinear operators over infinite-dimensional vector spaces that is purely algebraic. The topic becomes much more interesting when one brings in analytic notions by equipping the vector spaces with norms or inner products.

As is well known, it is easy to ascertain the continuity of linear operators between Banach spaces: $\Phi: \mathbb{V} \rightarrow \mathbb{W}$ is continuous if and only if it is bounded in the sense of $\|\Phi(v)\| \leq c\|v\|$ for some constant $c > 0$ and for all $v \in \mathbb{V}$, *i.e.* if and only if $\Phi \in \mathcal{B}(\mathbb{V}, \mathbb{W})$. We have slightly abused notation by not distinguishing the norms on different spaces and will continue to do so below. Almost exactly the same proof extends to multilinear operators on Banach spaces: $\Phi: \mathbb{V}_1 \times \cdots \times \mathbb{V}_d \rightarrow \mathbb{W}$ is continuous if and only if it is bounded in the sense of

$$\|\Phi(v_1, \dots, v_d)\| \leq c\|v_1\| \cdots \|v_d\|$$

for some constant $c > 0$ and for all $v_1 \in \mathbb{V}_1, \dots, v_d \in \mathbb{V}_d$ (Lang 1993, Chapter IV, Section 1), *i.e.* if and only if its spectral norm as defined in (3.19) is finite. So if $\mathbb{V}_1, \dots, \mathbb{V}_d$ are finite-dimensional, then Φ is automatically continuous; in this case it does not matter whether \mathbb{W} is finite-dimensional. We will write $\mathcal{B}(\mathbb{V}_1, \dots, \mathbb{V}_d; \mathbb{W})$ for the set of all bounded/continuous multilinear operators. Then $\mathcal{B}(\mathbb{V}_1, \dots, \mathbb{V}_d; \mathbb{W})$ is itself a Banach space when equipped with the spectral norm.

Example 3.7 (infinite-dimensional bilinear operators). Possibly the simplest continuous bilinear operator is the infinite-dimensional generalization of matrix–vector product $\mathbb{R}^{m \times n} \times \mathbb{R}^n \rightarrow \mathbb{R}^m$, $(A, v) \mapsto Av$, but where we instead have

$$\mathbf{M}: \mathcal{B}(\mathbb{V}, \mathbb{W}) \times \mathbb{V} \rightarrow \mathbb{W}, \quad (\Phi, v) \mapsto \Phi(v).$$

If \mathbb{V} and \mathbb{W} are Banach spaces equipped with norms $\|\cdot\|$ and $\|\cdot\|'$ respectively, then $\mathcal{B}(\mathbb{V}, \mathbb{W})$ is a Banach space equipped with the operator norm

$$\|\Phi\|'' = \sup_{v \neq 0} \frac{\|\Phi(v)\|'}{\|v\|}.$$

It is straightforward to see that the spectral norm of \mathbf{M} is then

$$\|\mathbf{M}\|_\sigma = \sup_{\Phi \neq 0, v \neq 0} \frac{\|\mathbf{M}(\Phi, v)\|'}{\|\Phi\|''\|v\|} = 1$$

and thus it is continuous. Next we will look at an actual Banach space of functions. Another quintessential bilinear operator is the convolution of two functions $f, g \in$

$L^1(\mathbb{R}^n)$, defined by

$$f * g(x) := \int_{\mathbb{R}^n} f(x - y)g(y) \, dy.$$

The result is an L^1 -function because $\|f * g\|_1 \leq \|f\|_1 \|g\|_1$. Thus we have a well-defined bilinear operator

$$B_* : L^1(\mathbb{R}^n) \times L^1(\mathbb{R}^n) \rightarrow L^1(\mathbb{R}^n), \quad (f, g) \mapsto f * g.$$

Since $\|B_*(f, g)\|_1 = \|f * g\|_1 \leq \|f\|_1 \|g\|_1$, we have

$$\|B_*\|_\sigma = \sup_{f, g \neq 0} \frac{\|f * g\|_1}{\|f\|_1 \|g\|_1} \leq 1,$$

and equality is attained by, say, choosing $f, g \in L^1(\mathbb{R}^n)$ to be non-negative and applying Fubini. So $\|B_*\|_\sigma = 1$ and in particular B_* is continuous.

For a continuous linear operator $\Phi \in \mathcal{B}(\mathbb{H}; \mathbb{H})$ and a continuous bilinear operator $B \in \mathcal{B}(\mathbb{H}, \mathbb{H}; \mathbb{H})$ on a separable Hilbert space \mathbb{H} with a countable orthonormal basis $\mathcal{B} = \{e_i : i \in \mathbb{N}\}$, Parseval's identity $f = \sum_{i=1}^\infty \langle f, e_i \rangle e_i$ gives us the following orthogonal expansions:

$$\begin{aligned} \Phi(f) &= \sum_{i=1}^\infty \sum_{j=1}^\infty \langle \Phi(e_i), e_j \rangle \langle f, e_i \rangle e_j, \\ B(f, g) &= \sum_{i=1}^\infty \sum_{j=1}^\infty \sum_{k=1}^\infty \langle B(e_i, e_j), e_k \rangle \langle f, e_i \rangle \langle g, e_j \rangle e_k \end{aligned}$$

for any $f, g \in \mathbb{H}$. Convergence of these infinite series is, as usual, in the norm induced by the inner product, and it follows from convergence that the hypermatrices representing Φ and B with respect to \mathcal{B} are l^2 -summable, that is,

$$(\langle \Phi(e_i), e_j \rangle)_{i,j=1}^\infty \in l^2(\mathbb{N} \times \mathbb{N}), \quad (\langle B(e_i, e_j), e_k \rangle)_{i,j,k=1}^\infty \in l^2(\mathbb{N} \times \mathbb{N} \times \mathbb{N})$$

(we discuss infinite-dimensional hypermatrices formally in Example 4.5). In fact the converse is also true. If we define a linear operator and a bilinear operator by

$$\begin{aligned} \Phi(f) &:= \sum_{i=1}^\infty \sum_{j=1}^\infty a_{ij} \langle f, e_j \rangle e_i, \\ B(f, g) &:= \sum_{i=1}^\infty \sum_{j=1}^\infty \sum_{k=1}^\infty b_{ijk} \langle f, e_i \rangle \langle g, e_j \rangle e_k \end{aligned}$$

for any $f, g \in \mathbb{H}$ and where

$$(a_{ij})_{i,j=1}^\infty \in l^2(\mathbb{N} \times \mathbb{N}), \quad (b_{ijk})_{i,j,k=1}^\infty \in l^2(\mathbb{N} \times \mathbb{N} \times \mathbb{N}),$$

then $\Phi \in \mathcal{B}(\mathbb{H}; \mathbb{H})$ and $B \in \mathcal{B}(\mathbb{H}, \mathbb{H}; \mathbb{H})$. So the continuity of these operators is determined by the l^2 -summability of the corresponding hypermatrices. What

is perhaps surprising is that this also holds to various extents for Banach spaces with ‘almost orthogonal bases’ and other growth conditions for the coefficient hypermatrices, as we will see below.

Going beyond elementary examples like M or B_* requires somewhat more background. The space of *Schwartz functions* is

$$S(\mathbb{R}^n) := \left\{ f \in C^\infty(\mathbb{R}^n) : \sup_{x \in \mathbb{R}^n} \left| x_1^{p_1} \cdots x_n^{p_n} \frac{\partial^{q_1 + \cdots + q_n} f}{\partial x_1^{q_1} \cdots \partial x_n^{q_n}} \right| < \infty, p_i, q_j \in \mathbb{N} \right\}.$$

The defining condition essentially states that all its derivatives decay rapidly to zero faster than any negative powers. Examples include compactly supported smooth functions and functions of the form $p(x) \exp(-x^T Ax)$, where p is a polynomial and $A \in \mathbb{S}_{++}^n$. While $S(\mathbb{R}^n)$ is not a Banach space, it is a dense subset of many common Banach spaces such as $L^p(\mathbb{R}^n)$, $p \in [1, \infty)$. Its continuous dual space,

$$S'(\mathbb{R}^n) := \mathcal{B}(S(\mathbb{R}^n); \mathbb{R}),$$

is the space of *tempered distributions*. It follows from Schwartz’s kernel theorem (see Example 4.2) that any continuous bilinear operator $B : S(\mathbb{R}^n) \times S(\mathbb{R}^n) \rightarrow S'(\mathbb{R}^n)$ must take the form

$$B(f, g)(x) = \int_{\mathbb{R}^n} \int_{\mathbb{R}^n} K(x, y, z) f(y) g(z) \, dy \, dz$$

for some tempered distribution $K \in S'(\mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R}^n)$ (Grafakos and Torres 2002a). It is a celebrated result of Frazier and Jawerth (1990) that if $\psi \in S(\mathbb{R}^n)$ has Fourier transform $\widehat{\psi}(\xi)$ vanishing outside the annulus $\pi/4 \leq |\xi| \leq \pi$ and bounded away from zero on a smaller annulus $\pi/4 + \varepsilon \leq |\xi| \leq \pi - \varepsilon$, then it is an *almost orthogonal wavelet* for $L^p(\mathbb{R}^n)$, $p \in (1, \infty)$. This means that if we set

$$\mathcal{B}_\psi := \{\psi_{k,v} : (k, v) \in \mathbb{Z} \times \mathbb{Z}^n\}, \quad \psi_{k,v}(x) := 2^{kn/2} \psi(2^k x - v),$$

then any $f \in L^p(\mathbb{R}^n)$ has an expansion

$$f = \sum_{(k,v) \in \mathbb{Z}^{n+1}} \langle f, \psi_{k,v} \rangle \psi_{k,v}$$

that converges in the L^p -norm, much like Parseval’s identity for a Hilbert space, even though we do not have a Hilbert space and \mathcal{B}_ψ is not an orthonormal basis (in the language of Example 4.47, \mathcal{B}_ψ is a tight wavelet frame with frame constant 1). Furthermore, if the matrix¹²

$$A : \mathbb{Z}^{n+1} \times \mathbb{Z}^{n+1} \rightarrow \mathbb{R}$$

satisfies an ‘almost diagonal’ growth condition that essentially says that $a_{(i,\lambda),(j,\mu)}$

¹² Note that A is a matrix in the sense of Example 4.5. Here the row and column indices take values in \mathbb{Z}^{n+1} . Likewise for the hypermatrix B later.

is small whenever (i, λ) and (j, μ) are far apart, then defining

$$\Phi_A(f) := \sum_{(i,\lambda) \in \mathbb{Z}^{n+1}} \sum_{(j,\mu) \in \mathbb{Z}^{n+1}} a_{(i,\lambda),(j,\mu)} \langle f, \psi_{i,\lambda} \rangle \psi_{j,\mu}$$

gives us a continuous linear operator $\Phi_A: L^p(\mathbb{R}^n) \rightarrow L^p(\mathbb{R}^n)$, $p \in (1, \infty)$. This extends to bilinear operators. If the 3-hypermatrix

$$B: \mathbb{Z}^{n+1} \times \mathbb{Z}^{n+1} \times \mathbb{Z}^{n+1} \rightarrow \mathbb{R}$$

satisfies an analogous ‘almost diagonal’ growth condition with $b_{(i,\lambda),(j,\mu),(k,\nu)}$ small whenever $(i, \lambda), (j, \mu), (k, \nu)$ are far apart from each other, then defining

$$\Phi_B(f, g) := \sum_{(i,\lambda) \in \mathbb{Z}^{n+1}} \sum_{(j,\mu) \in \mathbb{Z}^{n+1}} \sum_{(k,\nu) \in \mathbb{Z}^{n+1}} b_{(i,\lambda),(j,\mu),(k,\nu)} \langle f, \psi_{i,\lambda} \rangle \langle g, \psi_{j,\mu} \rangle \psi_{k,\nu}$$

gives us a continuous linear operator $\Phi_B: L^p(\mathbb{R}^n) \times L^q(\mathbb{R}^n) \rightarrow L^r(\mathbb{R}^n)$ whenever

$$\frac{1}{p} + \frac{1}{q} = \frac{1}{r}, \quad 1 < p, q, r < \infty.$$

The linear result is due to [Frazier and Jawerth \(1990\)](#) but our description is based on [Grafakos and Torres \(2002a, Theorem A\)](#), which also contains the bilinear result ([Grafakos and Torres 2002a, Theorem 1](#)). We refer the reader to these references for the exact statement of the ‘almost diagonal’ growth conditions.

Aside from convolution, the best-known continuous bilinear operator is probably the bilinear Hilbert transform

$$H(f, g)(x) := \lim_{\varepsilon \rightarrow 0} \int_{|y| > \varepsilon} f(x + y)g(x - y) \frac{dy}{y}, \tag{3.32}$$

being a bilinear extension of the Hilbert transform

$$H(f)(x) := \lim_{\varepsilon \rightarrow 0} \int_{|y| > \varepsilon} f(x - y) \frac{dy}{y}.$$

The latter, according to [Krantz \(2009, p. 15\)](#) ‘is, without question, the most important operator in analysis.’ While it is a standard result that the Hilbert transform is continuous as a linear operator $H: L^p(\mathbb{R}) \rightarrow L^p(\mathbb{R})$, $p \in (1, \infty)$, and we even know the exact value of its operator/spectral norm ([Grafakos 2014, Remark 5.1.8](#)),

$$\|H\|_\sigma = \begin{cases} \pi \tan\left(\frac{\pi}{2p}\right) & 1 < p \leq 2, \\ \pi \cot\left(\frac{\pi}{2p}\right) & 2 \leq p < \infty, \end{cases}$$

the continuity of its bilinear counterpart had been a long-standing open problem. It was resolved by [Lacey and Thiele \(1997, 1999\)](#), who showed that as a bilinear operator, $H: L^p(\mathbb{R}) \times L^q(\mathbb{R}) \rightarrow L^r(\mathbb{R})$ is continuous whenever

$$\frac{1}{p} + \frac{1}{q} = \frac{1}{r}, \quad 1 < p, q \leq \infty, \quad \frac{2}{3} < r < \infty,$$

that is, there exists $c > 0$ such that

$$\|\mathbf{H}(f, g)\|_r \leq c \|f\|_p \|g\|_q$$

for all $f \in L^p(\mathbb{R})$ and $g \in L^q(\mathbb{R})$. The special case $p = q = 2$, $r = 1$, open for more than thirty years, was known as the Calderón conjecture.

The study of infinite-dimensional multilinear operators along the above lines has become a vast undertaking, sometimes called multilinear harmonic analysis (Muscalu and Schlag 2013), with a multilinear Calderón–Zygmund theory (Grafakos and Torres 2002b) and profound connections to wavelets (Meyer and Coifman 1997) among its many cornerstones.

3.2. Decompositions of multilinear maps

For simplicity, we will assume a finite-dimensional setting in this section. We will see that every multilinear map can be constructed out of linear functionals and vectors. This is a simple observation but often couched in abstract terms under headings like the ‘universal factorization property’ or ‘universal mapping property’. Its simplicity notwithstanding, this observation is a fruitful one that is the basis of the notion of tensor rank.

In the context of definition ②, 0- and 1-tensors are building blocks, and there is little more that we may say about them. The definition really begins at $d = 2$. Let $\Phi: \mathbb{U} \rightarrow \mathbb{V}$ be a linear operator and let $\mathcal{B} = \{v_1, \dots, v_n\}$ be a basis of \mathbb{V} . Then, for any $u \in \mathbb{U}$, we have

$$\Phi(u) = \sum_{j=1}^n a_j v_j \tag{3.33}$$

for some $a_1, \dots, a_n \in \mathbb{R}$. Clearly, if we change u the coefficients a_1, \dots, a_n must in general change too, that is, they are functions of u and we ought to have written

$$\Phi(u) = \sum_{j=1}^n a_j(u) v_j \tag{3.34}$$

to indicate this. What kind of function is $a_i: \mathbb{U} \rightarrow \mathbb{R}$? It is easy to see if we take the linear functional $v_i^*: \mathbb{V} \rightarrow \mathbb{R}^*$ from the dual basis $\mathcal{B}^* = \{v_1^*, \dots, v_n^*\}$ and hit (3.34) on the left with it; then

$$v_i^*(\Phi(u)) = \sum_{j=1}^n a_j(u) v_i^*(v_j) = a_i(u)$$

by linearity of v_i^* and the fact that $v_i^*(v_j) = \delta_{ij}$. Since this holds for all $u \in \mathbb{U}$, we have

$$v_i^* \circ \Phi = a_i.$$

So $a_i: \mathbb{U} \rightarrow \mathbb{R}$ is a linear functional as v_i^* and Φ are both linear. Switching back

to our usual notation of denoting linear functionals as φ_i instead of a_i , we see that every linear operator $\Phi: \mathbb{U} \rightarrow \mathbb{V}$ takes the form

$$\Phi(u) = \sum_{j=1}^n \varphi_j(u)v_j$$

for some linear functionals $\varphi_1, \dots, \varphi_n \in \mathbb{V}^*$, that is, every linear operator is constructed out of linear functionals and vectors. Now observe that in order to get (3.34) we really did not need v_1, \dots, v_n to be a basis of \mathbb{V} : we just need v_1, \dots, v_r to be a basis of $\text{im}(\Phi)$ with $r = \dim \text{im}(\Phi) = \text{rank}(\Phi)$. Since r cannot be any smaller or else (3.34) would not hold, we must have

$$\text{rank}(\Phi) = \min \left\{ r : \Phi(u) = \sum_{i=1}^r \varphi_i(u)v_i \right\}.$$

The argument is similar for a bilinear functional $\beta: \mathbb{U} \times \mathbb{V} \rightarrow \mathbb{R}$ but also different enough to warrant going over. For any $u \in \mathbb{U}$, $\beta(u, \cdot): \mathbb{V} \rightarrow \mathbb{R}$ is a linear functional, so we must have

$$\beta(u, \cdot) = \sum_{i=1}^n a_i v_i^*$$

as $\mathcal{B}^* = \{v_1^*, \dots, v_n^*\}$ is a basis for \mathbb{V}^* . Since a_i depends on u , we write

$$\beta(u, \cdot) = \sum_{i=1}^n a_i(u)v_i^*.$$

Evaluating at v_j gives

$$\beta(u, v_j) = \sum_{i=1}^n a_i(u)v_i^*(v_j) = a_j(u)$$

for any $u \in \mathbb{U}$. So $a_j: \mathbb{U} \rightarrow \mathbb{R}$ is a linear functional. Since $v_i^*: \mathbb{V} \rightarrow \mathbb{R}$ is also a linear functional, switching to our usual notation for linear functionals, we conclude that a bilinear functional $\beta: \mathbb{U} \times \mathbb{V} \rightarrow \mathbb{R}$ must take the form

$$\beta(u, v) = \sum_{i=1}^n \varphi_i(u)\psi_i(v) \tag{3.35}$$

for some $\varphi_1, \dots, \varphi_n \in \mathbb{U}^*$ and $\psi_1, \dots, \psi_n \in \mathbb{V}^*$, that is, every bilinear functional is constructed out of linear functionals. In addition,

$$\text{rank}(\beta) = \min \left\{ r : \beta(u, v) = \sum_{i=1}^r \varphi_i(u)\psi_i(v) \right\}$$

defines a notation of rank for bilinear functionals.

For a bilinear operator $B: \mathbb{U} \times \mathbb{V} \rightarrow \mathbb{W}$ and any $u \in \mathbb{U}$ and $v \in \mathbb{V}$,

$$B(u, v) = \sum_{j=1}^p a_j w_j$$

with $\mathcal{C} = \{w_1, \dots, w_p\}$ a basis of \mathbb{W} , but the difference now is that a_j depends on both u and v and so is a function $a_j: \mathbb{U} \times \mathbb{V} \rightarrow \mathbb{R}$, that is,

$$B(u, v) = \sum_{j=1}^p a_j(u, v) w_j.$$

Hitting this equation on the left by w_i^* gives

$$w_i^*(B(u, v)) = a_i(u, v),$$

and since composing a bilinear operator and a linear functional $\mathbb{U} \times \mathbb{V} \xrightarrow{B} \mathbb{W} \xrightarrow{w_i^*} \mathbb{R}$ gives a bilinear functional, we conclude that a_i is a bilinear functional. Applying the result in the previous paragraph, a_i must take the form in (3.35), and changing to our notation for linear functionals we get

$$B(u, v) = \sum_{i=1}^n \sum_{j=1}^p \varphi_i(u) \psi_i(v) w_j = \sum_{k=1}^r \varphi_k(u) \psi_k(v) w_k, \quad (3.36)$$

where the last step is simply relabelling the indices, noting that both are sums of terms of the form $\varphi(u)\psi(v)w$, with $r = np$. The smallest r , that is,

$$\text{rank}(B) = \min \left\{ r: B(u, v) = \sum_{i=1}^r \varphi_i(u) \psi_i(v) w_i \right\}, \quad (3.37)$$

is called the *tensor rank* of B , a notion that may be traced to Hitchcock (1927, equations 2 and 2_a) and will play a critical role in the next section.

The same line of argument may be repeated on a trilinear functional $\tau: \mathbb{U} \times \mathbb{V} \times \mathbb{W} \rightarrow \mathbb{R}$ to show that they are just sums of products of linear functionals

$$\tau(u, v, w) = \sum_{i=1}^r \varphi_i(u) \psi_i(v) \theta_i(w) \quad (3.38)$$

and with it a corresponding notion of tensor rank. More generally, any d -linear map $\Phi: \mathbb{V}_1 \times \mathbb{V}_2 \times \dots \times \mathbb{V}_d \rightarrow \mathbb{W}$ is built up of linear functionals and vectors,

$$\Phi(v_1, \dots, v_d) = \sum_{i=1}^r \varphi_i(v_1) \psi_i(v_2) \dots \theta_i(v_d) w_i,$$

where the last w_i may be dropped if it is a d -linear functional with $\mathbb{W} = \mathbb{R}$. Consequently, we see that the ‘multilinearity’ in any multilinear map comes from that of

$$\mathbb{R}^d \ni (x_1, x_2, \dots, x_d) \mapsto x_1 x_2 \dots x_d \in \mathbb{R}.$$

Take (3.38) as illustration: where does the ‘trilinearity’ of τ come from? Say we look at the middle argument; then

$$\begin{aligned}\tau(u, \lambda v + \lambda' v', w) &= \sum_{i=1}^r \varphi_i(u) \psi_i(\lambda v + \lambda' v') \theta_i(w) \\ &= \sum_{i=1}^r \varphi_i(u) [\lambda \psi_i(v) + \lambda' \psi_i(v')] \theta_i(w) \\ &= \lambda \left[\sum_{i=1}^r \varphi_i(u) \psi_i(v) \theta_i(w) \right] + \lambda' \left[\sum_{i=1}^r \varphi_i(u) \psi_i(v') \theta_i(w) \right] \\ &= \lambda \tau(u, v, w) + \lambda' \tau(u, v', w).\end{aligned}$$

The reason why it is linear in the middle argument is simply a result of

$$x(\lambda y + \lambda' y')z = \lambda xyz + \lambda' xy'z,$$

which is in turn a result of the trilinearity of $(x, y, z) \mapsto xyz$. All ‘multilinearity’ in all multilinear maps arises in this manner. In addition, this ‘multilinearity’ may be ‘factored out’ of a multilinear map, and once we do that, whatever is left behind is a linear map. This is essentially the simple idea behind the universal factorization property of tensor products that we will formally state and discuss in Section 4.3.

3.3. Tensors in computations II: multilinearity

A basic utility of the multilinear map perspective of tensors is that it allows one to recognize 3-tensors in situations like the examples (i), (ii), (iii) on page 594, and in turn allows one to apply tensorial notions such as tensor ranks and tensor norms to analyse them. The most common 3-tensors are bilinear operators and this is the case in computations too, although we will see that trilinear functionals do also appear from time to time. When the bilinear operator is matrix multiplication, study of its tensor rank leads us to the fabled *exponent of matrix multiplication* ω . While most are aware that ω is important in the complexity of matrix–matrix products and matrix inversion, it actually goes far beyond. The complexity of computing *LU* decompositions, symmetric eigenvalue decompositions, determinants and characteristic polynomials, bases for null spaces, matrix sparsification – note that none of these are bilinear operators – all have asymptotic complexity either exactly ω or bounded by it (Bürgisser *et al.* 1997, Chapter 16). Determining the exact value of ω truly deserves the status of a Holy Grail problem in numerical linear algebra. Especially relevant to this article is the fact that ω is a tensorial notion.

Given three vector spaces $\mathbb{U}, \mathbb{V}, \mathbb{W}$, how could one construct a bilinear operator

$$\mathbf{B}: \mathbb{U} \times \mathbb{V} \rightarrow \mathbb{W}?$$

Taking a leaf from Section 3.2, a simple and natural way is to take a linear functional $\varphi: \mathbb{U} \rightarrow \mathbb{R}$, a linear functional $\psi: \mathbb{V} \rightarrow \mathbb{R}$, a vector $w \in \mathbb{W}$, and then define

$$B(u, v) = \varphi(u)\psi(v)w \tag{3.39}$$

for any $u \in \mathbb{U}$ and $v \in \mathbb{V}$; it is easy to check that this is bilinear. We will call a non-zero bilinear operator of such a form *rank-one*.

An important point to note is that in evaluating $B(u, v)$ in (3.39), only multiplication of variables matters, and for this $B(u, v)$ requires exactly one multiplication. Take for example $\mathbb{U} = \mathbb{V} = \mathbb{W} = \mathbb{R}^3$ and $\varphi(u) = u_1 + 2u_2 + 3u_3$, $\psi(v) = 2v_1 + 3v_2 + 4v_3$, $w = (3, 4, 5)$; then

$$B(u, v) = \begin{bmatrix} 3(u_1 + 2u_2 + 3u_3)(2v_1 + 3v_2 + 4v_3) \\ 4(u_1 + 2u_2 + 3u_3)(2v_1 + 3v_2 + 4v_3) \\ 5(u_1 + 2u_2 + 3u_3)(2v_1 + 3v_2 + 4v_3) \end{bmatrix}.$$

This appears to require far more than one multiplication, but multiplications such as $2u_2$ or $4v_3$ are all scalar multiplications, that is, one of the factors is a constant, and these are discounted.

This is the notion of *bilinear complexity* introduced in Strassen (1987) and is a very fruitful idea. Instead of trying to design algorithms that minimize all arithmetic operations at once, we focus on arguably the most expensive one, the multiplication of variables, as this is the part that cannot be hardwired or hardcoded. On the other hand, once we have fixed φ, ψ, w , the evaluation of these linear functionals can be implemented as specialized algorithms in hardware for maximum efficiency. One example is the discrete Fourier transform,

$$\begin{bmatrix} x'_0 \\ x'_1 \\ x'_2 \\ x'_3 \\ \vdots \\ x'_{n-1} \end{bmatrix} = \frac{1}{\sqrt{n}} \begin{bmatrix} 1 & 1 & 1 & 1 & \cdots & 1 \\ 1 & \omega & \omega^2 & \omega^3 & \cdots & \omega^{n-1} \\ 1 & \omega^2 & \omega^4 & \omega^6 & \cdots & \omega^{2(n-1)} \\ 1 & \omega^3 & \omega^6 & \omega^9 & \cdots & \omega^{3(n-1)} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \omega^{n-1} & \omega^{2(n-1)} & \omega^{3(n-1)} & \cdots & \omega^{(n-1)(n-1)} \end{bmatrix} \begin{bmatrix} x_0 \\ x_1 \\ x_2 \\ x_3 \\ \vdots \\ x_{n-1} \end{bmatrix}. \tag{3.40}$$

Fast Fourier transform is a specialized algorithm that gives us the value of x'_k by evaluating the linear functional $\varphi(x) = \sum_{j=0}^{n-1} \omega^{jk} x_j$ but not in the obvious way, and it is often built into software libraries or hardware. It is instructive to ask about the bilinear complexity of (3.40): the answer is in fact zero, as it involves only scalar multiplications with constants ω^{jk} .

By focusing only on variable multiplications, we may neatly characterize bilinear complexity in terms of its tensor rank in (3.37); in fact, bilinear complexity is tensor rank. Furthermore, upon finding the algorithms that are optimal in terms of variable multiplications – these are almost never unique – we may then seek among them the ones that are optimal in terms of scalar multiplications (*i.e.* sparsest), additions, numerical stability, communication cost, energy cost, *etc.* In fact, one may often bound the number of additions and scalar multiplications in terms of

the number of variable multiplications. For example, if an algorithm takes n^p variable multiplications, we may often show that it takes at most cn^p additions and scalar multiplications for some constant $c > 0$ as each variable multiplication in the algorithm is accompanied by at most c other additions and scalar multiplications; as a result the algorithm is still $O(n^p)$ even when we count all arithmetic operations. This will be the case for all problems considered in Examples 3.9 and 3.10, that is, the operation counts therein include all arithmetic operations. Henceforth, by ‘multiplication’ in the remainder of this section we will always mean ‘variable multiplication’ unless specified otherwise.

In general, a bilinear operator would not be rank-one as in (3.39). Nevertheless, as we saw in (3.36), it will always be decomposable into a sum of rank-one terms

$$B(u, v) = \sum_{i=1}^r \varphi_i(u) \psi_i(v) w_i \quad (3.41)$$

for some $r \in \mathbb{N}$ with the smallest possible r given by the tensor rank of B . Note that any decomposition of the form in (3.41) gives us an explicit algorithm for computing B with r multiplications, and thus $\text{rank}(B)$ gives us the bilinear complexity or least number of multiplications required to compute B . This relation between tensor rank and evaluation of bilinear operators first appeared in Strassen (1973).

As numerical computations go, there is no need to compute a quantity exactly. What if we just require the right-hand side (remember this gives an algorithm) of (3.41) to be an ε -approximation of the left-hand side? This leads to the notion of *border rank*:

$$\overline{\text{rank}}(B) = \min \left\{ r : B(u, v) = \lim_{\varepsilon \rightarrow 0^+} \sum_{i=1}^r \varphi_i^\varepsilon(u) \psi_i^\varepsilon(v) w_i^\varepsilon \right\}. \quad (3.42)$$

This was first proposed by Bini, Capovani, Romani and Lotti (1979) and Bini, Lotti and Romani (1980) and may be regarded as providing an algorithm (remembering that every such decomposition gives an algorithm)

$$B^\varepsilon(u, v) = \sum_{i=1}^r \varphi_i^\varepsilon(u) \psi_i^\varepsilon(v) w_i^\varepsilon$$

that approximates $B(u, v)$ up to any arbitrary ε -accuracy. While (3.42) relies on a limit, it may also be defined purely algebraically over any ring as in Knuth (1998, p. 522) and Bürgisser *et al.* (1997, Definition 15.19), which are in turn based on Bini (1980). Clearly $\overline{\text{rank}}(B) \leq \text{rank}(B)$, but one may wonder if there are indeed instances where the inequality is strict. There are explicit examples in Bini *et al.* (1979, 1980), Bürgisser *et al.* (1997) and Knuth (1998), but they do not reveal the familiar underlying phenomenon, namely that of a difference quotient converging

to a derivative (De Silva and Lim 2008):

$$\lim_{\varepsilon \rightarrow 0^+} \frac{(\varphi_1(u) + \varepsilon\varphi_2(u))(\psi_1(v) + \varepsilon\psi_2(v))(w_1 + \varepsilon w_2) - \varphi_1(u)\psi_1(v)w_1}{\varepsilon} = \varphi_2(u)\psi_2(v)w_1 + \varphi_1(u)\psi_2(v)w_1 + \varphi_1(u)\psi_1(v)w_2. \tag{3.43}$$

The left-hand side clearly has rank no more than two; one may show that as long as φ_1, φ_2 are not collinear, and likewise for ψ_1, ψ_2 and w_1, w_2 , the right-hand side of (3.43) must have rank three, that is, it defines a bilinear operator with rank three and border rank two.

Tensor rank and border rank are purely algebraic notions defined over any vector spaces, or even modules, which are generalizations of vector spaces that we will soon introduce. However, if $\mathbb{U}, \mathbb{V}, \mathbb{W}$ are norm spaces, we may introduce various notions of norms on bilinear operators $B: \mathbb{U} \times \mathbb{V} \rightarrow \mathbb{W}$. We will slightly abuse notation by denoting the norms on all three spaces by $\|\cdot\|$. Recall that for a linear functional $\psi: \mathbb{V} \rightarrow \mathbb{R}$, its *dual norm* is just

$$\|\psi\|_* := \sup_{v \neq 0} \frac{|\psi(v)|}{\|v\|}. \tag{3.44}$$

A continuous analogue of tensor rank in (3.37) may then be defined by

$$\|B\|_\nu := \inf \left\{ \sum_{i=1}^r \|\varphi_i\|_* \|\psi_i\|_* \|w_i\| : B(u, v) = \sum_{i=1}^r \varphi_i(u)\psi_i(v)w_i \right\} \tag{3.45}$$

and we call this a *tensor nuclear norm*. This defines a norm dual to the spectral norm in (3.19), which in this case becomes

$$\|B\|_\sigma = \sup_{u, v \neq 0} \frac{\|B(u, v)\|}{\|u\| \|v\|}. \tag{3.46}$$

We will argue later that the tensor nuclear norm, in an appropriate sense, quantifies the optimal numerical stability of computing B just as tensor rank quantifies bilinear complexity.

It will be instructive to begin from some low-dimensional examples where $\mathbb{U}, \mathbb{V}, \mathbb{W}$ are of dimensions two and three.

Example 3.8 (Gauss’s algorithm for complex multiplication). As is fairly well known, one may multiply a pair of complex numbers with three instead of the usual four real multiplications:

$$\begin{aligned} (a + bi)(c + di) &= (ac - bd) + i(bc + ad) \\ &= (ac - bd) + i[(a + b)(c + d) - ac - bd]. \end{aligned} \tag{3.47}$$

The latter algorithm is usually attributed to Gauss. Note that while Gauss’s algorithm uses three multiplications, it involves five additions/subtractions where the usual algorithm has only two, but in bilinear complexity we only care about multiplication of variables, which in this case are a, b, c, d . Observe that complex

multiplication $B_{\mathbb{C}}: \mathbb{C} \times \mathbb{C} \rightarrow \mathbb{C}$, $(z, w) \mapsto zw$ is an \mathbb{R} -bilinear map when we identify $\mathbb{C} \cong \mathbb{R}^2$. So we may regard $B_{\mathbb{C}}: \mathbb{R}^2 \times \mathbb{R}^2 \rightarrow \mathbb{R}^2$ and rewrite (3.47) as

$$B_{\mathbb{C}}\left(\begin{bmatrix} a \\ b \end{bmatrix}, \begin{bmatrix} c \\ d \end{bmatrix}\right) = \begin{bmatrix} ac - bd \\ bc + ad \end{bmatrix} = \begin{bmatrix} ac - bd \\ (a+b)(c+d) - ac - bd \end{bmatrix}.$$

The standard basis vectors in \mathbb{R}^2 ,

$$e_1 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \quad e_2 = \begin{bmatrix} 0 \\ 1 \end{bmatrix},$$

correspond to $1, i \in \mathbb{C}$ and they have dual basis $e_1^*, e_2^*: \mathbb{R}^2 \rightarrow \mathbb{R}$ given by

$$e_1^*\left(\begin{bmatrix} a \\ b \end{bmatrix}\right) = a, \quad e_2^*\left(\begin{bmatrix} a \\ b \end{bmatrix}\right) = b.$$

Then the standard algorithm is given by the decomposition

$$B_{\mathbb{C}}(z, w) = [e_1^*(z)e_1^*(w) - e_2^*(z)e_2^*(w)]e_1 + [e_1^*(z)e_2^*(w) + e_2^*(z)e_1^*(w)]e_2, \quad (3.48)$$

whereas Gauss's algorithm is given by the decomposition

$$B_{\mathbb{C}}(z, w) = [(e_1^* + e_2^*)(z)(e_1^* + e_2^*)(w)]e_2 + [e_1^*(z)e_1^*(w)](e_1 - e_2) - [e_2^*(z)e_2^*(w)](e_1 + e_2). \quad (3.49)$$

One may show that $\text{rank}(B_{\mathbb{C}}) = 3 = \overline{\text{rank}}(B_{\mathbb{C}})$, that is, Gauss's algorithm has optimal bilinear complexity whether in the exact or approximate sense. While using Gauss's algorithm for actual multiplication of complex *numbers* is pointless overkill, it is actually useful in practice (Higham 1992) as one may use it for the multiplication of complex *matrices*:

$$(A + iB)(C + iD) = (AC - BD) + i[(A + B)(C + D) - AC - BD] \quad (3.50)$$

for any $A + iB, C + iD \in \mathbb{C}^{n \times n}$ with $A, B, C, D \in \mathbb{R}^{n \times n}$.

For bilinear maps on two-dimensional vector spaces, Example 3.8 is essentially the only example. We might ask for, say, parallel evaluation of the standard inner product and the standard symplectic form on \mathbb{R}^2 , that is,

$$g(x, y) = x_1y_1 + x_2y_2 \quad \text{and} \quad \omega(x, y) = x_1y_2 - x_2y_1. \quad (3.51)$$

An algorithm similar to Gauss's algorithm would yield the result in three multiplications, and it is optimal going by either rank or border rank. For bilinear maps on three-dimensional vector spaces, a natural example is the bilinear operator $B_{\wedge}: \wedge^2(\mathbb{R}^n) \times \mathbb{R}^n \rightarrow \mathbb{R}^n$ given by the 3×3 skew-symmetric matrix–vector product

$$\begin{bmatrix} 0 & a & b \\ -a & 0 & c \\ -b & -c & 0 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} ay + bz \\ -ax + cz \\ -bx - cy \end{bmatrix}.$$

In this case $\text{rank}(B_\wedge) = 5 = \overline{\text{rank}}(B_\wedge)$; see [Ye and Lim \(2018a, Proposition 12\)](#) and [Krishna and Makam \(2018, Theorem 1.3\)](#). For a truly interesting example, one would have to look at bilinear maps on four-dimensional vector spaces and we shall do so next.

Example 3.9 (Strassen's algorithm for matrix multiplication). [Strassen \(1969\)](#) discovered an algorithm for inverting an $n \times n$ matrix in fewer than $5.64n^{\log_2 7}$ arithmetic operations (counting both additions and multiplications). This was a huge surprise at that time as there were results proving that one may not do this with fewer multiplications than the $n^3/3$ required by Gaussian elimination ([Kljuev and Kokovkin-Ščerbak 1965](#)). The issue is that such impossibility results invariably assume that one is limited to *row* and *column* operations. Strassen's algorithm, on the other hand, is based on *block* operations combined with a novel algorithm that multiplies two 2×2 matrices with seven multiplications:

$$\begin{aligned} & \begin{bmatrix} a_1 & a_2 \\ a_3 & a_4 \end{bmatrix} \begin{bmatrix} b_1 & b_2 \\ b_3 & b_4 \end{bmatrix} \\ &= \begin{bmatrix} a_1 b_1 + a_2 b_3 & \beta + \gamma + (a_1 + a_2 - a_3 - a_4) b_4 \\ \alpha + \gamma + a_4(b_2 + b_3 - b_1 - b_4) & \alpha + \beta + \gamma \end{bmatrix} \end{aligned}$$

with

$$\begin{aligned} \alpha &= (a_3 - a_1)(b_2 - b_4), & \beta &= (a_3 + a_4)(b_2 - b_1), \\ \gamma &= a_1 b_1 + (a_3 + a_4 - a_1)(b_1 - b_2 + b_4). \end{aligned}$$

As in the case of Gauss's algorithm, the saving of one multiplication comes at the cost of an increase in the number of additions/subtractions from eight to fifteen. In fact Strassen's original version had eighteen; the version presented here is the well-known but unpublished Winograd variant discussed in [Knuth \(1998, p. 500\)](#) and [Higham \(2002, equation 23.6\)](#). Recursively applying this algorithm to 2×2 block matrices produces an algorithm for multiplying $n \times n$ matrices with $O(n^{\log_2 7})$ multiplications. More generally, the bilinear operator defined by

$$M_{m,n,p} : \mathbb{R}^{m \times n} \times \mathbb{R}^{n \times p} \rightarrow \mathbb{R}^{m \times p}, \quad (A, B) \mapsto AB,$$

is called either the *matrix multiplication tensor* or *Strassen's tensor*. Every decomposition

$$M_{m,n,p}(A, B) = \sum_{i=1}^r \varphi_i(A) \psi_i(B) W_i$$

with linear functionals $\varphi_i : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$, $\psi_i : \mathbb{R}^{n \times p} \rightarrow \mathbb{R}$ and $W_i \in \mathbb{R}^{m \times p}$ gives us an algorithm for multiplying an $m \times n$ matrix by an $n \times p$ matrix that requires just r multiplications, with the smallest possible r given by $\text{rank}(M_{m,n,p})$. The *exponent of matrix multiplication* is defined to be

$$\omega := \inf\{p \in \mathbb{R} : \text{rank}(M_{n,n,n}) = O(n^p)\}.$$

Strassen's work showed that $\omega < \log_2 7 \approx 2.8073549$, and this has been improved over the years to $\omega < 2.3728596$ at the time of writing (Alman and Williams 2021). Recall that any linear functional $\varphi: \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ must take the form $\varphi(A) = \text{tr}(V^T A)$ for some matrix $V \in \mathbb{R}^{m \times n}$, a consequence of the Riesz representation theorem for an inner product space. For concreteness, when $m = n = p = 2$, Winograd's variant of Strassen's algorithm is given by

$$M_{2,2,2}(A, B) = \sum_{i=1}^7 \varphi_i(A) \psi_i(B) W_i,$$

where $\varphi_i(A) = \text{tr}(U_i^T A)$ and $\psi_i(B) = \text{tr}(V_i^T B)$ with

$$\begin{aligned} U_1 &= \begin{bmatrix} -1 & 0 \\ 1 & 1 \end{bmatrix}, & V_1 &= \begin{bmatrix} 1 & -1 \\ 0 & 1 \end{bmatrix}, & W_1 &= \begin{bmatrix} 0 & 1 \\ 1 & 1 \end{bmatrix}, \\ U_2 &= \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, & V_2 &= \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, & W_2 &= \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}, \\ U_3 &= \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, & V_3 &= \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}, & W_3 &= \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \\ U_4 &= \begin{bmatrix} 1 & 0 \\ -1 & 0 \end{bmatrix}, & V_4 &= \begin{bmatrix} 0 & -1 \\ 0 & 1 \end{bmatrix}, & W_4 &= \begin{bmatrix} 0 & 0 \\ 1 & 1 \end{bmatrix}, \\ U_5 &= \begin{bmatrix} 0 & 0 \\ 1 & 1 \end{bmatrix}, & V_5 &= \begin{bmatrix} -1 & 1 \\ 0 & 0 \end{bmatrix}, & W_5 &= \begin{bmatrix} 0 & 1 \\ 0 & 1 \end{bmatrix}, \\ U_6 &= \begin{bmatrix} 1 & 1 \\ -1 & -1 \end{bmatrix}, & V_6 &= \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}, & W_6 &= \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, \\ U_7 &= \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}, & V_7 &= \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}, & W_7 &= \begin{bmatrix} 0 & 0 \\ -1 & 0 \end{bmatrix}. \end{aligned}$$

The name 'exponent of matrix multiplication' is in retrospect a misnomer: ω should rightly be called the exponent of nearly all matrix computations, as we will see next.

Example 3.10 (beyond matrix multiplication). This example is a summary of the highly informative book by Bürgisser *et al.* (1997, Chapter 16), with a few more items drawn from Schönhage (1972/73), but stripped of the technical details to make the message clear, namely that ω pervades numerical linear algebra. Consider the following problems.¹³

- (i) *Inversion.* Given $A \in \text{GL}(n)$, find $A^{-1} \in \text{GL}(n)$.
- (ii) *Determinant.* Given $A \in \text{GL}(n)$, find $\det(A) \in \mathbb{R}$.
- (iii) *Null basis.* Given $A \in \mathbb{R}^{n \times n}$, find a basis $v_1, \dots, v_m \in \mathbb{R}^n$ of $\ker(A)$.

¹³ An upper Hessenberg matrix H is one where $h_{ij} = 0$ whenever $i > j + 1$ and $\text{nnz}(A) := \#\{(i, j): a_{ij} \neq 0\}$ counts the number of non-zero entries.

- (iv) *Linear system.* Given $A \in \text{GL}(n)$ and $b \in \mathbb{R}^n$, find $v \in \mathbb{R}^n$ so that $Av = b$.
- (v) *LU decomposition.* Given $A \in \mathbb{R}^{m \times n}$ of full rank, find permutation P , unit lower triangular $L \in \mathbb{R}^{m \times m}$, upper triangular $U \in \mathbb{R}^{m \times n}$ so that $PA = LU$.
- (vi) *QR decomposition.* Given $A \in \mathbb{R}^{n \times n}$, find orthogonal $Q \in \text{O}(n)$, upper triangular $U \in \mathbb{R}^{n \times n}$ so that $A = QR$.
- (vii) *Eigenvalue decomposition.* Given $A \in \mathbb{S}^n$, find $Q \in \text{O}(n)$ and a diagonal matrix $\Lambda \in \mathbb{R}^{n \times n}$ so that $A = Q\Lambda Q^T$.
- (viii) *Hessenberg decomposition.* Given $A \in \mathbb{R}^{n \times n}$, find $Q \in \text{O}(n)$ and an upper Hessenberg matrix $H \in \mathbb{R}^{n \times n}$ so that $A = QHQ^T$.
- (ix) *Characteristic polynomial.* Given $A \in \mathbb{R}^{n \times n}$, find $(a_0, \dots, a_{n-1}) \in \mathbb{R}^n$ so that $\det(xI - A) = x^n + a_{n-1}x^{n-1} + \dots + a_1x + a_0$.
- (x) *Sparsification.* Given $A \in \mathbb{R}^{n \times n}$ and $c \in [1, \infty)$, find $X, Y \in \text{GL}(n)$ so that $\text{nnz}(XAY^{-1}) \leq cn$.

For each of problems (i)–(x), if there is an algorithm that computes the $n \times n$ matrix product in $O(n^p)$ multiplications, then there is an algorithm that solves that problem in $O(n^p)$ or possibly $O(n^p \log n)$ arithmetic operations, *i.e.* inclusive of additions and scalar multiplications. By the definition of ω in Example 3.9, there is an algorithm for the $n \times n$ matrix product in $O(n^{\omega+\varepsilon})$ multiplications for any $\varepsilon > 0$. So the ‘exponents’ of problems (i)–(x), which may be properly defined (Bürgisser *et al.* 1997, Definition 16.1) even though they are not bilinear operators, are all equal to or bounded by ω . These results are known to be sharp, that is, one cannot do better than ω , except for problems (iv), (vi) and (viii). In particular, it is still an open problem whether one might solve a non-singular linear system with fewer than $O(n^\omega)$ arithmetic operations asymptotically (Strassen 1990).

To get an inkling of why these results hold, a key idea, mentioned in Example 3.9 and well known to practitioners of numerical linear algebra, is to avoid working with rows and columns and work with blocks instead, making recursive use of formulas for block factorization, inversion and determinant similar to

$$\begin{bmatrix} A & B \\ C & D \end{bmatrix} = \begin{bmatrix} I & 0 \\ CA^{-1} & I \end{bmatrix} \begin{bmatrix} A & B \\ 0 & S \end{bmatrix}, \quad \det \begin{bmatrix} A & B \\ C & D \end{bmatrix} = \det(A) \det(S),$$

$$\begin{bmatrix} A & B \\ C & D \end{bmatrix}^{-1} = \begin{bmatrix} A^{-1} + A^{-1}BS^{-1}CA^{-1} & -A^{-1}BS^{-1} \\ -S^{-1}CA^{-1} & S^{-1} \end{bmatrix},$$

where $S = D - CA^{-1}B$ is the Schur complement. For illustration, Strassen’s algorithm for matrix inversion (Strassen 1969) in $O(n^{\omega+\varepsilon})$ complexity computes

$$\begin{bmatrix} A & B \\ C & D \end{bmatrix}^{-1} = \begin{bmatrix} X_1 - X_3X_6X_2 & X_3X_6 \\ X_6X_2 & -X_6 \end{bmatrix}, \quad \begin{array}{ll} X_1 = A^{-1}, & X_2 = CX_1, \\ X_3 = X_1B, & X_4 = CX_3, \\ X_5 = X_4 - D, & X_6 = X_5^{-1}, \end{array}$$

with the inversion in X_1 and X_6 performed recursively using the same algorithm (Higham 2002, p. 449).

We would like to stress that the results in Examples 3.9 and 3.10, while certainly easy to state, took the combined efforts of a multitude of superb scholars over many decades to establish. We refer interested readers to Bürgisser *et al.* (1997, Sections 15.13 and 16.12), Knuth (1998, Section 4.6.4) and Landsberg (2017, Chapters 1–5) for further information on the work involved.

The examples above are primarily about the tensor rank of a bilinear operator over vector spaces. We next consider three variations on this main theme:

- nuclear/spectral norms instead of tensor rank,
- modules instead of vector spaces,
- trilinear functionals instead of bilinear operators.

An algorithm, however fast, is generally not practical if it lacks numerical stability, as the correctness of its output may no longer be guaranteed in finite precision. Fortunately, Strassen's algorithm in Example 3.9 is only slightly less stable than the standard algorithm for matrix multiplication, an unpublished result¹⁴ of Brent (1970) that has been reproduced and extended in Higham (2002, Theorems 23.2 and 23.3). Numerical stability is, however, a complicated issue that depends on many factors and on hardware architecture. Designing numerically stable algorithms is as much an art as it is a science, but the six *Higham guidelines* for numerical stability (Higham 2002, Section 1.18) capture the most salient aspects. The second guideline, to 'minimize the size of intermediate quantities relative to the final solution', together with the fifth are the two most unequivocal ones. We will discuss how such considerations naturally lead us to the definition of tensor nuclear norms in (3.45).

Example 3.11 (nuclear norm and numerical stability). As we have discussed, an algorithm for evaluating a bilinear operator is a decomposition of the form (3.41),

$$B(u, v) = \sum_{i=1}^r \varphi_i(u) \psi_i(v) w_i, \quad (3.52)$$

where r may or may not be minimal. An intermediate quantity has a natural interpretation as a rank-one term $\varphi_i(u) \psi_i(v) w_i$ among the summands and its size also has a natural interpretation as its norm. Note that such a rank-one bilinear operator has nuclear and spectral norms equal to each other and to $\|\varphi_i\|_* \|\psi_i\|_* \|w_i\|$. Hence the sum

$$\sum_{i=1}^r \|\varphi_i\|_* \|\psi_i\|_* \|w_i\|$$

¹⁴ For $C = AB \in \mathbb{R}^{n \times n}$, we have $\|C - \widehat{C}\| \leq c_n \|A\| \|B\| \varepsilon + O(\varepsilon^2)$ with $c_n = n^2$ (standard), $cn^{\log_2(12)}$ (Strassen), $c'n^{\log_2(18)}$ (Winograd) using $\|X\| = \max_{i,j=1,\dots,n} |x_{ij}|$.

measures the total size of the intermediate quantities in the algorithm (3.52) and its minimum value, given by the nuclear norm of B as defined in (3.45), provides a measure of optimal numerical stability in the sense of Higham’s second guideline. This was first discussed in Ye and Lim (2018a, Section 3.2). Using complex multiplication $B_{\mathbb{C}}$ in Example 3.8 for illustration, one may show that the nuclear norm (Friedland and Lim 2018, Lemma 6.1) is given by

$$\|B_{\mathbb{C}}\|_v = 4.$$

The standard algorithm (3.48) attains this minimum value:

$$\begin{aligned} &\|e_1^*\|_* \|e_1^*\|_* \|e_1\| + \|-e_2^*\|_* \|e_2^*\|_* \|e_1\| + \|e_1^*\|_* \|e_2^*\|_* \|e_2\| \\ &+ \|e_2^*\|_* \|e_1^*\|_* \|e_2\| = 4 = \|B_{\mathbb{C}}\|_v \end{aligned}$$

but not Gauss’s algorithm (3.49), which has total size

$$\begin{aligned} &\|e_1^* + e_2^*\|_* \|e_1^* + e_2^*\|_* \|e_2\| + \|e_1^*\|_* \|e_1^*\|_* \|e_1 - e_2\| \\ &+ \|-e_2^*\|_* \|e_2^*\|_* \|e_1 + e_2\| = 2(1 + \sqrt{2}) > \|B_{\mathbb{C}}\|_v. \end{aligned}$$

So Gauss’s algorithm is faster (by bilinear complexity) but less stable (by Higham’s second guideline) than the standard algorithm. One might ask if it is possible to have an algorithm for complex multiplication that is optimal in both measures. It turns out that the decomposition¹⁵

$$\begin{aligned} B_{\mathbb{C}}(z, w) = &\frac{4}{3} \left(\left[\frac{\sqrt{3}}{2} e_1^*(z) + \frac{1}{2} e_2^*(z) \right] \left[\frac{\sqrt{3}}{2} e_1^*(w) + \frac{1}{2} e_2^*(w) \right] \left[\frac{1}{2} e_1 + \frac{\sqrt{3}}{2} e_2 \right] \right. \\ &+ \left[\frac{\sqrt{3}}{2} e_1^*(z) - \frac{1}{2} e_2^*(z) \right] \left[\frac{\sqrt{3}}{2} e_1^*(w) - \frac{1}{2} e_2^*(w) \right] \left[\frac{1}{2} e_1 - \frac{\sqrt{3}}{2} e_2 \right] \\ &\left. - e_2^*(z) e_2^*(w) e_1 \right) \end{aligned} \tag{3.53}$$

gives an algorithm that attains both $\text{rank}(B_{\mathbb{C}})$ and $\|B_{\mathbb{C}}\|_v$. In conventional notation, this algorithm multiplies two complex numbers via

$$\begin{aligned} &(a + bi)(c + di) \\ &= \frac{1}{2} \left[\left(a + \frac{1}{\sqrt{3}} b \right) \left(c + \frac{1}{\sqrt{3}} d \right) + \left(a - \frac{1}{\sqrt{3}} b \right) \left(c - \frac{1}{\sqrt{3}} d \right) - \frac{8}{3} bd \right] \\ &\quad + \frac{i\sqrt{3}}{2} \left[\left(a + \frac{1}{\sqrt{3}} b \right) \left(c + \frac{1}{\sqrt{3}} d \right) - \left(a - \frac{1}{\sqrt{3}} b \right) \left(c - \frac{1}{\sqrt{3}} d \right) \right]. \end{aligned}$$

We stress that the nuclear norm is solely a measure of stability in the sense of Higham’s second guideline. Numerical stability is too complicated an issue to

¹⁵ We have gone to some lengths to avoid the tensor product \otimes in this section, preferring to defer it to Section 4. The decompositions (3.48), (3.49), (3.53) would be considerably neater if expressed in tensor product form, giving another impetus for definition ③.

be adequately captured by a single number. For instance, from the perspective of cancellation errors, our algorithm above also suffers from the same issue pointed out in Higham (2002, Section 23.2.4) for Gauss's algorithm. By choosing $z = w$ and $b = \sqrt{3}/a$, our algorithm computes

$$\frac{1}{2} \left[\left(a + \frac{1}{a} \right)^2 + \left(a - \frac{1}{a} \right)^2 - \frac{8}{a^2} \right] + \frac{i\sqrt{3}}{2} \left[\left(a + \frac{1}{a} \right)^2 - \left(a - \frac{1}{a} \right)^2 \right] =: x + iy.$$

There will be cancellation error in the computed real part \hat{x} when $|a|$ is small and likewise in the computed imaginary part \hat{y} when $|a|$ is large. Nevertheless, as discussed in Higham (2002, Section 23.2.4), the algorithm is still stable in the weaker sense of having acceptably small $|x - \hat{x}|/|z|$ and $|y - \hat{y}|/|z|$ even if $|x - \hat{x}|/|x|$ or $|y - \hat{y}|/|y|$ might be large.

As is the case for Example 3.8, the algorithm for complex multiplication above is useful only when applied to complex matrices. When $A, B, C, D \in \mathbb{R}^{n \times n}$, the algorithms in Examples 3.8 and 3.11 provide substantial savings when used to multiply $A + iB, C + iD \in \mathbb{C}^{n \times n}$. This gives a good reason for extending multilinear maps and tensors to *modules* (Lang 2002), *i.e.* vector spaces whose field of scalars is replaced by other more general rings. Formally, if R is a ring with multiplicative identity 1 (which we will assume henceforth), an R -module \mathbb{V} is a commutative group under a group operation denoted $+$ and has a scalar multiplication operation denoted \cdot such that, for all $a, b \in R$ and $v, w \in \mathbb{V}$,

- (i) $a \cdot (v + w) = a \cdot v + a \cdot w$,
- (ii) $(a + b) \cdot v = a \cdot v + b \cdot v$,
- (iii) $(ab) \cdot v = a \cdot (b \cdot v)$,
- (iv) $1 \cdot v = v$.

Clearly, when $R = \mathbb{R}$ or \mathbb{C} or any other field, a module just becomes a vector space. What is useful about the notion is that it allows us to include rings of objects we would not normally consider as 'scalars'. For example, in (3.47) we regard \mathbb{C} as a two-dimensional vector space over \mathbb{R} , but in (3.50) we regard $\mathbb{C}^{n \times n}$ as a two-dimensional¹⁶ module over $\mathbb{R}^{n \times n}$. So in the latter the 'scalars' are actually matrices, *i.e.* $R = \mathbb{R}^{n \times n}$. When we consider block matrix operations on square matrices such as those on page 624, we are implicitly doing linear algebra over the ring $R = \mathbb{R}^{n \times n}$, which is not even commutative.

Many standard notions in linear and multilinear algebra carry through from vector spaces to modules with little or no change. For example, the multilinear maps and multilinear functionals of Definitions 3.1 and 3.3 apply verbatim to modules with the field of scalars \mathbb{R} replaced by any ring R . In other words, the notion of tensors in the sense of definition ② applies equally well over modules.

¹⁶ Strictly speaking, the terminology over modules should be *length* instead of dimension.

We will discuss three examples of multilinear maps over modules: tensor fields, fast integer multiplication algorithms and cryptographic multilinear maps.

When people speak of tensors in physics or geometry, they often really mean a tensor field. As a result one may find statements of the tensor transformation rules that bear little resemblance to our version. The next two examples are intended to describe the analogues of definitions ① and ② for a tensor field and show how they fit into the narrative of this article. Also, outside pure mathematics, defining a tensor field is by far the biggest reason for considering multilinear maps over modules.

Example 3.12 (modules and tensor fields). A *tensor field* is – roughly speaking – a tensor-valued function on a manifold. Let M be a smooth manifold and $C^\infty(M) := \{f: M \rightarrow \mathbb{R}: f \text{ is smooth}\}$. Note that $C^\infty(M)$ is a ring, as products and linear combinations of smooth real-valued functions are smooth (henceforth we drop the word ‘smooth’ as all functions and fields in this example are assumed smooth). Thus $R = C^\infty(M)$ may play the role of the ring of scalars. A 0-tensor field is just a function in $C^\infty(M)$. A contravariant 1-tensor field is a vector field, *i.e.* a function f whose value at $x \in M$ is a vector in the tangent space at x , denoted $\mathbb{T}_x(M)$. A covariant 1-tensor field is a covector field, *i.e.* a function f whose value at $x \in M$ is a vector in the cotangent space at x , denoted $\mathbb{T}_x^*(M)$. Here $\mathbb{T}_x(M)$ is a vector space and $\mathbb{T}_x^*(M)$ is its dual vector space, in which case a d -tensor field of contravariant order p and covariant order $d - p$ is simply a function φ whose value at $x \in M$ is such a d -tensor, that is, by Definition 3.3, a multilinear functional

$$\varphi(x): \underbrace{\mathbb{T}_x^*(M) \times \cdots \times \mathbb{T}_x^*(M)}_{p \text{ copies}} \times \underbrace{\mathbb{T}_x(M) \times \cdots \times \mathbb{T}_x(M)}_{d-p \text{ copies}} \rightarrow \mathbb{R}.$$

This seems pretty straightforward; the catch is that φ is not a function in the usual sense, which has a fixed domain and a fixed codomain, but the codomain of φ depends on the point x where it is evaluated. So we have only defined a tensor field at a point $x \in M$ but we still need a way to ‘glue’ all these pointwise definitions together. The customary way to do this is via coordinate charts and transition maps, but an alternative is to simply define the *tangent bundle*

$$\mathbb{T}(M) := \{(x, v): x \in M, v \in \mathbb{T}_x(M)\}$$

and *cotangent bundle*

$$\mathbb{T}^*(M) := \{(x, \varphi): x \in M, \varphi \in \mathbb{T}_x^*(M)\}$$

and observe that these are $C^\infty(M)$ -modules with scalar product given by pointwise multiplication of real-valued functions with vector/covector fields. A d -tensor field of contravariant order p and covariant order $d - p$ is then defined to be the multilinear functional

$$\varphi: \underbrace{\mathbb{T}^*(M) \times \cdots \times \mathbb{T}^*(M)}_{p \text{ copies}} \times \underbrace{\mathbb{T}(M) \times \cdots \times \mathbb{T}(M)}_{d-p \text{ copies}} \rightarrow C^\infty(M).$$

Note that this is a multilinear functional in the sense of modules: the ‘scalars’ are drawn from a ring $R = C^\infty(M)$ and the ‘vector spaces’ are the R -modules $\mathbb{T}(M)$ and $\mathbb{T}^*(M)$. This is another upside of defining tensors via definition ②; it may be easily extended to include tensor fields.

What about definition ①? The first thing to note is that not every result in linear algebra over vector spaces carries over to modules. An example is the notion of basis. While some modules do have a basis – for example, when we speak of $\mathbb{C}^{n \times n}$ as a two-dimensional $\mathbb{R}^{n \times n}$ -module, it is with the basis $\{1, i\}$ in mind – others such as the $C^\infty(M)$ -module of d -tensor fields may not have a basis when $d \geq 1$. An explicit and well-known example is given by vector fields on the 2-sphere S^2 , which as a $C^\infty(S^2)$ -module does not have a basis because of the hairy ball theorem (Passman 1991, Proposition 17.7). Consequently, given a d -tensor field on M , the idea that one may just choose bases and represent it as d -dimensional hypermatrix with entries in $C^\infty(M)$ is false even when $d = 1$. There is, however, one important exception, namely when $M = \mathbb{R}^n$, in which case, given a d -tensor field φ , it is possible to write down a d -hypermatrix

$$A(v) = [\varphi_{ij\dots k}(v)]_{i,j,\dots,k=1}^n \in R^{n \times n \times \dots \times n}$$

with entries in $R = C^\infty(\mathbb{R}^n)$ that holds for all $v \in \mathbb{R}^n$. The analogue of the transformation rules is a bit more complicated since we are now allowed a *non-linear change of coordinates* $v' = F(v)$ as opposed to merely a linear change of basis as in Section 2.2. Here the change-of-coordinates function $F: N(v) \rightarrow \mathbb{R}^n$ is any smooth function defined on a neighbourhood $N(v) \subseteq \mathbb{R}^n$ of v that is *locally invertible*, that is, the derivative $DF(v)$ as defined in Example 3.2 is an invertible linear map in a neighbourhood of v . This is sometimes called a *curvilinear* change of coordinates. The analogue of tensor transformation rule (2.12) for a d -tensor field on \mathbb{R}^n is then

$$A(v') = (DF(v)^{-1}, \dots, DF(v)^{-1}, DF(v)^\top, \dots, DF(v)^\top) \cdot A(v). \tag{3.54}$$

If we have $v' = Xv$ for some $X \in GL(n)$, then $DF(v) = X$ and we recover (2.12). More generally, if we write $v = (v_1, \dots, v_n)$ and $v' = (v'_1, \dots, v'_n)$, then $DF(v)$ and $DF(v)^{-1}$ may be written down as Jacobian matrices,

$$DF(v) = \begin{bmatrix} \frac{\partial v'_1}{\partial v_1} & \dots & \frac{\partial v'_1}{\partial v_n} \\ \vdots & \ddots & \vdots \\ \frac{\partial v'_n}{\partial v_1} & \dots & \frac{\partial v'_n}{\partial v_n} \end{bmatrix}, \quad DF(v)^{-1} = \begin{bmatrix} \frac{\partial v_1}{\partial v'_1} & \dots & \frac{\partial v_1}{\partial v'_n} \\ \vdots & \ddots & \vdots \\ \frac{\partial v_n}{\partial v'_1} & \dots & \frac{\partial v_n}{\partial v'_n} \end{bmatrix},$$

where the latter is a consequence of the inverse function theorem: $DF(v)^{-1} = DF^{-1}(v')$ for $v' = F(v)$.

Take, for instance, Cartesian coordinates $v = (x, y, z)$ and spherical coordinates $v' = (r, \theta, \varphi)$ on \mathbb{R}^3 . Then

$$\begin{aligned} r &= \sqrt{x^2 + y^2 + z^2}, & x &= r \sin \theta \cos \varphi, \\ \theta &= \arctan(\sqrt{x^2 + y^2}/z), & y &= r \sin \theta \sin \varphi, \\ \varphi &= \arctan(y/x), & z &= r \cos \theta, \end{aligned} \quad (3.55)$$

and

$$DF(v) = \frac{\partial(x, y, z)}{\partial(r, \theta, \varphi)} = \begin{bmatrix} \frac{x}{r} & \frac{y}{r} & \frac{z}{r} \\ \frac{xz}{r^2\sqrt{x^2 + y^2}} & \frac{yz}{r^2\sqrt{x^2 + y^2}} & \frac{-(x^2 + y^2)}{r^2\sqrt{x^2 + y^2}} \\ \frac{-y}{x^2 + y^2} & \frac{x}{x^2 + y^2} & 0 \end{bmatrix},$$

$$DF(v)^{-1} = \frac{\partial(r, \theta, \varphi)}{\partial(x, y, z)} = \begin{bmatrix} \sin \theta \cos \varphi & r \cos \theta \cos \varphi & -r \sin \theta \sin \varphi \\ \sin \theta \sin \varphi & r \cos \theta \sin \varphi & r \sin \theta \cos \varphi \\ \cos \theta & -r \sin \theta & 0 \end{bmatrix}.$$

Note that either of the last two matrices may be expressed solely in terms of x, y, z or solely in terms of r, θ, φ ; the forms above are chosen for convenience. The transformation rule (3.54) then allows us to transform a hypermatrix in x, y, z to one in r, θ, φ or *vice versa*.

Example 3.13 (tensor fields over manifolds). Since a smooth manifold M locally resembles \mathbb{R}^n , we may take (3.54) as the *local* transformation rule for a tensor field φ . It describes how, for a point $x \in M$, a hypermatrix $A(v(x))$ representing φ in one system of local coordinates $v: N(x) \rightarrow \mathbb{R}^n$ on a neighbourhood $N(x) \subseteq M$ is related to another hypermatrix $A(v'(x))$ representing φ in another system of local coordinates $v': N'(x) \rightarrow \mathbb{R}^n$. As is the case for the tensor transformation rule (2.12), one may use the tensor field transformation rule (3.54) to ascertain whether or not a physical or mathematical object represented locally by a hypermatrix with entries in $C^\infty(M)$ is a tensor field; this is how one would usually show that connection forms or Christoffel symbols are not tensor fields (Simmonds 1994, pp. 64–65).

In geometry, tensor fields play a central role, as geometric structures are nearly always tensor fields (exceptions such as Finsler metrics tend to be less studied). The most common ones include Riemannian metrics, which are symmetric 2-tensor fields $g: \mathbb{T}(M) \times \mathbb{T}(M) \rightarrow C^\infty(M)$, and symplectic forms, which are alternating 2-tensor fields $\omega: \mathbb{T}(M) \times \mathbb{T}(M) \rightarrow C^\infty(M)$, toy examples of which we have seen in (3.51). As we saw in the previous example, the change-of-coordinates maps for tensors are invertible linear maps, but for tensor fields they are locally invertible linear maps; these are called *diffeomorphisms*, *Riemannian isometries*,

symplectomorphisms, etc., depending on what geometric structures they preserve. The analogues of the finite-dimensional matrix groups in (2.14) then become infinite-dimensional Lie groups such as $\text{Diff}(M)$, $\text{Iso}(M, g)$, $\text{Symp}(M, \omega)$, *etc.*

Although beyond the scope of our article, there are two tensor fields that are too important to be left completely unmentioned: the Ricci curvature tensor and the Riemann curvature tensor. Without introducing additional materials, a straightforward way to define them is to use a special system of local coordinates called *Riemannian normal coordinates* (Chern, Chen and Lam 1999, Chapter 5). For a point $x \in M$, an n -dimensional Riemannian manifold, we choose normal coordinates $(x_1, \dots, x_n): N(x) \rightarrow \mathbb{R}^n$ in a neighbourhood $N(x) \subseteq M$ with x at the origin 0. Then the Riemannian metric g takes values in \mathbb{S}_{++}^n and we have

$$r_{ijkl}(x) = \frac{3}{2} \frac{\partial^2 g_{ij}}{\partial x_k \partial x_l}(0), \quad \bar{r}_{ij}(x) = \frac{3}{\sqrt{\det(g)}} \frac{\partial^2 \sqrt{\det(g)}}{\partial x_k \partial x_l}(0).$$

The Riemann and Ricci curvature tensors at x , respectively, are the quadrilinear and bilinear functionals

$$R(x): \mathbb{T}_x(M) \times \mathbb{T}_x(M) \times \mathbb{T}_x(M) \times \mathbb{T}_x(M) \rightarrow \mathbb{R}, \quad \bar{R}(x): \mathbb{T}_x(M) \times \mathbb{T}_x(M) \rightarrow \mathbb{R},$$

given by

$$R(x)(u, v, w, z) = \sum_{i,j,k,l=1}^n r_{ijkl}(x) u_i v_j w_k z_l, \quad \bar{R}(x)(u, v) = \sum_{i=1}^n \sum_{j=1}^n \bar{r}_{ij}(x) u_i v_j.$$

These are tensor fields – note that the coefficients depend on x – so $R(x)$ will generally be a different multilinear functional at different points $x \in M$. The Ricci curvature tensor will make a brief appearance in Example 4.33 on separation of variables for PDEs. Riemann curvature, being a 4-tensor, is difficult to handle as is, but when M is embedded in Euclidean space, it appears implicitly in the form of a 2-tensor called the Weingarten map (Bürgisser and Cucker 2013, Chapter 21) or second fundamental form (Niyogi, Smale and Weinberger 2008), whose eigenvalues, called *principal curvatures*, give us condition numbers. There are other higher-order tensor fields in geometry (Dodson and Poston 1991) such as the torsion tensor, the Nijenhuis tensor ($d = 3$) and the Weyl curvature tensor ($d = 4$), all of which are unfortunately beyond our scope.

In physics, it is probably fair to say that (i) most physical fields are tensor fields and (ii) most tensors are tensor fields. For (i), while there are important exceptions such as spinor fields, the most common fields such as temperature, pressure and Higgs fields are scalar fields; electric, magnetic and flow velocity fields are vector fields; the Cauchy stress tensor, Einstein tensor and Faraday electromagnetic tensor are 2-tensor fields; higher-order tensor fields are rarer in physics but there are also examples such as the Cotton tensor (García, Hehl, Heinicke and Macías 2004) and

the Lanczos tensor (Roberts 1995), both 3-tensor fields. The last five named tensors also make the case for (ii): a ‘tensor’ in physics almost always means a tensor field of order two or (more rarely) higher. We will describe the Cauchy stress tensor and mention a few higher-order tensors related to it in Example 4.8.

The above examples are analytic in nature but the next two will be algebraic. They show why it is useful to consider bilinear and more generally multilinear maps for modules over $\mathbb{Z}/m\mathbb{Z}$, the ring of integers modulo m .

Example 3.14 (modules and integer multiplication). Trivially, integer multiplication $B: \mathbb{Z} \times \mathbb{Z} \rightarrow \mathbb{Z}$, $(a, b) \mapsto ab$ is a bilinear map over the \mathbb{Z} -module \mathbb{Z} , but this is not the relevant module structure that one exploits in fast integer multiplication algorithms. Instead they are based primarily on two key ideas. The first idea is that integers (assumed unsigned) may be represented as polynomials,

$$a = \sum_{i=0}^{p-1} a_i \theta^i =: a(\theta) \quad \text{and} \quad b = \sum_{j=0}^{p-1} b_j \theta^j =: b(\theta)$$

for some number base θ , and the product has coefficients given by convolutions,

$$ab = \sum_{k=0}^{2p-2} c_k \theta^k =: c(\theta) \quad \text{with} \quad c_k = \sum_{i=0}^k a_i b_{k-i}.$$

Let $n = 2p - 1$ and pad the vectors of coefficients with enough zeros so that we may consider (a_0, \dots, a_{n-1}) , (b_0, \dots, b_{n-1}) , (c_0, \dots, c_{n-1}) on an equal footing. The second idea is to use the discrete Fourier transform¹⁷ (DFT) for some root of unity ω to perform the convolution,

$$\begin{bmatrix} a'_0 \\ a'_1 \\ a'_2 \\ \vdots \\ a'_{n-1} \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & \omega & \omega^2 & \cdots & \omega^{n-1} \\ 1 & \omega^2 & \omega^4 & \cdots & \omega^{2(n-1)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \omega^{n-1} & \omega^{2(n-1)} & \cdots & \omega^{(n-1)(n-1)} \end{bmatrix} \begin{bmatrix} a_0 \\ a_1 \\ a_2 \\ \vdots \\ a_{n-1} \end{bmatrix},$$

$$\begin{bmatrix} b'_0 \\ b'_1 \\ b'_2 \\ \vdots \\ b'_{n-1} \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & \omega & \omega^2 & \cdots & \omega^{n-1} \\ 1 & \omega^2 & \omega^4 & \cdots & \omega^{2(n-1)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \omega^{n-1} & \omega^{2(n-1)} & \cdots & \omega^{(n-1)(n-1)} \end{bmatrix} \begin{bmatrix} b_0 \\ b_1 \\ b_2 \\ \vdots \\ b_{n-1} \end{bmatrix},$$

¹⁷ A slight departure from (3.40) is that we dropped the $1/\sqrt{n}$ coefficient from our DFT and instead put a $1/n$ with our inverse DFT to avoid surds.

$$\begin{bmatrix} c_0 \\ c_1 \\ c_2 \\ \vdots \\ c_{n-1} \end{bmatrix} = \frac{1}{n} \begin{bmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & \omega^{-1} & \omega^{-2} & \cdots & \omega^{-(n-1)} \\ 1 & \omega^{-2} & \omega^{-4} & \cdots & \omega^{-2(n-1)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \omega^{-(n-1)} & \omega^{-2(n-1)} & \cdots & \omega^{-(n-1)(n-1)} \end{bmatrix} \begin{bmatrix} a'_0 b'_0 \\ a'_1 b'_1 \\ a'_2 b'_2 \\ \vdots \\ a'_{n-1} b'_{n-1} \end{bmatrix},$$

taking advantage of the following well-known property: a Fourier transform \mathcal{F} turns convolution product $*$ into pointwise product \cdot and the inverse Fourier transform turns it back, that is,

$$a * b = \mathcal{F}^{-1}(\mathcal{F}(a) \cdot \mathcal{F}(b)).$$

Practical considerations inform the way we choose θ and ω . We choose $\theta = 2^s$ so that $a_i, b_i \in \{0, 1, 2, \dots, 2^s - 1\}$ are s -bit binary numbers that can be readily handled. We choose ω to be a $2p$ th root of unity in a ring $\mathbb{Z}/m\mathbb{Z}$ with $m > c_k$ for all $k = 0, 1, \dots, n-1$. The $2p$ th root of unity ensures that the powers $1, \omega, \dots, \omega^{n-1}$ are distinct and $m > c_k$ prevents any ‘carrying’ when computing c_k . It turns out that such a choice is always possible with, say, $m = 2^{3d} + 1$. Practical considerations aside, the key ideas are to convert integer multiplication to a bilinear operator,

$$\begin{aligned} B_1: (\mathbb{Z}/2^s\mathbb{Z})[\theta] \times (\mathbb{Z}/2^s\mathbb{Z})[\theta] &\rightarrow (\mathbb{Z}/2^s\mathbb{Z})[\theta], \\ (a(\theta), b(\theta)) &\mapsto a(\theta)b(\theta), \end{aligned}$$

followed by a Fourier conversion into another bilinear operator,

$$\begin{aligned} B_2: (\mathbb{Z}/m\mathbb{Z})^n \times (\mathbb{Z}/m\mathbb{Z})^n &\rightarrow (\mathbb{Z}/m\mathbb{Z})^n, \\ ((a_0, \dots, a_{n-1}), (b_0, \dots, b_{n-1})) &\mapsto (a_0 b_0, \dots, a_{n-1} b_{n-1}). \end{aligned}$$

In the former, the univariate polynomial ring $(\mathbb{Z}/2^s\mathbb{Z})[\theta]$ is a $\mathbb{Z}/2^s\mathbb{Z}$ -module with θ regarded as the indeterminate of the polynomials. In the latter $(\mathbb{Z}/m\mathbb{Z})^n$ is a $\mathbb{Z}/m\mathbb{Z}$ -module, in fact it is the direct sum of n copies of $\mathbb{Z}/m\mathbb{Z}$. The algorithms of Karatsuba and Ofman (1962), Cook and Aanderaa (1969), Toom (1963), Schönhage and Strassen (1971) and Fürer (2009) are all variations of one or both of these two ideas: incorporating a divide-and-conquer strategy, computing the DFT with fast Fourier transform, replacing the DFT over $\mathbb{Z}/m\mathbb{Z}$ with one over \mathbb{C} , etc. The recent breakthrough of Harvey and van der Hoeven (2021) that led to an $O(n \log n)$ algorithm for n -bit integer multiplications was achieved with a multidimensional variation: (i) replacing the univariate polynomials $(\mathbb{Z}/2^s\mathbb{Z})[\theta]$ with an R -module of multivariate polynomials $R[\theta_1, \dots, \theta_d]$ with a slightly more complicated coefficient ring R , and (ii) replacing the one-dimensional DFT with a d -dimensional DFT,

$$a'(\phi_1, \phi_2, \dots, \phi_d) = \sum_{\theta_1=0}^{n_1} \cdots \sum_{\theta_d=0}^{n_d} \omega_1^{\phi_1 \theta_1} \omega_2^{\phi_2 \theta_2} \cdots \omega_d^{\phi_d \theta_d} a(\theta_1, \theta_2, \dots, \theta_d),$$

and we will discuss some tensorial features of such multidimensional transforms in Example 4.46. Their resulting algorithm is not only the fastest ever but the

fastest possible. As in the case of matrix multiplication, the implications of an $O(n \log n)$ integer multiplication algorithm stretch far beyond integer arithmetic. Among other things, we may use the algorithm of [Harvey and van der Hoeven \(2021\)](#) to compute to n -bit precision x/y or $\sqrt[n]{x}$ in time $O(n \log n)$ and e^x or π in time $O(n \log^2 n)$, for real inputs x, y ([Brent and Zimmermann 2011](#)).

While the previous example exploits the computational efficiency of multilinear maps over modules, the next one exploits their computational intractability.

Example 3.15 (cryptographic multilinear maps). Suppose Alice and Bob want to generate a (secure) common password they may both use over the (insecure) internet. One way they may do so is to pick a large prime number p and pick a primitive root of unity $g \in (\mathbb{Z}/p\mathbb{Z})^\times$, that is, g generates the *multiplicative* group of integers modulo p in the sense that every non-zero $x \in \mathbb{Z}/p\mathbb{Z}$ may be expressed as $x = g^a$ (group theory notation) or $x \equiv g^a \pmod{p}$ (number theory notation) for some $a \in \mathbb{Z}$. Alice will pick a secret $a \in \mathbb{Z}$ and send g^a publicly to Bob; Bob will pick a secret $b \in \mathbb{Z}$ and send g^b publicly to Alice. Alice, knowing the value of a , may compute $g^{ab} = (g^b)^a$ from the g^b she received from Bob, and Bob, knowing the value of b , may compute $g^{ab} = (g^a)^b$ from the g^a he received from Alice. They now share the secure password g^{ab} . This is the renowned Diffie–Hellman key exchange. The security of the version described is based on the intractability of the discrete log problem: determining the value $a = \log_g(g^a)$ from g^a and p is believed to be intractable. Although the problem has a well-known polynomial-time quantum algorithm ([Shor 1994](#)) and has recently been shown to be quasi-polynomial-time in n ([Kleinjung and Wesolowski 2021](#)) for a finite field \mathbb{F}_{p^n} when p is fixed (note that for $n = 1$, $\mathbb{F}_p = \mathbb{Z}/p\mathbb{Z}$) the technology required for the former is still in its infancy, whereas the latter does not apply in our case where complexity is measured in terms of p and not n (for us, $n = 1$ always but $p \rightarrow \infty$).

Now observe that $(\mathbb{Z}/p\mathbb{Z})^\times$ is a commutative group under the group operation of multiplication modulo p , and it is a \mathbb{Z} -module as we may check that it satisfies the properties on page 627: for any $x, y \in \mathbb{Z}/p\mathbb{Z}$ and $a, b \in \mathbb{Z}$,

- (i) $(xy)^a = x^a y^a$,
- (ii) $x^{(a+b)} = x^a x^b$,
- (iii) $x^{(ab)} = (x^b)^a$,
- (iv) $x^1 = x$.

Furthermore, the Diffie–Hellman key exchange is a \mathbb{Z} -bilinear map

$$\mathbf{B}: (\mathbb{Z}/p\mathbb{Z})^\times \times (\mathbb{Z}/p\mathbb{Z})^\times \rightarrow (\mathbb{Z}/p\mathbb{Z})^\times, \quad (g^a, g^b) \mapsto g^{ab}$$

since, for any $\lambda, \lambda' \in \mathbb{Z}$ and $g^a, g^b \in (\mathbb{Z}/p\mathbb{Z})^\times$,

$$\mathbf{B}(g^{\lambda a + \lambda' a'}, g^b) = g^{(\lambda a + \lambda' a')b} = (g^{ab})^\lambda (g^{a'b})^{\lambda'} = \mathbf{B}(g^a, g^b)^\lambda \mathbf{B}(g^{a'}, g^b)^{\lambda'}$$

and likewise for the second argument. That the notation is written multiplicatively

with coefficients appearing in the power is immaterial; if anything it illustrates the power of abstract algebra in recognizing common structures across different scenarios. While one may express everything in additive notation by taking discrete logs whenever necessary, the notation g^a serves as a useful mnemonic: anything appearing in the power is hard to extract, while using additive notation means having to constantly remind ourselves that extracting a from ag and λa is intractable for the former and trivial for the latter.

What if $d + 1$ different parties need to establish a common password (say, in a 1000-participant Zoom session)? In principle one may successively apply the two-party Diffie–Hellman key exchange $d + 1$ times with the $d + 1$ parties each doing $d + 1$ exponentiations, which is obviously expensive. One may reduce the number of exponentiations to $\log_2(d + 1)$ with a more sophisticated protocol, but it was discovered in Joux (2004) that with some mild assumptions a bilinear map like the one above already allows us to generate a tripartite password in just one round since

$$B(g^a, g^b)^c = B(g^a, g^c)^b = B(g^b, g^c)^a = B(g, g)^{abc}.$$

Boneh and Silverberg (2003) generalized and formalized this idea as a *cryptographic multilinear map* to allow a one-round $(d + 1)$ -partite password. Let p be a prime and let G_1, \dots, G_d, G be cyclic groups of p elements, written multiplicatively.¹⁸ Then, as in the $d = 2$ case, these groups are \mathbb{Z} -modules and we may consider a d -linear map over \mathbb{Z} -modules:

$$\Phi: G_1 \times \dots \times G_d \rightarrow G.$$

The only property that matters is the following consequence of multilinearity:

$$\Phi(g_1^{a_1}, g_2^{a_2}, \dots, g_d^{a_d}) = \Phi(g_1, g_2, \dots, g_d)^{a_1 a_2 \dots a_d}$$

for any $g_1 \in G_1, \dots, g_d \in G_d$ and $a_1, \dots, a_d \in \mathbb{Z}$. Slightly abusing notation, we will write 1 for the identity element in all groups. To exclude trivialities, we assume that Φ is non-degenerate, that is, if $\Phi(g_1, \dots, g_d) = 1$, then we must have $g_i = 1$ for some i (the converse is always true). For a non-degenerate multilinear map Φ to be *cryptographic*, it needs to have two additional properties:

- (i) the discrete log problem in each of G_1, \dots, G_d is intractable,
- (ii) $\Phi(g_1, \dots, g_d)$ is efficiently computable for any $g_1 \in G_1, \dots, g_d \in G_d$.

Note that these assumptions imply that the discrete log problem in G must also be intractable. Given, say, g_1 and $g_1^{a_1}$, since $\Phi(g_1, 1, \dots, 1)$ and $\Phi(g_1^{a_1}, 1, \dots, 1)$

¹⁸ Of course they are all isomorphic to each other, but in actual cryptographic applications, it matters how the groups are realized: one might be an elliptic curve over a finite field and another a cyclic subgroup of the Jacobian of a hyperelliptic curve. Also, an explicit isomorphism may not be easy to identify or compute in practice.

may be efficiently computed, if we could efficiently solve the discrete log problem $\Phi(g_1^{a_1}, 1, \dots, 1) = \Phi(g_1, 1, \dots, 1)^{a_1}$ in G to get a_1 , we would have solved the discrete log problem in G_1 efficiently.

Suppose we have a cryptographic d -linear map, $\Phi: G \times \dots \times G \rightarrow G$, where we have assumed $G_1 = \dots = G_d = G$ for simplicity. For $d + 1$ parties to generate a common password, the i th party just needs to pick a password a_i , perform a single exponentiation to get g^{a_i} , and broadcast g^{a_i} to the other parties, who will each do likewise so that every party now has $g^{a_1}, \dots, g^{a_{d+1}}$. With these and the password a_i , the i th party will now compute

$$\Phi(g^{a_1}, \dots, g^{a_{i-1}}, g^{a_{i+1}}, \dots, g^{a_{d+1}})^{a_i} = \Phi(g, \dots, g)^{a_1 \dots a_{d+1}},$$

which will serve as the common password for the $d + 1$ parties. Note that each of them would have arrived at $\Phi(g, \dots, g)^{a_1 \dots a_{d+1}}$ in a different way with their own password but their results are guaranteed to be equal as a consequence of multilinearity. There are several candidates for such cryptographic multilinear maps (Gentry, Gorbunov and Halevi 2015, Garg, Gentry and Halevi 2013) and a variety of cryptographic applications that go beyond multipartite key exchange (Boneh and Silverberg 2003).

We return to multilinear maps over familiar vector spaces. The next two examples are about trilinear functionals.

Example 3.16 (trilinear functionals and self-concordance). The definition of self-concordance is usually stated over \mathbb{R}^n . Let $f: \Omega \rightarrow \mathbb{R}$ be a convex C^3 -function on an open convex subset $\Omega \subseteq \mathbb{R}^n$. Then f is said to be *self-concordant* at $x \in \Omega$ if

$$|D^3 f(x)(h, h, h)| \leq 2[D^2 f(x)(h, h)]^{3/2} \quad (3.56)$$

for all $h \in \mathbb{R}^n$ (Nesterov and Nemirovskii 1994). As we discussed in Example 3.2, for any fixed $x \in \Omega$, the higher derivatives $D^2 f(x)$ and $D^3 f(x)$ in this case are bilinear and trilinear functionals on \mathbb{R}^n given by

$$[D^2 f(x)](h, h) = \sum_{i,j=1}^n \frac{\partial^2 f(x)}{\partial x_i \partial x_j} h_i h_j,$$

$$[D^3 f(x)](h, h, h) = \sum_{i,j,k=1}^n \frac{\partial^3 f(x)}{\partial x_i \partial x_j \partial x_k} h_i h_j h_k.$$

The affine invariance (Nesterov and Nemirovskii 1994, Proposition 2.1.1) of self-concordance implies that self-concordance is a tensorial property in the sense of definition ①. For the convergence and complexity analysis of interior point methods, it goes hand in hand with the affine invariance of Newton's method in Example 2.15. Such analysis in turn allows one to establish the celebrated result that a convex optimization problem may be solved to arbitrary ε -accuracy in polynomial time using interior point methods if it has a self-concordant barrier function whose

first and second derivatives may be evaluated in polynomial time (Nesterov and Nemirovskii 1994, Chapter 6). These conditions are satisfied for many common problems including linear programming, convex quadratic programming, second-order cone programming, semidefinite programming and geometric programming (Boyd and Vandenberghe 2004). Contrary to popular belief, polynomial-time solvability to ε -accuracy is not guaranteed by convexity alone: copositive and complete positive programming are convex optimization problems but both are known to be NP-hard (Murty and Kabadi 1987, Dickinson and Gijben 2014).

By Example 3.12, we may view D^2f and D^3f as covariant tensor fields on the manifold Ω (any open subset of a manifold is itself a manifold) with

$$D^2f(x): \mathbb{T}_x(\Omega) \times \mathbb{T}_x(\Omega) \rightarrow \mathbb{R}, \quad D^3f(x): \mathbb{T}_x(\Omega) \times \mathbb{T}_x(\Omega) \times \mathbb{T}_x(\Omega) \rightarrow \mathbb{R}$$

for any fixed $x \in \Omega$. While this tensor field perspective is strictly speaking unnecessary, it helps us formulate (3.56) concretely in situations when we are not working over \mathbb{R}^n . For instance, among the aforementioned optimization problems, semidefinite, complete positive and copositive programming require that we work over the space of symmetric matrices \mathbb{S}^n with Ω given respectively by the following open convex cones:

$$\begin{aligned} \mathbb{S}_{++}^n &= \{A \in \mathbb{S}^n : x^\top Ax > 0, 0 \neq x \in \mathbb{R}^n\} = \{BB^\top \in \mathbb{S}^n : B \in \text{GL}(n)\}, \\ \mathbb{S}_{+++}^n &= \{BB^\top \in \mathbb{S}^n : B \in \text{GL}(n) \cap \mathbb{R}_+^{n \times n}\}, \\ \mathbb{S}_{+++}^{n*} &= \{A \in \mathbb{S}^n : x^\top Ax > 0, 0 \neq x \in \mathbb{R}_+^n\}. \end{aligned}$$

In all three cases we have $\mathbb{T}_X(\Omega) = \mathbb{S}^n$ for all $X \in \Omega$, equipped with the usual trace inner product. To demonstrate self-concordance for the log barrier function $f(X) = -\log \det(X)$, Examples 3.4 and 3.6 give us

$$\begin{aligned} D^2f(X)(H, H) &= \text{tr}(H^\top [\nabla^2 f(X)](H)) = \text{tr}(HX^{-1}HX^{-1}), \\ D^3f(X)(H, H, H) &= \text{tr}(H^\top [\nabla^3 f(X)](H, H)) = -2 \text{tr}(HX^{-1}HX^{-1}HX^{-1}), \end{aligned}$$

and it follows from Cauchy–Schwarz that

$$|D^3f(X)(H, H, H)| \leq 2\|HX^{-1}\|^3 = 2[D^2f(X)(H, H)]^{3/2},$$

as required. For comparison, take the *inverse barrier* function $g(X) = \text{tr}(X^{-1})$ which, as a convex function that blows up near the boundary of \mathbb{S}_{++}^n and has easily computable gradient and Hessian as we saw in Example 3.4, appears to be a perfect barrier function for semidefinite programming. Nevertheless, using the derivatives calculated in Example 3.4,

$$\begin{aligned} D^2g(X)(H, H) &= 2 \text{tr}(HX^{-1}HX^{-2}), \\ D^3g(X)(H, H, H) &= -6 \text{tr}(HX^{-1}HX^{-1}HX^{-2}), \end{aligned}$$

and since $6|h|^3/x^4 > 2(2h^2/x^3)^{3/2}$ as $x \rightarrow 0^+$, (3.56) will not be satisfied when X is near singular. Thus the inverse barrier function for \mathbb{S}_{++}^n is not self-concordant.

What about \mathbb{S}_{+++}^n , the completely positive cone, and its dual cone \mathbb{S}_{+++}^{n*} , the copositive cone? It turns out that it is possible to construct convex self-concordant barrier functions for these (Nesterov and Nemirovskii 1994, Theorem 2.5.1), but the problem is that the gradients and Hessians of these barrier functions are not computable in polynomial time.

The last example, though more of a curiosity, is a personal favourite of the author (Friedland, Lim and Zhang 2019).

Example 3.17 (spectral norm and Grothendieck's inequality). One of the most fascinating inequalities in matrix and operator theory is the following. For any $A \in \mathbb{R}^{m \times n}$, there exists a constant $K_G > 0$ such that

$$\max_{\|x_i\|=\|y_j\|=1} \sum_{i=1}^m \sum_{j=1}^n a_{ij} \langle x_i, y_j \rangle \leq K_G \max_{|\varepsilon_i|=|\delta_j|=1} \sum_{i=1}^m \sum_{j=1}^n a_{ij} \varepsilon_i \delta_j, \quad (3.57)$$

with the maximum on the left taken over vectors $x_1, \dots, x_m, y_1, \dots, y_n \in \mathbb{R}^p$ of unit 2-norm and that on the right taken over $\varepsilon_1, \dots, \varepsilon_m, \delta_1, \dots, \delta_n \in \{-1, +1\}$. Here $p \geq m+n$ can be arbitrarily large. What is remarkable is that the constant K_G is *universal*, i.e. independent of m, n, p or the matrix A . It is christened the Grothendieck constant after the discoverer of the inequality (Grothendieck 1953), which has found widespread applications in combinatorial optimization, complexity theory and quantum information (Pisier 2012). The exact value of the constant is unknown but there are excellent bounds (Davie 1985, Krivine 1979):

$$1.67696 \leq K_G \leq 1.78221.$$

Nevertheless the nature of the Grothendieck constant remains somewhat mysterious. One clue is that since the constant is independent of the matrix A and obviously also independent of the $x_i, y_j, \varepsilon_i, \delta_j$, as these are dummy variables that are maximized over, whatever underlying object the Grothendieck constant measures must have nothing to do with any of the quantities appearing in (3.57). We will show that this object is in fact the Strassen matrix multiplication tensor $M_{m,n,p}$ in Example 3.9. First, observe that the right-hand side of (3.57) is just the matrix $(\infty, 1)$ -norm,

$$\|A\|_{\infty,1} = \max_{v \neq 0} \frac{\|Av\|_1}{\|v\|_\infty} = \max_{\varepsilon_i, \delta_j = \pm 1} \sum_{i=1}^m \sum_{j=1}^n a_{ij} \varepsilon_i \delta_j,$$

which is in fact a reason why the inequality is so useful, the $(\infty, 1)$ -norm being ubiquitous in combinatorial optimization but NP-hard to compute, and the left-hand side of (3.57) being readily computable via semidefinite programming. Writing x_i and y_j , respectively, as columns and rows of matrices $X = [x_1, \dots, x_m] \in \mathbb{R}^{p \times m}$

and $Y = [y_1^\top, \dots, y_n^\top]^\top \in \mathbb{R}^{n \times p}$, we next observe that the constraints on the left-hand side of (3.57) may be expressed in terms of their $(1, 2)$ -norm and $(2, \infty)$ -norm:

$$\|X\|_{1,2} = \max_{v \neq 0} \frac{\|Xv\|_2}{\|v\|_1} = \max_{i=1, \dots, m} \|x_i\|_2,$$

$$\|Y\|_{2,\infty} = \max_{v \neq 0} \frac{\|Yv\|_\infty}{\|v\|_2} = \max_{j=1, \dots, n} \|y_j\|_2,$$

namely, as $\|X\|_{1,2} \leq 1$ and $\|Y\|_{2,\infty} \leq 1$. Lastly, we observe that for the standard inner product on \mathbb{R}^P ,

$$\sum_{i=1}^m \sum_{j=1}^n a_{ij} \langle x_i, y_j \rangle = \text{tr}(XAY),$$

and thus (3.57) may be written as

$$\max_{A, X, Y \neq 0} \frac{\text{tr}(XAY)}{\|A\|_{\infty,1} \|X\|_{1,2} \|Y\|_{2,\infty}} \leq K_G.$$

As we saw in (3.19), the expression on the left is precisely the spectral norm of the trilinear functional

$$\tau_{m,n,p}: \mathbb{R}^{m \times n} \times \mathbb{R}^{n \times p} \times \mathbb{R}^{p \times m} \rightarrow \mathbb{R}, \quad (A, B, C) \mapsto \text{tr}(ABC),$$

if we equip the three vector spaces of matrices with the matrix $(1, 2)$, $(\infty, 1)$ and $(2, \infty)$ -norms respectively. We shall denote this norm as $\|\cdot\|_{1,2,\infty}$. The last step is to observe that as 3-tensors, the trilinear functional $\tau_{m,n,p}$ and the matrix multiplication tensor

$$\mathbb{M}_{m,n,p}: \mathbb{R}^{m \times n} \times \mathbb{R}^{n \times p} \rightarrow \mathbb{R}^{m \times p}, \quad (A, B) \mapsto AB$$

are one and the same, assuming we identify $(\mathbb{R}^{p \times n})^* = \mathbb{R}^{n \times p}$. One way to see this is that the same type of tensor transformation rule applies:

$$(XAY^{-1})(YBZ^{-1}) = X(AB)Z^{-1},$$

$$\text{tr}((XAY^{-1})(YBZ^{-1})(ZCX^{-1})) = \text{tr}(ABC)$$

for any $(X, Y, Z) \in \text{GL}(m) \times \text{GL}(n) \times \text{GL}(p)$. As we have alluded, recognizing that two different multilinear maps correspond to the same tensor without reference to the transformation rule is one benefit of definition ③, which we will discuss immediately after this example. Hence we have shown that

$$K_G = \sup_{m,n,p \in \mathbb{N}} \|\mathbb{M}_{m,n,p}\|_{1,2,\infty}.$$

In particular, the Grothendieck constant is an asymptotic value of the spectral norm of $\mathbb{M}_{m,n,p}$ just as the exponent of matrix multiplication ω in Example 3.9 is an asymptotic value of the tensor rank of $\mathbb{M}_{m,n,p}$.

4. Tensors via tensor products

In the last two sections we have alluded to various inadequacies of definitions ① and ②. The modern definition of tensors, *i.e.* definition ③, remedies them by regarding a tensor as an element of a tensor product of vector spaces (or, more generally, modules). We may view the development of the notion of tensor via these three definitions as a series of improvements. Definition ① leaves the tensor unspecified and instead just describes its change-of-basis theorems as transformation rules. Definition ② fixes this by supplying a multilinear map as a candidate for the unspecified object but the problem becomes one of oversupply: with increasing order d , there is an exponential multitude of different d -linear maps for each transformation rule. Definition ③ fixes both issues by providing an object for definition ① that at the same time also neatly classifies the multilinear maps in definition ② by the transformation rules they satisfy. In physics, the perspective may be somewhat different. Instead of viewing them as a series of pedagogical improvements, definition ① was discovered alongside (most notably) general relativity, with definition ② as its addendum, and definition ③ was discovered alongside (most notably) quantum mechanics. Both remain useful in different ways for different purposes and each is used independently of the other.

One downside of definition ③ is that it relies on the tensor product construction and/or the universal factorization property; both have a reputation of being abstract. We will see that the reputation is largely undeserved; with the proper motivations they are far easier to appreciate than, say, definition ①. One reason for its ‘abstract’ reputation is that the construction is often cast in a way that lends itself to vast generalizations. In fact, tensor products of Hilbert spaces (Reed and Simon 1980), modules (Lang 2002), vector bundles (Milnor and Stasheff 1974), operator algebras (Takesaki 2002), representations (Fulton and Harris 1991), sheaves (Hartshorne 1977), cohomology rings (Hatcher 2002), *etc.*, have become foundational materials covered in standard textbooks, with more specialized tensor product constructions, such as those of Banach spaces (Ryan 2002), distributions (Trèves 2006), operads (Dunn 1988) or more generally objects in any tensor category (Etingof *et al.* 2015), readily available in various monographs. This generality is a feature, not a bug. In our (much more modest) context, it allows us to define tensor products of norm spaces and inner product spaces, which in turn allows us to define norms and inner products for tensors, to view multivariate functions as tensors, an enormously useful perspective in computation, and to identify ‘separation of variables’ as the common underlying thread in a disparate array of well-known algorithms. In physics, the importance of tensor products cannot be over-emphasized: it is one of the fundamental postulates of quantum mechanics (Nielsen and Chuang 2000, Section 2.2.8), the source behind many curious quantum phenomena that lie at the heart of the subject, and is indispensable in technologies such as quantum computing and quantum cryptography.

Considering its central role, we will discuss three common approaches to constructing tensor products, increasing in abstraction:

- (i) via tensor products of function spaces,
- (ii) via tensor products of more general vector spaces,
- (iii) via the universal factorization property.

As in the case of the three definitions of tensor, each of these constructions is useful in its own way and each may be taken to be a variant of definition ③. We will motivate each construction with concrete examples but we defer all examples of computational relevance to Section 4.4. The approach of defining a tensor as an element of a tensor product of vector spaces likely first appeared in the first French edition of Bourbaki (1998) and is now standard in graduate algebra texts (Dummit and Foote 2004, Lang 2002, Vinberg 2003). It has also caught on in physics (Geroch 1985) and in statistics (McCullagh 1987). For further historical information we refer the reader to Conrad (2018) and the last chapter of Kostrikin and Manin (1997).

4.1. Tensor products of function spaces

This is the most intuitive of our three tensor product constructions as it reflects the way we build functions of more variables out of functions of fewer variables. In many situations, especially analytic ones, this tensor-as-a-multivariate-function view suffices and one need go no further.

For any set X , the set of all real-valued functions on X ,

$$\mathbb{R}^X := \{f : X \rightarrow \mathbb{R}\},$$

is always a vector space under pointwise addition and scalar multiplication of functions, that is, by defining

$$(a_1 f_1 + a_2 f_2)(x) := a_1 f_1(x) + a_2 f_2(x) \quad \text{for each } x \in X.$$

Note that the left-hand side is a real-valued function whose value at x is defined by the right-hand side, which involves only additions and multiplications of real numbers $a_1, a_2, f_1(x), f_2(x)$. For any finite set X , $\dim \mathbb{R}^X = \#X$, the cardinality of X .

It turns out that the converse is also true: any vector space \mathbb{V} may be regarded as a function space \mathbb{R}^X for some X . It is easiest to see this for finite-dimensional vector spaces, so we first assume that $\dim \mathbb{V} = n \in \mathbb{N}$. In this case \mathbb{V} has a basis $\mathcal{B} = \{v_1, \dots, v_n\}$ and every element $v \in \mathbb{V}$ may be written as

$$v = a_1 v_1 + \dots + a_n v_n \quad \text{for some unique } a_1, \dots, a_n \in \mathbb{R}. \quad (4.1)$$

So we simply take $X = \mathcal{B}$ and regard $v \in \mathbb{V}$ as the function $f_v \in \mathbb{R}^X$ defined by

$$f_v : X \rightarrow \mathbb{R}, \quad f_v(v_i) = a_i, \quad i = 1, \dots, n.$$

The map $\mathbb{V} \rightarrow \mathbb{R}^{\mathcal{B}}, v \mapsto f_v$ is linear and invertible and so the two vector spaces are isomorphic. The function f_v depends on our choice of basis, and any two such functions f_v and f'_v corresponding to different bases \mathcal{B} and \mathcal{B}' will be related via the contravariant 1-tensor transformation rule if we represent f_v as a vector in \mathbb{R}^n with coordinates given by its values

$$[f_v]_{\mathcal{B}} := \begin{bmatrix} f(v_1) \\ \vdots \\ f(v_n) \end{bmatrix} = \begin{bmatrix} a_1 \\ \vdots \\ a_n \end{bmatrix} \in \mathbb{R}^n$$

and do likewise for f'_v . This is of course just restating (3.2) and (3.3) in slightly different language. As long as we assume that every vector space \mathbb{V} has a basis \mathcal{B} , or, equivalently, the axiom of choice (Blass 1984), the above construction for $\mathbb{V} \cong \mathbb{R}^{\mathcal{B}}$ in principle applies verbatim to infinite-dimensional vector spaces, as having a unique finite linear combination of the form (4.1) is the very definition of a basis. When used in this sense, such a basis for a vector space is called a Hamel basis.

Since any vector space is isomorphic to a function space, by defining tensor products on function spaces we define tensor products on all vector spaces. This is the most straightforward approach to defining tensor products. In this regard, a d -tensor is just a real-valued function $f(x_1, \dots, x_d)$ of d variables. For two sets X and Y , we define the *tensor product* $\mathbb{R}^X \otimes \mathbb{R}^Y$ of \mathbb{R}^X and \mathbb{R}^Y to be the subspace of $\mathbb{R}^{X \times Y}$ comprising all functions that can be written as a finite sum of product of two univariate functions, one of x and another of y :

$$f(x, y) = \sum_{i=1}^r \varphi_i(x)\psi_i(y). \tag{4.2}$$

For univariate functions $\varphi: X \rightarrow \mathbb{R}$ and $\psi: Y \rightarrow \mathbb{R}$, we will write $\varphi \otimes \psi: X \times Y \rightarrow \mathbb{R}$ for the bivariate function defined by

$$(\varphi \otimes \psi)(x, y) = \varphi(x)\psi(y) \quad \text{for all } x \in X, y \in Y.$$

Such a function is called a *separable function*, and when \otimes is used in this sense, we will refer to it a *separable product*; see Figure 4.1 for a simple example. In notation,

$$\mathbb{R}^X \otimes \mathbb{R}^Y := \left\{ f \in \mathbb{R}^{X \times Y} : f = \sum_{i=1}^r \varphi_i \otimes \psi_i, \varphi_i \in \mathbb{R}^X, \psi_i \in \mathbb{R}^Y, r \in \mathbb{N} \right\}.$$

Note that $\mathbb{R}^{X \times Y} = \{f: X \times Y \rightarrow \mathbb{R}\}$ is the set of all bivariate \mathbb{R} -valued functions on $X \times Y$, separable or not. In reality, when \mathbb{V} is infinite-dimensional, its Hamel basis \mathcal{B} is often uncountably infinite and impossible to write down explicitly for many common infinite-dimensional vector spaces (e.g. any complete topological vector space), and thus this approach of constructing tensor products by identifying a vector space \mathbb{V} with $\mathbb{R}^{\mathcal{B}}$ is not practically useful for most infinite-dimensional




Figure 4.1. Separable function $f(x, y) = \varphi(x)\psi(y)$ with $\varphi(x) = 1 + \sin(180x)$ and $\psi(y) = \exp(-(y - 3/2)^2)$.

vector spaces. Furthermore, if we take this route we will inevitably have to discuss a ‘change of Hamel bases’, a tricky undertaking that has never been done, as far as we know. As such, tensor products of infinite-dimensional vector spaces are almost always discussed in the context of *topological tensor products*, which we will describe in Example 4.1.

If X and Y are finite sets, then in fact we have

$$\mathbb{R}^X \otimes \mathbb{R}^Y = \mathbb{R}^{X \times Y} \quad (4.3)$$

as, clearly,

$$\dim(\mathbb{R}^X \otimes \mathbb{R}^Y) = \#X \cdot \#Y = \#(X \times Y) = \dim \mathbb{R}^{X \times Y},$$

and so every function $f: X \times Y \rightarrow \mathbb{R}$ can be expressed in the form (4.2) for some finite r . In other words, elements of $\mathbb{R}^X \otimes \mathbb{R}^Y$ give all functions of two variables $(x, y) \in X \times Y$. If X and Y are infinite sets, then (4.3) is demonstrably false. Nevertheless, function spaces over infinite sets almost always have additional properties imposed (*e.g.* continuity, differentiability, integrability, *etc.*) and/or carry additional structures (*e.g.* metrics, norms, inner products, completeness in the

topology induced by these, *etc.*). If we constrain our functions in some manner (instead of allowing them to be completely arbitrary) and limit our attention to some smaller subspace $\mathcal{F}(X) \subseteq \mathbb{R}^X$, then it turns out that the analogue of (4.3) often holds, that is,

$$\mathcal{F}(X) \otimes \mathcal{F}(Y) = \mathcal{F}(X \times Y) \tag{4.4}$$

for the right function class \mathcal{F} and the right interpretation of \otimes , as we will see in the next two examples. For infinite-dimensional vector spaces, there are many ways to define (topological) tensor products, and the ‘right tensor product’ is usually regarded as the one that gives us (4.4).

Example 4.1 (tensor products in infinite dimensions). We write $\mathbb{R}[x_1, \dots, x_m]$ for the set of all multivariate polynomials in m variables, with no restriction on degrees. This is one of the few infinite-dimensional vector spaces that has a countable Hamel basis, namely the set of all monic monomials

$$\mathcal{B} = \{x_1^{d_1} x_2^{d_2} \dots x_m^{d_m} : d_1, d_2, \dots, d_m \in \mathbb{N} \cup \{0\}\}.$$

For multivariate polynomials, it is clearly true that¹⁹

$$\mathbb{R}[x_1, \dots, x_m] \otimes \mathbb{R}[y_1, \dots, y_n] = \mathbb{R}[x_1, \dots, x_m, y_1, \dots, y_n], \tag{4.5}$$

as polynomials are sums of finitely many monomials and monomials are always separable, for example

$$7x_1^2 x_2 y_2^3 y_3 - 6x_2^4 y_1 y_2^5 y_3 = (7x_1^2 x_2) \cdot (y_2^3 y_3) + (-6x_2^4) \cdot (y_1 y_2^5 y_3).$$

But (4.4) is also clearly false in general for other infinite-dimensional vector spaces. Take continuous real-valued functions for illustration. While $\sin(x + y) = \sin x \cos y + \cos x \sin y$ and $\log(xy) = \log x + \log y$, we can never have

$$\sin(xy) = \sum_{i=1}^r \varphi_i(x)\psi_i(y) \quad \text{or} \quad \log(x - y) = \sum_{i=1}^r \varphi_i(x)\psi_i(y)$$

for any continuous functions φ_i, ψ_i and finite $r \in \mathbb{N}$. But if we allow $r \rightarrow \infty$, then

$$\sin(xy) = \sum_{n=0}^{\infty} \frac{(-1)^n}{(2n + 1)!} x^{2n+1} y^{2n+1},$$

and by Taylor’s theorem, $\sin(xy)$ can be approximated to arbitrary accuracy by sums of separable functions. Likewise,

$$\log(x - y) = \log(-y) - \frac{x}{y} - \frac{x^2}{2y^2} - \frac{x^3}{3y^3} - \dots - \frac{x^n}{ny^n} + O(x^{n+1})$$

¹⁹ One may extend (4.5) to monomials involving arbitrary real powers, *i.e.* Laurent polynomials, posynomials, signomials, *etc.*, as they are all *finite* sums of monomials.

as $x \rightarrow 0$, or, more relevant for our later purposes,

$$\log(x - y) = \log(x) - \frac{y}{x} - \frac{y^2}{2x^2} - \frac{y^3}{3x^3} - \cdots - \frac{y^n}{nx^n} + O\left(\frac{1}{x^{n+1}}\right),$$

as $x \rightarrow \infty$. In other words, we may approximate $\log(x - y)$ by sums of separable functions when x is very small or very large; the latter will be important when we discuss the fast multipole method in Example 4.45.

Given that (4.5) holds for polynomials, the Stone–Weierstrass theorem gives an indication that we might be able to extend (4.4) to other infinite-dimensional spaces, say, the space of continuous functions $C(X)$ on some nice domain X . However, we will need to relax the definition of tensor product to allow for limits, *i.e.* taking *completion* with respect to some appropriate choice of norm. For instance, we may establish, for $1 \leq p < \infty$,

$$C(X) \widehat{\otimes} C(Y) = C(X \times Y), \quad L^p(X) \widehat{\otimes} L^p(Y) = L^p(X \times Y),$$

where X and Y are locally compact Hausdorff spaces in the left equality and σ -finite measure spaces in the right equality (Light and Cheney 1985, Corollaries 1.14 and 1.52). The *topological tensor product* $C(X) \widehat{\otimes} C(Y)$ above refers to the completion of

$$C(X) \otimes C(Y) = \left\{ f \in C(X \times Y) : f = \sum_{i=1}^r \varphi_i \otimes \psi_i, \varphi_i \in C(X), \psi_i \in C(Y), r \in \mathbb{N} \right\}$$

with respect to some *cross-norm*, *i.e.* a norm $\|\cdot\|$ that is multiplicative on rank-one elements,

$$\|\varphi \otimes \psi \otimes \cdots \otimes \theta\| = \|\varphi\| \|\psi\| \cdots \|\theta\|,$$

like those in (1.6), (1.7), (3.45), (3.46); we will discuss cross-norms in greater detail in Examples 4.17 and 4.18 for Banach and Hilbert spaces. As we have seen, without taking completion, $C(X) \otimes C(Y)$ will be a strictly smaller subset of $C(X \times Y)$. The same goes for $L^p(X) \widehat{\otimes} L^p(Y)$, but the completion is with respect to some other cross-norm dependent on the value of p .

There many important infinite-dimensional vector spaces that are not Banach (and thus not Hilbert) spaces. The topology of these spaces is usually defined by families of *seminorms*. Fortunately, the construction outlined above also works with such spaces to give topological tensor products that satisfy (4.4). If we limit our consideration of X and Y to subsets of Euclidean spaces, then for spaces of rapidly decaying Schwartz functions (see Example 3.7), smooth functions, compactly supported smooth functions and holomorphic functions, we have (Trèves 2006, Theorem 51.6)

$$\begin{aligned} S(X) \widehat{\otimes} S(Y) &= S(X \times Y), & C^\infty(X) \widehat{\otimes} C^\infty(Y) &= C^\infty(X \times Y), \\ C_c^\infty(X) \widehat{\otimes} C_c^\infty(Y) &= C_c^\infty(X \times Y), & H(X) \widehat{\otimes} H(Y) &= H(X \times Y), \end{aligned}$$

where X and Y , respectively, are \mathbb{R}^m and \mathbb{R}^n (Schwartz), open subsets of \mathbb{R}^m and \mathbb{R}^n

(smooth), compact subsets of \mathbb{R}^m and \mathbb{R}^n (compactly supported) and open subsets of \mathbb{C}^m and \mathbb{C}^n (holomorphic).

Similar results have been established for more complicated function spaces such as Sobolev and Besov spaces (Sickel and Ullrich 2009) or extended to more general domains such as having X and Y be smooth manifolds or Stein manifolds (Akbarov 2003). The bottom line is that while it is not always true that all (X, Y) -variate functions are separable products of X -variate functions and Y -variate functions – L^∞ -functions being a notable exception – (4.4) holds true for many common function spaces and many common types of domains when \otimes is interpreted as an appropriate topological tensor product $\widehat{\otimes}$.

To keep our promise of simplicity, our discussion in Example 4.1 is necessarily terse as the technical details are quite involved. Nevertheless, the readers will get a fairly good idea of the actual construction for the special cases of Banach and Hilbert spaces in Examples 4.17 and 4.18. We will briefly mention another example.

Example 4.2 (tensor products of distributions). Strictly speaking, the construction of tensor products via sums of separable functions gives us *contravariant* tensors, as we will see more clearly in Example 4.5. Nevertheless, we may get covariant tensors by taking the continuous dual of the function spaces in Example 4.1. For a given \mathbb{V} equipped with a norm, this is the set of all linear functionals $\varphi \in \mathbb{V}^*$ whose dual norm in (3.44) is finite. For a function space $\mathcal{F}(X)$, the continuous dual space will be denoted $\mathcal{F}'(X)$ and it is a subspace of the algebraic dual space $\mathcal{F}(X)^*$. Elements of $\mathcal{F}'(X)$ are called *distributions* or generalized functions with the best-known example being the Dirac delta ‘function’, intuitively defined by

$$\int_{-\infty}^{\infty} f(x)\delta(x) dx = f(0),$$

and rigorously as the continuous linear functional $\delta: C_c^\infty(\mathbb{R}) \rightarrow \mathbb{R}$, $\varphi \rightarrow \varphi(0)$, i.e. given by evaluation at zero. The continuous duals of $C_c^\infty(X)$, $S(X)$, $C^\infty(X)$ and $H(X)$ are denoted by $D'(X)$ (distributions), $S'(X)$ (tempered distributions), $E'(X)$ (compactly supported distributions) and $H'(X)$ (analytic functionals) respectively, and one may obtain the following dual analogues (Trèves 2006, Theorem 51.7):

$$\begin{aligned} S'(X) \widehat{\otimes} S'(Y) &= S'(X \times Y), & E'(X) \widehat{\otimes} E'(Y) &= E'(X \times Y), \\ D'(X) \widehat{\otimes} D'(Y) &= D'(X \times Y), & H'(X) \widehat{\otimes} H'(Y) &= H'(X \times Y), \end{aligned} \quad (4.6)$$

with X and Y as in Example 4.1. These apparently innocuous statements conceal some interesting facts. While every linear operator $\Phi: \mathbb{V} \rightarrow \mathbb{W}$ and bilinear functional $\beta: \mathbb{V} \times \mathbb{W} \rightarrow \mathbb{R}$ on finite-dimensional vector spaces may be represented by a matrix $A \in \mathbb{R}^{m \times n}$, as we discussed in Section 3.1, the continuous analogue on infinite-dimensional spaces is clearly false. For instance, not every continuous

linear operator $\Phi: L^2(\mathbb{R}) \rightarrow L^2(\mathbb{R})$ may be expressed as

$$[\Phi(f)](x) = \int_{-\infty}^{\infty} K(x, y)f(y) dy \quad (4.7)$$

with an integral kernel K in place of the matrix A , as the identity operator is already a counterexample. This is the impetus behind the idea of a *reproducing kernel Hilbert space* that we will discuss in Example 4.20. What (4.6) implies is that this is true for the function spaces in Example 4.1. Any continuous linear operator $\Phi: S(X) \rightarrow S'(Y)$ has a unique kernel $K \in S'(X \times Y)$ and likewise with C_c^∞ , C^∞ , H in place of S (Trèves 2006, Chapters 50 and 51). For bilinear functionals, (4.6) implies that for these spaces, say, $\beta: C_c^\infty(X) \times C_c^\infty(Y) \rightarrow \mathbb{R}$, we may decompose it into

$$\beta = \sum_{i=1}^{\infty} a_i \varphi_i \otimes \psi_i \quad (4.8)$$

with $(a_i)_{i=1}^{\infty} \in l^1(\mathbb{N})$, and $\varphi_i \in D'(X)$, $\psi_i \in D'(Y)$ satisfying $\lim_{i \rightarrow \infty} \varphi_i = 0 = \lim_{i \rightarrow \infty} \psi_i$, among other properties (Schaefer and Wolff 1999, Theorem 9.5). These results are all consequences of Schwartz's kernel theorem, which, alongside Mercer's kernel theorem in Example 4.20, represent different approaches to resolving the falsity of (4.7). While Mercer restricted to smaller spaces (reproducing kernel Hilbert spaces), Schwartz expanded to larger ones (spaces of distributions), but both are useful. Among other things, the former allows for the φ_i and ψ_i in (4.8) to be functions instead of distributions, from which we obtain feature maps for support vector machines (Scholköpfung and Smola 2002, Steinwart and Christmann 2008); on the other hand the latter would give us distributional solutions for PDEs (Trèves 2006, Theorem 52.6).

The discussion up to this point extends to any d sets X_1, X_2, \dots, X_d , either by repeated application of the $d = 2$ case or directly:

$$\begin{aligned} & \mathbb{R}^{X_1} \otimes \mathbb{R}^{X_2} \otimes \dots \otimes \mathbb{R}^{X_d} \\ & := \left\{ f \in \mathbb{R}^{X_1 \times X_2 \times \dots \times X_d} : f = \sum_{i=1}^r \varphi_i \otimes \psi_i \otimes \dots \otimes \theta_i, \right. \\ & \quad \left. \varphi_i \in \mathbb{R}^{X_1}, \psi_i \in \mathbb{R}^{X_2}, \dots, \theta_i \in \mathbb{R}^{X_d}, i = 1, \dots, r \in \mathbb{N} \right\}. \quad (4.9) \end{aligned}$$

Each summand is a separable function defined by

$$(\varphi \otimes \psi \otimes \dots \otimes \theta)(x_1, x_2, \dots, x_d) := \varphi(x_1)\psi(x_2)\dots\theta(x_d) \quad (4.10)$$

for $x_1 \in X_1, x_2 \in X_2, \dots, x_d \in X_d$, and

$$\mathbb{R}^{X_1 \times \dots \times X_d} = \{f: X_1 \times \dots \times X_d \rightarrow \mathbb{R}\}$$

is the set of d -variate real-valued functions on $X_1 \times \dots \times X_d$. We will next look at some concrete examples of separable functions.

Example 4.3 (multivariate normal and quantum harmonic oscillator). The quintessential example of an n -variate separable function is a Gaussian:

$$f(x) = \exp(x^*Ax + b^*x + c)$$

for $A \in \mathbb{C}^{n \times n}$, $b \in \mathbb{C}^n$, $c \in \mathbb{C}$, which is separable under a change of variables $x \mapsto V^*x$ where the columns of $V \in U(n)$ are an orthonormal eigenbasis of $(A + A^*)/2$. Usually one requires that $(A + A^*)/2$ is negative definite and b is purely imaginary to ensure that it is an L^2 -function. This is a singularly important separable function.

In statistics, it arises as the probability density, moment generating and characteristic functions of a normal random variable $X \sim N(\mu, \Sigma)$, with the last of the three given by

$$\varphi_X(x) = \exp\left(i\mu^\top x - \frac{1}{2}x^\top \Sigma x\right).$$

This gives the unique distribution all of whose cumulant d -tensors, *i.e.* the d th derivative as in Example 3.2 of $\log \varphi_X$, vanish when $d \geq 3$, a fact from which the central limit theorem follows almost as an afterthought (McCullagh 1987, Section 2.7). It is an undisputed fact that this is the single most important probability distribution in statistics.

In physics, it arises as Gaussian wave functions for the quantum harmonic oscillator. For a spinless particle in \mathbb{R}^3 of mass μ subjected to the potential

$$V(x, y, z) = \frac{1}{2}\mu(\omega_x^2 x^2 + \omega_y^2 y^2 + \omega_z^2 z^2),$$

if we assume that V is isotropic, *i.e.* $\omega_x = \omega_y = \omega_z = \omega$, then the states are given by

$$\psi_{m,n,p}(x, y, z) = \left(\frac{\beta^2}{\pi}\right)^{3/4} \frac{H_m(\beta x)H_n(\beta y)H_p(\beta z)}{\sqrt{2^{m+n+p} m! n! p!}} \exp\left[\frac{\beta^2}{2}(x^2 + y^2 + z^2)\right],$$

where $\beta^2 := \mu\omega/\hbar$ and H_n denotes a Hermite polynomial of degree n (Cohen-Tannoudji, Diu and Laloë 2020a, Chapter VII, Complement B). The ground state, *i.e.* $m = n = p = 0$, is a Gaussian but note that even excited states, *i.e.* $m + n + p > 0$, are separable functions. If one subscribes to the dictum that ‘physics is that subset of human experience which can be reduced to coupled harmonic oscillators’ (attributed to Michael Peskin), then this has a status not unlike that of the normal distribution in statistics.

As in the case $d = 2$, when X_1, \dots, X_d are finite sets, we have

$$\mathbb{R}^{X_1} \otimes \dots \otimes \mathbb{R}^{X_d} = \mathbb{R}^{X_1 \times \dots \times X_d}, \quad (4.11)$$

that is, d -tensors in $\mathbb{R}^{X_1} \otimes \dots \otimes \mathbb{R}^{X_d}$ are just functions of d variables $(x_1, \dots, x_d) \in X_1 \times \dots \times X_d$. In particular, the dimension of the tensor space is

$$\dim(\mathbb{R}^{X_1} \otimes \dots \otimes \mathbb{R}^{X_d}) = n_1 \dots n_d,$$

where $n_i := \#X_i$, $i = 1, \dots, d$. As in the $d = 2$ case, constructing a tensor product of infinite-dimensional vector spaces $\mathbb{V}_1, \dots, \mathbb{V}_d$ by picking Hamel bases $\mathcal{B}_1, \dots, \mathcal{B}_d$ and then defining

$$\mathbb{V}_1 \otimes \cdots \otimes \mathbb{V}_d := \mathbb{R}^{\mathcal{B}_1} \otimes \cdots \otimes \mathbb{R}^{\mathcal{B}_d} \quad (4.12)$$

is rarely practical, with multivariate polynomials just about the only exception: Clearly,

$$\begin{aligned} &\mathbb{R}[x_1, \dots, x_m] \otimes \mathbb{R}[y_1, \dots, y_n] \otimes \cdots \otimes \mathbb{R}[z_1, \dots, z_p] \\ &= \mathbb{R}[x_1, \dots, x_m, y_1, \dots, y_n, \dots, z_1, \dots, z_p], \end{aligned}$$

just as in (4.5). However, if $\widehat{\otimes}$ is interpreted to be an appropriate topological tensor product, then repeated applications of the results in Example 4.1 lead us to equalities like

$$L^2(X_1) \widehat{\otimes} L^2(X_2) \widehat{\otimes} \cdots \widehat{\otimes} L^2(X_d) = L^2(X_1 \times X_2 \times \cdots \times X_d). \quad (4.13)$$

We will furnish more details in Examples 4.17 and 4.18.

We now attempt to summarize the discussions up to this point with a definition that is intended to capture all examples in this section. Note that it is difficult to be more precise without restricting the scope.

Definition 4.4 (tensors as multivariate functions). Let $\mathcal{F}_i(X_i) \subseteq \mathbb{R}^{X_i}$ be a subspace of real-valued functions on X_i , $i = 1, \dots, d$. The topological tensor product of $\mathcal{F}_1(X_1), \dots, \mathcal{F}_d(X_d)$ is a subspace

$$\mathcal{F}_1(X_1) \widehat{\otimes} \cdots \widehat{\otimes} \mathcal{F}_d(X_d) \subseteq \mathbb{R}^{X_1 \times \cdots \times X_d}$$

that comprises finite sums of real-valued separable functions $\varphi_1 \otimes \cdots \otimes \varphi_d$ with $\varphi_i \in \mathcal{F}_i(X_i)$, $i = 1, \dots, d$, and their limits with respect to some topology. A d -tensor is an element of this space.

Some explanations are in order. Definition 4.4 includes the special cases of functions on finite sets and multivariate polynomials by simply taking the topology to be the discrete topology. It is not limited to norm topology, as numerous tensor product constructions require seminorms (Trèves 2006) or quasinorms (Sickel and Ullrich 2009). The definition allows for tensor products of different types of function spaces so as to accommodate different regularities in different arguments. For example, a classical solution for the fluid velocity v in the Navier–Stokes equation (Fefferman 2006)

$$\frac{\partial v_i}{\partial t} + \sum_{j=1}^3 \frac{\partial v_i}{\partial x_j} v_j = -\frac{1}{\rho} \frac{\partial p}{\partial x_i} + \nu \sum_{j=1}^3 \frac{\partial^2 v_i}{\partial x_j^2} + f_i, \quad i = 1, 2, 3, \quad (4.14)$$

may be regarded as a tensor $v \in C^2(\mathbb{R}^3) \widehat{\otimes} C^1[0, \infty) \otimes \mathbb{R}^3$, that is, twice continuously differentiable in its spatial arguments, once in its temporal argument, with a discrete third argument capturing the fact that it is a three-dimensional vector field. We will have more to say about the last point in Example 4.30.

Definition 4.4 allows us to view numerous types of multivariate functions as tensors. The variables involved may be continuous or discrete or a mix of both. The functions may represent anything from solutions to PDEs to target functions in machine learning. In short, Definition 4.4 is extremely versatile and is by far the most common manifestation of tensors.

We will reflect on how Definition 4.4 relates to definitions ① and ② by way of two examples.

Example 4.5 (hypermatrices). A hypermatrix has served as our ‘multi-indexed object’ in definition ① but may now be properly defined as an element of $\mathbb{R}^{X_1 \times \dots \times X_d}$ when X_1, \dots, X_d are finite or countable discrete sets. For any $n \in \mathbb{N}$, we will write

$$[n] := \{1, 2, \dots, n\}.$$

We begin from $d = 1$. While \mathbb{R}^n is usually taken as a Cartesian product $\mathbb{R}^n = \mathbb{R} \times \dots \times \mathbb{R}$, it is often fruitful to regard it as the set of functions

$$\mathbb{R}^{[n]} = \{f: [n] \rightarrow \mathbb{R}\}. \quad (4.15)$$

Any n -tuple $(a_1, \dots, a_n) \in \mathbb{R}^n$ defines a function $f: [n] \rightarrow \mathbb{R}$ with $f(i) = a_i$, $i = 1, \dots, n$, and any function $f \in \mathbb{R}^{[n]}$ defines an n -tuple $(f(1), \dots, f(n)) \in \mathbb{R}^n$. So we may as well assume that $\mathbb{R}^n \equiv \mathbb{R}^{[n]}$, given that any difference is superficial. There are two advantages. The first is that it allows us to speak of sequence spaces in the same breath, as the space of all sequences is then

$$\mathbb{R}^{\mathbb{N}} = \{f: \mathbb{N} \rightarrow \mathbb{R}\} = \{(a_i)_{i=1}^{\infty} : a_i \in \mathbb{R}\},$$

and this in turn allows us to treat various Banach spaces of sequences $l^p(\mathbb{N})$, $c(\mathbb{N})$, $c_0(\mathbb{N})$ as special cases of their function spaces counterpart $L^p(\Omega)$, $C(\Omega)$, $C_0(\Omega)$. This perspective is of course common knowledge in real and functional analysis.

Given that there is no such thing as a ‘two-way Cartesian product’, a second advantage is that it allows us to define matrix spaces via

$$\mathbb{R}^{m \times n} := \mathbb{R}^{[m] \times [n]} = \{f: [m] \times [n] \rightarrow \mathbb{R}\}, \quad (4.16)$$

where we regard $(a_{ij})_{i,j=1}^{m,n}$ as shorthand for the function with $f(i, j) = a_{ij}$, $i \in [m]$, $j \in [n]$. Why not just say that it is a two-dimensional array? The answer is that an ‘array’ is an undefined term in mathematics; at best we may view it as a type of data structure, and as soon as we try to define it rigorously, we will end up with essentially (4.16). Again this perspective is known; careful treatments of matrices in linear algebra such as Berberian (2014, Definition 4.1.3) or Brualdi (1992) would define a matrix in the same vein. An infinite-dimensional variation gives us double sequences or equivalently infinite-dimensional matrices in $l^2(\mathbb{N} \times \mathbb{N})$. These are essential in Heisenberg’s matrix mechanics, the commutation relation

$PQ - QP = i\hbar I$ being false for finite-dimensional matrices (trace zero on the left, non-zero on the right), requiring infinite-dimensional ones instead:

$$Q = \sqrt{\frac{\hbar}{2\nu\omega}} \begin{bmatrix} 0 & \sqrt{1} & 0 & 0 & \cdots \\ \sqrt{1} & 0 & \sqrt{2} & 0 & \cdots \\ 0 & \sqrt{2} & 0 & \sqrt{3} & \cdots \\ 0 & 0 & \sqrt{3} & 0 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix},$$

$$P = \sqrt{\frac{\hbar\nu\omega}{2}} \begin{bmatrix} 0 & -i\sqrt{1} & 0 & 0 & \cdots \\ i\sqrt{1} & 0 & -i\sqrt{2} & 0 & \cdots \\ 0 & i\sqrt{2} & 0 & -i\sqrt{3} & \cdots \\ 0 & 0 & i\sqrt{3} & 0 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}.$$

P and Q are in fact the position and momentum operators for the harmonic oscillator in Example 4.3 restricted to one dimension with all y - and z -terms dropped (Jordan 2005, Chapter 18).

It is straightforward to generalize. Let $n_1, \dots, n_d \in \mathbb{N}$. Then

$$\mathbb{R}^{n_1 \times \cdots \times n_d} := \mathbb{R}^{[n_1] \times \cdots \times [n_d]} = \{f: [n_1] \times \cdots \times [n_d] \rightarrow \mathbb{R}\}, \quad (4.17)$$

that is, a d -dimensional hypermatrix or more precisely an $n_1 \times \cdots \times n_d$ hypermatrix is a real-valued function on $[n_1] \times \cdots \times [n_d]$. We introduce the shorthand $(a_{i_1 \dots i_d})_{i_1, \dots, i_d=1}^{n_1, \dots, n_d}$ for a hypermatrix with $f(i_1, \dots, i_d) = a_{i_1 \dots i_d}$. We say that a hypermatrix is *hypercubical* when $n_1 = \cdots = n_d$. These terms originated in the Russian mathematical literature, notably in Gelfand, Kapranov and Zelevinsky (1992, 1994), although the notion itself has appeared implicitly as coefficients of multilinear forms in Cayley (1845), alongside the term ‘hyperdeterminant’. The names ‘multidimensional matrix’ (Sokolov 1972) or ‘spatial matrix’ (Sokolov 1960) have also been used for hypermatrices. We would have preferred simply ‘matrix’ or ‘ d -matrix’ instead of ‘hypermatrix’ or ‘ d -hypermatrix’, but the conventional view of the term ‘matrix’ as exclusively two-dimensional has become too deeply entrenched to permit such a change in nomenclature.

We will allow for a further generalization of the above, namely, X_1, \dots, X_d may be any discrete sets.

- Instead of restricting i_1, \dots, i_d to ordinal variables with values in $[n_k] = \{1, \dots, n_k\}$, we allow for them to be nominal, e.g. $\{\uparrow, \downarrow\}$, $\{A, C, G, T\}$, or irreducible representations of some group (see Example 2.16).
- The discrete set may have additional structures, e.g. they might be graphs or finite groups or combinatorial designs.
- We permit sets of countably infinite cardinality, e.g. \mathbb{N} or \mathbb{Z} (see Example 3.7).

Readers may have noted that we rarely invoke a d -hypermatrix with $d \geq 3$ in actual examples. The reason is that whatever we can do with a d -hypermatrix $A \in \mathbb{R}^{n_1 \times \cdots \times n_d}$ we can do without or do better with $f: [n_1] \times \cdots \times [n_d] \rightarrow \mathbb{R}$. If anything, hypermatrices give a misplaced sense of familiarity by furnishing a false analogy with standard two-dimensional matrices; they are nothing alike, as we explain next.

In conjunction with specific bases, that is, when a d -hypermatrix is a coordinate representation of a d -tensor, the values $a_{i_1 \dots i_d}$ convey information *about those basis vectors* such as near orthogonality (Example 3.7) or incidence (below). But on their own, the utility of d -hypermatrices is far from that of the usual two-dimensional matrices, one of the most powerful and universal mathematical tools. While there is a dedicated calculus²⁰ for working with two-dimensional matrices – a rich collection of matrix operations, functions, invariants, decompositions, canonical forms, *etc.*, that are consistent with the tensor transformation rules and most of them readily computable – there is no equivalent for d -hypermatrices. There are no canonical forms when $d \geq 3$ (Landsberg 2012, Chapter 10), no analogue of matrix–matrix product when d is odd (Example 2.7) and thus no analogue of matrix inverse, and almost every calculation is computationally intractable even for small examples (see Section 5.2); here ‘no’ is in the sense of mathematically proved to be non-existent. The most useful aspect of hypermatrices is probably in supplying some analogues of basic block operations, notably in Strassen’s laser method (Bürgisser *et al.* 1997, Chapter 15), or of row and column operations, notably in various variations of slice rank (Lovett 2019), although there is nothing as sophisticated as the kind of block matrix operations on page 624.

Usually when we recast a problem in the form of a matrix problem, we are on our way to a solution: matrices are the tool that gets us there. The same is not true with hypermatrices. For instance, while we may easily capture the adjacency structure of a d -hypergraph $G = (V, E)$, $E \subseteq \binom{V}{d}$, with a hypermatrix $f: V \times \cdots \times V \rightarrow \mathbb{R}$,

$$f(i_1, \dots, i_d) := \begin{cases} 1 & \{i_1, \dots, i_d\} \in E, \\ 0 & \{i_1, \dots, i_d\} \notin E, \end{cases}$$

the analogy with spectral graph theory stops here; the hypermatrix view (Friedman 1991, Friedman and Wigderson 1995) does not get us anywhere close to what is possible in the $d = 2$ case. Almost every mathematically sensible problem involving hypermatrices is a challenge, even for $d = 3$ and very small dimensions such as $4 \times 4 \times 4$. For example, we know that the $m \times n$ matrices of rank not more than r are precisely the ones with vanishing $(r + 1) \times (r + 1)$ minors. Resolving the equivalent problem for $4 \times 4 \times 4$ hypermatrices of rank not more than 4 requires that we first reduce it to a series of questions about matrices and then throw every tool in the matrix arsenal at them (Friedland 2013, Friedland and Gross 2012).

²⁰ In the sense of a system of calculation and reasoning.

The following resolves a potential point of confusion that the observant reader might have already noticed. In Definition 3.3, a d -tensor is a multilinear functional, that is,

$$\varphi: \mathbb{V}_1 \times \cdots \times \mathbb{V}_d \rightarrow \mathbb{R} \quad (4.18)$$

satisfying (3.18), but in Definition 4.4 a d -tensor is merely a multivariate function

$$f: X_1 \times \cdots \times X_d \rightarrow \mathbb{R} \quad (4.19)$$

that is not required to be multilinear, a requirement that in fact makes no sense as the sets X_1, \dots, X_d may not be vector spaces. One might think that (4.18) is a special case of (4.19) as a multilinear functional is a special case of a multivariate function, but this would be incorrect: they are different types of tensors. By Definition 3.3, φ is a covariant tensor whereas, as we pointed out in Example 4.2, f is a contravariant tensor. It suffices to explain this issue over finite dimensions and so we will assume that $\mathbb{V}_1, \dots, \mathbb{V}_d$ are finite-dimensional vector spaces and X_1, \dots, X_d are finite sets. We will use slightly different notation below to avoid having to introduce more indices.

Example 4.6 (multivariate versus multilinear). The relation between (4.18) and (4.19) is subtler than meets the eye. Up to a choice of bases, we can construct a unique φ from any f and *vice versa*. One direction is easy. Given vector spaces $\mathbb{U}, \mathbb{V}, \dots, \mathbb{W}$, choose any bases

$$\mathcal{A} = \{u_1, \dots, u_m\}, \quad \mathcal{B} = \{v_1, \dots, v_n\}, \dots, \mathcal{C} = \{w_1, \dots, w_p\}.$$

Then any multilinear functional $\varphi: \mathbb{U} \times \mathbb{V} \times \cdots \times \mathbb{W} \rightarrow \mathbb{R}$ gives us a real-valued function

$$f_\varphi: \mathcal{A} \times \mathcal{B} \times \cdots \times \mathcal{C} \rightarrow \mathbb{R}, \quad (u_i, v_j, \dots, w_k) \mapsto \varphi(u_i, v_j, \dots, w_k).$$

The other direction is a bit more involved. For any finite set X and a point $x \in X$, we define the delta function $\delta_x \in \mathbb{R}^X$ given by

$$\delta_x: X \rightarrow \mathbb{R}, \quad \delta_x(x') = \begin{cases} 1 & x = x', \\ 0 & x \neq x', \end{cases} \quad (4.20)$$

for all $x' \in X$, and the *point evaluation* linear functional $\varepsilon_x \in (\mathbb{R}^X)^*$ given by

$$\varepsilon_x: \mathbb{R}^X \rightarrow \mathbb{R}, \quad \varepsilon_x(f) = f(x)$$

for all $f \in \mathbb{R}^X$. These are bases of their respective spaces and for any set we have the following:

$$X = \{x_1, \dots, x_m\}, \quad \mathbb{R}^X = \text{span}\{\delta_{x_1}, \dots, \delta_{x_m}\}, \quad (\mathbb{R}^X)^* = \text{span}\{\varepsilon_{x_1}, \dots, \varepsilon_{x_m}\}.$$

Given any d sets,

$$X = \{x_1, \dots, x_m\}, \quad Y = \{y_1, \dots, y_n\}, \dots, Z = \{z_1, \dots, z_p\},$$

a real-valued function $f: X \times Y \times \dots \times Z \rightarrow \mathbb{R}$ gives us a multilinear functional

$$\varphi_f: \mathbb{R}^X \times \mathbb{R}^Y \times \dots \times \mathbb{R}^Z \rightarrow \mathbb{R}$$

defined by

$$\varphi_f = \sum_{i=1}^m \sum_{j=1}^n \cdots \sum_{k=1}^p f(x_i, y_j, \dots, z_k) \varepsilon_{x_i} \otimes \varepsilon_{y_j} \otimes \cdots \otimes \varepsilon_{z_k}.$$

Evaluating φ_f on basis vectors, we get

$$\varphi_f(\delta_{x_i}, \delta_{y_j}, \dots, \delta_{z_k}) = f(x_i, y_j, \dots, z_k),$$

taken together with the earlier definition of f_φ ,

$$f_\varphi(u_i, v_j, \dots, w_k) = \varphi(u_i, v_j, \dots, w_k),$$

and we obtain

$$\varphi_{f_\varphi} = \varphi, \quad f_{\varphi_f} = f.$$

Roughly speaking, a real-valued multivariate function is a tensor in the sense of Definition 3.3 because one may view it as a linear combination of point evaluation linear functionals. Such point evaluation functionals are the basis behind numerical quadrature (Example 4.48).

The next example is intended to serve as further elaboration for Examples 4.3, 4.5 and 4.6, and as partial impetus for Sections 4.2 and 4.3.

Example 4.7 (quantum spin and covalent bonds). As we saw in Example 4.3, the quantum state of a spinless particle like the Gaussian wave function $\psi_{m,n,p}$ is an element of the Hilbert space²¹ $L^2(\mathbb{R}^3)$. But even for a single particle, tensor products come in handy when we need to discuss particles with spin, which are crucial in chemistry. The simplest but also the most important case is when a quantum particle has spin $-\frac{1}{2}, \frac{1}{2}$, called a *spin-half particle* for short, and its quantum state is

$$\Psi \in L^2(\mathbb{R}^3) \otimes \mathbb{C}^2. \quad (4.21)$$

Going by Definition 4.4, this is a function $\Psi: \mathbb{R}^3 \times \{-\frac{1}{2}, \frac{1}{2}\} \rightarrow \mathbb{C}$, where $\Psi(x, \sigma)$ is an $L^2(\mathbb{R}^3)$ -integrable function in the argument x for each $\sigma = -\frac{1}{2}$ or $\sigma = \frac{1}{2}$. It is

²¹ To avoid differentiability and integrability issues, we need $\psi_{m,n,p}$ to be in the Schwartz space $S(\mathbb{R}^3) \subseteq L^2(\mathbb{R}^3)$ if we take a rigged Hilbert space approach (de la Madrid 2005) or the Sobolev space $H^2(\mathbb{R}^3) \subseteq L^2(\mathbb{R}^3)$ if we take an unbounded self-adjoint operators approach (Teschl 2014, Chapter 7), but we disregard these to keep things simple.

also not hard to see that when we have a tensor product of an infinite-dimensional space with a finite-dimensional one, every element must be a *finite* sum of separable terms, so Ψ in (4.21) may be expressed as

$$\Psi = \sum_{i=1}^r \psi_i \otimes \chi_i \quad \text{or} \quad \Psi(x, \sigma) = \sum_{i=1}^r \psi_i(x) \chi_i(\sigma) \quad (4.22)$$

with $\psi_i \in L^2(\mathbb{R}^3)$ and $\chi_i: \{-\frac{1}{2}, \frac{1}{2}\} \rightarrow \mathbb{C}$. In physics parlance, Ψ is the *total wave function* of the particle and ψ_i and χ_i are its spatial and spin wave functions; the variables that Ψ depends on are called *degrees of freedom*, with those describing position and momentum called *external* degrees of freedom and others like spin called *internal* degrees of freedom (Cohen-Tannoudji *et al.* 2020a, Chapter II, Section F).

This also illustrates why, as we discussed in Example 4.5, it is often desirable to view \mathbb{C}^n as the set of complex-valued functions on some finite sets of n elements. Here the finite set is $\{-\frac{1}{2}, \frac{1}{2}\}$ and \mathbb{C}^2 in (4.21) is just shorthand for the *spin state space*

$$\mathbb{C}^{\{-1/2, 1/2\}} = \left\{ \chi: \left\{ -\frac{1}{2}, \frac{1}{2} \right\} \rightarrow \mathbb{C} \right\}.$$

The domain $\mathbb{R}^3 \times \{-\frac{1}{2}, \frac{1}{2}\}$ is called the *position–spin space* (Pauli 1980) and plays a role in the Pauli exclusion principle. While two classical particles cannot simultaneously occupy the exact same location in \mathbb{R}^3 , two quantum particles can as long as they have different spins, as that means they are occupying different locations in $\mathbb{R}^3 \times \{-\frac{1}{2}, \frac{1}{2}\}$. A consequence is that two electrons with opposite spins can occupy the same molecular orbital, and when they do, we have a covalent bond in the molecule. The Pauli exclusion principle implies the converse: if two electrons occupy the same molecular orbital, then they must have opposite spins. We will see that this is a consequence of the antisymmetry of the total wave function.

In fact, ‘two electrons with opposite spins occupying the same molecular orbital’ is the quantum mechanical definition of a covalent bond in chemistry. When this happens, quantum mechanics mandates that we may not speak of each electron individually but need to consider both as a single entity described by a single wave function

$$\Psi \in (L^2(\mathbb{R}^3) \otimes \mathbb{C}^2) \widehat{\otimes} (L^2(\mathbb{R}^3) \otimes \mathbb{C}^2), \quad (4.23)$$

with each copy of $L^2(\mathbb{R}^3) \otimes \mathbb{C}^2$ associated with one of the electrons. The issue with writing the tensor product in the form (4.23) is that Ψ will not be a finite sum of separable terms since it is not a tensor product of two infinite-dimensional spaces. As we will see in Section 4.3, we may arrange the factors in a tensor product in arbitrary order and obtain isomorphic spaces

$$(L^2(\mathbb{R}^3) \otimes \mathbb{C}^2) \widehat{\otimes} (L^2(\mathbb{R}^3) \otimes \mathbb{C}^2) \cong L^2(\mathbb{R}^3) \widehat{\otimes} L^2(\mathbb{R}^3) \otimes \mathbb{C}^2 \otimes \mathbb{C}^2, \quad (4.24)$$

but this is also obvious from Definition 4.4 by observing that for finite sums

$$\sum_{i=1}^r \psi_i(x)\chi_i(\sigma)\varphi_i(y)\xi_i(\tau) = \sum_{i=1}^r \psi_i(x)\varphi_i(y)\chi_i(\sigma)\xi_i(\tau),$$

and then taking completion. Since $L^2(\mathbb{R}^3) \widehat{\otimes} L^2(\mathbb{R}^3) = L^2(\mathbb{R}^3 \times \mathbb{R}^3)$ by (4.13) and $\mathbb{C}^2 \otimes \mathbb{C}^2 = \mathbb{C}^{2 \times 2}$ by Example 4.5, again the latter is shorthand for

$$\mathbb{C}^{\{-1/2, 1/2\} \times \{-1/2, 1/2\}} = \left\{ \chi : \left\{ \left(-\frac{1}{2}, -\frac{1}{2} \right), \left(-\frac{1}{2}, \frac{1}{2} \right), \left(\frac{1}{2}, -\frac{1}{2} \right), \left(\frac{1}{2}, \frac{1}{2} \right) \right\} \rightarrow \mathbb{C} \right\},$$

we deduce that the state space in (4.24) is $L^2(\mathbb{R}^6) \otimes \mathbb{C}^{2 \times 2}$. The total wave function for two electrons may now be expressed as a finite sum of separable terms

$$\Psi(x, y, \sigma, \tau) = \sum_{i=1}^r \varphi_i(x, y)\chi_i(\sigma, \tau),$$

nicely separated into external degrees of freedom $(x, y) \in \mathbb{R}^3 \times \mathbb{R}^3$ and internal ones (σ, τ) as before. Henceforth we will write $\Psi(x, \sigma; y, \tau)$ if we view Ψ as an element of the space on the left of (4.24) and $\Psi(x, y, \sigma, \tau)$ if we view it as an element of the space on the right. The former emphasizes that Ψ is a wave function of two electrons, one with degrees of freedom (x, σ) and the other (y, τ) ; the latter emphasizes the separation of the wave function into external and internal degrees of freedom.

Electrons are *fermions*, which are quantum particles with antisymmetric wave functions, that is,

$$\Psi(x, \sigma; y, \tau) = -\Psi(y, \tau; x, \sigma). \tag{4.25}$$

For simplicity, suppose we have a separable wave function

$$\Psi(x, \sigma; y, \tau) = \Psi(x, y, \sigma, \tau) = \varphi(x, y)\chi(\sigma, \tau).$$

Aside from their spins, identical electrons in the same orbital are indistinguishable and we must have $\varphi(x, y) = \varphi(y, x)$, so the antisymmetry (4.25) must be a consequence of $\chi(\sigma, \tau) = -\chi(\tau, \sigma)$. As in Example 4.6, a basis for \mathbb{C}^2 is given by the delta functions $\delta_{-1/2}$ and $\delta_{1/2}$ and thus a basis for $\mathbb{C}^{2 \times 2}$ is given by

$$\delta_{-1/2} \otimes \delta_{-1/2}, \quad \delta_{-1/2} \otimes \delta_{1/2}, \quad \delta_{1/2} \otimes \delta_{-1/2}, \quad \delta_{1/2} \otimes \delta_{1/2}.$$

We may thus express

$$\chi = a\delta_{-1/2} \otimes \delta_{-1/2} + b\delta_{-1/2} \otimes \delta_{1/2} + c\delta_{1/2} \otimes \delta_{-1/2} + d\delta_{1/2} \otimes \delta_{1/2},$$

and if $\chi(\sigma, \tau) = -\chi(\tau, \sigma)$ for all σ, τ , then it implies that $a = d = 0$ and $b = -c$. Choosing $b = 1/\sqrt{2}$ to get unit norms, we have

$$\chi = \frac{1}{\sqrt{2}}(\delta_{-1/2} \otimes \delta_{1/2} - \delta_{1/2} \otimes \delta_{-1/2}).$$

In other words, the two electrons in the same molecular orbital must have opposite spins – either $(-\frac{1}{2}, \frac{1}{2})$ or $(\frac{1}{2}, -\frac{1}{2})$.

The discussion above may be extended to more particles ($d > 2$) and more spin values ($s > \frac{1}{2}$). When we have d quantum particles with spins $-\frac{1}{2}, \frac{1}{2}$, the total wave function satisfies

$$\Psi \in (L^2(\mathbb{R}^3) \otimes \mathbb{C}^2) \widehat{\otimes} \dots \widehat{\otimes} (L^2(\mathbb{R}^3) \otimes \mathbb{C}^2) \cong L^2(\mathbb{R}^{3d}) \otimes \mathbb{C}^{2 \times \dots \times 2} \tag{4.26}$$

via the same argument. Again we may group the degrees of freedom either by particles or by external/internal, that is,

$$\Psi(x_1, \sigma_1; x_2, \sigma_2; \dots; x_d, \sigma_d) = \Psi(x_1, x_2, \dots, x_d, \sigma_1, \sigma_2, \dots, \sigma_d),$$

depending on whether we want to view Ψ as an element of the left or right space in (4.26), with antisymmetry usually expressed in the former,

$$\Psi(x_1, \sigma_1; \dots; x_d, \sigma_d) = (-1)^{\text{sgn}(\pi)} \Psi(x_{\pi(1)}, \sigma_{\pi(1)}; \dots; x_{\pi(d)}, \sigma_{\pi(d)}) \tag{4.27}$$

for all $\pi \in \mathfrak{S}_d$, and separability usually expressed in the latter,

$$\Psi(x_1, \dots, x_d, \sigma_1, \dots, \sigma_d) = \sum_{i=1}^r \psi_i(x_1, \dots, x_d) \chi_i(\sigma_1, \dots, \sigma_d).$$

As before, $\mathbb{C}^{2 \times \dots \times 2}$ is to be regarded as

$$\mathbb{C}^{\{-1/2, 1/2\} \times \dots \times \{-1/2, 1/2\}} = \left\{ \chi: \left\{ -\frac{1}{2}, \frac{1}{2} \right\}^n \rightarrow \mathbb{C} \right\}.$$

We may easily extend our discussions to a spin- s quantum particle for any

$$s \in \left\{ 0, \frac{1}{2}, 1, \frac{3}{2}, 2, \frac{5}{2}, \dots \right\},$$

in which case its state space would be

$$L^2(\mathbb{R}^3) \otimes \mathbb{C}^{2s+1},$$

where \mathbb{C}^{2s+1} is shorthand for the space of functions $\chi: \{-s, -s+1, \dots, s-1, s\} \rightarrow \mathbb{C}$. The particle is called a fermion or a *boson* depending on whether s is a half-integer or an integer. The key difference is that the total wave function of d identical fermions is antisymmetric as in (4.27) whereas that of d identical bosons is *symmetric*, that is,

$$\Psi(x_1, \sigma_1; \dots; x_d, \sigma_d) = \Psi(x_{\pi(1)}, \sigma_{\pi(1)}; \dots; x_{\pi(d)}, \sigma_{\pi(d)}) \tag{4.28}$$

for all $\pi \in \mathfrak{S}_d$. Collectively, (4.27) and (4.28) are the modern form of the Pauli exclusion principle. Even more generally, we may consider d quantum particles with different spins s_1, \dots, s_d and state space

$$(L^2(\mathbb{R}^3) \otimes \mathbb{C}^{2s_1+1}) \widehat{\otimes} \dots \widehat{\otimes} (L^2(\mathbb{R}^3) \otimes \mathbb{C}^{2s_d+1}) \cong L^2(\mathbb{R}^{3d}) \otimes \mathbb{C}^{(2s_1+1) \times \dots \times (2s_d+1)}.$$

We recommend the very lucid exposition in [Faddeev and Yakubovskii \(2009\)](#), Sections 46–48) and [Takhtajan \(2008, Chapter 4\)](#) for further information.

For most purposes in applied and computational mathematics, $s = 1/2$, that is, the spin-half case is literally the only one that matters, as particles that constitute ordinary matter (electrons, protons, neutrons, quarks, *etc.*) are all spin-half fermions. Spin-zero (*e.g.* Higgs) and spin-one bosons (*e.g.* W and Z, photons, gluons) are already the realm of particle physics, and higher-spin particles, whether fermions or bosons, tend to be hypothetical (*e.g.* gravitons, gravitinos).

In [Example 4.7](#) we have implicitly relied on the following three background assumptions for a physical system in quantum mechanics.

- (a) The state space of a system is a Hilbert space.
- (b) The state space of a composite system is the tensor product of the state spaces of the component systems.
- (c) The state space of a composite system of identical particles lies either in the symmetric or alternating tensor product of the state spaces of the component systems.

These assumptions are standard; see [Cohen-Tannoudji *et al.* \(2020a\)](#), Chapter III, Section B1), [Nielsen and Chuang \(2000\)](#), Section 2.2.8) and [Cohen-Tannoudji, Diu and Laloë \(2020b\)](#), Chapter XIV, Section C).

[Example 4.7](#) shows the simplicity of [Definition 4.4](#). If our quantum system has yet other degrees of freedom, say, flavour or colour, we just include the corresponding wave functions as factors in the separable product. For instance, the total wave function of a proton or neutron is a finite sum of separable terms of the form ([Griffiths 2008](#), Section 5.6.1)

$$\psi_{\text{spatial}} \otimes \psi_{\text{spin}} \otimes \psi_{\text{flavour}} \otimes \psi_{\text{colour}}.$$

Nevertheless, we would like more flexibility than [Definition 4.4](#) provides. For instance, instead of writing the state of a spin-half particle as a sum of real-valued functions in [\(4.22\)](#), we might prefer to view it as a \mathbb{C}^2 -valued L^2 -vector field on \mathbb{R}^3 indexed by spin:

$$\Psi(x) = \begin{bmatrix} \psi_{-1/2}(x) \\ \psi_{1/2}(x) \end{bmatrix} \quad (4.29)$$

with

$$\|\Psi\|^2 = \int_{\mathbb{R}^3} |\psi_{-1/2}(x)|^2 dx + \int_{\mathbb{R}^3} |\psi_{1/2}(x)|^2 dx < \infty.$$

In other words, we want to allow $L^2(\mathbb{R}^3; \mathbb{C}^2)$ to be a possible interpretation for $L^2(\mathbb{R}^3) \otimes \mathbb{C}^2$. The next two sections will show how to accomplish this.

4.2. Tensor products of abstract vector spaces

Most readers' first encounter with scalars and vectors will likely take the following forms:

- a scalar is an object that only has magnitude,
- a vector is an object with both magnitude and direction.

In this section we will construct tensors of arbitrary order starting from these. For a vector v in this sense, we will denote its magnitude by $\|v\|$ and its direction by \widehat{v} . Note that here $\|\cdot\|$ does not need to be a norm, just some notion of magnitude. Likewise we will write $|a| \in [0, \infty)$ for the magnitude of a scalar a and $\text{sgn}(a)$ for its sign if it has one.

If we go along this line of reasoning, what should the next object be? One possible answer is that it ought to be an object with a magnitude and *two* directions; this is called a *dyad* and we have in fact encountered it in Section 3.1. Henceforth it is straightforward to generalize to objects with a magnitude and an arbitrary number of directions:

OBJECT	PROPERTY
scalar a	has magnitude $ a $
vector v	has magnitude $\ v\ $ and a direction \widehat{v}
dyad $v \otimes w$	has magnitude $\ v \otimes w\ $ and two directions \widehat{v}, \widehat{w}
triad $u \otimes v \otimes w$	has magnitude $\ u \otimes v \otimes w\ $ and three directions $\widehat{u}, \widehat{v}, \widehat{w}$
\vdots	\vdots
d -ad $u \otimes v \otimes \cdots \otimes w$	has magnitude $\ u \otimes v \otimes \cdots \otimes w\ $ and d directions $\widehat{u}, \widehat{v}, \dots, \widehat{w}$.

Such objects are called *polyads* and have been around for over a century; the first definition of tensor rank was in fact in terms of polyads (Hitchcock 1927). We have slightly modified the notation by adding the tensor product symbol \otimes between vectors to bring it in line with modern treatments in algebra, which we will describe later. Note that the \otimes in $v \otimes w$ is solely used as a delimiter; we could have written it as vw as in Hitchcock (1927) and Morse and Feshbach (1953) or $|v\rangle|w\rangle$ or $|v, w\rangle$ in Dirac notation (Cohen-Tannoudji *et al.* 2020a, Chapter II, Section F2c). Indeed, v and w may not have coordinates and there is no 'formula for $v \otimes w$ '. Also, the order matters as v and w may denote vectors of different nature, so in general

$$v \otimes w \neq w \otimes v. \quad (4.30)$$

The reason for this would be clarified with Example 4.8. Used in this sense, \otimes is called an *abstract tensor product* or *dyadic product* in the older literature (Chou and Pagano 1992, Morse and Feshbach 1953, Tai 1997).

With hindsight, we see that the extension of a scalar in the sense of an object with magnitude, and a vector in the sense of one with magnitude and direction, is

not a tensor but a rank-one tensor. To get tensors that are not rank-one we will have to look at the arithmetic of dyads and polyads. As we have discussed, each dyad is essentially a placeholder for three pieces of information:

(magnitude, first direction, second direction)

or, in our notation ($\|v \otimes w\|, \widehat{v}, \widehat{w}$). We are trying to devise a consistent system of arithmetic for such objects that (a) preserves their information content, (b) merges information whenever possible, and (c) is consistent with scalar and vector arithmetic.

Scalar products are easy. The role of scalars is that they scale vectors, as reflected in its name. We expect the same for dyads. If any of the vectors in a dyad is scaled by a , we expect its magnitude to be scaled by a but the two directions to remain unchanged. So we require

$$(av) \otimes w = v \otimes (aw) =: av \otimes w, \quad (4.31)$$

where the last term is defined to be the common value of the first two. This seemingly innocuous property is the main reason tensor products, not direct sums, are used to combine quantum state spaces. In quantum mechanics, a quantum state is described not so much by a vector v but the entire one-dimensional subspace spanned by v ; thus (4.31) ensures that when combining two quantum states in the form of two one-dimensional subspaces, it does not matter which (non-zero) vector in the subspace we pick to represent the state. On the other hand, for direct sums,

$$(av) \oplus w \neq v \oplus (aw) \neq a(v \oplus w) \quad (4.32)$$

in general; for example, with $v = (1, 0) \in \mathbb{R}^2$, $w = 1 \in \mathbb{R}^1$ and $a = 2 \in \mathbb{R}$, we get $(2, 0, 1) \neq (1, 0, 2) \neq (2, 0, 2)$.

To ensure consistency with the usual scalar and vector arithmetic, one also needs the assumption that scalar multiplication is always distributive and associative in the following sense:

$$(a + b)v \otimes w = av \otimes w + bv \otimes w, \quad (ab)v \otimes w = a(bv \otimes w).$$

Addition of dyads is trickier. If a_1, a_2 are scalars, then $a_1 + a_2$ is a scalar; if v_1, v_2 are vectors, then $v_1 + v_2$ is a vector according to either the parallelogram law of addition (for geometrical or physical vectors) or the axioms of vector spaces (for abstract vectors). But for dyads $v_1 \otimes w_1$ and $v_2 \otimes w_2$, their sum $v_1 \otimes w_1 + v_2 \otimes w_2$ is no longer a dyad; in general there is no way to simplify the sum further and it should be regarded as an object with two magnitudes and four directions:²²

(magnitude 1, first direction 1, second direction 1)

& (magnitude 2, first direction 2, second direction 2).

²² We could of course add corresponding vectors $v_1 + v_2$ and $w_1 + w_2$, but that leads to direct sums of vector spaces, which is inappropriate because of (4.32).

Just as the \otimes is used as a delimiter where order matters, the $+$ in the sum of dyads is used as a different type of delimiter where order does not matter:

$$v_1 \otimes w_1 + v_2 \otimes w_2 = v_2 \otimes w_2 + v_1 \otimes w_1,$$

that is, $+$ is commutative. We also want $+$ to be associative so that we may unambiguously define, for any finite $r \in \mathbb{N}$, a sum of r dyads

$$v_1 \otimes w_1 + v_2 \otimes w_2 + \cdots + v_r \otimes w_r$$

that we will call a *dyadic*. To ensure consistency with the usual scalar and vector arithmetic, we assume throughout this article that addition $+$ is always commutative and associative. This assumption is also perfectly reasonable as a dyadic is just a placeholder for the information contained in its dyad summands:

$$\begin{aligned} &(\text{magnitude } 1, \text{ first direction } 1, \text{ second direction } 1) \\ &\& (\text{magnitude } 2, \text{ first direction } 2, \text{ second direction } 2) \& \cdots \\ &\& (\text{magnitude } r, \text{ first direction } r, \text{ second direction } r). \end{aligned}$$

There is one scenario where a sum can be simplified and the directional information in two dyads merged. Whenever $v_1 = v_2$ or $w_1 = w_2$, we want

$$\begin{aligned} v \otimes w_1 + v \otimes w_2 &= v \otimes (w_1 + w_2), \\ v_1 \otimes w + v_2 \otimes w &= (v_1 + v_2) \otimes w, \end{aligned} \tag{4.33}$$

and these may be applied to combine any pair of summands in a dyadic that share a common factor. The following example provides further motivation for why we would like these arithmetical properties to hold.

Example 4.8 (stress tensor). In physics, a dyadic is essentially a vector whose components are themselves vectors. The standard example is the *stress* tensor, sometimes called the Cauchy stress tensor to distinguish it from other related models. The stress σ at a point in space²³ has three components in the directions given by unit vectors e_x, e_y, e_z :

$$\sigma = \sigma_x e_x + \sigma_y e_y + \sigma_z e_z. \tag{4.34}$$

Now the deviation from a usual linear combination of vectors is that each of the coefficients $\sigma_x, \sigma_y, \sigma_z$ is not a scalar but a vector. This is the nature of stress: in every direction, the stress in that direction has a normal component in that direction and two shear components in the plane perpendicular to it. For example,

$$\sigma_x = \sigma_{xx} e_x + \sigma_{yx} e_y + \sigma_{zx} e_z, \tag{4.35}$$

the component σ_{xx} in the direction of e_x is called *normal stress* whereas the two components σ_{yx} and σ_{zx} in the plane perpendicular to e_x , *i.e.* $\text{span}\{e_y, e_z\}$, is

²³ We assume ‘space’ here means \mathbb{R}^3 with coordinate axes labelled x, y, z . The older literature would use i, j, k instead of e_x, e_y, e_z .




Figure 4.2. Depiction of the stress tensor. The cube represents a single point in \mathbb{R}^3 and should be regarded as infinitesimally small. In particular, the origins of the three coordinate frames on the surface of the cube are the same point.

called *shear stress*. Here $\sigma_{xx}, \sigma_{yx}, \sigma_{zx}$ are scalars and (4.35) is an honest linear combination. It is easiest to represent σ pictorially, as in Borg (1990, Figure 4.10) or (Chou and Pagano 1992, Figure 1.5), adapted here as Figure 4.2, where the point is blown up into a cube to show the fine details. The coefficients $\sigma_x, \sigma_y, \sigma_z$ are called *stress vectors* in the directions e_x, e_y, e_z respectively. As Figure 4.2 indicates, each stress vector has a normal component pointing in the same direction as itself and two shear components in the plane perpendicular to it:

$$\begin{aligned} \sigma_y &= \sigma_{xy}e_x + \sigma_{yy}e_y + \sigma_{zy}e_z, \\ \sigma_z &= \sigma_{xz}e_x + \sigma_{yz}e_y + \sigma_{zz}e_z. \end{aligned}$$

Since the ‘coefficients’ $\sigma_x, \sigma_y, \sigma_z$ in the ‘linear combination’ in (4.34) are vectors, to conform to modern notation we just insert \otimes :

$$\sigma = \sigma_x \otimes e_x + \sigma_y \otimes e_y + \sigma_z \otimes e_z.$$

If we then plug the expressions for $\sigma_x, \sigma_y, \sigma_z$ into that of σ and use (4.33), we get

$$\begin{aligned} \sigma &= \sigma_{xx}e_x \otimes e_x + \sigma_{yx}e_y \otimes e_x + \sigma_{zx}e_z \otimes e_x \\ &\quad + \sigma_{xy}e_x \otimes e_y + \sigma_{yy}e_y \otimes e_y + \sigma_{zy}e_z \otimes e_y \\ &\quad + \sigma_{xz}e_x \otimes e_z + \sigma_{yz}e_y \otimes e_z + \sigma_{zz}e_z \otimes e_z. \end{aligned} \tag{4.36}$$

Without the \otimes , this is the expression of stress as a dyadic in Morse and Feshbach (1953, pp. 70–71). This also explains why we want (4.30), that is, \otimes should be non-commutative. For instance, $e_z \otimes e_y$, the z -shear component for the y -normal

direction, and $e_z \otimes e_y$, the y -shear component for the z -normal direction mean completely different things.

If we set a basis $\mathcal{B} = \{e_x, e_y, e_z\}$, we may represent σ as a 3×3 matrix, that is,

$$\Sigma := [\sigma]_{\mathcal{B}} = \begin{bmatrix} \sigma_{xx} & \sigma_{xy} & \sigma_{xz} \\ \sigma_{yx} & \sigma_{yy} & \sigma_{yz} \\ \sigma_{zx} & \sigma_{zy} & \sigma_{zz} \end{bmatrix}$$

with normal stresses on the diagonal and shear stresses on the off-diagonal. If we have a different basis $\mathcal{B}' = \{e'_x, e'_y, e'_z\}$ with a different representation

$$\Sigma' := [\sigma]_{\mathcal{B}'} = \begin{bmatrix} \sigma'_{xx} & \sigma'_{xy} & \sigma'_{xz} \\ \sigma'_{yx} & \sigma'_{yy} & \sigma'_{yz} \\ \sigma'_{zx} & \sigma'_{zy} & \sigma'_{zz} \end{bmatrix},$$

then by plugging the change-of-basis relations

$$\begin{cases} e'_x = c_{xx}e_x + c_{yx}e_y + c_{zx}e_z, \\ e'_y = c_{xy}e_x + c_{yy}e_y + c_{zy}e_z, \\ e'_z = c_{xz}e_x + c_{yz}e_y + c_{zz}e_z, \end{cases} \quad C := \begin{bmatrix} c_{xx} & c_{xy} & c_{xz} \\ c_{yx} & c_{yy} & c_{yz} \\ c_{zx} & c_{zy} & c_{zz} \end{bmatrix}$$

into

$$\begin{aligned} \sigma &= \sigma'_{xx}e'_x \otimes e'_x + \sigma'_{yx}e'_y \otimes e'_x + \sigma'_{zx}e'_z \otimes e'_x \\ &\quad + \sigma'_{xy}e'_x \otimes e'_y + \sigma'_{yy}e_y \otimes e'_y + \sigma'_{zy}e'_z \otimes e'_y \\ &\quad + \sigma'_{xz}e'_x \otimes e'_z + \sigma'_{yz}e'_y \otimes e'_z + \sigma'_{zz}e'_z \otimes e'_z, \end{aligned}$$

applying (4.33) and comparing with (4.36), we get

$$\begin{bmatrix} \sigma_{xx} & \sigma_{xy} & \sigma_{xz} \\ \sigma_{yx} & \sigma_{yy} & \sigma_{yz} \\ \sigma_{zx} & \sigma_{zy} & \sigma_{zz} \end{bmatrix} = \begin{bmatrix} c_{xx} & c_{xy} & c_{xz} \\ c_{yx} & c_{yy} & c_{yz} \\ c_{zx} & c_{zy} & c_{zz} \end{bmatrix} \begin{bmatrix} \sigma'_{xx} & \sigma'_{xy} & \sigma'_{xz} \\ \sigma'_{yx} & \sigma'_{yy} & \sigma'_{yz} \\ \sigma'_{zx} & \sigma'_{zy} & \sigma'_{zz} \end{bmatrix} \begin{bmatrix} c_{xx} & c_{yx} & c_{zx} \\ c_{xy} & c_{yy} & c_{zy} \\ c_{xz} & c_{yz} & c_{zz} \end{bmatrix}.$$

By the table on page 574, the fact that any two coordinate representations of Cauchy stress σ satisfy the transformation rule $\Sigma' = C^{-1}\Sigma C^{-T}$ says that it is a contravariant 2-tensor. This provides yet another example why we do not want to identify a 2-tensor with the matrix that represents it: stripped of the basis vectors, the physical connotations of (4.34) and (4.35) that we discussed earlier are irretrievably lost.

What we have described above is the behaviour of stress at a single point in space, assumed to be \mathbb{R}^3 , and in Cartesian coordinates. As we mentioned in Example 3.13, stress is a tensor field and e_x, e_y, e_z really form a basis of the tangent space $\mathbb{T}_v(\mathbb{R}^3)$ at the point $v = (x, y, z) \in \mathbb{R}^3$, and we could have used any coordinates. Note that our discussion above merely used stress in a purely nominal way: our focus is on the dyadic that represents stress. The same discussion applies *mutatis mutandis* to any contravariant 2-tensors describing inertia, polarization, strain, tidal force, viscosity, etc. (Borg 1990). We refer the readers to Borg (1990, Chapter 4), Chou

and Pagano (1992, Chapter 1) and Irgens (2019, Chapter 2) for the actual physical details regarding stress, including how to derive (4.34) from first principles.

Stress is a particularly vital notion, as many tensors describing physical properties in various constitutive equations (Hartmann 1984, Table 1) are defined as derivatives with respect to it. In fact, this is one way in which tensors of higher order arise in physics – as *second*-order derivatives of real-valued functions with respect to first- and second-order tensors. For instance, in Cartesian coordinates, the piezoelectric tensor, the piezo-magnetic tensor and the elastic tensor are represented as hypermatrices $D, Q \in \mathbb{R}^{3 \times 3 \times 3}$ and $S \in \mathbb{R}^{3 \times 3 \times 3 \times 3}$, where

$$d_{ijk} = -\frac{\partial^2 G}{\partial \sigma_{ij} \partial e_k}, \quad q_{ijk} = -\frac{\partial^2 G}{\partial \sigma_{ij} \partial h_k}, \quad s_{ijkl} = -\frac{\partial^2 G}{\partial \sigma_{ij} \partial \sigma_{kl}},$$

and $i, j, k, l = 1, 2, 3$. Here $G = G(\Sigma, E, H, T)$ is the *Gibbs potential*, a real-valued function of the second-order stress tensor Σ , the first-order tensors E and H representing electric and magnetic field respectively, and the zeroth-order tensor T representing temperature (Hartmann 1984, Section 3).

It is straightforward to extend the above construction of a dyadic to arbitrary dimensions and arbitrary vectors in abstract vector spaces \mathbb{U} and \mathbb{V} . A dyadic is a ‘linear combination’ of vectors in \mathbb{V} ,

$$\sigma = \sigma_1 v_1 + \cdots + \sigma_n v_n,$$

whose coefficients σ_j are vectors in \mathbb{U} :

$$\sigma_j = \sigma_{1j} u_1 + \cdots + \sigma_{mj} u_j, \quad j = 1, \dots, n.$$

Thus by (4.31) and (4.33) we have

$$\sigma = \sum_{i=1}^m \sum_{j=1}^n \sigma_{ij} u_i v_j = \sum_{i=1}^m \sum_{j=1}^n \sigma_{ij} u_i \otimes v_j$$

in old (pre- \otimes) and modern notation respectively. Note that the coefficients σ_{ij} are now scalars and a dyadic is an honest linear combination of dyads with scalar coefficients. We denote the set of all such dyadics as $\mathbb{U} \otimes \mathbb{V}$.

We may recursively extend the construction to arbitrary order. With a third vector space \mathbb{W} , a *tridic* is a ‘linear combination’ of vectors in \mathbb{W} ,

$$\tau = \tau_1 w_1 + \cdots + \tau_p w_p,$$

whose coefficients τ_k are dyadics in $\mathbb{U} \otimes \mathbb{V}$:

$$\tau_k = \sum_{i=1}^m \sum_{j=1}^n \tau_{ijk} u_i v_j = \sum_{i=1}^m \sum_{j=1}^n \tau_{ijk} u_i \otimes v_j, \quad k = 1, \dots, p.$$

Thus we have

$$\tau = \sum_{i=1}^m \sum_{j=1}^n \sum_{k=1}^p \tau_{ijk} u_i v_j w_k = \sum_{i=1}^m \sum_{j=1}^n \sum_{k=1}^p \tau_{ijk} u_i \otimes v_j \otimes w_k$$

in old and modern notation respectively. Henceforth we will ditch the old notation. Again a triadic is an honest linear combination of triads and the set of all triadics will be denoted $\mathbb{U} \otimes \mathbb{V} \otimes \mathbb{W}$. Strictly speaking, we have constructed $(\mathbb{U} \otimes \mathbb{V}) \otimes \mathbb{W}$ and an analogous construction considering ‘linear combinations’ of dyadics in $\mathbb{V} \otimes \mathbb{W}$ with coefficients in \mathbb{U} would give us $\mathbb{U} \otimes (\mathbb{V} \otimes \mathbb{W})$, but it does not matter for us as we have implicitly imposed that

$$(u \otimes v) \otimes w = u \otimes (v \otimes w) =: u \otimes v \otimes w \quad (4.37)$$

for all $u \in \mathbb{U}$, $v \in \mathbb{V}$, $w \in \mathbb{W}$.

The observant reader may have noticed that aside from (4.37) we have also implicitly imposed the third-order analogues of (4.31) and (4.33) in the construction above. For completeness we will state them formally but in a combined form:

$$\begin{aligned} (\lambda u + \lambda' u') \otimes v \otimes w &= \lambda u \otimes v \otimes w + \lambda' u' \otimes v \otimes w, \\ u \otimes (\lambda v + \lambda' v') \otimes w &= \lambda u \otimes v \otimes w + \lambda' u \otimes v' \otimes w, \\ u \otimes v \otimes (\lambda w + \lambda' w') &= \lambda u \otimes v \otimes w + \lambda' u \otimes v \otimes w' \end{aligned} \quad (4.38)$$

for all vectors $u, u' \in \mathbb{U}$, $v, v' \in \mathbb{V}$, $w, w' \in \mathbb{W}$ and scalars λ, λ' .

The construction described above makes the set $\mathbb{U} \otimes \mathbb{V} \otimes \mathbb{W}$ into an algebraic object with scalar product, tensor addition $+$ and tensor product \otimes interacting in a consistent manner according to the algebraic rules (4.37) and (4.38). We will call it a *tensor product* of \mathbb{U} , \mathbb{V} , \mathbb{W} , or an *abstract tensor product* if we need to emphasize that it refers to this particular construction. Observe that this is an extremely general construction.

- (i) It does not depend what we call ‘scalars’ so long as we may add and multiply them, that is, this construction works for any arbitrary modules \mathbb{U} , \mathbb{V} , \mathbb{W} over a ring R .
- (ii) It does not require \mathbb{U} , \mathbb{V} , \mathbb{W} to be finite-dimensional (vector spaces) or finitely generated (modules); the whole construction is about specifying how linear combinations behave under \otimes , and there is no need for elements to be linear combinations of some basis or generating set.
- (iii) It does not call for separate treatments of covariant and mixed tensors; the definition is agnostic to having some or all of \mathbb{U} , \mathbb{V} , \mathbb{W} replaced by their dual spaces \mathbb{U}^* , \mathbb{V}^* , \mathbb{W}^* .

It almost goes without saying that the construction can be readily extended to arbitrary order d . We may consider ‘linear combinations’ of dyadics with dyadic coefficients to get $d = 4$, ‘linear combinations’ of triadics with dyadic coefficients to get $d = 5$, etc. But looking at the end results of the $d = 2$ and $d = 3$ constructions,

if we have d vector spaces $\mathbb{U}, \mathbb{V}, \dots, \mathbb{W}$, a more direct way is to simply consider the set of all d -adics, *i.e.* finite sums of d -ads,

$$\mathbb{U} \otimes \mathbb{V} \otimes \dots \otimes \mathbb{W} := \left\{ \sum_{i=1}^r u_i \otimes v_i \otimes \dots \otimes w_i : u_i \in \mathbb{U}, v_i \in \mathbb{V}, \dots, w_i \in \mathbb{W}, r \in \mathbb{N} \right\}, \tag{4.39}$$

and decree that \otimes is associative, $+$ is associative and commutative, and \otimes is distributive over $+$ in the sense of

$$\begin{aligned} (\lambda u + \lambda' u') \otimes v \otimes \dots \otimes w &= \lambda u \otimes v \otimes \dots \otimes w + \lambda' u' \otimes v \otimes \dots \otimes w, \\ u \otimes (\lambda v + \lambda' v') \otimes \dots \otimes w &= \lambda u \otimes v \otimes \dots \otimes w + \lambda' u \otimes v' \otimes \dots \otimes w, \\ &\vdots \\ u \otimes v \otimes \dots \otimes (\lambda w + \lambda' w') &= \lambda u \otimes v \otimes \dots \otimes w + \lambda' u \otimes v \otimes \dots \otimes w' \end{aligned} \tag{4.40}$$

for all vectors $u, u' \in \mathbb{U}, v, v' \in \mathbb{V}, \dots, w, w' \in \mathbb{W}$, and scalars λ, λ' .

We have used the old polyadic terminology to connect our discussions to the older engineering and physics literature, and to show that the modern abstract construction²⁴ is firmly rooted in concrete physical considerations. We will henceforth ditch the obsolete names d -ads and d -adics and use their modern names, rank-one d -tensors and d -tensors. Note that these arithmetic rules also allow us to define, for any $T \in \mathbb{U} \otimes \mathbb{V} \otimes \dots \otimes \mathbb{W}$ and any $T' \in \mathbb{U}' \otimes \mathbb{V}' \otimes \dots \otimes \mathbb{W}'$, an element

$$T \otimes T' \in \mathbb{U} \otimes \mathbb{V} \otimes \dots \otimes \mathbb{W} \otimes \mathbb{U}' \otimes \mathbb{V}' \otimes \dots \otimes \mathbb{W}'. \tag{4.41}$$

We emphasize that $\mathbb{U} \otimes \mathbb{V} \otimes \dots \otimes \mathbb{W}$ does not mean the set of rank-one tensors but the set of finite sums of them as in (4.39). This mistake is common enough to warrant a displayed equation to highlight the pitfall:

$$\mathbb{U} \otimes \mathbb{V} \otimes \dots \otimes \mathbb{W} \neq \{u \otimes v \otimes \dots \otimes w : u \in \mathbb{U}, v \in \mathbb{V}, \dots, w \in \mathbb{W}\}, \tag{4.42}$$

unless all but one of $\mathbb{U}, \mathbb{V}, \dots, \mathbb{W}$ are one-dimensional.

For easy reference, we state a special case of the above abstract tensor product construction, but where we have both vector spaces and dual spaces, as another definition for tensors.

Definition 4.9 (tensors as elements of tensor spaces). Let $p \leq d$ be non-negative integers and let $\mathbb{V}_1, \dots, \mathbb{V}_d$ be vector spaces. A *tensor of contravariant order p and covariant order $d - p$* is an element

$$T \in \mathbb{V}_1 \otimes \dots \otimes \mathbb{V}_p \otimes \mathbb{V}_{p+1}^* \otimes \dots \otimes \mathbb{V}_d^*.$$

²⁴ This construction is equivalent to the one in Lang (2002, Chapter XVI, Section 1), likely adapted from Bourbaki (1998, Chapter II, Section 3), and widely used in pure mathematics. Take the free module generated by all rank-one elements $u \otimes v \otimes \dots \otimes w$ and quotient it by the submodule generated by elements of the forms $(\lambda u + \lambda' u') \otimes v \otimes \dots \otimes w - \lambda u \otimes v \otimes \dots \otimes w - \lambda' u' \otimes v \otimes \dots \otimes w$, *etc.*, to ensure that (4.40) holds.

The set above is as defined in (4.39) and is called a *tensor product* of vector spaces or a *tensor space* for short. The tensor T is said to be of *type* $(p, d - p)$ and of *order* d . A tensor of type $(d, 0)$, i.e. $T \in \mathbb{V}_1 \otimes \cdots \otimes \mathbb{V}_d$, is called a contravariant d -tensor, and a tensor of type $(0, d)$, i.e. $T \in \mathbb{V}_1^* \otimes \cdots \otimes \mathbb{V}_d^*$ is called a covariant d -tensor.

Vector spaces have bases and a tensor space $\mathbb{V}_1 \otimes \cdots \otimes \mathbb{V}_d$ has a *tensor product basis* given by

$$\mathcal{B}_1 \otimes \mathcal{B}_2 \otimes \cdots \otimes \mathcal{B}_d := \{u \otimes v \otimes \cdots \otimes w : u \in \mathcal{B}_1, v \in \mathcal{B}_2, \dots, w \in \mathcal{B}_d\}, \quad (4.43)$$

where \mathcal{B}_i is any basis of \mathbb{V}_i , $i = 1, \dots, d$. This is the *only* way to get a basis of rank-one tensors on a tensor space. If $\mathcal{B}_1 \otimes \cdots \otimes \mathcal{B}_d$ is a basis for $\mathbb{V}_1 \otimes \cdots \otimes \mathbb{V}_d$, then $\mathcal{B}_1, \dots, \mathcal{B}_d$ must be bases of $\mathbb{V}_1, \dots, \mathbb{V}_d$, respectively. More generally, for mixed tensor spaces, $\mathcal{B}_1 \otimes \cdots \otimes \mathcal{B}_p \otimes \mathcal{B}_{p+1}^* \otimes \cdots \otimes \mathcal{B}_d^*$ is a basis for $\mathbb{V}_1 \otimes \cdots \otimes \mathbb{V}_p \otimes \mathbb{V}_{p+1}^* \otimes \cdots \otimes \mathbb{V}_d^*$, where \mathcal{B}^* denotes the dual basis of \mathcal{B} . There is no need to assume finite dimension; these bases may well be uncountable. If the vector spaces are finite-dimensional, then (4.43) implies that

$$\dim(\mathbb{V}_1 \otimes \cdots \otimes \mathbb{V}_d) = \dim(\mathbb{V}_1) \cdots \dim(\mathbb{V}_d). \quad (4.44)$$

Aside from (4.31), relation (4.44) is another reason the tensor product is the proper operation for combining quantum states of different systems: if a quantum system is made up of superpositions (i.e. linear combinations) of m distinct (i.e. linearly independent) states and another is made up of superpositions of n distinct states, then it is physically reasonable for the combined system to be made up of superpositions of mn distinct states (Nielsen and Chuang 2000, p. 94). As a result, we expect the combination of an m -dimensional quantum state space and an n -dimensional quantum state space to be an mn -dimensional quantum state space, and by (4.44), the tensor product fits the bill perfectly. This of course is hand-waving and assumes finite-dimensionality. For a fully rigorous justification, we refer readers to Aerts and Daubechies (1978, 1979a,b).

Note that there is no contradiction between (4.42) and (4.43): in the former the objects are vector spaces; in the latter they are bases of vector spaces. The tensor product symbol \otimes is used in at least a dozen different ways, but fortunately there is little cause for confusion as its meaning is almost always unambiguous from the context.

One downside of Definition 4.4 in Section 4.1 is that it requires us to convert everything into real-valued functions before we can discuss tensor products. What we gain from the abstraction in Definition 4.9 is generality: the construction above allows us to form tensor products of any objects *as is*, so long as they belong to some vector space (or module).

Example 4.10 (concrete tensor products). As we mentioned on page 659, when $\mathbb{U}, \mathbb{V}, \mathbb{W}$ are abstract vector spaces, there is no further ‘meaning’ one may attach to \otimes , just that it satisfies arithmetical properties such as (4.38). But with abstraction

comes generality: it also means that any associative product operation that satisfies (4.38) may play the role of \otimes . As soon as we pick specific $\mathbb{U}, \mathbb{V}, \mathbb{W}$ – Euclidean spaces, function spaces, distribution spaces, measure spaces, *etc.* – we may start taking tensor products of these spaces with concrete formulas for \otimes . We will go over the most common examples.

(i) Outer products of vectors $a = (a_1, \dots, a_m) \in \mathbb{R}^m, b = (b_1, \dots, b_n) \in \mathbb{R}^n$:

$$a \otimes b := ab^T = \begin{bmatrix} a_1 b_1 & \cdots & a_1 b_n \\ \vdots & \ddots & \vdots \\ a_m b_1 & \cdots & a_m b_n \end{bmatrix} \in \mathbb{R}^{m \times n}.$$

(ii) Kronecker products of matrices $A \in \mathbb{R}^{m \times n}, B \in \mathbb{R}^{p \times q}$:

$$A \otimes B := \begin{bmatrix} a_{11} B & \cdots & a_{1n} B \\ \vdots & \ddots & \vdots \\ a_{m1} B & \cdots & a_{mn} B \end{bmatrix} \in \mathbb{R}^{mp \times nq}.$$

(iii) Separable products of functions $f: X \rightarrow \mathbb{R}, g: Y \rightarrow \mathbb{R}$:

$$f \otimes g: X \times Y \rightarrow \mathbb{R}, \quad (x, y) \mapsto f(x)g(y).$$

(iv) Separable products of kernels $K: X \times X' \rightarrow \mathbb{R}, H: Y \times Y' \rightarrow \mathbb{R}$:

$$K \otimes H: (X \times X') \times (Y \times Y') \rightarrow \mathbb{R}, \\ ((x, x'), (y, y')) \mapsto K(x, x')H(y, y').$$

Hence these apparently different products of different objects may all be regarded as tensor products. In particular, associativity (4.37) is satisfied and we may extend these products to an arbitrary number of factors. Readers who recall Example 4.5 would see that (i) is a special case of (iii) with $X = [m], Y = [n]$, and (ii) is a special case of (iv) with $X = [m], X' = [n], Y = [p], Y' = [q]$. Nevertheless, (i) and (ii) may also be independently derived from Definition 4.9. More generally (iii) recovers Definition 4.4.

The outer product in (i) comes from Definition 4.9. Take two vectors represented in, say, Cartesian and spherical coordinates:

$$u = a_1 e_x + a_2 e_y + a_3 e_z, \quad v = b_1 e_r + b_2 e_\theta + b_3 e_\phi \quad (4.45)$$

respectively. Consider the contravariant 2-tensor $u \otimes v$. It follows from (4.38) that

$$u \otimes v = a_1 b_1 e_x \otimes e_r + a_1 b_2 e_x \otimes e_\theta + a_1 b_3 e_x \otimes e_\phi \\ + a_2 b_1 e_y \otimes e_r + a_2 b_2 e_y \otimes e_\theta + a_2 b_3 e_y \otimes e_\phi \\ + a_3 b_1 e_z \otimes e_r + a_3 b_2 e_z \otimes e_\theta + a_3 b_3 e_z \otimes e_\phi.$$

The coefficients are captured by the outer product of their coefficients, that is,

$$a \otimes b = \begin{bmatrix} a_1 b_1 & a_1 b_2 & a_1 b_3 \\ a_2 b_1 & a_2 b_2 & a_2 b_3 \\ a_3 b_1 & a_3 b_2 & a_3 b_3 \end{bmatrix} = [u \otimes v]_{\mathcal{B}_1 \otimes \mathcal{B}_2},$$

which is a (hyper)matrix representation of $u \otimes v$ in the tensor product basis

$$\begin{aligned} \mathcal{B}_1 \otimes \mathcal{B}_2 := \{ & e_x \otimes e_r, e_x \otimes e_\theta, e_x \otimes e_\phi, e_y \otimes e_r, \\ & e_y \otimes e_\theta, e_y \otimes e_\phi, e_z \otimes e_r, e_z \otimes e_\theta, e_z \otimes e_\phi \} \end{aligned} \quad (4.46)$$

of the two bases $\mathcal{B}_1 = \{e_x, e_y, e_z\}$ and $\mathcal{B}_2 = \{e_r, e_\theta, e_\phi\}$. If we have another pair of vectors

$$u' = a'_1 e_x + a'_2 e_y + a'_3 e_z, \quad v' = b'_1 e_r + b'_2 e_\theta + b'_3 e_\phi,$$

then $u \otimes v + u' \otimes v'$, or more generally $\lambda u \otimes v + \lambda' u' \otimes v'$ for any scalars λ, λ' , obeys

$$[\lambda u \otimes v + \lambda' u' \otimes v']_{\mathcal{B}_1 \otimes \mathcal{B}_2} = \lambda a \otimes b + \lambda' a' \otimes b'.$$

In other words, the arithmetic of 2-tensors in $\mathbb{V} \otimes \mathbb{W}$ mirrors the arithmetic of rank-one matrices in $\mathbb{R}^{m \times n}$, $m = \dim \mathbb{W}$ and $n = \dim \mathbb{V}$. It is easy to see that this statement extends to higher order, that is, the arithmetic of d -tensors in $\mathbb{U} \otimes \mathbb{V} \otimes \cdots \otimes \mathbb{W}$ mirrors the arithmetic of rank-one d -hypermatrices in $\mathbb{R}^{m \times n \times \cdots \times p}$:

$$\left[\sum_{i=1}^r \lambda_i u_i \otimes v_i \otimes \cdots \otimes w_i \right]_{\mathcal{B}_1 \otimes \mathcal{B}_2 \otimes \cdots \otimes \mathcal{B}_d} = \sum_{i=1}^r \lambda_i [u_i]_{\mathcal{B}_1} \otimes [v_i]_{\mathcal{B}_2} \otimes \cdots \otimes [w_i]_{\mathcal{B}_d}.$$

Note that the tensor products on the left of the equation are abstract tensor products and those on the right are outer products.

A few words of caution are in order whenever hypermatrices are involved. Firstly, hypermatrices are dependent on the bases: if we change u in (4.45) to spherical coordinates with the same numerical values,

$$u = a_1 e_r + a_2 e_\theta + a_3 e_\phi,$$

then $\lambda a \otimes b + \lambda' a' \otimes b'$, perfectly well-defined as a (hyper)matrix, is a meaningless thing to compute as we are adding coordinate representations in different bases. Secondly, hypermatrices are oblivious to covariance and contravariance: if we change u in (4.45) to a dual vector with exactly the same coordinates,

$$u = a_1 e_x^* + a_2 e_y^* + a_3 e_z^*,$$

then again $\lambda a \otimes b + \lambda' a' \otimes b'$ is meaningless as we are adding coordinate representations of incompatible objects in different vector spaces. Both are of course obvious points but they are often obfuscated when higher-order tensors are brought into the picture.

The Kronecker product above may also be viewed as another manifestation of Definition 4.9, but a more fruitful way is to deduce it as a tensor product of

linear operators that naturally follows from Definition 4.9. This is important and sufficiently interesting to warrant separate treatment.

Example 4.11 (Kronecker product). Given linear operators

$$\Phi_1: \mathbb{V}_1 \rightarrow \mathbb{W}_1 \quad \text{and} \quad \Phi_2: \mathbb{V}_2 \rightarrow \mathbb{W}_2,$$

forming the tensor products $\mathbb{V}_1 \otimes \mathbb{V}_2$ and $\mathbb{W}_1 \otimes \mathbb{W}_2$ automatically²⁵ gives us a linear operator

$$\Phi_1 \otimes \Phi_2: \mathbb{V}_1 \otimes \mathbb{V}_2 \rightarrow \mathbb{W}_1 \otimes \mathbb{W}_2$$

defined on rank-one elements by

$$\Phi_1 \otimes \Phi_2(v_1 \otimes v_2) := \Phi_1(v_1) \otimes \Phi_2(v_2)$$

and extended linearly to all elements of $\mathbb{V}_1 \otimes \mathbb{V}_2$, which are all finite linear combinations of rank-one elements. The way it is defined, $\Phi_1 \otimes \Phi_2$ is clearly unique and we will call it the Kronecker product of linear operators Φ_1 and Φ_2 . It extends easily to an arbitrary number of linear operators via

$$\Phi_1 \otimes \cdots \otimes \Phi_d(v_1 \otimes \cdots \otimes v_d) := \Phi_1(v_1) \otimes \cdots \otimes \Phi_d(v_d), \tag{4.47}$$

and the result obeys (4.40):

$$\begin{aligned} (\lambda\Phi_1 + \lambda'\Phi'_1) \otimes \Phi_2 \otimes \cdots \otimes \Phi_d &= \lambda\Phi_1 \otimes \Phi_2 \otimes \cdots \otimes \Phi_d + \lambda'\Phi'_1 \otimes \Phi_2 \otimes \cdots \otimes \Phi_d, \\ \Phi_1 \otimes (\lambda\Phi_2 + \lambda'\Phi'_2) \otimes \cdots \otimes \Phi_d &= \lambda\Phi_1 \otimes \Phi_2 \otimes \cdots \otimes \Phi_d + \lambda'\Phi_1 \otimes \Phi'_2 \otimes \cdots \otimes \Phi_d, \\ &\vdots \\ \Phi_1 \otimes \Phi_2 \otimes \cdots \otimes (\lambda\Phi_d + \lambda'\Phi'_d) &= \lambda\Phi_1 \otimes \Phi_2 \otimes \cdots \otimes \Phi_d + \lambda'\Phi_1 \otimes \Phi_2 \otimes \cdots \otimes \Phi'_d. \end{aligned}$$

In other words, the Kronecker product defines a tensor product in the sense of Definition 4.9 on the space of linear operators,

$$\mathcal{L}(\mathbb{V}_1; \mathbb{W}_1) \otimes \cdots \otimes \mathcal{L}(\mathbb{V}_d; \mathbb{W}_d),$$

that for finite-dimensional vector spaces equals

$$\mathcal{L}(\mathbb{V}_1 \otimes \cdots \otimes \mathbb{V}_d; \mathbb{W}_1 \otimes \cdots \otimes \mathbb{W}_d).$$

In addition, as linear operators they may be composed and Moore–Penrose-inverted, and have adjoints and ranks, images and null spaces, all of which work in tandem with the Kronecker product:

$$(\Phi_1 \otimes \cdots \otimes \Phi_d)(\Psi_1 \otimes \cdots \otimes \Psi_d) = \Phi_1\Psi_1 \otimes \cdots \otimes \Phi_d\Psi_d, \tag{4.48}$$

$$(\Phi_1 \otimes \cdots \otimes \Phi_d)^\dagger = \Phi_1^\dagger \otimes \cdots \otimes \Phi_d^\dagger, \tag{4.49}$$

$$(\Phi_1 \otimes \cdots \otimes \Phi_d)^* = \Phi_1^* \otimes \cdots \otimes \Phi_d^*, \tag{4.50}$$

²⁵ This is called *functoriality* in category theory. The fact that the tensor product of vector spaces in Definition 4.9 gives a tensor product of linear operators on these spaces says that the tensor product in Definition 4.9 is functorial.

$$\text{rank}(\Phi_1 \otimes \cdots \otimes \Phi_d) = \text{rank}(\Phi_1) \cdots \text{rank}(\Phi_d), \quad (4.51)$$

$$\text{im}(\Phi_1 \otimes \cdots \otimes \Phi_d) = \text{im}(\Phi_1) \otimes \cdots \otimes \text{im}(\Phi_d). \quad (4.52)$$

Observe that the \otimes on the left of (4.52) is the Kronecker product of operators whereas that on the right is the tensor product of vector spaces. For null spaces, we have

$$\ker(\Phi_1 \otimes \Phi_2) = \ker(\Phi_1) \otimes \mathbb{V}_2 + \mathbb{V}_1 \otimes \ker(\Phi_2),$$

when $d = 2$ and more generally

$$\begin{aligned} \ker(\Phi_1 \otimes \cdots \otimes \Phi_d) &= \ker(\Phi_1) \otimes \mathbb{V}_2 \otimes \cdots \otimes \mathbb{V}_d + \mathbb{V}_1 \otimes \ker(\Phi_2) \otimes \cdots \otimes \mathbb{V}_d \\ &\quad + \cdots + \mathbb{V}_1 \otimes \mathbb{V}_2 \otimes \cdots \otimes \ker(\Phi_d). \end{aligned}$$

Therefore injectivity and surjectivity are preserved by taking Kronecker products. If $\mathbb{V}_i = \mathbb{W}_i$, $i = 1, \dots, d$, then the eigenpairs of $\Phi_1 \otimes \cdots \otimes \Phi_d$ are exactly those given by

$$(\lambda_1 \lambda_2 \cdots \lambda_d, v_1 \otimes v_2 \otimes \cdots \otimes v_d), \quad (4.53)$$

where (λ_i, v_i) is an eigenpair of Φ_i , $i = 1, \dots, d$. All of these are straightforward consequences of (4.47) (Berberian 2014, Section 13.2).

The Kronecker product of matrices simply expresses how the matrix representing $\Phi_1 \otimes \cdots \otimes \Phi_d$ is related to those representing Φ_1, \dots, Φ_d . If we pick bases \mathcal{B}_i for \mathbb{V}_i and \mathcal{C}_i for \mathbb{W}_i , then each linear operator Φ_i has a matrix representation

$$[\Phi_i]_{\mathcal{B}_i, \mathcal{C}_i} = X_i \in \mathbb{R}^{m_i \times n_i}$$

as in (3.5), $i = 1, \dots, d$, and

$$[\Phi_1 \otimes \cdots \otimes \Phi_d]_{\mathcal{B}_1 \otimes \cdots \otimes \mathcal{B}_d, \mathcal{C}_1 \otimes \cdots \otimes \mathcal{C}_d} = X_1 \otimes \cdots \otimes X_d \in \mathbb{R}^{m_1 m_2 \cdots m_d \times n_1 n_2 \cdots n_d}.$$

Note that the tensor products on the left of the equation are Kronecker products as defined in (4.47) and those on the right are matrix Kronecker products as defined in Example 4.10(ii), applied d times in any order (order does not matter as \otimes is associative). The tensor product of d bases is as defined in (4.43).

Clearly the matrix Kronecker product inherits the properties (4.48)–(4.53), and when $m_i = n_i$, the eigenvalue property in particular gives

$$\begin{aligned} \text{tr}(X_1 \otimes \cdots \otimes X_d) &= \text{tr}(X_1) \cdots \text{tr}(X_d), \\ \det(X_1 \otimes \cdots \otimes X_d) &= \det(X_1)^{p_1} \cdots \det(X_d)^{p_d}, \end{aligned} \quad (4.54)$$

where $p_i = (n_1 n_2 \cdots n_d) / n_i$, $i = 1, \dots, d$. Incidentally these last two properties also apply to Kronecker products of operators with coordinate independent definitions of trace (see Example 4.28) and determinant (Berberian 2014, Section 7.2).

Next we will see that the multilinear matrix multiplication in (2.6) and the matrix Kronecker product in Example 4.10(ii) are intimately related.

Example 4.12 (multilinear matrix multiplication revisited). There is one caveat to obtaining the neat formula for the matrix Kronecker product in Example 4.10(ii) from the operator Kronecker product in Example 4.11, namely, we must choose the *lexicographic order* for the tensor product bases $\mathcal{B}_1 \otimes \cdots \otimes \mathcal{B}_d$ and $\mathcal{C}_1 \otimes \cdots \otimes \mathcal{C}_d$, recalling that ‘basis’ in this article always means ordered basis. This is the most common and intuitive way to totally order a Cartesian product of totally ordered sets. For our purposes, we just need to define it for $[n_1] \times [n_2] \times \cdots \times [n_d]$ since our tensor product bases and hypermatrices are all indexed by this. Let

$$(i_1, \dots, i_d), (j_1, \dots, j_d) \in [n_1] \times \cdots \times [n_d].$$

The lexicographic order $<$ is defined as

$$(i_1, i_2, \dots, i_d) < (j_1, j_2, \dots, j_d) \quad \text{if and only if} \quad \begin{cases} i_1 = j_1, \dots, i_{k-1} = j_{k-1}, \\ i_k < j_k \text{ for some } k \in [d]. \end{cases}$$

Take $d = 2$, for instance. If $\mathcal{B}_1 = \{u_1, u_2\}$, $\mathcal{B}_2 = \{v_1, v_2, v_3\}$, we have to order

$$\mathcal{B}_1 \otimes \mathcal{B}_2 = \{u_1 \otimes v_1, u_1 \otimes v_2, u_1 \otimes v_3, u_2 \otimes v_1, u_2 \otimes v_2, u_2 \otimes v_3\}$$

and not, say, $\{u_1 \otimes v_1, u_2 \otimes v_1, u_1 \otimes v_2, u_2 \otimes v_2, u_1 \otimes v_3, u_2 \otimes v_3\} =: \mathcal{A}$. While this is a trifling point, the latter ordering is unfortunately what is used in the *vec* operation common in mathematical software where one concatenates the n columns of an $m \times n$ matrix into a vector of dimension mn . The coordinate representation of the 2-tensor

$$T = a_{11}u_1 \otimes v_1 + a_{12}u_1 \otimes v_2 + a_{13}u_1 \otimes v_3 + a_{21}u_2 \otimes v_1 + a_{22}u_2 \otimes v_2 + a_{23}u_2 \otimes v_3$$

in $\mathcal{B}_1 \otimes \mathcal{B}_2$ (lexicographic) and \mathcal{A} (non-lexicographic) order are

$$[T]_{\mathcal{B}_1 \otimes \mathcal{B}_2} = \begin{bmatrix} a_{11} \\ a_{12} \\ a_{13} \\ a_{21} \\ a_{22} \\ a_{23} \end{bmatrix} \in \mathbb{R}^6, \quad [T]_{\mathcal{A}} = \begin{bmatrix} a_{11} \\ a_{21} \\ a_{12} \\ a_{22} \\ a_{13} \\ a_{23} \end{bmatrix} = \text{vec} \left(\begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \end{bmatrix} \right) \in \mathbb{R}^6$$

respectively, where *vec* denotes the usual ‘column-concatenated’ operation stacking the columns of an $m \times n$ matrix to obtain a vector of dimension mn . To see why we favour lexicographic order, define the vector space isomorphism

$$\text{vec}_\ell: \mathbb{R}^{n_1 \times n_2 \times \cdots \times n_d} \rightarrow \mathbb{R}^{n_1 n_2 \cdots n_d}$$

that takes a d -hypermatrix $A = [a_{i_1 i_2 \dots i_d}]$ to a vector

$$\text{vec}_\ell(A) = (a_{11\dots 1}, a_{11\dots 2}, \dots, a_{n_1 n_2 \dots n_d}) \in \mathbb{R}^{n_1 n_2 \cdots n_d},$$

whose coordinates are those in A ordered lexicographically. For any matrices

$X_1 \in \mathbb{R}^{m_1 \times n_1}, \dots, X_d \in \mathbb{R}^{m_d \times n_d}$, it is routine to check that

$$\text{vec}_\ell((X_1, \dots, X_d) \cdot A) = (X_1 \otimes \dots \otimes X_d) \text{vec}_\ell(A). \tag{4.55}$$

In other words, the lexicographic vec_ℓ operator takes multilinear matrix multiplication to the matrix Kronecker product. They are the same operation represented on different but isomorphic vector spaces $\mathbb{R}^{n_1 \times \dots \times n_d}$ and $\mathbb{R}^{n_1 \dots n_d}$, respectively. Incidentally, vec_ℓ is an example of a ‘forgetful map’: it forgets the hypermatrix structure. This loss of information is manifested in that, given (X_1, \dots, X_d) , we can compute $X_1 \otimes \dots \otimes X_d$ but we cannot uniquely recover (X_1, \dots, X_d) from their matrix Kronecker product.

For $d = 2$, vec_ℓ is just the ‘row-concatenated vec operation’, *i.e.* $\text{vec}_\ell(A) = \text{vec}(A^\top)$ for any $A \in \mathbb{R}^{m \times n}$, and we have

$$\text{vec}_\ell(XAY^\top) = (X \otimes Y) \text{vec}_\ell(A), \quad \text{vec}(XAY^\top) = (Y^\top \otimes X) \text{vec}(A).$$

The neater form for vec_ℓ , especially when extended to higher order, is why we prefer lexicographic ordering.

As we discussed in Example 2.7, there is no formula that extends the usual matrix–matrix multiplication to 3-hypermatrices or indeed any d -hypermatrices for odd $d \in \mathbb{N}$. Such a formula, however, exists for any even d , thanks to the Kronecker product.

Example 4.13 (‘multiplying higher-order tensors’ revisited). Consider the following $2d$ -hypermatrices:

$$A \in \mathbb{R}^{m_1 \times n_1 \times m_2 \times n_2 \times \dots \times m_d \times n_d}, \quad B \in \mathbb{R}^{n_1 \times p_1 \times n_2 \times p_2 \times \dots \times n_d \times p_d}.$$

Their product, denoted AB to be consistent with the usual matrix–matrix product, is the $2d$ -hypermatrix $C \in \mathbb{R}^{m_1 \times p_1 \times m_2 \times p_2 \times \dots \times m_d \times p_d}$, where

$$c_{i_1 k_1 \dots i_d k_d} := \sum_{j_1=1}^{n_1} \dots \sum_{j_d=1}^{n_d} a_{i_1 j_1 \dots i_d j_d} b_{j_1 k_1 \dots j_d k_d}.$$

This clearly extends the matrix–matrix product, which is just the case when $d = 1$. We have in fact already encountered Example 4.11 in a different form: this is (4.48), the composition of two Kronecker products of d linear operators. To view it in hypermatrix form, we just use the fact that $\mathcal{L}(\mathbb{V}; \mathbb{W}) = \mathbb{V}^* \otimes \mathbb{W}$. Indeed, we may also view it as a tensor contraction of

$$T \in (\mathbb{U}_1^* \otimes \mathbb{V}_1) \otimes \dots \otimes (\mathbb{U}_d^* \otimes \mathbb{V}_d), \quad T' \in (\mathbb{V}_1^* \otimes \mathbb{W}_1) \otimes \dots \otimes (\mathbb{V}_d^* \otimes \mathbb{W}_d),$$

from which we obtain

$$\langle T, T' \rangle \in (\mathbb{U}_1^* \otimes \mathbb{W}_1) \otimes \dots \otimes (\mathbb{U}_d^* \otimes \mathbb{W}_d).$$

This is entirely within expectation as the first fundamental theorem of invariant theory, the result that prevents the existence of a d -hypermatrix–hypermatrix product for odd d , also tells us that essentially the only way to multiply tensors without

increasing order is via contractions. This product is well-defined for $2d$ -tensors and does not depend on bases: in terms of Kronecker products of matrices,

$$\begin{aligned} & ((X_1 A_1 Y_1^{-1}) \otimes \cdots \otimes (X_d A_d Y_d^{-1})) ((Y_1 B_1 Z_1^{-1}) \otimes \cdots \otimes (Y_d B_d Z_d^{-1})) \\ &= (X_1 A_1 B_1 Z_1^{-1}) \otimes \cdots \otimes (X_d A_d B_d Z_d^{-1}), \end{aligned}$$

and in terms of hypermatrix–hypermatrix products,

$$\begin{aligned} & ((X_1, Y_1^{-1}, \dots, X_d, Y_d^{-1}) \cdot A) ((Y_1, Z_1^{-1}, \dots, Y_d, Z_d^{-1}) \cdot B) \\ &= (X_1, Z_1^{-1}, \dots, X_d, Z_d^{-1}) \cdot (AB), \end{aligned}$$

that is, they satisfy the higher-order analogue of (2.2). Frankly, we do not see any advantage in formulating such a product as $2d$ -hypermatrices, but an abundance of disadvantages.

Just as it is possible to deduce the tensor transformation rules in definition ① from definition ②, we can do likewise with definition ③ in the form of Definition 4.9.

Example 4.14 (tensor transformation rules revisited). As stated on page 665, the tensor product construction in this section does not require that $\mathbb{U}, \mathbb{V}, \dots, \mathbb{W}$ have bases, that is, they could be modules, which do not have bases in general (those with bases are called *free modules*). Nevertheless, in the event where they do have bases, their change-of-basis theorems would lead us directly to the transformation rules discussed in Section 2.

We first remind the reader of a simple notion, discussed in standard linear algebra textbooks such as Berberian (2014, Section 3.9) and Friedberg *et al.* (2003, Section 2.6) but often overlooked. Any linear operator $\Phi: \mathbb{V} \rightarrow \mathbb{W}$ induces a *transpose* linear operator on the dual spaces defined as

$$\Phi^T: \mathbb{W}^* \rightarrow \mathbb{V}^*, \quad \Phi(\varphi) := \varphi \circ \Phi$$

for any linear functional $\varphi \in \mathbb{W}^*$. Note that the composition $\varphi \circ \Phi: \mathbb{V} \rightarrow \mathbb{R}$ is indeed a linear functional in \mathbb{V}^* . The reason for its name is that

$$[\Phi]_{\mathcal{B}, \mathcal{C}} = A \in \mathbb{R}^{m \times n} \quad \text{if and only if} \quad [\Phi^T]_{\mathcal{C}^*, \mathcal{B}^*} = A^T \in \mathbb{R}^{n \times m}.$$

One may show that Φ is injective if and only if Φ^T is surjective, and Φ is surjective if and only if Φ^T is injective.²⁶ So if $\Phi: \mathbb{V} \rightarrow \mathbb{W}$ is invertible, then so is $\Phi^T: \mathbb{W}^* \rightarrow \mathbb{V}^*$ and its inverse is a linear operator,

$$\Phi^{-T}: \mathbb{V}^* \rightarrow \mathbb{W}^*.$$

Another name for ‘invertible linear operator’ is vector space isomorphism, especially when used in the following context. Any basis $\mathcal{B} = \{v_1, \dots, v_n\}$ of \mathbb{V}

²⁶ Not a typo. Injectivity and surjectivity are in fact dual notions in this sense.

gives us two vector space isomorphisms, namely (i) $\Phi_{\mathcal{B}}: \mathbb{V} \rightarrow \mathbb{R}^n$ that takes $v = a_1v_1 + \dots + a_nv_n \in \mathbb{V}$ to its coordinate representation in \mathcal{B} , and (ii) $\Phi_{\mathcal{B}^*}: \mathbb{V}^* \rightarrow \mathbb{R}^n$ that takes $\varphi = b_1v_1^* + \dots + b_nv_n^* \in \mathbb{V}^*$ to its coordinate representation in \mathcal{B}^* :

$$\Phi_{\mathcal{B}}(v) = \begin{bmatrix} a_1 \\ \vdots \\ a_n \end{bmatrix} \in \mathbb{R}^n, \quad \Phi_{\mathcal{B}^*}(\varphi) = \begin{bmatrix} b_1 \\ \vdots \\ b_n \end{bmatrix} \in \mathbb{R}^n, \tag{4.56}$$

and as expected

$$\Phi_{\mathcal{B}^*} = \Phi_{\mathcal{B}}^{-T}.$$

This is of course just paraphrasing what we have already discussed in (3.2) and thereabouts. Note that we do not distinguish between \mathbb{R}^n and its dual space $(\mathbb{R}^n)^*$.

Given vector spaces $\mathbb{V}_1, \dots, \mathbb{V}_d$ and $\mathbb{W}_1, \dots, \mathbb{W}_d$ and invertible linear operators $\Phi_i: \mathbb{V}_i \rightarrow \mathbb{W}_i, i = 1, \dots, d$, take the Kronecker product of $\Phi_1, \dots, \Phi_p, \Phi_{p+1}^{-T}, \dots, \Phi_d^{-T}$ as defined in (4.47) of the previous example. Then the linear operator

$$\begin{aligned} &\Phi_1 \otimes \dots \otimes \Phi_p \otimes \Phi_{p+1}^{-T} \otimes \dots \otimes \Phi_d^{-T}: \mathbb{V}_1 \otimes \dots \otimes \mathbb{V}_p \otimes \mathbb{V}_{p+1}^* \otimes \dots \otimes \mathbb{V}_d^* \\ &\rightarrow \mathbb{W}_1 \otimes \dots \otimes \mathbb{W}_p \otimes \mathbb{W}_{p+1}^* \otimes \dots \otimes \mathbb{W}_d^* \end{aligned}$$

must also be invertible by (4.49).

For each $i = 1, \dots, d$, let \mathcal{B}_i and \mathcal{B}'_i be two different bases of \mathbb{V}_i ; let $\mathbb{W}_i = \mathbb{R}^{n_i}$ with $n_i = \dim \mathbb{V}_i$ and

$$\Phi_i := \Phi_{\mathcal{B}_i}, \quad \Psi_i := \Phi_{\mathcal{B}'_i}$$

as in (4.56). Since we assume $\mathbb{R}^n = (\mathbb{R}^n)^*$, we have

$$\mathbb{W}_1 \otimes \dots \otimes \mathbb{W}_p \otimes \mathbb{W}_{p+1}^* \otimes \dots \otimes \mathbb{W}_d^* \cong \mathbb{R}^{n_1} \otimes \dots \otimes \mathbb{R}^{n_d} \cong \mathbb{R}^{n_1 \times \dots \times n_d}, \tag{4.57}$$

where the last \cong is a consequence of using Example 4.10(i). With these choices, the vector space isomorphisms

$$\begin{aligned} &\Phi_1 \otimes \dots \otimes \Phi_p \otimes \Phi_{p+1}^{-T} \otimes \dots \otimes \Phi_d^{-T}: \\ &\quad \mathbb{V}_1 \otimes \dots \otimes \mathbb{V}_p \otimes \mathbb{V}_{p+1}^* \otimes \dots \otimes \mathbb{V}_d^* \rightarrow \mathbb{R}^{n_1 \times \dots \times n_d}, \\ &\Psi_1 \otimes \dots \otimes \Psi_p \otimes \Psi_{p+1}^{-T} \otimes \dots \otimes \Psi_d^{-T}: \\ &\quad \mathbb{V}_1 \otimes \dots \otimes \mathbb{V}_p \otimes \mathbb{V}_{p+1}^* \otimes \dots \otimes \mathbb{V}_d^* \rightarrow \mathbb{R}^{n_1 \times \dots \times n_d} \end{aligned}$$

give us hypermatrix representations of d -tensors with respect to $\mathcal{B}_1, \dots, \mathcal{B}_d$ and $\mathcal{B}'_1, \dots, \mathcal{B}'_d$ respectively. Let $T \in \mathbb{V}_1 \otimes \dots \otimes \mathbb{V}_p \otimes \mathbb{V}_{p+1}^* \otimes \dots \otimes \mathbb{V}_d^*$ and

$$\begin{aligned} &\Phi_1 \otimes \dots \otimes \Phi_p \otimes \Phi_{p+1}^{-T} \otimes \dots \otimes \Phi_d^{-T}(T) = A \in \mathbb{R}^{n_1 \times \dots \times n_d}, \\ &\Psi_1 \otimes \dots \otimes \Psi_p \otimes \Psi_{p+1}^{-T} \otimes \dots \otimes \Psi_d^{-T}(T) = A' \in \mathbb{R}^{n_1 \times \dots \times n_d}. \end{aligned}$$

Then the relation between the hypermatrices A and A' is given by

$$A = \Phi_1 \Psi_1^{-1} \otimes \cdots \otimes \Phi_p \Psi_p^{-1} \otimes (\Phi_{p+1} \Psi_{p+1}^{-1})^{-\top} \otimes \cdots \otimes (\Phi_d \Psi_d^{-1})^{-\top}(A'),$$

where we have used (4.48) and (4.49). Note that each $\Phi_i \Psi_i^{-1}: \mathbb{R}^{n_i} \rightarrow \mathbb{R}^{n_i}$ is an invertible linear operator and so it must be given by $\Phi_i \Psi_i^{-1}(v) = X_i v$ for some $X_i \in GL(n_i)$. Using (4.55), the above relation between A and A' in terms of multilinear matrix multiplication is just

$$A = (X_1, \dots, X_p, X_{p+1}^{-\top}, \dots, X_d^{-\top}) \cdot A',$$

which gives us the tensor transformation rule (2.9). To get the isomorphism in (4.57), we had to identify \mathbb{R}^n with $(\mathbb{R}^n)^*$, and this is where we lost all information pertaining to covariance and contravariance; a hypermatrix cannot perfectly represent a tensor, which is one reason why we need to look at the transformation rules to ascertain the tensor.

The example above essentially discusses change of basis without any mention of bases. In pure mathematics, sweeping such details under the rug is generally regarded as a good thing; in applied and computational mathematics, such details are usually unavoidable and we will work them out below.

Example 4.15 (representing tensors as hypermatrices). Let $\mathcal{B}_1 = \{u_1, \dots, u_m\}$, $\mathcal{B}_2 = \{v_1, \dots, v_n\}, \dots, \mathcal{B}_d = \{w_1, \dots, w_p\}$ be any bases of finite-dimensional vector spaces $\mathbb{U}, \mathbb{V}, \dots, \mathbb{W}$. Then any d -tensor $T \in \mathbb{U} \otimes \mathbb{V} \otimes \cdots \otimes \mathbb{W}$ may be expressed as

$$T = \sum_{i=1}^m \sum_{j=1}^n \cdots \sum_{k=1}^p a_{ij\dots k} u_i \otimes v_j \otimes \cdots \otimes w_k, \tag{4.58}$$

and the d -hypermatrix of coefficients

$$[T]_{\mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_d} := [a_{ij\dots k}]_{i,j,\dots,k=1}^{m,n,\dots,p} = A \in \mathbb{R}^{m \times n \times \cdots \times p}$$

is said to *represent* or to be a *coordinate representation* of T with respect to $\mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_d$. That (4.58) always holds is easily seen as follows. First it holds for a rank-one tensor $T = u \otimes v \otimes \cdots \otimes w$ as we may express u, v, \dots, w as a linear combination of basis vectors and apply (4.40) to get

$$\begin{aligned} & (a_1 u_1 + \cdots + a_m u_m) \otimes (b_1 v_1 + \cdots + b_n v_n) \otimes \cdots \otimes (c_1 w_1 + \cdots + c_p w_p) \\ &= \sum_{i=1}^m \sum_{j=1}^n \cdots \sum_{k=1}^p a_i b_j \cdots c_k u_i \otimes v_j \otimes \cdots \otimes w_k, \end{aligned} \tag{4.59}$$

which has the form in (4.58). Since an arbitrary T is simply a finite sum of rank-one

tensors as in (4.39), the distributivity in (4.40) again ensures that

$$\begin{aligned} & \lambda \sum_{i=1}^m \sum_{j=1}^n \cdots \sum_{k=1}^p a_{ij\dots k} u_i \otimes v_j \otimes \cdots \otimes w_k \\ & \quad + \lambda' \sum_{i=1}^m \sum_{j=1}^n \cdots \sum_{k=1}^p b_{ij\dots k} u_i \otimes v_j \otimes \cdots \otimes w_k \\ & = \sum_{i=1}^m \sum_{j=1}^n \cdots \sum_{k=1}^p (\lambda a_{ij\dots k} + \lambda' b_{ij\dots k}) u_i \otimes v_j \otimes \cdots \otimes w_k. \end{aligned} \quad (4.60)$$

By (4.59) and (4.60), any T must be expressible in the form (4.58). For a different choice of bases $\mathcal{B}'_1, \mathcal{B}'_2, \dots, \mathcal{B}'_d$ we obtain a different hypermatrix representation $A' \in \mathbb{R}^{m \times n \times \cdots \times p}$ for T that is related to A by a multilinear matrix multiplication as in Example 4.14.

Note that (4.59) and (4.60), respectively, give the formulas for outer product and linear combination:

$$\begin{aligned} (a_1, \dots, a_m) \otimes (b_1, \dots, b_n) \otimes \cdots \otimes (c_1, \dots, c_p) &= [a_i b_j \cdots c_k], \\ \lambda [a_{ij\dots k}] + \lambda' [b_{ij\dots k}] &= [\lambda a_{ij\dots k} + \lambda' b_{ij\dots k}]. \end{aligned}$$

These also follow from our definition of hypermatrices as real-valued functions in Example 4.5.

If $T' \in \mathbb{U}' \otimes \mathbb{V}' \otimes \cdots \otimes \mathbb{W}'$ is a d' -tensor, expressed as in (4.58) as

$$T' = \sum_{i'=1}^{m'} \sum_{j'=1}^{n'} \cdots \sum_{k'=1}^{p'} b_{i'j'\dots k'} u'_{i'} \otimes v'_{j'} \otimes \cdots \otimes w'_{k'},$$

with bases of $\mathbb{U}', \mathbb{V}', \dots, \mathbb{W}'$

$$\mathcal{B}'_1 = \{u'_1, \dots, u'_{m'}\}, \quad \mathcal{B}'_2 = \{v'_1, \dots, v'_{n'}\}, \dots, \mathcal{B}'_{d'} = \{w'_1, \dots, w'_{p'}\},$$

then the outer product of T and T' in (4.41) may be expressed as

$$T \otimes T' = \sum_{i,j,\dots,k,i',j',\dots,k'=1}^{m,n,\dots,p,m',n',\dots,p'} a_{ij\dots k} b_{i'j'\dots k'} u_i \otimes v_j \otimes \cdots \otimes w_k \otimes u'_{i'} \otimes v'_{j'} \otimes \cdots \otimes w'_{k'},$$

which is a $(d + d')$ -tensor. In terms of their hypermatrix representations,

$$A \otimes B = [a_{ij\dots k} b_{i'j'\dots k'}]_{i,j,\dots,k,i',j',\dots,k'=1}^{m,n,\dots,p,m',n',\dots,p'} \in \mathbb{R}^{m \times n \times \cdots \times p \times m' \times n' \times \cdots \times p'}.$$

It is straightforward to show that with respect to the Kronecker product or multi-linear matrix multiplication,

$$\begin{aligned}
 & [\Phi_1 \otimes \cdots \otimes \Phi_d(T)] \otimes [\Psi_1 \otimes \cdots \otimes \Psi_{d'}(T')] \\
 & \quad = \Phi_1 \otimes \cdots \otimes \Phi_d \otimes \Psi_1 \otimes \cdots \otimes \Psi_{d'}(T \otimes T'), \\
 & [(X_1, \dots, X_d) \cdot A] \otimes [(Y_1, \dots, Y_{d'}) \cdot B] \\
 & \quad = (X_1, \dots, X_d, Y_1, \dots, Y_{d'}) \cdot A \otimes B,
 \end{aligned}$$

with notation as in Examples 4.11, 4.12 and 4.14.

The above discussion extends verbatim to covariant and mixed tensors by replacing some or all of these vector spaces with their duals and bases with dual bases. It also extends to separable Hilbert spaces with orthonormal bases $\mathcal{B}_1, \dots, \mathcal{B}_d$ and $A \in l^2(\mathbb{N}^d)$ an infinite-dimensional hypermatrix as discussed in Example 4.5. Nevertheless we do not recommend working exclusively with such hypermatrices. The reason is simple: there is a better representation of the tensor T sitting in plain sight, namely (4.58), which contains not just all information in the hypermatrix A but also the bases $\mathcal{B}_1, \dots, \mathcal{B}_d$. Indeed, the representation of a tensor in (4.58), coupled with the arithmetic rules for manipulating \otimes like (4.40), is more or less the standard method for working with tensors when bases are involved. Take $d = 2$ for illustration. We can tell at a glance from

$$\beta = \sum_{i=1}^m \sum_{j=1}^n a_{ij} u_i^* \otimes v_j^*, \quad \Phi = \sum_{i=1}^m \sum_{j=1}^n a_{ij} u_i \otimes v_j^*$$

that β is a bilinear functional and Φ a linear operator (the dyad case is discussed in Example 4.8), information that is lost if one just looks at the matrix $A = [a_{ij}]$. Furthermore, the u_i and v_j are sometimes more important than the a_{ij} . We might have

$$\beta = \varepsilon_{-1/\sqrt{3}} \otimes \varepsilon_{-1/\sqrt{3}} + \varepsilon_{-1/\sqrt{3}} \otimes \varepsilon_{1/\sqrt{3}} + \varepsilon_{1/\sqrt{3}} \otimes \varepsilon_{-1/\sqrt{3}} + \varepsilon_{1/\sqrt{3}} \otimes \varepsilon_{1/\sqrt{3}}, \tag{4.61}$$

where the a_{ij} are all ones, and all the important information is in the basis vectors, which are the point evaluation functionals we encountered in Example 4.6, that is, $\varepsilon_x(f) = f(x)$ for any $f: [-1, 1] \rightarrow \mathbb{R}$. The decomposition (4.61) is in fact a four-point Gauss quadrature formula on the domain $[-1, 1] \times [-1, 1]$. We will revisit this in Example 4.48.

We next discuss tensor rank in light of Definition 4.9.

Example 4.16 (multilinear rank and tensor rank). For any d vector spaces $\mathbb{U}, \mathbb{V}, \dots, \mathbb{W}$ and any subspaces $\mathbb{U}' \subseteq \mathbb{U}, \mathbb{V}' \subseteq \mathbb{V}, \dots, \mathbb{W}' \subseteq \mathbb{W}$, it follows from (4.39) that

$$\mathbb{U}' \otimes \mathbb{V}' \otimes \cdots \otimes \mathbb{W}' \subseteq \mathbb{U} \otimes \mathbb{V} \otimes \cdots \otimes \mathbb{W}.$$

Given a non-zero d -tensor $T \in \mathbb{U} \otimes \mathbb{V} \otimes \cdots \otimes \mathbb{W}$, one might ask for a smallest tensor subspace that contains T , that is,

$$\begin{aligned} \mu\text{rank}(T) := \min\{(\dim \mathbb{U}', \dim \mathbb{V}', \dots, \dim \mathbb{W}') \in \mathbb{N}^d : \\ T \in \mathbb{U}' \otimes \mathbb{V}' \otimes \cdots \otimes \mathbb{W}', \mathbb{U}' \subseteq \mathbb{U}, \mathbb{V}' \subseteq \mathbb{V}, \dots, \mathbb{W}' \subseteq \mathbb{W}\}. \end{aligned} \quad (4.62)$$

This gives the *multilinear rank* of T , first studied in Hitchcock (1927, Theorem II). Let $\mathbb{U}', \mathbb{V}', \dots, \mathbb{W}'$ be subspaces that attain the minimum in (4.62) and let $\{u_1, \dots, u_p\}, \{v_1, \dots, v_q\}, \dots, \{w_1, \dots, w_r\}$ be any respective bases. Then, by (4.58), we get a multilinear rank decomposition (Hitchcock 1927, Corollary I):

$$T = \sum_{i=1}^p \sum_{j=1}^q \cdots \sum_{k=1}^r c_{ij\dots k} u_i \otimes v_j \otimes \cdots \otimes w_k. \quad (4.63)$$

For a given T , the vectors on the right of (4.63) are not unique but the subspaces $\mathbb{U}', \mathbb{V}', \dots, \mathbb{W}'$ that they span are unique, and thus $\mu\text{rank}(T)$ is well-defined.²⁷ One may view these subspaces as generalizations of the row and column spaces of a matrix and multilinear rank as a generalization of the row and column ranks. Note that the coefficient hypermatrix $C = [c_{ij\dots k}] \in \mathbb{R}^{p \times q \times \cdots \times r}$ is not unique either, but two such coefficient hypermatrices C and C' must be related by the tensor transformation rule (2.9), as we saw in Example 4.15.

The decomposition in (4.63) is a sum of rank-one tensors, and if we simply use this as impetus and define

$$\begin{aligned} \text{rank}(T) := \min \left\{ r \in \mathbb{N} : T = \sum_{i=1}^r u_i \otimes v_i \otimes \cdots \otimes w_i, u_1, \dots, u_r \in \mathbb{U}, \right. \\ \left. v_1, \dots, v_r \in \mathbb{V}, \dots, w_1, \dots, w_r \in \mathbb{W} \right\}, \end{aligned}$$

we obtain *tensor rank* (Hitchcock 1927, equations 2 and 2_a), which we have encountered in a more restricted form in (3.37). For completeness, we define

$$\mu\text{rank}(T) := (0, 0, \dots, 0) \Leftrightarrow \text{rank}(T) := 0 \Leftrightarrow T = 0.$$

Note that

$$\mu\text{rank}(T) = (1, 1, \dots, 1) \Leftrightarrow \text{rank}(T) = 1 \Leftrightarrow T = u \otimes v \otimes \cdots \otimes w,$$

that is, a rank-one tensor has rank one irrespective of the rank we use. Also, we did not exclude infinite-dimensional vector spaces from our discussion. Note that by construction $\mathbb{U} \otimes \mathbb{V} \otimes \cdots \otimes \mathbb{W}$ comprises *finite* linear combinations of rank-one tensors, and so by its definition, $\mathbb{U} \otimes \mathbb{V} \otimes \cdots \otimes \mathbb{W}$ is the set of all tensors of *finite rank*. Furthermore, whether or not a tensor has finite rank is independent of the rank we use.

²⁷ While \mathbb{N}^d is only partially ordered, this shows that the minimum in (4.62) is unique.

Vector spaces become enormously more interesting when equipped with inner products and norms; it is the same with tensor spaces. The tensor product construction leading to Definition 4.9 allows us to incorporate them easily.

Example 4.17 (tensor products of inner product and norm spaces). We let $\mathbb{U}, \mathbb{V}, \dots, \mathbb{W}$ be d vector spaces equipped with either norms

$$\|\cdot\|_1: \mathbb{U} \rightarrow [0, \infty), \|\cdot\|_2: \mathbb{V} \rightarrow [0, \infty), \dots, \|\cdot\|_d: \mathbb{W} \rightarrow [0, \infty)$$

or inner products

$$\langle \cdot, \cdot \rangle_1: \mathbb{U} \times \mathbb{U} \rightarrow \mathbb{R}, \langle \cdot, \cdot \rangle_2: \mathbb{V} \times \mathbb{V} \rightarrow \mathbb{R}, \dots, \langle \cdot, \cdot \rangle_d: \mathbb{W} \times \mathbb{W} \rightarrow \mathbb{R}.$$

We would like to define a norm or an inner product on the tensor space $\mathbb{U} \otimes \mathbb{V} \otimes \dots \otimes \mathbb{W}$ as defined in (4.39). Note that we do not assume that these vector spaces are finite-dimensional.

The inner product is easy. First define it on rank-one tensors,

$$\langle u \otimes v \otimes \dots \otimes w, u' \otimes v' \otimes \dots \otimes w' \rangle := \langle u, u' \rangle_1 \langle v, v' \rangle_2 \dots \langle w, w' \rangle_d \quad (4.64)$$

for all $u, u' \in \mathbb{U}$, $v, v' \in \mathbb{V}$, \dots , $w, w' \in \mathbb{W}$. Then extend it bilinearly, that is, decree that for rank-one tensors S, T, S', T' and scalars λ, λ' ,

$$\begin{aligned} \langle \lambda S + \lambda' S', T \rangle &:= \lambda \langle S, T \rangle + \lambda' \langle S', T \rangle, \\ \langle S, \lambda T + \lambda' T' \rangle &:= \lambda \langle S, T \rangle + \lambda' \langle S, T' \rangle. \end{aligned}$$

As $\mathbb{U} \otimes \mathbb{V} \otimes \dots \otimes \mathbb{W}$ comprises finite linear combinations of rank-one tensors by definition, this defines $\langle \cdot, \cdot \rangle$ on the whole tensor space. It is then routine to check that the axioms of inner product are satisfied by $\langle \cdot, \cdot \rangle$. This construction applies verbatim to other bilinear functionals such as the Lorentzian scalar product in Example 2.4.

Norms are slightly trickier because there are multiple ways to define them on tensor spaces. We will restrict ourselves to the three most common constructions. If we have an inner product as above, then the norm induced by the inner product,

$$\|T\|_F := \sqrt{\langle S, T \rangle},$$

is called the Hilbert–Schmidt norm. Note that by virtue of (4.64), the Hilbert–Schmidt norm would satisfy

$$\|u \otimes v \otimes \dots \otimes w\| = \|u\|_1 \|v\|_2 \dots \|w\|_d, \quad (4.65)$$

a property that we will require of all tensor norms. As we mentioned in Example 4.1, such norms are called cross-norms. We have to extend (4.65) so that $\|\cdot\|$ is defined on not just rank-one tensors but on all finite sums of these. Naively we might just define it to be the sum of the norms of its rank-one summands, but that does not work as we get different values with different sums of rank-one tensors. Fortunately,

taking the infimum over all possible finite sums

$$\|T\|_v := \inf \left\{ \sum_{i=1}^r \|u_i\|_1 \|v_i\|_2 \cdots \|w_i\|_d : T = \sum_{i=1}^r u_i \otimes v_i \otimes \cdots \otimes w_i \right\} \tag{4.66}$$

fixes the issue and gives us a norm. This is the *nuclear norm*, special cases of which we have encountered in (3.45) and in Examples 3.11 and 4.1.

There is an alternative. We may also take the supremum over all rank-one tensors in the dual tensor space:

$$\|T\|_\sigma := \sup \left\{ \frac{|\varphi \otimes \psi \otimes \cdots \otimes \theta(T)|}{\|\varphi\|_{1^*} \|\psi\|_{2^*} \cdots \|\theta\|_{d^*}} : \varphi \in \mathbb{U}^*, \psi \in \mathbb{V}^*, \dots, \theta \in \mathbb{W}^* \right\}, \tag{4.67}$$

where $\|\cdot\|_{i^*}$ is the dual norm of $\|\cdot\|_i$ as defined by (3.44), $i = 1, \dots, d$. This is the *spectral norm*, special cases of which we have encountered in (3.19) and (3.46) and in Example 3.17. There are many more cross-norms (Diestel *et al.* 2008, Chapter 4) but these three are the best known. The nuclear and spectral norms are also special in that any cross-norm $\|\cdot\|$ must satisfy

$$\|T\|_\sigma \leq \|T\| \leq \|T\|_v \quad \text{for all } T \in \mathbb{U} \otimes \mathbb{V} \otimes \cdots \otimes \mathbb{W}; \tag{4.68}$$

and conversely any norm $\|\cdot\|$ that satisfies (4.68) must be a cross-norm (Ryan 2002, Proposition 6.1). In this sense, the nuclear and spectral norms, respectively, are the largest and smallest cross-norms. If \mathbb{V} is a Hilbert space, then by the Riesz representation theorem, a linear functional $\varphi: \mathbb{V} \rightarrow \mathbb{R}$ takes the form $\varphi = \langle v, \cdot \rangle$ for some $v \in \mathbb{V}$ and the Hilbert space norm on \mathbb{V} and its dual norm may be identified. Thus, if $\mathbb{U}, \mathbb{V}, \dots, \mathbb{W}$ are Hilbert spaces, then the spectral norm in (4.67) takes the form

$$\|T\|_\sigma = \sup \left\{ \frac{|\langle T, u \otimes v \otimes \cdots \otimes w \rangle|}{\|u\|_1 \|v\|_2 \cdots \|w\|_d} : u \in \mathbb{U}, v \in \mathbb{V}, \dots, w \in \mathbb{W} \right\}. \tag{4.69}$$

The definition of inner products on tensor spaces fits perfectly with other tensorial notions such as the tensor product basis and the Kronecker product discussed in Example 4.11. If $\mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_d$ are orthonormal bases for $\mathbb{U}, \mathbb{V}, \dots, \mathbb{W}$, then the tensor product basis $\mathcal{B}_1 \otimes \mathcal{B}_2 \otimes \cdots \otimes \mathcal{B}_d$ as defined in (4.43) is an orthonormal basis for $\mathbb{U} \otimes \mathbb{V} \otimes \cdots \otimes \mathbb{W}$. Here we have no need to assume finite dimension or separability: these orthonormal bases may be uncountable. If $\Phi_1: \mathbb{U} \rightarrow \mathbb{U}, \dots, \Phi_d: \mathbb{W} \rightarrow \mathbb{W}$ are orthogonal linear operators in the sense that

$$\langle \Phi_1(u), \Phi_1(u') \rangle_1 = \langle u, u' \rangle_1, \dots, \langle \Phi_d(w), \Phi_d(w') \rangle_d = \langle w, w' \rangle_d$$

for all $u, u' \in \mathbb{U}, \dots, w, w' \in \mathbb{W}$, then the Kronecker product $\Phi_1 \otimes \cdots \otimes \Phi_d: \mathbb{U} \otimes \cdots \otimes \mathbb{W} \rightarrow \mathbb{U} \otimes \cdots \otimes \mathbb{W}$ is an orthogonal linear operator, that is,

$$\langle \Phi_1 \otimes \cdots \otimes \Phi_d(S), \Phi_1 \otimes \cdots \otimes \Phi_d(T) \rangle = \langle S, T \rangle$$

for all $S, T \in \mathbb{U} \otimes \cdots \otimes \mathbb{W}$. For Kronecker products of operators defined on norm spaces, we will defer the discussion to the next example.

Let U, V, \dots, W be finite-dimensional inner product spaces. Let $\mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_d$ be orthonormal bases. Then, as we saw in Example 4.15, a d -tensor in $T \in U \otimes V \otimes \dots \otimes W$ may be represented by a d -hypermatrix $A \in \mathbb{R}^{m \times n \times \dots \times p}$, in which case the expressions for norms and inner product may be given in terms of A :

$$\|A\|_F = \left(\sum_{i=1}^m \sum_{j=1}^n \dots \sum_{k=1}^p |a_{ij\dots k}|^2 \right)^{1/2}, \quad \langle A, B \rangle = \sum_{i=1}^m \sum_{j=1}^n \dots \sum_{k=1}^p a_{ij\dots k} b_{ij\dots k},$$

$$\|A\|_v = \inf \left\{ \sum_{i=1}^r |\lambda_i| : A = \sum_{i=1}^r \lambda_i u_i \otimes v_i \otimes \dots \otimes w_i, \|u\| = \|v\| = \dots = \|w\| = 1 \right\},$$

$$\|A\|_\sigma = \sup \{ |\langle A, u \otimes v \otimes \dots \otimes w \rangle| : \|u\| = \|v\| = \dots = \|w\| = 1 \}.$$

Here the nuclear and spectral norms are in fact dual norms (not true in general)

$$\|A\|_{v^*} = \sup_{\|B\|_v \leq 1} |\langle A, B \rangle| = \|A\|_\sigma \tag{4.70}$$

and thus satisfy

$$|\langle A, B \rangle| \leq \|A\|_\sigma \|B\|_v.$$

For $d = 2$, these inner product and norms become the usual trace inner product $\langle A, B \rangle = \text{tr}(A^T B)$ and spectral, Frobenius, nuclear norms of matrices:

$$\|A\|_v = \sum_{i=1}^r \sigma_i(A), \quad \|A\|_F = \left(\sum_{i=1}^r \sigma_i(A)^2 \right)^{1/2}, \quad \|A\|_\sigma = \max_{i=1, \dots, r} \sigma_i(A),$$

where $\sigma_i(A)$ denotes the i th singular value of $A \in \mathbb{R}^{m \times n}$ and $r = \text{rank}(A)$.

We emphasize that the discussion in the above paragraph requires the bases in Example 4.15 to be orthonormal. Nevertheless, the values of the inner product and norms do not depend on which orthonormal bases we choose. In the terminology of Section 2 they are invariant under multilinear matrix multiplication by $(X_1, X_2, \dots, X_d) \in O(m) \times O(n) \times \dots \times O(p)$, or equivalently, they are defined on Cartesian tensors. More generally, if $\langle X_i v, X_i v \rangle_i = \langle v, v \rangle_i, i = 1, \dots, d$, then

$$\langle (X_1, X_2, \dots, X_d) \cdot A, (X_1, X_2, \dots, X_d) \cdot B \rangle = \langle A, B \rangle,$$

and thus $\|(X_1, X_2, \dots, X_d) \cdot A\|_F = \|A\|_F$; if $\|X_i v\|_i = \|v\|_i, i = 1, \dots, d$, then

$$\|(X_1, X_2, \dots, X_d) \cdot A\|_v = \|A\|_v, \quad \|(X_1, X_2, \dots, X_d) \cdot A\|_\sigma = \|A\|_\sigma.$$

This explains why, when discussing definition ① in conjunction with inner products or norms, we expect the change-of-basis matrices in the transformation rules to preserve these inner products or norms.

Inner product and norm spaces become enormously more interesting when they are completed into Hilbert and Banach spaces. The study of cross-norms was in fact started by Grothendieck (1955) and Schatten (1950) in order to define tensor products of Banach spaces, and this has grown into a vast subject (Defant and Floret

1993, Diestel *et al.* 2008, Light and Cheney 1985, Ryan 2002, Trèves 2006) that we are unable to survey at any reasonable level of detail. The following example is intended to convey an idea of how cross-norms allow one to complete the tensor spaces in a suitable manner.

Example 4.18 (tensor product of Hilbert and Banach spaces). We revisit our discussion in Example 4.1, properly defining topological tensor products $\widehat{\otimes}$ using continuous and integrable functions as illustrations.

- (i) If X and Y are compact Hausdorff topological spaces, then by the Stone–Weierstrass theorem, $C(X) \otimes C(Y)$ is a dense subset of $C(X \times Y)$ with respect to the uniform norm

$$\|f\|_\infty = \sup_{(x,y) \in X \times Y} |f(x,y)|. \tag{4.71}$$

- (ii) If X and Y are σ -finite measure spaces, then by the Fubini–Tonelli theorem, $L^1(X) \otimes L^1(Y)$ is a dense subset of $L^1(X \times Y)$ with respect to the L^1 -norm

$$\|f\|_1 = \int_{X \times Y} |f(x,y)| \, dx \, dy. \tag{4.72}$$

Here a tensor product of vector spaces is as defined in (4.39), and as we just saw in (4.17), there are several ways to equip it with a norm and with respect to any norm we may complete it (*i.e.* by adding the limits of all Cauchy sequences) to obtain Banach and Hilbert spaces out of norm and inner product spaces. The completed space depends on the choice of norms; with a judicious choice, we get

$$C(X) \widehat{\otimes}_\sigma C(Y) = C(X \times Y), \quad L^1(X) \widehat{\otimes}_\nu L^1(Y) = L^1(X \times Y), \tag{4.73}$$

as was first discovered in Grothendieck (1955). Here $\widehat{\otimes}_\sigma$ and $\widehat{\otimes}_\nu$ denote completion in the spectral and nuclear norm respectively and are called the *injective tensor product* and *projective tensor product* respectively. To be clear, the first equality in (4.73) says that if we equip $C(X) \otimes C(Y)$ with the spectral norm (4.67) and complete it to obtain $C(X) \widehat{\otimes}_\sigma C(Y)$, then the resulting space is $C(X \times Y)$ equipped with the uniform norm (4.71), and likewise for the second equality. In particular, (4.73) also tells us that the uniform and spectral norms are equal on $C(X \times Y)$, and likewise for the L^1 -norm in (4.72) and nuclear norm in (4.66). For $f \in L^1(X \times Y)$,

$$\begin{aligned} & \int_{X \times Y} |f(x,y)| \, dx \, dy \\ &= \inf \left\{ \sum_{i=1}^r \int_X |\varphi_i(x)| \, dx \int_Y |\psi_i(y)| \, dy : \varphi_i \in L^1(X), \psi_i \in L^1(Y), r \in \mathbb{N} \right\}, \end{aligned}$$

a fact that can be independently verified by considering simple functions²⁸ and taking limits.

²⁸ Finite linear combinations of characteristic functions of measurable sets.

These are examples of topological tensor products that involve completing the (algebraic) tensor product in (4.39) with respect to a choice of cross-norm to obtain a complete topological vector space. These also suggest why it is desirable to have a variety of different cross-norms, and with each a different topological tensor product, as the ‘right’ cross-norm to choose for a class of functions $\mathcal{F}(X)$ is usually the one that gives us $\mathcal{F}(X) \widehat{\otimes} \mathcal{F}(Y) = \mathcal{F}(X \times Y)$ as we discussed after (4.4). For example, to get the corresponding result for L^2 -functions, we have to use the Hilbert–Schmidt norm

$$L^2(X) \widehat{\otimes}_F L^2(Y) = L^2(X \times Y). \tag{4.74}$$

Essentially the proof relies on the fact that the completion of $L^2(X) \otimes L^2(Y)$ with respect to the Hilbert–Schmidt norm is the closure of $L^2(X) \otimes L^2(Y)$ as a subspace of $L^2(X \times Y)$, and as the orthogonal complement of $L^2(X) \otimes L^2(Y)$ is zero, its closure is the whole space (Light and Cheney 1985, Theorem 1.39). Nevertheless, such results are not always possible: there is no cross-norm that will complete $L^\infty(X) \otimes L^\infty(Y)$ into $L^\infty(X \times Y)$ for all σ -finite X and Y (Light and Cheney 1985, Theorem 1.53). Also, we should add that a ‘right’ cross-norm that guarantees (4.73) may be less interesting than a ‘wrong’ cross-norm that gives us a new tensor space. For instance, had we used $\widehat{\otimes}_v$ to form a tensor product of $C(X)$ and $C(Y)$, we would have obtained a smaller subset:

$$C(X) \widehat{\otimes}_v C(Y) \subsetneq C(X \times Y).$$

This smaller tensor space of continuous functions on $X \times Y$, more generally the tensor space $C(X_1) \widehat{\otimes}_v \cdots \widehat{\otimes}_v C(X_n)$, is called the *Varopoulos algebra* and it turns out to be very interesting and useful in harmonic analysis (Varopoulos 1965, 1967).

A point worth highlighting is the difference between $L^2(X) \otimes L^2(Y)$ and $L^2(X) \widehat{\otimes}_F L^2(Y)$. While the former contains only *finite* sums of separable functions

$$\sum_{i=1}^r f_i \otimes g_i$$

with $f_i \in L^2(X)$, $g_i \in L^2(Y)$, the latter includes all convergent infinite series of separable functions

$$\sum_{i=1}^\infty f_i \otimes g_i := \lim_{r \rightarrow \infty} \sum_{i=1}^r f_i \otimes g_i,$$

that is, the right-hand side converges to some limit in $L^2(X \times Y)$ in the Hilbert–Schmidt norm $\|\cdot\|_F$. The equality in (4.74) says that every function in $L^2(X \times Y)$ is given by such a limit, but if we had taken completion with respect to some other norms such as $\|\cdot\|_v$ or $\|\cdot\|_\sigma$, their topological tensor products $L^2(X) \widehat{\otimes}_v L^2(Y)$ or $L^2(X) \widehat{\otimes}_\sigma L^2(Y)$ would in general be smaller or larger than $L^2(X \times Y)$ respectively. Whatever the choice of $\widehat{\otimes}$, the (algebraic) tensor product $\mathbb{U} \otimes \mathbb{V} \otimes \cdots \otimes \mathbb{W}$ should be regarded as the subset of all finite-rank tensors in the topological tensor product $\mathbb{U} \widehat{\otimes} \mathbb{V} \widehat{\otimes} \cdots \widehat{\otimes} \mathbb{W}$.

The tensor product of Hilbert spaces invariably refers to the topological tensor product with respect to the Hilbert–Schmidt norm because the result is always a Hilbert space. In particular, the $\widehat{\otimes}$ in (4.13) should be interpreted as $\widehat{\otimes}_F$. However, as we pointed out above, other topological tensor products with respect to other cross-norms may also be very interesting.

Example 4.19 (trace-class, Hilbert–Schmidt, compact operators). For a separable Hilbert space \mathbb{H} with inner product $\langle \cdot, \cdot \rangle$ and induced norm $\| \cdot \|$, completing $\mathbb{H} \otimes \mathbb{H}^*$ with respect to the nuclear, Hilbert–Schmidt and spectral norms, we obtain different types of bounded linear operators on \mathbb{H} as follows:

$$\begin{aligned} \text{trace-class} \quad \mathbb{H} \widehat{\otimes}_v \mathbb{H}^* &= \left\{ \Phi \in \mathcal{B}(\mathbb{H}) : \sum_{i \in I} \sum_{j \in I} |\langle \Phi(e_i), f_j \rangle| < \infty \right\}, \\ \text{Hilbert–Schmidt} \quad \mathbb{H} \widehat{\otimes}_F \mathbb{H}^* &= \left\{ \Phi \in \mathcal{B}(\mathbb{H}) : \sum_{i \in I} \|\Phi(e_i)\|^2 < \infty \right\}, \\ \text{compact} \quad \mathbb{H} \widehat{\otimes}_\sigma \mathbb{H}^* &= \left\{ \Phi \in \mathcal{B}(\mathbb{H}) : \begin{array}{l} X \subseteq \mathbb{H} \text{ bounded} \\ \Rightarrow \overline{\Phi(X)} \subseteq \mathbb{H} \text{ compact} \end{array} \right\}. \end{aligned}$$

The series convergence is understood to mean ‘for some orthonormal bases $\{e_i : i \in I\}$ and $\{f_i : i \in I\}$ of \mathbb{H} ’, although for trace-class operators the condition could be simplified to $\sum_{i \in I} \langle \Phi(e_i), e_j \rangle < \infty$ provided that ‘for some’ is replaced by ‘for all’. See Schaefer and Wolff (1999, p. 278) for the trace-class result and Trèves (2006, Theorems 48.3) for the compact result.

If \mathbb{H} is separable, then such operators are characterized by their having a *Schmidt decomposition*:

$$\Phi = \sum_{i=1}^{\infty} \sigma_i u_i \otimes v_i^* \quad \text{or} \quad \Phi(x) = \sum_{i=1}^{\infty} \sigma_i \langle v_i, x \rangle u_i \quad \text{for all } x \in \mathbb{H}, \tag{4.75}$$

where $\{u_i : i \in \mathbb{N}\}$ and $\{v_i : i \in \mathbb{N}\}$ are orthonormal sets, $\sigma_i \geq 0$, and

$$\sum_{i=1}^{\infty} \sigma_i < \infty, \quad \sum_{i=1}^{\infty} \sigma_i^2 < \infty, \quad \lim_{i \rightarrow \infty} \sigma_i = 0$$

depending on whether they are trace-class, Hilbert–Schmidt or compact, respectively (Reed and Simon 1980, Chapter VI, Sections 5–6). Note that \mathbb{H} and \mathbb{H}^* are naturally isomorphic by the Riesz representation theorem and we have identified v_i^* with $\langle v_i, \cdot \rangle$. As expected, we also have

$$\|\Phi\|_v = \sum_{i=1}^{\infty} \sigma_i, \quad \|\Phi\|_F = \left(\sum_{i=1}^{\infty} \sigma_i^2 \right)^{1/2}, \quad \|\Phi\|_\sigma = \sup_{i \in \mathbb{N}} \sigma_i.$$

In other words, this is the infinite-dimensional version of the relation between the various norms and matrix singular values in Example 4.17, and the Schmidt decomposition is an infinite-dimensional generalization of singular value decomposition.

Unlike the finite-dimensional case, where one may freely speak of the nuclear, Frobenius and spectral norms of any matrix, for infinite-dimensional \mathbb{H} we have

$$\mathbb{H} \otimes \mathbb{H}^* \subsetneq \mathbb{H} \widehat{\otimes}_\nu \mathbb{H}^* \subsetneq \mathbb{H} \widehat{\otimes}_F \mathbb{H}^* \subsetneq \mathbb{H} \widehat{\otimes}_\sigma \mathbb{H}^*, \quad \|\Phi\|_\sigma \leq \|\Phi\|_F \leq \|\Phi\|_\nu,$$

and the inclusions are strict, for example, a compact operator Φ may have $\|\Phi\|_\nu = \infty$. By our discussion at the end of Example 4.18, $\mathbb{H} \otimes \mathbb{H}^*$ is the subset of finite-rank operators in any of these larger spaces. The inequality relating the three norms is a special case of (4.68) as $\|\cdot\|_F$ is a cross-norm. On $\mathbb{H} \widehat{\otimes}_\nu \mathbb{H}^*$, nuclear and spectral norms are dual norms as in the finite-dimensional case (4.70):

$$\|\Phi\|_{\nu^*} = \sup_{\|\Psi\|_\nu \leq 1} |\langle \Phi, \Psi \rangle_F| = \|\Phi\|_\sigma,$$

where the inner product will be defined below in (4.78).

The reason for calling $\Phi \in \mathbb{H} \widehat{\otimes}_\nu \mathbb{H}^*$ trace-class is that its *trace* is well-defined and always finite:

$$\text{tr}(\Phi) := \sum_{i=1}^\infty \langle \Phi(e_i), e_i \rangle$$

for any orthonormal basis $\{e_i : i \in \mathbb{N}\}$ of \mathbb{H} . As trace is independent of the choice of orthonormal basis, it is truly a tensorial notion. Moreover $\text{tr}(|\Phi|) = \|\Phi\|_\nu$, where $|\Phi| := \sqrt{\Phi^* \Phi}$. To see why the expression above is considered a trace, we may, as discussed in Example 4.15, represent Φ as an infinite-dimensional matrix with respect to the chosen orthonormal basis,

$$\Phi = \sum_{i=1}^\infty \sum_{j=1}^\infty a_{ij} e_i \otimes e_j^*, \tag{4.76}$$

where $A = (a_{ij})_{i,j=1}^\infty \in l^2(\mathbb{N}^2)$, and now observe that

$$\text{tr}(\Phi) = \sum_{k=1}^\infty \langle \Phi(e_k), e_k \rangle = \sum_{k=1}^\infty \left\langle \sum_{i=1}^\infty a_{ik} e_i, e_k \right\rangle = \sum_{k=1}^\infty a_{kk}. \tag{4.77}$$

One may show that an operator is trace-class if and only if it is a product of two Hilbert–Schmidt operators. A consequence is that

$$\langle \Phi, \Psi \rangle_F := \text{tr}(\Phi^* \Psi) \tag{4.78}$$

is always finite and defines an inner product on $\mathbb{H} \widehat{\otimes}_F \mathbb{H}^*$ that gives $\|\cdot\|_F$ as its induced norm:

$$\langle \Phi, \Phi \rangle_F = \text{tr}(\Phi^* \Phi) = \|\Phi\|_F^2.$$

While $(\mathbb{H} \widehat{\otimes}_\nu \mathbb{H}^*, \|\cdot\|_\nu)$, $(\mathbb{H} \widehat{\otimes}_F \mathbb{H}^*, \|\cdot\|_F)$, $(\mathbb{H} \widehat{\otimes}_\sigma \mathbb{H}^*, \|\cdot\|_\sigma)$ are all Banach spaces, only $\mathbb{H} \widehat{\otimes}_F \mathbb{H}^*$ is a Hilbert space with the inner product in (4.78). So a topological tensor product of Hilbert spaces may or may not be a Hilbert space: it depends on the choice of cross-norms. Also, the choice of norm is critical in determining

whether we have a Banach space: for instance, $\mathbb{H} \widehat{\otimes}_\nu \mathbb{H}^*$ is complete with respect to $\|\cdot\|_\nu$ but not $\|\cdot\|_\sigma$.

Taking topological tensor products of infinite-dimensional spaces has some unexpected properties. When \mathbb{H} is infinite-dimensional, the identity operator $I_{\mathbb{H}} \in \mathcal{B}(\mathbb{H})$ is not in $\mathbb{H} \widehat{\otimes} \mathbb{H}^*$ regardless of whether we use $\widehat{\otimes}_\nu$, $\widehat{\otimes}_F$ or $\widehat{\otimes}_\sigma$, as $I_{\mathbb{H}}$ is non-compact (and thus neither Hilbert–Schmidt nor trace-class). Hence $\mathbb{H} \widehat{\otimes} \mathbb{H}^* \neq \mathcal{B}(\mathbb{H})$, no matter which notion of topological tensor product $\widehat{\otimes}$ we use. Incidentally, the strong and weak duals of $\mathcal{B}(\mathbb{H})$, *i.e.* the set of linear functionals on $\mathcal{B}(\mathbb{H})$ continuous under the strong and weak operator topologies, are equal to each other and to $\mathbb{H} \otimes \mathbb{H}^*$, *i.e.* the set of finite-rank operators (Dunford and Schwartz 1988, Chapter IV, Section 1). Taking the (continuous) dual of the compact operators gives us the trace-class operators

$$(\mathbb{H} \widehat{\otimes}_\sigma \mathbb{H}^*)^* \cong \mathbb{H} \widehat{\otimes}_\nu \mathbb{H}^*, \quad (4.79)$$

so taking dual can change the topological tensor product.

We have framed our discussion in terms of $\mathbb{H} \widehat{\otimes} \mathbb{H}^*$ as this is the most common scenario, but it may be extended to other situations.

Different Hilbert spaces. It applies verbatim to $\mathbb{H}_1 \widehat{\otimes} \mathbb{H}_2^*$ as operators in $\mathcal{B}(\mathbb{H}_1; \mathbb{H}_2)$.

Covariant and contravariant. We covered the mixed case $\mathbb{H}_1 \widehat{\otimes} \mathbb{H}_2^*$, but the covariant case $\mathbb{H}_1^* \widehat{\otimes} \mathbb{H}_2^*$ and contravariant case $\mathbb{H}_1 \widehat{\otimes} \mathbb{H}_2$ follow from replacing \mathbb{H}_1 with \mathbb{H}_1^* or \mathbb{H}_2 with \mathbb{H}_2^* .

Banach spaces. Note that (4.66) and (4.67) are defined without any reference to inner products, so for $\mathbb{H}_1 \widehat{\otimes}_\sigma \mathbb{H}_2^*$ and $\mathbb{H}_1 \widehat{\otimes}_\nu \mathbb{H}_2^*$ we do not need \mathbb{H}_1 and \mathbb{H}_2 to be Hilbert spaces; compact and trace-class operators may be defined for any pair of Banach spaces, although the latter are usually called nuclear operators in this context.

Higher order. One may define order- $d \geq 3$ analogues of bounded, compact, Hilbert–Schmidt, trace-class operators in a straightforward manner, but corresponding results are more difficult; one reason is that Schmidt decomposition (4.75) no longer holds (Bényi and Torres 2013, Cobos, Kühn and Peetre 1992, 1999).

We remind the reader that many of the topological vector spaces considered in Examples 4.1 and 4.2, such as $S(X)$, $C^\infty(X)$, $C_c^\infty(X)$, $H(X)$, $S'(X)$, $E'(X)$, $D'(X)$ and $H'(X)$, are not Banach spaces (and thus not Hilbert spaces) but so-called *nuclear spaces*. Nevertheless, similar ideas apply to yield topological tensor products. In fact nuclear spaces have the nice property that topological tensor products with respect to the nuclear and spectral norms, *i.e.* $\widehat{\otimes}_\nu$ and $\widehat{\otimes}_\sigma$, are always equal and thus independent of the choice of cross-norms.

We mention two applications of Example 4.19 in machine learning (Mercer kernels) and quantum mechanics (density matrices).

Example 4.20 (Mercer kernels). Hilbert–Schmidt operators, in the form of Mercer kernels (Mercer 1909), are of special significance to machine learning. Here $\mathbb{H} = L^2(X)$ for a set X of machine learning significance (where a training set is sampled from) that has some appropriate measure (usually Borel) and topology (usually locally compact). For any L^2 -integral kernel $K: X \times X \rightarrow \mathbb{R}$, i.e. $K \in L^2(X \times X) = L^2(X) \widehat{\otimes}_{\mathbb{F}} L^2(X)$, the integral operator $\Phi: L^2(X) \rightarrow L^2(X)$,

$$\Phi(f)(x) = \int_X K(x, y)f(y) dy,$$

is Hilbert–Schmidt; in fact every Hilbert–Schmidt operator on $L^2(X)$ arises in this manner and we always have $\|\Phi\|_{\mathbb{F}} = \|K\|_{L^2}$ (Conway 1990, p. 267). By our discussion above, Φ has a Schmidt decomposition as in (4.75):

$$\Phi = \sum_{i=1}^{\infty} \sigma_i \varphi_i \otimes \psi_i^*, \tag{4.80}$$

where $\psi^*(f) = \langle \psi, f \rangle = \int_X \psi(x)f(x) dx$ by the Riesz representation. Hence

$$K = \sum_{i=1}^{\infty} \sigma_i \varphi_i \otimes \psi_i, \tag{4.81}$$

and we see that the difference between (4.80) and (4.81) is just one of covariance and contravariance: every $K \in L^2(X) \widehat{\otimes}_{\mathbb{F}} L^2(X)$ gives us a $\Phi \in L^2(X) \widehat{\otimes}_{\mathbb{F}} L^2(X)^*$ and *vice versa*. Mercer’s kernel theorem is essentially the statement that if K is continuous and symmetric positive semidefinite, i.e. $[K(x_i, x_j)]_{i,j=1}^n \in \mathbb{S}_+^n$ for any $x_1, \dots, x_n \in X$ and any $n \in \mathbb{N}$, then φ_i and ψ_i may be chosen to be equal in (4.80) and (4.81). In this case we obtain a *feature map*:

$$F: X \rightarrow l^2(\mathbb{N}), \quad x \mapsto (\sqrt{\sigma_i} \varphi_i(x))_{i=1}^{\infty},$$

and in this context $l^2(\mathbb{N})$ is called the *feature space*. It follows from (4.81) that $K(x, y) = \langle F(x), F(y) \rangle_{l^2(\mathbb{N})}$. Assuming that $\sigma_i > 0$ for all $i \in \mathbb{N}$, the subspace of $L^2(X)$ defined by

$$L_K^2(X) := \left\{ \sum_{i=1}^{\infty} a_i \varphi_i \in L^2(X) : (a_i / \sqrt{\sigma_i})_{i=1}^{\infty} \in l^2(\mathbb{N}) \right\},$$

equipped with the inner product

$$\left\langle \sum_{i=1}^{\infty} a_i \varphi_i, \sum_{i=1}^{\infty} b_i \varphi_i \right\rangle_K := \sum_{i=1}^{\infty} \frac{a_i b_i}{\sigma_i} = a^T \Sigma^{-1} b,$$

is called the *reproducing kernel Hilbert space* associated with the kernel K (Cucker and Smale 2002). The last expression is given in terms of infinite-dimensional matrices, as described in Example 4.5, i.e. $a, b \in l^2(\mathbb{N})$ and $\Sigma \in l^2(\mathbb{N} \times \mathbb{N})$, to show its connection with the finite-dimensional case.




Figure 4.3. Depiction of the feature map $F: X \rightarrow l^2(\mathbb{N})$ in the context of support-vector machines.

These notions are used to great effect in machine learning (Hofmann, Schölkopf and Smola 2008, Schölkopf and Smola 2002, Steinwart and Christmann 2008). For instance, as depicted in Figure 4.3, the feature map allows one to unfurl a complicated set of points x_1, \dots, x_n in a lower-dimensional space X into a simpler set of points $F(x_1), \dots, F(x_n)$, say, one that may be partitioned with a hyperplane, in an infinite-dimensional feature space $l^2(\mathbb{N})$.

The previous example is about symmetric positive semidefinite Hilbert–Schmidt operators; the next one will be about Hermitian positive semidefinite trace-class operators.

Example 4.21 (density operators). These are sometimes called density matrices, and are important alike in quantum physics (Blum 1996), quantum chemistry (Davidson 1976) and quantum computing (Nielsen and Chuang 2000). Let \mathbb{H} be a Hilbert space, assumed separable for notational convenience, and let $\langle \cdot, \cdot \rangle$ be its inner product. A trace-class operator $\rho \in \mathbb{H} \widehat{\otimes}_v \mathbb{H}^*$ is called a density operator if it is Hermitian $\rho^* = \rho$, positive semidefinite $|\rho| = \rho$, and of unit trace $\text{tr}(\rho) = 1$ (or, equivalently, unit nuclear norm as $\text{tr}(\rho) = \|\rho\|_v$ in this case). By the same considerations as in Example 4.20, its Schmidt decomposition may be chosen to have $u_i = v_i$, that is,

$$\rho = \sum_{i=1}^{\infty} \sigma_i v_i \otimes v_i^*, \quad \text{tr}(\rho) = \sum_{i=1}^{\infty} \sigma_i = 1. \quad (4.82)$$

A key reason density operators are important is that the postulates of quantum mechanics may be reformulated in terms of them (Nielsen and Chuang 2000, Section 2.4.2). Instead of representing a quantum state as a one-dimensional

subspace of \mathbb{H} spanned by a non-zero vector $v \in \mathbb{H}$, we let it be represented by the projection operator onto that subspace, that is,

$$v \otimes v^*: \mathbb{H} \rightarrow \mathbb{H}, \quad w \mapsto v^*(w)v = \langle v, w \rangle v,$$

where we have assumed that v has unit norm (if not, normalize). This is evidently a density operator and such rank-one density operators are called *pure states*. An advantage is that one may now speak of *mixed states*: they are infinite convex linear combination of pure states as in (4.82), *i.e.* the density operators. Note that this is not possible with states represented as vectors in \mathbb{H} : a linear combination of two vectors would just be another vector in \mathbb{H} . There is a well-known condition for telling pure states apart from mixed states without involving a Schmidt decomposition, namely, ρ is pure if and only if $\text{tr}(\rho^2) = 1$ and ρ is mixed if and only if $\text{tr}(\rho^2) < 1$. Consequently, the value of $\text{tr}(\rho^2)$ is called the *purity*.

In this density operator formulation, a quantum system is completely described by a single density operator, and for d quantum systems described by density operators ρ_1, \dots, ρ_d , the composite system is described by their Kronecker product $\rho_1 \otimes \dots \otimes \rho_d$ as defined in Example 4.11. The reader should compare these with (a) and (b) on page 658. A Kronecker product of density operators remains a density operator because of (4.50), (4.53) and the fact that (4.54) holds for trace-class operators:

$$\text{tr}(\rho_1 \otimes \dots \otimes \rho_d) = \text{tr}(\rho_1) \cdots \text{tr}(\rho_d).$$

We end with a word about the role of compact operators $\mathbb{H} \widehat{\otimes}_\sigma \mathbb{H}^*$ in this discussion. For any density operator ρ , the linear functional $\omega_\rho: \mathcal{B}(\mathbb{H}) \rightarrow \mathbb{C}$, $\Phi \mapsto \text{tr}(\rho\Phi)$ satisfies $\omega_\rho(I_{\mathbb{H}}) = 1$ and $\omega_\rho(|\Phi|) \geq 0$ for all $\Phi \in \mathcal{B}(\mathbb{H})$. A linear functional ω with the last two properties is called a *positive linear functional*, abbreviated $\omega \geq 0$. It turns out that every positive linear functional on the space of compact operators is of the form ω_ρ for some density operator ρ (Landsman 2017, Corollaries 4.14 and 4.15):

$$\{\omega \in (\mathbb{H} \widehat{\otimes}_\sigma \mathbb{H}^*)^*: \omega \geq 0\} \cong \{\rho \in \mathbb{H} \widehat{\otimes}_\nu \mathbb{H}^*: \rho \text{ density operator}\},$$

a very useful addendum to (4.79). With this observation, we may now use positive linear functionals to represent quantum states. The advantage is that it permits ready generalization: $\mathbb{H} \widehat{\otimes}_\sigma \mathbb{H}^*$ has the structure of a C^* -algebra and one may replace it with other C^* -algebra \mathcal{A} and represent quantum states as positive linear functionals $\omega: \mathcal{A} \rightarrow \mathbb{C}$.

The evolution of the definition of quantum states from vectors in Hilbert spaces to density operators to positive linear functionals on C^* -algebras is not unlike the three increasingly sophisticated definitions of tensors or the three increasingly abstract definitions of tensor products. As in the case of tensors and tensor products, each definition of quantum states is useful in its own way, and all three remain in use.

Our discussion of tensor product will not be complete without mentioning the tensor algebra. We will make this the last example of this section.

Example 4.22 (tensor algebra). Let \mathbb{V} be a vector space and $v \in \mathbb{V}$. We introduce the shorthand

$$\mathbb{V}^{\otimes d} := \overbrace{\mathbb{V} \otimes \cdots \otimes \mathbb{V}}^{d \text{ copies}}, \quad v^{\otimes d} := \overbrace{v \otimes \cdots \otimes v}^{d \text{ copies}}$$

for any $d \in \mathbb{N}$. We also define $\mathbb{V}^{\otimes 0} := \mathbb{R}$ and $v^{\otimes 0} := 1$. There is a risk of construing erroneous relations from notation like this. We caution that $\mathbb{V}^{\otimes d}$ is *not* the set of tensors of the form $v^{\otimes d}$, as we pointed out in (4.42), but neither is it the set of linear combinations of such tensors; $\mathbb{V}^{\otimes d}$ contains all finite linear combinations of any $v_1 \otimes \cdots \otimes v_d$, including but not limited to those of the form $v^{\otimes d}$. As $\mathbb{V}^{\otimes d}$, $d = 0, 1, 2, \dots$, are all vector spaces, we may form their direct sum to obtain an infinite-dimensional vector space:

$$\mathbb{T}(\mathbb{V}) = \bigoplus_{k=0}^{\infty} \mathbb{V}^{\otimes k} = \mathbb{R} \oplus \mathbb{V} \oplus \mathbb{V}^{\otimes 2} \oplus \mathbb{V}^{\otimes 3} \oplus \cdots, \quad (4.83)$$

called the *tensor algebra* of \mathbb{V} . Those unfamiliar with direct sums may simply regard $\mathbb{T}(\mathbb{V})$ as the set of all *finite* sums of the tensors of any order:

$$\mathbb{T}(\mathbb{V}) = \left\{ \sum_{k=0}^d T_k : T_k \in \mathbb{V}^{\otimes k}, d \in \mathbb{N} \right\}. \quad (4.84)$$

What does it mean to add a tensor $T \in \mathbb{V}^{\otimes 2}$ to another tensor $T' \in \mathbb{V}^{\otimes 5}$, *i.e.* of different orders? This is exactly like the addition of polyads we discussed at the beginning of this section. Namely there is no formula for actually adding T to T' , a sum like $T + T'$ cannot be further simplified, it is just intended as a place holder for the two tensors T and T' , we could have written instead $(0, 0, T, 0, 0, T', 0, \dots)$, and indeed another possible way to view $\mathbb{T}(\mathbb{V})$ is that it is the set of all sequences with finitely many non-zero terms:

$$\{(T_k)_{k=0}^{\infty} : T_k \in \mathbb{V}^{\otimes k}, T_k = 0 \text{ for all but finitely many } k\}.$$

However, we prefer the form in (4.84). The reason $\mathbb{T}(\mathbb{V})$ is called an algebra is that it inherits a product operation given by the product of tensors in (4.41):

$$\mathbb{V}^{\otimes d} \times \mathbb{V}^{\otimes d'} \ni (T, T') \mapsto T \otimes T' \in \mathbb{V}^{\otimes(d+d')},$$

which can be extended to $\mathbb{T}(\mathbb{V})$ by

$$\left(\sum_{j=0}^d T_j \right) \otimes \left(\sum_{k=0}^{d'} T'_k \right) := \sum_{j=0}^d \sum_{k=0}^{d'} T_j \otimes T'_k.$$

Nothing we have discussed up to this point involves a basis of \mathbb{V} , so it all applies to modules as well. If we now take advantage of the fact that \mathbb{V} is a vector space and further assume that it is finite-dimensional, then \mathbb{V} has a basis $\mathcal{B} = \{e_1, \dots, e_n\}$

and we may represent any element in $T(V)$ as

$$\sum_{k=0}^d \sum_{i_1, i_2, \dots, i_k=1}^n a_{i_1 i_2 \dots i_k} e_{i_1} \otimes e_{i_2} \otimes \dots \otimes e_{i_k}$$

with $d \in \mathbb{N}$. Another way to represent this is as a *non-commutative polynomial*

$$f(X_1, X_2, \dots, X_n) = \sum_{k=0}^d \sum_{i_1, i_2, \dots, i_k=1}^n a_{i_1 i_2 \dots i_k} X_{i_1} X_{i_2} \dots X_{i_k},$$

in n non-commutative variables X_1, X_2, \dots, X_n . Here non-commutative means $X_i X_j \neq X_j X_i$ if $i \neq j$, *i.e.* just like (4.30), and we require other arithmetic rules for manipulating X_1, X_2, \dots, X_n to mirror those in (4.40). The set of all such polynomials is usually denoted $\mathbb{R}\langle X_1, \dots, X_n \rangle$. Hence, as algebras,

$$T(V) \cong \mathbb{R}\langle X_1, \dots, X_n \rangle, \tag{4.85}$$

where the isomorphism is given by the map that takes $e_{i_1} \otimes e_{i_2} \otimes \dots \otimes e_{i_k} \mapsto X_{i_1} X_{i_2} \dots X_{i_k}$. Such a representation also applies to any d -tensor in $V^{\otimes d}$ as a special case, giving us a *homogeneous* non-commutative polynomial

$$\sum_{i_1, i_2, \dots, i_d=1}^n a_{i_1 i_2 \dots i_d} X_{i_1} X_{i_2} \dots X_{i_d},$$

where all terms have degree d . This is a superior representation compared to merely representing the tensor as a hypermatrix $(a_{i_1 i_2 \dots i_d})_{i_1, i_2, \dots, i_d=1}^n \in \mathbb{R}^{n \times n \times \dots \times n}$ because, like usual polynomials, we can take derivatives and integrals of non-commutative polynomials and evaluate them on matrices, that is, we can take $A_1, A_2, \dots, A_n \in \mathbb{R}^{m \times m}$ and plug them into $f(X_1, X_2, \dots, X_n)$ to get $f(A_1, A_2, \dots, A_n) \in \mathbb{R}^{m \times m}$. This last observation alone leads to a rich subject, often called ‘non-commutative sums-of-squares’, that has many engineering applications (Helton and Putinar 2007).

What if we need to speak of an infinite sum of tensors of different orders? This is not just of theoretical interest; as we will see in Example 4.45, a multipole expansion is such an infinite sum. A straightforward way to do this would be to replace the direct sums in (4.83) with direct products; the difference, if the reader recalls, is that while the elements of a direct sum are zero for all but a finite number of summands, those of a direct product may contain an infinite number of non-zero summands. If there are just a finite number of vector spaces V_1, \dots, V_d , the direct sum $V_1 \oplus \dots \oplus V_d$ and the direct product $V_1 \times \dots \times V_d$ are identical, but if we have infinitely many vector spaces, a direct product is much larger than a direct sum. For instance, a direct sum of countably many *two*-dimensional vector spaces $\bigoplus_{k \in \mathbb{N}} V_k$ has countable dimension whereas a direct product $\prod_{k \in \mathbb{N}} V_k$ has uncountable dimension.

Nevertheless, instead of introducing new notation for the direct product of $\mathbb{V}^{\otimes k}$, $k = 0, 1, 2, \dots$, we may simply consider the dual vector space of $T(\mathbb{V})$,

$$T(\mathbb{V})^* = \prod_{k=0}^{\infty} \mathbb{V}^{*\otimes k} = \left\{ \sum_{k=0}^{\infty} \varphi_k : \varphi_k \in \mathbb{V}^{*\otimes k} \right\}, \tag{4.86}$$

noting that taking duals turns a direct sum into a direct product. We call this the *dual tensor algebra*. While it is true that $(\mathbb{V}^{\otimes k})^* = \mathbb{V}^{*\otimes k}$ – this follows from the ‘calculus of tensor products’ that we will present in Example 4.31 – we emphasize that the dual tensor algebra of any non-zero vector space \mathbb{V} is not the same as the tensor algebra of its dual space \mathbb{V}^* , that is,

$$T(\mathbb{V})^* \neq T(\mathbb{V}^*).$$

Indeed, the dual vector space of any infinite-dimensional vector space must have dimension strictly larger than that of the vector space (Jacobson 1975, Chapter IX, Section 5). In this case, if \mathbb{V} is finite-dimensional, then by what we discussed earlier, $T(\mathbb{V}^*)$ has countable dimension whereas $T(\mathbb{V})^*$ has uncountable dimension. Although what we have defined in (4.86) is the direct product of $\mathbb{V}^{*\otimes k}$, $k = 0, 1, 2, \dots$, the direct product of $\mathbb{V}^{\otimes k}$, $k = 0, 1, 2, \dots$, is simply given by the dual tensor algebra of the dual space:

$$T(\mathbb{V}^*)^* = \prod_{k=0}^{\infty} \mathbb{V}^{\otimes k} = \left\{ \sum_{k=0}^{\infty} T_k : T_k \in \mathbb{V}^{\otimes k} \right\}.$$

Just as $T(\mathbb{V})$ may be regarded as the ring of non-commutative polynomials, the same arguments that led to (4.85) also give

$$T(\mathbb{V}^*)^* \cong \mathbb{R}\langle\langle X_1, \dots, X_n \rangle\rangle,$$

that is, $T(\mathbb{V}^*)^*$ may be regarded as the ring of non-commutative power series.

One issue with these constructions is that taking direct sums results in a space $T(\mathbb{V})$ that is sometimes too small whereas taking direct products leads to a space $T(\mathbb{V}^*)^*$ that is often too big. However, these constructions are purely algebraic; if \mathbb{V} is equipped with a norm or inner product, we can often use that to obtain a tensor algebra of the desired size. Suppose \mathbb{V} has an inner product $\langle \cdot, \cdot \rangle$; then as in Example 4.17, defining

$$\langle u \otimes v \otimes \dots \otimes w, u' \otimes v' \otimes \dots \otimes w' \rangle := \langle u, u' \rangle \langle v, v' \rangle \dots \langle w, w' \rangle$$

for any $u \otimes v \otimes \dots \otimes w, u' \otimes v' \otimes \dots \otimes w' \in \mathbb{V}^{\otimes d}$, any $d \in \mathbb{N}$, and extending bilinearly to all of $T(\mathbb{V})$ makes it into an inner product space. We will denote the inner products on $\mathbb{V}^{\otimes d}$ and $T(\mathbb{V})$ by $\langle \cdot, \cdot \rangle$ and the corresponding induced norms by $\| \cdot \|$; there is no ambiguity since they are all consistently defined. Taking completion gives us a Hilbert space

$$\widehat{T}(\mathbb{V}) := \left\{ \sum_{k=0}^{\infty} T_k : T_k \in \mathbb{V}^{\otimes k}, \sum_{k=0}^{\infty} \|T_k\|^2 < \infty \right\}$$

with inner product

$$\left\langle \sum_{k=0}^{\infty} T_k, \sum_{k=0}^{\infty} T'_k \right\rangle := \sum_{k=0}^{\infty} \langle T_k, T'_k \rangle.$$

Alternatively, one may also describe $\widehat{T}(\mathbb{V})$ as the *Hilbert space direct sum* of the Hilbert spaces $\mathbb{V}^{\otimes k}$, $k = 0, 1, 2, \dots$. If \mathcal{B} is an orthonormal basis of \mathbb{V} , then the tensor product basis (4.43), denoted $\mathcal{B}^{\otimes k}$, is an orthonormal basis on $\mathbb{V}^{\otimes k}$, and $\bigcup_{k=0}^{\infty} \mathcal{B}^{\otimes k}$ is a countable orthonormal basis on $\widehat{T}(\mathbb{V})$, i.e. it is a separable Hilbert space. This can be used in a ‘tensor trick’ for linearizing a non-linear function with convergent Taylor series. For example, for any $x, y \in \mathbb{R}^n$,

$$\begin{aligned} \exp(\langle x, y \rangle) &= \sum_{k=0}^{\infty} \frac{1}{k!} \langle x, y \rangle^k = \sum_{k=0}^{\infty} \frac{1}{k!} \langle x^{\otimes k}, y^{\otimes k} \rangle \\ &= \left\langle \sum_{k=0}^{\infty} \frac{1}{\sqrt{k!}} x^{\otimes k}, \sum_{k=0}^{\infty} \frac{1}{\sqrt{k!}} y^{\otimes k} \right\rangle = \langle S(x), S(y) \rangle, \end{aligned}$$

where the map

$$S: \mathbb{R}^n \rightarrow \widehat{T}(\mathbb{R}^n), \quad S(x) = \sum_{k=0}^{\infty} \frac{1}{\sqrt{k!}} x^{\otimes k}$$

is well-defined since

$$\|S(x)\|^2 = \sum_{k=0}^{\infty} \left\| \frac{1}{\sqrt{k!}} x^{\otimes k} \right\|^2 = \sum_{k=0}^{\infty} \frac{1}{k!} \|x\|^{2k} = \exp(\|x\|^2) < \infty.$$

This was used, with \sin in place of \exp , in an elegant proof (Krivine 1979) of Grothendieck’s inequality that also yields

$$K_G \leq \frac{\pi}{2 \log(1 + \sqrt{2})} \approx 1.78221,$$

the upper bound we saw in Example 3.17. This is in fact the best explicit upper bound for the Grothendieck constant over \mathbb{R} , although it is now known that it is not sharp.


4.3. Tensors via the universal factorization property

We now come to the most abstract and general way to define tensor products and thus tensors: the universal factorization property. It is both a potent tool and an illuminating concept. Among other things, it unifies almost every notion of tensors we have discussed thus far. For instance, an immediate consequence is that multilinear maps (Definitions 3.1 and 3.3) and sums of separable functions (Definitions 4.4 and 4.9) are one and the same notion, unifying definitions ② and ③.

The straightforward way to regard the universal factorization property is that it is a property of multilinear maps. It says that the multilinearity in any multilinear map

Φ is the result of a special multilinear map σ_{\otimes} that is an omnipresent component of all multilinear maps: $\Phi = F \circ \sigma_{\otimes}$ – once σ_{\otimes} is factored out of Φ , what remains, *i.e.* the component F , is just a linear map. Any multilinearity contained in Φ is all due to σ_{\otimes} .

However, because of its universal nature, we may turn the property on its head and use it to define the tensor product: a tensor product is whatever satisfies the universal factorization property. More precisely, whatever vector space one can put in the box below that makes the diagram



commute for all vector spaces \mathbb{W} with Φ multilinear and F linear is defined to be a tensor product of $\mathbb{V}_1, \dots, \mathbb{V}_d$. One may show that this defines the tensor product (and the map σ) uniquely up to vector space isomorphism.


We now supply the details. The special multilinear map in question is $\sigma_{\otimes}: \mathbb{V}_1 \times \cdots \times \mathbb{V}_d \rightarrow \mathbb{V}_1 \otimes \cdots \otimes \mathbb{V}_d$,

$$\sigma_{\otimes}(v_1, \dots, v_d) = v_1 \otimes \cdots \otimes v_d, \tag{4.87}$$

called the *Segre map* or Segre embedding.²⁹ Here the tensor space $\mathbb{V}_1 \otimes \cdots \otimes \mathbb{V}_d$ is as defined in Section 4.2 and (4.87) defines a multilinear map by virtue of (4.40).

Universal factorization property. Every multilinear map is the composition of a linear operator with the Segre map.

This is also called the universal mapping property, or the universal property of the tensor product. Formally, it says that if Φ is a multilinear map from $\mathbb{V}_1 \times \cdots \times \mathbb{V}_d$ into a vector space \mathbb{W} , then Φ determines a unique linear map F_{Φ} from $\mathbb{V}_1 \otimes \cdots \otimes \mathbb{V}_d$ into \mathbb{W} , that makes the following diagram commute:



or, equivalently, Φ can be factored as

$$\Phi = F_{\Phi} \circ \sigma_{\otimes}, \tag{4.89}$$

where F_{Φ} is a ‘linear component’ dependent on Φ and σ_{\otimes} is a ‘multilinear component’ independent of Φ . The component σ_{\otimes} is universal, *i.e.* common to all

²⁹ There is no standard notation. It has been variously denoted as σ (Harris 1995), as φ (Bourbaki 1998, Lang 2002), as \otimes (Conrad 2018) and as Seg (Landsberg 2012).

multilinear maps. This can be stated in less precise but more intuitive terms as follows.

- (a) All the ‘multilinearity’ in a multilinear map Φ may be factored out of Φ , leaving behind only the ‘linearity’ that is encapsulated in F_Φ .
- (b) The ‘multilinearity’ extracted from any multilinear map is identical, that is, all multilinear maps are multilinear because they contain a copy of the Segre map σ_\otimes .

Actually, σ_\otimes does depend on d and so we should have said it is universal for d -linear maps. An immediate consequence of the universal factorization property is that

$$\mathcal{M}^d(\mathbb{V}_1, \dots, \mathbb{V}_d; \mathbb{W}) = \mathcal{L}(\mathbb{V}_1 \otimes \dots \otimes \mathbb{V}_d; \mathbb{W}).$$

In principle one could use the universal factorization property to avoid multilinear maps entirely and discuss only linear maps, but of course there is no reason to do that: as we saw in Section 3, multilinearity is a very useful notion in its own right. The universal factorization property should be viewed as a way to move back and forth between the multilinear realm and the linear one. While we have assumed contravariant tensors for simplicity, the discussions above apply verbatim to mixed tensor spaces $\mathbb{V}_1 \otimes \dots \otimes \mathbb{V}_p \otimes \mathbb{V}_{p+1}^* \otimes \dots \otimes \mathbb{V}_d^*$ and with modules in place of vector spaces.

Pure mathematicians accustomed to commutative diagrams would swear by (4.88) but it expresses the same thing as (4.89), which is more palatable to applied and computational mathematicians as it reminds us of matrix factorizations like Example 2.13, where we factor a matrix $A \in \mathbb{R}^{m \times n}$ into $A = QR$ with an orthogonal component $Q \in O(m)$ that captures the ‘orthogonality’ in A , leaving behind a triangular component R . There is one big difference, though: both Q and R depend on A , but in (4.89), whatever our choice of Φ , the multilinear component will always be σ_\otimes . The next three examples will bear witness to this intriguing property.

Example 4.23 (trilinear functionals). The universal factorization property explains why (3.9) and (4.38) look alike. A trilinear functional $\tau: \mathbb{U} \times \mathbb{V} \times \mathbb{W} \rightarrow \mathbb{R}$ can be factored as

$$\tau = F_\tau \circ \sigma_\otimes \tag{4.90}$$

and the multilinearity of σ_\otimes in (4.38) accounts for that of τ in (3.9). For example, take the first equation in (3.9). We may view it as arising from

$$\begin{aligned} \tau(\lambda u + \lambda' u', v, w) &= F_\tau(\sigma_\otimes(\lambda u + \lambda' u', v, w)) && \text{by (4.90),} \\ &= F_\tau((\lambda u + \lambda' u') \otimes v \otimes w) && \text{by (4.87),} \\ &= F_\tau(\lambda u \otimes v \otimes w + \lambda' u' \otimes v \otimes w) && \text{by (4.38),} \\ &= \lambda F_\tau(u \otimes v \otimes w) + \lambda' F_\tau(u' \otimes v \otimes w) && F_\tau \text{ is linear,} \\ &= \lambda F_\tau(\sigma_\otimes(u, v, w)) + \lambda' F_\tau(\sigma_\otimes(u', v, w)) && \text{by (4.87),} \\ &= \lambda \tau(u, v, w) + \lambda' \tau(u', v, w) && \text{by (4.90),} \end{aligned}$$

and similarly for the second and third equations in (3.9). More generally, if we reread Section 3.2 with the hindsight of this section, we will realize that it is simply a discussion of the universal factorization property without invoking the \otimes symbol.

We next see how the universal factorization property ties together the three matrix products that have made an appearance in this article.

Example 4.24 (matrix, Hadamard and Kronecker products). Let us consider the standard matrix product in (2.5) and Hadamard product in (2.4) on 2×2 matrices. Let $\mu, \eta: \mathbb{R}^{2 \times 2} \times \mathbb{R}^{2 \times 2} \rightarrow \mathbb{R}^{2 \times 2}$ be, respectively,

$$\begin{aligned} \mu \left(\begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix}, \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix} \right) &= \begin{bmatrix} a_{11}b_{11} + a_{12}b_{21} & a_{11}b_{12} + a_{12}b_{22} \\ a_{21}b_{11} + a_{22}b_{21} & a_{21}b_{12} + a_{22}b_{22} \end{bmatrix}, \\ \eta \left(\begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix}, \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix} \right) &= \begin{bmatrix} a_{11}b_{11} & a_{12}b_{12} \\ a_{21}b_{21} & a_{22}b_{22} \end{bmatrix}. \end{aligned}$$

They look nothing alike and, as discussed in Section 2, are of entirely different natures. But the universal factorization property tells us that, by virtue of the fact that both are bilinear operators from $\mathbb{R}^{2 \times 2} \times \mathbb{R}^{2 \times 2}$ to $\mathbb{R}^{2 \times 2}$, the ‘bilinearity’ in μ and η is the same, encapsulated in the Segre map $\sigma_{\otimes}: \mathbb{R}^{2 \times 2} \times \mathbb{R}^{2 \times 2} \rightarrow \mathbb{R}^{2 \times 2} \otimes \mathbb{R}^{2 \times 2} \cong \mathbb{R}^{4 \times 4}$,

$$\begin{aligned} \sigma_{\otimes} \left(\begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix}, \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix} \right) &= \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \otimes \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix} \\ &= \begin{bmatrix} a_{11}b_{11} & a_{11}b_{12} & a_{12}b_{11} & a_{12}b_{12} \\ a_{11}b_{21} & a_{11}b_{22} & a_{12}b_{21} & a_{12}b_{22} \\ a_{21}b_{11} & a_{21}b_{12} & a_{22}b_{11} & a_{22}b_{12} \\ a_{21}b_{21} & a_{21}b_{22} & a_{22}b_{21} & a_{22}b_{22} \end{bmatrix}, \end{aligned}$$

i.e. the Kronecker product in Example 4.10(ii). The difference between μ and η is due entirely to their linear components $F_{\mu}, F_{\eta}: \mathbb{R}^{4 \times 4} \rightarrow \mathbb{R}^{2 \times 2}$,

$$\begin{aligned} F_{\mu} \left(\begin{bmatrix} c_{11} & c_{12} & c_{13} & c_{14} \\ c_{21} & c_{22} & c_{23} & c_{24} \\ c_{31} & c_{32} & c_{33} & c_{34} \\ c_{41} & c_{42} & c_{43} & c_{44} \end{bmatrix} \right) &= \begin{bmatrix} c_{11} + c_{23} & c_{12} + c_{24} \\ c_{31} + c_{43} & c_{32} + c_{44} \end{bmatrix}, \\ F_{\eta} \left(\begin{bmatrix} c_{11} & c_{12} & c_{13} & c_{14} \\ c_{21} & c_{22} & c_{23} & c_{24} \\ c_{31} & c_{32} & c_{33} & c_{34} \\ c_{41} & c_{42} & c_{43} & c_{44} \end{bmatrix} \right) &= \begin{bmatrix} c_{11} & c_{14} \\ c_{41} & c_{44} \end{bmatrix}, \end{aligned}$$

which are indeed very different. As a sanity check, the reader may like to verify from these formulas that

$$\mu = F_{\mu} \circ \sigma_{\otimes}, \quad \eta = F_{\eta} \circ \sigma_{\otimes}.$$

More generally, the same argument shows that any products we define on 2×2 matrices must share something in common, namely the Kronecker product that

takes a pair of 2×2 matrices to a 4×4 matrix. The Kronecker product accounts for the ‘bilinearity’ of all conceivable matrix products.

We will include an infinite-dimensional example for variety.

Example 4.25 (convolution I). We recall the convolution bilinear operator

$$*: L^1(\mathbb{R}^n) \times L^1(\mathbb{R}^n) \rightarrow L^1(\mathbb{R}^n), \quad (f, g) \mapsto f * g$$

from Example 3.7. As we saw in Example 4.18, $L^1(\mathbb{R}^n) \otimes L^1(\mathbb{R}^n) = L^1(\mathbb{R}^n \times \mathbb{R}^n)$ when we interpret \otimes as $\widehat{\otimes}_v$, the topological tensor product with respect to the nuclear norm, and so the Segre map is

$$\sigma_{\otimes}: L^1(\mathbb{R}^n) \times L^1(\mathbb{R}^n) \rightarrow L^1(\mathbb{R}^n \times \mathbb{R}^n), \quad (f, g) \mapsto f \otimes g,$$

where, as a reminder, $f \otimes g(x, y) = f(x)g(y)$, where \otimes is interpreted as the separable product. Commutative diagrams like (4.88) are useful when we have three or more maps between multiple spaces, as they allow us to assemble all maps into a single picture, which helps to keep track of what gets mapped where. Here we have

$$\begin{array}{ccc} L^1(\mathbb{R}^n) \times L^1(\mathbb{R}^n) & \xrightarrow{\sigma_{\otimes}} & L^1(\mathbb{R}^n \times \mathbb{R}^n) \\ & \searrow * & \downarrow F_* \\ & & L^1(\mathbb{R}^n) \end{array}$$

and we would like to find the unique linear map F_* that makes the diagram commute, i.e. $f * g = F_*(f \otimes g)$. Writing out this expression, we see that $F_* = G \circ H$ is a composition of the two linear maps defined by

$$\begin{aligned} G: L^1(\mathbb{R}^n \times \mathbb{R}^n) &\rightarrow L^1(\mathbb{R}^n), & G(K)(x) &:= \int_{\mathbb{R}^n} K(x, y) \, dy, \\ H: L^1(\mathbb{R}^n \times \mathbb{R}^n) &\rightarrow L^1(\mathbb{R}^n \times \mathbb{R}^n), & H(K)(x, y) &:= K(x - y, y) \end{aligned}$$

for all $K \in L^1(\mathbb{R}^n \times \mathbb{R}^n)$.

Strictly speaking, the space $\mathbb{V}_1 \otimes \cdots \otimes \mathbb{V}_d$ has already been fixed by Definition 4.9. In order to allow some flexibility in (4.88), we will have to allow for $\mathbb{V}_1 \otimes \cdots \otimes \mathbb{V}_d$ to be replaced by any vector space \mathbb{T} and σ_{\otimes} by any multilinear map

$$\sigma_{\mathbb{T}}: \mathbb{V}_1 \times \cdots \times \mathbb{V}_d \rightarrow \mathbb{T}$$

as long as the universal factorization property holds for this choice of \mathbb{T} and $\sigma_{\mathbb{T}}$, that is,

$$\begin{array}{ccc} \mathbb{V}_1 \times \cdots \times \mathbb{V}_d & \xrightarrow{\sigma_{\mathbb{T}}} & \mathbb{T} \\ & \searrow \text{multilinear } \Phi & \downarrow \text{linear } F_{\Phi} \\ & & \mathbb{W} \end{array} \tag{4.91}$$

Definition 4.26 (tensors via universal factorization). Let \mathbb{T} be a vector space and let $\sigma_{\mathbb{T}}: \mathbb{V}_1 \times \cdots \times \mathbb{V}_d \rightarrow \mathbb{T}$ be a multilinear map with the universal factorization property, that is, for any vector space \mathbb{W} and any multilinear map $\Phi: \mathbb{V}_1 \times \cdots \times \mathbb{V}_d \rightarrow \mathbb{W}$, there is a unique linear map $F_{\Phi}: \mathbb{T} \rightarrow \mathbb{W}$, called the *linearization* of Φ , such that the factorization $\Phi = F_{\Phi} \circ \sigma_{\mathbb{T}}$ holds. Then \mathbb{T} is unique up to vector space isomorphisms and is called a *tensor space* or more precisely a *tensor product* of the vector spaces $\mathbb{V}_1, \dots, \mathbb{V}_d$. An element of \mathbb{T} is called a *d-tensor*.

As usual, one may easily incorporate covariance and contravariance by replacing some of the vector spaces with their dual spaces.

In Example 4.23, $\mathbb{T} = \mathbb{U} \otimes \mathbb{V} \otimes \mathbb{W}$, as defined in Definition 4.9, $\sigma_{\mathbb{T}} = \sigma_{\otimes}$, as defined in (4.87), and we recover (4.88). But for the subsequent two examples we have in fact already used Definition 4.26 implicitly. In Example 4.24, $\mathbb{T} = \mathbb{R}^{4 \times 4}$ and $\sigma_{\mathbb{T}}: \mathbb{R}^{2 \times 2} \times \mathbb{R}^{2 \times 2} \rightarrow \mathbb{T}$ takes a pair of 2×2 matrices to their Kronecker product. In Example 4.25, $\mathbb{T} = L^1(\mathbb{R}^n \times \mathbb{R}^n)$ and $\sigma_{\mathbb{T}}: L^1(\mathbb{R}^n) \times L^1(\mathbb{R}^n) \rightarrow \mathbb{T}$ takes a pair of L^1 -functions to their separable product.

Definition 4.26 is intended to capture as many notions of tensor products as possible. The Segre map $\sigma_{\mathbb{T}}$ may be an abstract tensor product of vectors, an outer product of vectors in \mathbb{R}^n , a Kronecker product of operators or matrices, a separable product of functions, etc.; the \otimes in (4.91) could be abstract tensor products of vector spaces, Kronecker products of operator spaces, topological tensor products of topological vector spaces, etc. In the last case we require the maps $\Phi, \sigma_{\mathbb{T}}, F_{\Phi}$ to be also *continuous*; without this requirement the topological tensor product just becomes the algebraic tensor product. Recall that continuity takes a simple form for a multilinear map, being equivalent to the statement that its spectral norm as defined in (3.19) is finite, that is, continuous multilinear maps are precisely bounded multilinear maps. Let us verify that the maps in Example 4.25 are continuous.

Example 4.27 (convolution II). We follow the same notation in Example 4.25 except that we will write $B_*: L^1(\mathbb{R}^n) \times L^1(\mathbb{R}^n) \rightarrow L^1(\mathbb{R}^n)$, $(f, g) \mapsto f * g$ for convolution as a bilinear map. We show that $\|B_*\|_{\sigma} = \|\sigma_{\otimes}\|_{\sigma} = \|F_*\|_{\sigma} = 1$ and thus these maps are continuous. First recall from Example 3.7 that $\|B_*\|_{\sigma} = 1$. Next, since $\|f \otimes g\|_1 = \|f\|_1 \|g\|_1$, we have

$$\|\sigma_{\otimes}\|_{\sigma} = \sup_{f, g \neq 0} \frac{\|f \otimes g\|_1}{\|f\|_1 \|g\|_1} = 1.$$

Lastly, for $K \in L^1(\mathbb{R}^n \times \mathbb{R}^n)$,

$$\|F_*(K)\|_1 = \int_{\mathbb{R}^n} \int_{\mathbb{R}^n} |K(x - y, y)| \, dy \, dx = \|K\|_1$$

by Fubini and a change of variables. Thus we get

$$\|F_*\|_{\sigma} = \sup_{K \neq 0} \frac{\|F_*(K)\|_1}{\|K\|_1} = 1.$$

Note that even though we use the *spectral* norm to ascertain continuity of the linear and multilinear maps between spaces, the tensor product space above and in Example 4.25 is formed with $\widehat{\otimes}_\nu$, the topological tensor product with respect to the *nuclear* norm. Indeed, $\widehat{\otimes}_\nu$ is an amazing tensor product in that for any norm spaces U, V, W , if we have the commutative diagram

$$\begin{array}{ccc}
 U \times V & \xrightarrow{\sigma_\otimes} & U \widehat{\otimes}_\nu V \\
 & \searrow \Phi & \downarrow F_\Phi \\
 & & W
 \end{array}$$

then the bilinear map $\Phi: U \times V \rightarrow W$ is continuous if and only if its linearization $F_\Phi: U \widehat{\otimes}_\nu V \rightarrow W$ is continuous (Defant and Floret 1993, Chapter 3).

We will now examine a non-example to see that Definition 4.26 is not omnipotent. The tensor product of Hilbert spaces is an important exception that does not satisfy the universal factorization property. The following example is adapted from Garrett (2010).

Example 4.28 (no universal factorization for Hilbert spaces). We begin with a simple observation. For a finite-dimensional vector space V , the map defined by $\beta: V \times V^* \rightarrow \mathbb{R}, (v, \varphi) \mapsto \varphi(v)$ is clearly bilinear. So we may apply universal factorization to get

$$\begin{array}{ccc}
 V \times V^* & \xrightarrow{\sigma_\otimes} & V \otimes V^* \\
 & \searrow \beta & \downarrow \text{tr} \\
 & & \mathbb{R}
 \end{array}$$

The linear map that we obtained, $\text{tr}: V \otimes V^* \rightarrow \mathbb{R}$, is in fact a coordinate-free way to define trace. To see this, note that if we choose $V = \mathbb{R}^n$, then for any $a \in \mathbb{R}^n$ and $b^T \in \mathbb{R}^{1 \times n} = \mathbb{R}^{n*}$,

$$\beta(a, b^T) = b^T a = \text{tr}(b^T a) = \text{tr}(ab^T) = \text{tr}(a \otimes b).$$

This observation is key to our subsequent discussion. We will see that when V is replaced by an infinite-dimensional Hilbert space and we require all maps to be continuous, then there is no continuous linear map that can take the place of trace.

Let H be a separable infinite-dimensional Hilbert space. As discussed in Examples 4.18 and 4.19, the tensor product of H and H^* is interpreted to be the topological tensor product $H \widehat{\otimes}_F H^*$ with respect to the Hilbert–Schmidt norm $\|\cdot\|_F$. So we must have $\sigma_\otimes: H \times H^* \rightarrow H \widehat{\otimes}_F H^*, (v, \varphi) \mapsto v \otimes \varphi$. Take $W = \mathbb{C}$ and consider the continuous bilinear functional $\beta: H \times H^* \rightarrow \mathbb{C}, (v, \varphi) \mapsto \varphi(v)$. We claim that there is no continuous linear map $F_\beta: H \widehat{\otimes}_F H^* \rightarrow \mathbb{C}$ such that $\beta = F_\beta \circ \sigma_\otimes$, that is,

we cannot make the diagram

$$\begin{array}{ccc} \mathbb{H} \times \mathbb{H}^* & \xrightarrow{\sigma_{\otimes}} & \mathbb{H} \widehat{\otimes}_F \mathbb{H}^* \\ & \searrow \beta & \downarrow F_{\beta} ? \\ & & \mathbb{C} \end{array}$$

commute. To this end, suppose F_{β} exists. Take any orthonormal basis $\{e_i : i \in \mathbb{N}\}$. Then

$$\beta(e_i, e_j^*) = \delta_{ij}, \quad \sigma_{\otimes}(e_i, e_j^*) = e_i \otimes e_j^*,$$

and since $\beta = F_{\beta} \circ \sigma_{\otimes}$, we must have

$$F_{\beta}(e_i \otimes e_j^*) = \delta_{ij}.$$

Now take any $\Phi \in \mathbb{H} \widehat{\otimes}_F \mathbb{H}^*$ and express it as in (4.76). Then since F_{β} is continuous and linear, we have

$$F_{\beta}(\Phi) = \sum_{k=1}^{\infty} a_{kk} = \text{tr}(\Phi),$$

that is, F_{β} must be the trace as defined in (4.77). This is a contradiction as trace is unbounded on $\mathbb{H} \widehat{\otimes}_F \mathbb{H}^*$; just take

$$\Phi = \sum_{k=1}^{\infty} k^{-1} e_k \otimes e_k^*,$$

which is Hilbert–Schmidt as $\|\Phi\|_F^2 = \sum_{k=1}^{\infty} k^{-2} < \infty$ but $\text{tr}(\Phi) = \sum_{k=1}^{\infty} k^{-1} = \infty$. Hence no such map exists.

A plausible follow-up question is that instead of $\widehat{\otimes}_F$, why not take topological tensor products with respect to the nuclear norm $\widehat{\otimes}_v$? As we mentioned at the end of Example 4.27, this automatically guarantees continuity of all maps. The problem, as we pointed out in Example 4.19, is that the resulting tensor space, *i.e.* the trace-class operators $\mathbb{H} \widehat{\otimes}_v \mathbb{H}^*$, is not a Hilbert space. In fact, one may show that no such Hilbert space exists. Suppose there is a Hilbert space \mathbb{T} and a Segre map $\sigma_{\mathbb{T}} : \mathbb{H} \times \mathbb{H}^* \rightarrow \mathbb{T}$ so that the universal factorization property in (4.91) holds. Then we may choose $\mathbb{W} = \mathbb{H} \widehat{\otimes}_F \mathbb{H}^*$ and $\Phi = \sigma_{\otimes}$ to get

$$\begin{array}{ccc} \mathbb{H} \times \mathbb{H}^* & \xrightarrow{\sigma_{\mathbb{T}}} & \mathbb{T} \\ & \searrow \sigma_{\otimes} & \downarrow F_{\otimes} \\ & & \mathbb{H} \widehat{\otimes}_F \mathbb{H}^* \end{array}$$

from which one may deduce (see Garrett 2010 for details) that if F_{\otimes} is a continuous linear map, then

$$\langle \cdot, \cdot \rangle_{\mathbb{T}} = c \langle \cdot, \cdot \rangle_F$$

for some $c > 0$, that is, F_{\otimes} is up to a constant multiple an isometry (Hilbert space isomorphism). So \mathbb{T} and $\mathbb{H} \widehat{\otimes}_{\mathbb{F}} \mathbb{H}^*$ are essentially the same Hilbert space up to scaling and thus the same trace argument above shows that \mathbb{T} does not satisfy the universal factorization property.

While the goal of this example is to demonstrate that Hilbert spaces generally do not satisfy the universal factorization property, there is a point worth highlighting in this construction. As we saw at the beginning, if we do not care about continuity, then everything goes through without a glitch: the *evaluation bilinear functional* $\beta: \mathbb{V} \times \mathbb{V}^* \rightarrow \mathbb{R}, (v, \varphi) \mapsto \varphi(v)$ satisfies the universal factorization property

$$\begin{array}{ccc} \mathbb{V} \times \mathbb{V}^* & \xrightarrow{\sigma_{\otimes}} & \mathbb{V} \otimes \mathbb{V}^* \\ & \searrow \beta & \downarrow \text{tr} \\ & & \mathbb{R} \end{array}$$

with the linear functional $\text{tr}: \mathbb{V} \otimes \mathbb{V}^* \rightarrow \mathbb{R}$ given by

$$\text{tr} \left(\sum_{i=1}^r v_i \otimes \varphi_i \right) = \sum_{i=1}^r \varphi_i(v_i)$$

for any $r \in \mathbb{N}$. In the language of Definition 4.26, trace is the linearization of the evaluation bilinear functional, a definition adopted in both algebra (Lang 2002, Proposition 5.7) and analysis (Ryan 2002, Section 1.3). Note that here \mathbb{V} can be any vector space, not required to be finite-dimensional or complete or normed, and the maps involved are not required to be continuous. The trace as defined above is always a finite sum and thus finite-valued.

The next example clarifies an occasional ambiguity in the definition of higher-order derivatives of vector-valued functions.

Example 4.29 (linearizing higher-order derivatives). Recall from Example 3.2 that the d th derivative of $f: \Omega \rightarrow \mathbb{W}$ at a point $v \in \Omega \subseteq \mathbb{V}$ is a multilinear map denoted $D^d f(v)$ as in (3.23). Applying the universal factorization property

$$\begin{array}{ccc} \overbrace{\mathbb{V} \times \cdots \times \mathbb{V}}^{d \text{ copies}} & \xrightarrow{\sigma_{\otimes}} & \overbrace{\mathbb{V} \otimes \cdots \otimes \mathbb{V}}^{d \text{ copies}} \\ & \searrow D^d f(v) & \downarrow \partial^d f(v) \\ & & \mathbb{W} \end{array}$$

we obtain a linear map on a tensor space,

$$\partial^d f(v): \mathbb{V}^{\otimes d} \rightarrow \mathbb{W},$$

that may be regarded as an alternative candidate for the d th derivative of f . In the language of Definition 4.26, $\partial^d f(v)$ is the linearization of $D^d f(v)$. While we have

denoted them differently for easy distinction, it is not uncommon to see $D^d f(v)$ and $\partial^d f(v)$ used interchangeably, sometimes in the same sentence.

Take the log barrier $f: \mathbb{S}_{++}^n \rightarrow \mathbb{R}, f(X) = -\log \det(X)$ discussed in Examples 3.6 and 3.16. We have

$$\begin{aligned} D^2 f(X): \mathbb{S}^n \times \mathbb{S}^n &\rightarrow \mathbb{R}, & (H_1, H_2) &\mapsto \text{tr}(X^{-1}H_1X^{-1}H_2), \\ \sigma_{\otimes}: \mathbb{S}^n \times \mathbb{S}^n &\rightarrow \mathbb{S}^n \otimes \mathbb{S}^n, & (H_1, H_2) &\mapsto H_1 \otimes H_2, \end{aligned}$$

where the latter denotes the Kronecker product as discussed in Example 4.11. The universal factorization property

$$\begin{array}{ccc} \mathbb{S}^n \times \mathbb{S}^n & \xrightarrow{\sigma_{\otimes}} & \mathbb{S}^n \otimes \mathbb{S}^n \\ & \searrow D^d f(X) & \downarrow \partial^d f(X) \\ & & \mathbb{R} \end{array}$$

gives us

$$\partial^2 f(X): \mathbb{S}^n \otimes \mathbb{S}^n \rightarrow \mathbb{R}, \quad (H_1, H_2) = \text{vec}(X^{-1})^\top (H_1 \otimes H_2) \text{vec}(X^{-1}).$$

Note that here it does not matter whether we use vec or vec_ℓ , since the matrices involved are all symmetric.

Taylor’s theorem for a vector-valued function $f: \Omega \subseteq \mathbb{V} \rightarrow \mathbb{W}$ in Example 3.2 may alternatively be expressed in the form

$$\begin{aligned} f(v) = f(v_0) &+ [\partial f(v_0)](v - v_0) + \frac{1}{2}[\partial^2 f(v_0)](v - v_0)^{\otimes 2} + \dots \\ &\dots + \frac{1}{d!}[\partial^d f(v_0)](v - v_0)^{\otimes d} + R(v - v_0), \end{aligned}$$

where $v_0 \in \Omega$ and $\|R(v - v_0)\|/\|v - v_0\|^d \rightarrow 0$ as $v \rightarrow v_0$. Here the ‘Taylor coefficients’ are linear maps $\partial^d f(v): \mathbb{V}^{\otimes d} \rightarrow \mathbb{W}$. For the important special case $\mathbb{W} = \mathbb{R}$, *i.e.* a real-valued function $f: \mathbb{V} \rightarrow \mathbb{R}$, we have that $\partial^d f(v): \mathbb{V}^{\otimes d} \rightarrow \mathbb{R}$ is a linear functional. So the Taylor coefficients are covariant d -tensors

$$\partial^d f(v) \in \mathbb{V}^{*\otimes d}.$$

By our discussions in Example 4.22 and assuming that $f \in C^\infty(\Omega)$, we obtain an element of the dual tensor algebra

$$\partial f(v) := \sum_{k=0}^{\infty} \frac{1}{k!} \partial^k f(v) \in \mathbb{T}(\mathbb{V})^*$$

that we will call the tensor Taylor series.

A slight variation is when \mathbb{V} is equipped with an inner product $\langle \cdot, \cdot \rangle$. We may use the Riesz representation theorem to write, for $f \in C^{d+1}(\Omega)$,

$$f(v) = f(v_0) + \langle \partial f(v_0), v - v_0 \rangle + \frac{1}{2} \langle \partial^2 f(v_0), (v - v_0)^{\otimes 2} \rangle + \dots \\ \dots + \frac{1}{d!} \langle \partial^d f(v_0), (v - v_0)^{\otimes d} \rangle + R(v - v_0), \tag{4.92}$$

where the inner products on $\mathbb{V}^{\otimes d}$ are defined as in (4.64). Here the Taylor coefficients are contravariant d -tensors $\partial^d f(v) \in \mathbb{V}^{\otimes d}$. If f is furthermore analytic, we may rewrite this as an inner product between maps into the tensor algebra:

$$f(v) = \left\langle \sum_{k=0}^{\infty} \frac{1}{k!} \partial^k f(v_0), \sum_{k=0}^{\infty} (v - v_0)^{\otimes k} \right\rangle = \langle \partial f(v_0), S(v - v_0) \rangle, \tag{4.93}$$

where the maps $\partial f: \Omega \rightarrow \widehat{\mathbb{T}}(\mathbb{V})$ and $S: \Omega \rightarrow \widehat{\mathbb{T}}(\mathbb{V})$ are given by

$$\partial f(v) = \sum_{k=0}^{\infty} \frac{1}{k!} \partial^k f(v), \quad S(v) = \sum_{k=0}^{\infty} v^{\otimes k}.$$

The tensor geometric series $S(v)$ is clearly a well-defined element of $\widehat{\mathbb{T}}(\mathbb{V})$ whenever $\|v\| < 1$ and the tensor Taylor series $\partial f(v)$ is well-defined as long as $\|\partial^k f(v)\| \leq B$ for some uniform bound $B > 0$. We will see this in action when we discuss multipole expansion in Example 4.45, which is essentially (4.93) applied to $f(v) = 1/\|v\|$.

Another utility of the universal factorization property is in defining various vector-valued objects, *i.e.* taking values in a vector space \mathbb{V} . The generality here is deliberate, intended to capture a myriad of possibilities for \mathbb{V} , which could refer to anything from \mathbb{C} as a two-dimensional vector space over \mathbb{R} to an infinite-dimensional function space, a space of matrices or hypermatrices, a Hilbert or Banach or Fréchet space, a space of linear or multilinear operators, a tensor space itself, *etc.*

Example 4.30 (vector-valued objects). For any real vector space \mathbb{V} and any set X , we denote the set of all functions taking values in \mathbb{V} by

$$\mathbb{V}^X := \{f: X \rightarrow \mathbb{V}\}.$$

This is a real vector space in the obvious way: $\lambda f + \lambda' f'$ is defined to be the function whose value at $x \in X$ is given by the vector $\lambda f(x) + \lambda' f'(x) \in \mathbb{V}$. It is also consistent with our notation \mathbb{R}^X for the vector space of all real-valued functions on X in Section 4.1. Consider the bilinear map

$$\Phi: \mathbb{R}^X \times \mathbb{V} \rightarrow \mathbb{V}^X, \quad (f, v) \mapsto f \cdot v,$$

where, for any fixed $f \in \mathbb{R}^X$ and $v \in \mathbb{V}$, the function $f \cdot v$ is defined to be

$f \cdot v: X \rightarrow \mathbb{V}, x \mapsto f(x)v$. Applying the universal factorization property, we have

$$\begin{array}{ccc} \mathbb{R}^X \times \mathbb{V} & \xrightarrow{\sigma_\otimes} & \mathbb{R}^X \otimes \mathbb{V} \\ & \searrow \Phi & \downarrow F_\Phi \\ & & \mathbb{V}^X \end{array} \quad (4.94)$$

with the linearization of Φ given by

$$F_\Phi \left(\sum_{i=1}^r f_i \otimes v_i \right) = \sum_{i=1}^r f_i \cdot v_i.$$

It is straightforward (Ryan 2002, Section 1.5) to show that F_Φ gives an isomorphism $\mathbb{R}^X \otimes \mathbb{V} \cong \mathbb{V}^X$. The standard practice is to identify the two spaces, that is, replace \cong with $=$ and take $f \otimes v$ to mean $f \cdot v$ as defined above.

In reality we are often more interested in various subspaces $\mathcal{F}(X) \subseteq \mathbb{R}^X$ whose elements satisfy properties such as continuity, differentiability, integrability, summability, *etc.*, and we would like to have a corresponding space of vector-valued functions $\mathcal{F}(X; \mathbb{V}) \subseteq \mathbb{V}^X$ that also have these properties. If \mathbb{V} is finite-dimensional, applying the universal factorization property (4.94) with the subspaces $\mathcal{F}(X)$ and $\mathcal{F}(X; \mathbb{V})$ in place of \mathbb{R}^X and \mathbb{V}^X gives us

$$\mathcal{F}(X; \mathbb{V}) = \mathcal{F}(X) \otimes \mathbb{V},$$

and this already leads to some useful constructions.

Extension of scalars. Any real vector space \mathbb{U} has a *complexification* $\mathbb{U}_{\mathbb{C}}$ with elements given by

$$u_1 \otimes 1 + u_2 \otimes i, \quad u_1, u_2 \in \mathbb{U},$$

and scalar multiplication by $a + bi \in \mathbb{C}$ given by

$$(a + ib)(u_1 \otimes 1 + u_2 \otimes i) = (au_1 - bu_2) \otimes 1 + (bu_1 + au_2) \otimes i.$$

This construction is just

$$\mathbb{W} \otimes \mathbb{C} = \mathbb{W}_{\mathbb{C}}.$$

Polynomial matrices. In mechanical engineering, one of the simplest and most fundamental system is the mass–spring–damper model described by

$$q(\lambda) = \lambda^2 M + \lambda C + K,$$

where $M, C, K \in \mathbb{R}^{n \times n}$ represent mass, damping and stiffness respectively. This a polynomial matrix, which may be viewed as either a matrix of polynomials or a polynomial with matrix coefficients. This construction is just

$$\mathbb{R}[\lambda] \otimes \mathbb{R}^{n \times n} = (\mathbb{R}[\lambda])^{n \times n} = (\mathbb{R}^{n \times n})[\lambda].$$

C^k - and L^p -vector fields. In (4.14) we viewed the solution to a system of partial differential equations as an element of $C^2(\mathbb{R}^3) \widehat{\otimes} C^1[0, \infty) \otimes \mathbb{R}^3$, and in (4.29) we viewed an L^2 -vector field as an element of $L^2(\mathbb{R}^3) \otimes \mathbb{C}^2$. These are consequences of

$$C^k(\Omega) \otimes \mathbb{R}^n = C^k(\Omega; \mathbb{R}^n) \quad \text{and} \quad L^p(\Omega) \otimes \mathbb{C}^n = L^p(\Omega; \mathbb{C}^n)$$

where the latter spaces refer to the sets of vector-valued functions

$$f = (f_1, \dots, f_n): \Omega \rightarrow \mathbb{R}^n \text{ or } \mathbb{C}^n$$

whose components f_i are k -times continuously differentiable or L^p -integrable on Ω respectively.

On the other hand, if \mathbb{V} is an infinite-dimensional topological vector space, then the results depend on the choice of topological tensor product and the algebraic construction above will need to be adapted to address issues of convergence. We present a few examples involving Banach-space-valued functions, mostly drawn from Ryan (2002, Chapters 2 and 3). In the following, \mathbb{B} denotes a Banach space with norm $\|\cdot\|$.

Absolutely summable sequences. This is a special case of the L^1 -functions below but worth a separate statement:

$$l^1(\mathbb{N}) \widehat{\otimes}_v \mathbb{B} = l^1(\mathbb{N}; \mathbb{B}) = \left\{ (x_i)_{i=1}^\infty : x_i \in \mathbb{B}, \sum_{i=1}^\infty \|x_i\| < \infty \right\}.$$

If we use $\widehat{\otimes}_\sigma$ in place of $\widehat{\otimes}_v$, we obtain a larger subspace, as we will see next.

Unconditionally summable sequences. Absolute summability implies unconditional summability but the converse is generally false when \mathbb{B} is infinite-dimensional. We have

$$l^1(\mathbb{N}) \widehat{\otimes}_\sigma \mathbb{B} = \left\{ (x_i)_{i=1}^\infty : x_i \in \mathbb{B}, \sum_{i=1}^\infty \varepsilon_i x_i < \infty \text{ for any } \varepsilon_i = \pm 1 \right\}.$$

The last condition is also equivalent to $\sum_{i=1}^\infty x_{\sigma(i)} < \infty$ for any bijection $\sigma: \mathbb{N} \rightarrow \mathbb{N}$.

Integrable functions. Let Ω be σ -finite. Then

$$L^1(\Omega) \widehat{\otimes}_v \mathbb{B} = L^1(\Omega; \mathbb{B}) = \{f: \Omega \rightarrow \mathbb{B} : \|f\|_1 < \infty\}.$$

This is called the Lebesgue–Bochner space. Here the L^1 -norm is defined as $\|f\|_1 = \int_\Omega \|f(x)\| dx$. In fact the proof of the second half of (4.73) is via

$$L^1(X) \widehat{\otimes}_v L^1(Y) = L^1(X; L^1(Y)) = L^1(X \times Y).$$

Continuous functions. Let Ω be compact Hausdorff. Then

$$C(\Omega) \widehat{\otimes}_\sigma \mathbb{B} = C(\Omega; \mathbb{B}) = \{f: \Omega \rightarrow \mathbb{B} : f \text{ continuous}\}.$$

Here continuity is with respect to $\|f\|_\infty = \sup_{x \in \Omega} \|f(x)\|$. In fact the proof of the first half of (4.73) is via

$$C(X) \widehat{\otimes}_\sigma C(Y) = C(X; C(Y)) = C(X \times Y).$$

Infinite-dimensional diagonal matrices. For finite-dimensional square matrices, the subspace of diagonal matrices $\text{diag}(\mathbb{R}^{n \times n}) := \text{span}\{e_i \otimes e_i : i = 1, \dots, n\} \cong \mathbb{R}^n$ is independent of any other structure on $\mathbb{R}^{n \times n}$. For infinite-dimensional matrices, the closed linear span $\overline{\text{span}}\{e_i \otimes e_i : i \in \mathbb{N}\}$ depends on both the choice of norms and the choice of tensor products:

$$\text{diag}(l^p(\mathbb{N}) \widehat{\otimes}_\nu l^p(\mathbb{N})) \cong \begin{cases} l^1(\mathbb{N}) & 1 \leq p \leq 2, \\ l^{p/2}(\mathbb{N}) & 2 < p < \infty, \\ c_0(\mathbb{N}) & p = \infty. \end{cases}$$

We refer readers to Holub (1970) for other cases with $\text{diag}(l^p(\mathbb{N}) \widehat{\otimes}_\nu l^q(\mathbb{N}))$, or $\widehat{\otimes}_\sigma$ in place of $\widehat{\otimes}_\nu$, and to Arias and Farmer (1996) for the higher-order case $l^{p_1}(\mathbb{N}) \otimes \dots \otimes l^{p_d}(\mathbb{N})$ with $d > 2$.

Partial traces. For separable Hilbert spaces \mathbb{H}_1 and \mathbb{H}_2 , the partial traces are the continuous linear maps

$$\begin{aligned} \text{tr}_1 : (\mathbb{H}_1 \widehat{\otimes}_F \mathbb{H}_2) \widehat{\otimes}_\nu (\mathbb{H}_1 \widehat{\otimes}_F \mathbb{H}_2)^* &\rightarrow \mathbb{H}_2 \widehat{\otimes}_F \mathbb{H}_2^*, \\ \text{tr}_2 : (\mathbb{H}_1 \widehat{\otimes}_F \mathbb{H}_2) \widehat{\otimes}_\nu (\mathbb{H}_1 \widehat{\otimes}_F \mathbb{H}_2)^* &\rightarrow \mathbb{H}_1 \widehat{\otimes}_F \mathbb{H}_1^*. \end{aligned}$$

Note that the domain of these maps are trace-class operators on the Hilbert space $\mathbb{H}_1 \widehat{\otimes}_F \mathbb{H}_2$ by Example 4.19. Choose any orthonormal basis on \mathbb{H}_1 to obtain a Hilbert space isomorphism $\mathbb{H}_1 \widehat{\otimes}_F \mathbb{H}_1^* \cong l^2(\mathbb{N} \times \mathbb{N})$. Then

$$\begin{aligned} (\mathbb{H}_1 \widehat{\otimes}_F \mathbb{H}_2) \widehat{\otimes}_\nu (\mathbb{H}_1 \widehat{\otimes}_F \mathbb{H}_2)^* &\cong (\mathbb{H}_1 \widehat{\otimes}_F \mathbb{H}_1^*) \widehat{\otimes}_\nu (\mathbb{H}_2 \widehat{\otimes}_F \mathbb{H}_2^*) \\ &\cong l^2(\mathbb{N} \times \mathbb{N}) \widehat{\otimes}_\nu (\mathbb{H}_2 \widehat{\otimes}_F \mathbb{H}_2^*) \end{aligned}$$

and the last space is the set of infinite matrices of the form

$$\begin{bmatrix} \Phi_{11} & \Phi_{12} & \cdots & \Phi_{1j} & \cdots \\ \Phi_{21} & \Phi_{22} & \cdots & \Phi_{2j} & \cdots \\ \vdots & \vdots & \ddots & \vdots & \\ \Phi_{i1} & \Phi_{i2} & \cdots & \Phi_{ij} & \cdots \\ \vdots & \vdots & & \vdots & \ddots \end{bmatrix}, \tag{4.95}$$

where each $\Phi_{ij} : \mathbb{H}_2 \rightarrow \mathbb{H}_2^*$ is a Hilbert–Schmidt operator for any $i, j \in \mathbb{N}$. Given a trace-class operator $\Phi : \mathbb{H}_1 \widehat{\otimes}_F \mathbb{H}_2 \rightarrow \mathbb{H}_1 \widehat{\otimes}_F \mathbb{H}_2$, we may express it in the form (4.95) and then define the first partial trace as

$$\text{tr}_1(\Phi) := \sum_{i=1}^\infty \Phi_{ii}.$$

The other partial trace tr_2 may be similarly defined by reversing the roles of \mathbb{H}_1 and \mathbb{H}_2 . Partial traces are an indispensable tool for working with the density operators discussed in Example 4.21; we refer readers to Cohen-Tannoudji *et al.* (2020a, Chapter III, Complement E, Section 5b) and Nielsen and Chuang (2000, Sections 2.4.3) for more information.

Our discussion about partial traces contains quite a bit of hand-waving; we will justify some of it with the next example, where we discuss some properties of tensor products that we have used liberally above.

Example 4.31 (calculus of tensor products). The universal factorization property is particularly useful for establishing other properties (Greub 1978, Chapter I) of the tensor product operation \otimes such as how it interacts with itself,

$$U \otimes V \cong V \otimes U, \quad U \otimes (V \otimes W) \cong (U \otimes V) \otimes W \cong U \otimes V \otimes W,$$

with direct sum \oplus ,

$$U \otimes (V \oplus W) \cong (U \otimes V) \oplus (U \otimes W), \quad (U \oplus V) \otimes W \cong (U \otimes W) \oplus (V \otimes W),$$

and with intersection \cap ,

$$(V \otimes W) \cap (V' \otimes W') = (V \cap V') \otimes (W \cap W'),$$

as well as how it interacts with linear and multilinear maps,

$$\mathcal{L}(U \otimes V; W) \cong \mathcal{L}(U; \mathcal{L}(V; W)) \cong \mathcal{M}^2(U, V; W),$$

with duality $V^* = \mathcal{L}(V; \mathbb{R})$ a special case,

$$(V \otimes W)^* \cong V^* \otimes W^*, \quad V^* \otimes W \cong \mathcal{L}(V; W),$$

and with the Kronecker product,

$$\mathcal{L}(V; W) \otimes \mathcal{L}(V'; W') \cong \mathcal{L}(V \otimes V'; W \otimes W').$$

Collectively, these properties form a system of calculus for manipulating tensor products of vector spaces. When pure mathematicians speak of *multilinear algebra*, this is often what they have in mind, that is, the subject is less about manipulating individual tensors and more about manipulating whole spaces of tensors. Many, but not all, of these properties may be extended to other vector-space-like objects such as modules and vector bundles, or to vector spaces with additional structures such as metrics, products and topologies. One needs to exercise some caution in making such extensions. For instance, if V and W are infinite-dimensional Hilbert spaces, then

$$V^* \otimes W \not\cong \mathcal{B}(V; W),$$

no matter what notion of tensor product or topological tensor product one uses for \otimes , that is, in this context $\mathcal{B}(V; W)$ is not the infinite-dimensional analogue of $\mathcal{L}(V; W)$. As we saw in Examples 4.19 and 4.28, the proper interpretation of \otimes for Hilbert spaces is more subtle.

To establish these properties, the ‘uniqueness up to vector space isomorphisms’ in Definition 4.26 is the tool of choice. This says that if we have two tensor spaces \mathbb{T} and \mathbb{T}' , both satisfying the universal factorization property for a multilinear map Φ ,

$$\begin{array}{ccc} \mathbb{V}_1 \times \cdots \times \mathbb{V}_d & \xrightarrow{\sigma_{\mathbb{T}}} & \mathbb{T} \\ & \searrow \Phi & \downarrow F_{\Phi} \\ & & \mathbb{W} \end{array} \qquad \begin{array}{ccc} \mathbb{V}_1 \times \cdots \times \mathbb{V}_d & \xrightarrow{\sigma_{\mathbb{T}'}} & \mathbb{T}' \\ & \searrow \Phi & \downarrow F'_{\Phi} \\ & & \mathbb{W} \end{array}$$

then we must have $\mathbb{T} \cong \mathbb{T}'$ as vector spaces. For each of the above properties, say, $\mathbb{U} \otimes \mathbb{V} \cong \mathbb{V} \otimes \mathbb{U}$, we plug in the vector spaces on both sides of the purported isomorphism \cong as our \mathbb{T} and \mathbb{W} and make a judicious choice for Φ , say, $\mathbb{U} \times \mathbb{V} \rightarrow \mathbb{V} \otimes \mathbb{U}$, $(u, v) \mapsto v \otimes u$, to obtain F_{Φ} as the required isomorphism.

While we have said early on in (4.30) that $u \otimes v \neq v \otimes u$, note that here we did not say $\mathbb{U} \otimes \mathbb{V} = \mathbb{V} \otimes \mathbb{U}$ but $\mathbb{U} \otimes \mathbb{V} \cong \mathbb{V} \otimes \mathbb{U}$. We sometimes *identify* isomorphic spaces, *i.e.* replace \cong with $=$, but this depends on the context. In Navier–Stokes (4.14), we do not want to say $C^2(\mathbb{R}^3) \widehat{\otimes} C^1[0, \infty) = C^1[0, \infty) \widehat{\otimes} C^2(\mathbb{R}^3)$, since the order of the variables in a solution $v(x, y, z, t)$ carries a specific meaning. On the other hand, it does not matter whether we treat the quantum states of a composite system of two particles as $\mathbb{H}_1 \widehat{\otimes} \mathbb{H}_2$ or $\mathbb{H}_2 \widehat{\otimes} \mathbb{H}_1$: either the particles are distinguishable with \mathbb{H}_1 and \mathbb{H}_2 distinct and the difference becomes merely a matter of which particle we label as first and which as second (completely arbitrary); or, if they are indistinguishable, then $\mathbb{H}_1 = \mathbb{H}_2$ and the question does not arise.

The reader may remember from our discussions in Section 3.1 about a major shortcoming of definition ②, that for each type of tensor there are many different types of multilinear maps. Definition 4.26 neatly classifies them in that we may use the calculus of \otimes described above to reduce any type of multilinear map into a tensor of type $(p, d - p)$ for some p and d . A noteworthy special case is the multilinear functionals in Definition 3.3 where we have

$$\mathcal{M}(\mathbb{V}_1^*, \dots, \mathbb{V}_p^*, \mathbb{V}_{p+1}, \dots, \mathbb{V}_d; \mathbb{R}) \cong \mathbb{V}_1 \otimes \cdots \otimes \mathbb{V}_p \otimes \mathbb{V}_{p+1}^* \otimes \cdots \otimes \mathbb{V}_d^*.$$

For instance, recall the matrix–matrix product in Example 3.9 and triple product trace in Example 3.17, respectively,

$$M_{m,n,p}(A, B) = AB, \quad \tau_{m,n,p}(A, B, C) = \text{tr}(ABC),$$

for $A \in \mathbb{R}^{m \times n}$, $B \in \mathbb{R}^{n \times p}$, $C \in \mathbb{R}^{p \times m}$. Since

$$\begin{aligned} M_{m,n,p} &\in \mathcal{M}(\mathbb{R}^{m \times n}, \mathbb{R}^{n \times p}; \mathbb{R}^{m \times p}) \cong (\mathbb{R}^{m \times n})^* \otimes (\mathbb{R}^{n \times p})^* \otimes \mathbb{R}^{m \times p}, \\ \tau_{m,n,p} &\in \mathcal{M}(\mathbb{R}^{m \times n}, \mathbb{R}^{n \times p}, \mathbb{R}^{p \times m}; \mathbb{R}) \cong (\mathbb{R}^{m \times n})^* \otimes (\mathbb{R}^{n \times p})^* \otimes (\mathbb{R}^{p \times m})^*, \end{aligned}$$

and $(\mathbb{R}^{p \times m})^* \cong \mathbb{R}^{m \times p}$, the bilinear operator $M_{m,n,p}$ and the trilinear functional $\tau_{m,n,p}$ may be regarded as elements of the same tensor space. Checking their values on the standard bases shows that they are in fact the same tensor. More

generally, the $(d - 1)$ -linear map and d -linear functional given by

$$\begin{aligned} M_{n_1, n_2, \dots, n_d}(A_1, A_2, \dots, A_{d-1}) &= A_1 A_2 \cdots A_{d-1}, \\ \tau_{n_1, n_2, \dots, n_d}(A_1, A_2, \dots, A_d) &= \text{tr}(A_1 A_2 \cdots A_d), \end{aligned}$$

for matrices

$$A_1 \in \mathbb{R}^{n_1 \times n_2}, A_2 \in \mathbb{R}^{n_2 \times n_3}, \dots, A_{d-1} \in \mathbb{R}^{n_{d-1} \times n_d}, A_d \in \mathbb{R}^{n_d \times n_1},$$

are one and the same tensor. The tensor $\tau_{n_1, n_2, \dots, n_d}$ will appear again in Example 4.44 in the form of the *matrix product states* tensor network.

Definition 4.26 was originally due to Bourbaki (1998) and, as is typical of their style, provides the most general and abstract definition of a tensor, stated in language equally abstruse. As a result it is sometimes viewed with trepidation by students attempting to learn tensor products from standard algebra texts, although it has also become the prevailing definition of a tensor in modern mathematics (Dummit and Foote 2004, Kostrikin and Manin 1997, Lang 2002, Vinberg 2003). As we saw in this section, however, it is a simple and practical idea.

4.4. Tensors in computations III: separability

The notion of separability is central to the utility of definition ③ in its various forms. We begin with its best-known manifestation: the separation-of-variables technique.

We will first present an abstract version to show the tensor perspective and then apply it to PDEs, finite difference schemes and integro-differential equations.

Example 4.32 (separation of variables: abstract). Let $\Phi: \mathbb{V}_1 \otimes \cdots \otimes \mathbb{V}_d \rightarrow \mathbb{V}_1 \otimes \cdots \otimes \mathbb{V}_d$ be a linear operator of the form

$$\Phi = \Phi_1 \otimes I_2 \otimes \cdots \otimes I_d + I_1 \otimes \Phi_2 \otimes \cdots \otimes I_d + \cdots + I_1 \otimes I_2 \otimes \cdots \otimes \Phi_d, \tag{4.96}$$

where I_d is the identity operator on \mathbb{V}_d and \otimes is the Kronecker product in Example 4.11. The separation-of-variables technique essentially transforms a homogeneous linear system into a collection of eigenproblems:

$$\Phi(v_1 \otimes v_2 \otimes \cdots \otimes v_d) = 0 \quad \longrightarrow \quad \begin{cases} \Phi_1(v_1) = \lambda_1 v_1, \\ \Phi_2(v_2) = \lambda_2 v_2, \\ \vdots \\ \Phi_d(v_d) = -(\lambda_1 + \cdots + \lambda_{d-1})v_d, \end{cases} \tag{4.97}$$

and Φ being linear, any sum, linear combination or, under the right conditions, even integral of $v_1 \otimes v_2 \otimes \cdots \otimes v_d$ is also a solution. The constants $\lambda_1, \dots, \lambda_{d-1}$ are called separation constants. The technique relies only on one easy fact about tensor products: for any non-zero $v \in \mathbb{V}$ and $w \in \mathbb{W}$,

$$v \otimes w = v' \otimes w' \quad \Rightarrow \quad v = \lambda v', \quad w = \lambda^{-1} w' \tag{4.98}$$

for some non-zero $\lambda \in \mathbb{R}$. More generally, for any d non-zero vectors $v_1 \in \mathbb{V}_1, v_2 \in \mathbb{V}_2, \dots, v_d \in \mathbb{V}_d$, we have

$$v_1 \otimes v_2 \otimes \dots \otimes v_d = v'_1 \otimes v'_2 \otimes \dots \otimes v'_d$$

$$\Rightarrow v_1 = \lambda_1 v'_1, v_2 = \lambda_2 v'_2, \dots, v_d = \lambda_d v'_d, \quad \lambda_1 \lambda_2 \dots \lambda_d = 1,$$

although this more general version is not as useful for separation of variables, which just requires repeated applications of (4.98).

Take $d = 3$ for illustration. Dropping subscripts to avoid clutter, we have

$$(\Phi \otimes I \otimes I + I \otimes \Psi \otimes I + I \otimes I \otimes \Theta)(u \otimes v \otimes w) = 0 \tag{4.99}$$

or equivalently $\Phi(u) \otimes v \otimes w + u \otimes \Psi(v) \otimes w + u \otimes v \otimes \Theta(w) = 0$. Since $\Phi(u) \otimes (v \otimes w) = u \otimes [-\Psi(v) \otimes w - v \otimes \Theta(w)]$, applying (4.98) we get

$$\Phi(u) = \lambda u, \quad v \otimes w = -\lambda^{-1}[\Psi(v) \otimes w + v \otimes \Theta(w)].$$

Rearranging the second equation, $\Psi(v) \otimes w = v \otimes [-\Theta(w) - \lambda w]$, and applying (4.98) again, we get

$$\Psi(v) = \mu v, \quad \Theta(w) = -(\mu + \lambda)w.$$

Thus we have transformed (4.99) into three eigenproblems:

$$\begin{cases} \Phi(u) = \lambda u, \\ \Psi(v) = \mu v, \\ \Theta(w) = (-\mu - \lambda)w. \end{cases}$$

The technique applies widely. The operators $\Phi_i: \mathbb{V}_i \rightarrow \mathbb{V}_i$ involved can be any operator in any coordinates. While they are most commonly differential operators such as

$$\Phi_\theta(f) = \frac{1}{r^2 \sin^2 \phi} \frac{\partial^2 f}{\partial \theta^2}, \quad \Phi_\sigma(f) = \frac{1}{\sigma^2 + \tau^2} \frac{\partial^2 f}{\partial \sigma^2}$$

in Example 4.33, they may well be finite difference or integro-differential operators,

$$\Phi_k(u_{k,n}) = u_{k-1,n} + (1 - 2r)u_{k,n} + ru_{k+1,n}, \quad \Phi_x(f) = a \frac{\partial^2 f}{\partial x^2} + b \int_0^x f,$$

as in Examples 4.34 and 4.35 respectively. The important thing is that they must take the form³⁰ in (4.96): there should not be any ‘cross-terms’ like $\Phi_1 \otimes I_2 \otimes \Phi_3$, *e.g.* $\partial^2 f / \partial x_1 \partial x_3$, because if so then the first variable and third variable will not be separable, as we will see.

³⁰ One sometimes hears of ‘Kronecker sums’ in the $d = 2$ case: $\Phi_1 \oplus \Phi_2 := \Phi_1 \otimes I_2 + I_1 \otimes \Phi_2$. This has none of the usual characteristics of a sum: it is not associative, not commutative, and $\Phi_1 \oplus \Phi_2 \oplus \Phi_3 \neq \Phi_1 \otimes I_2 \otimes I_3 + I_1 \otimes \Phi_2 \otimes I_3 + I_1 \otimes I_2 \otimes \Phi_3$ regardless of the order \oplus is performed. As such we avoid using \oplus in this sense.

The abstract generality of Example 4.32 makes the separation-of-variables technique appear trivial, although as we will see in Examples 4.34 and 4.35, it also makes the technique widely applicable. Once we restrict to specific classes of problems like PDEs, there are much deeper results about the technique. Nevertheless, elementary treatments tend to focus on auxiliary issues such as Sturm–Liouville theory and leave out what we think are the most important questions: Why and when does the technique work? Fortunately these questions have been thoroughly investigated and addressed in Koornwinder (1980), Kalnins, Kress and Miller (2018), Miller (1977) and Schöbel (2015). The explanation is intimately related to tensors but in the sense of tensor fields in Example 3.13: the metric tensor, Ricci tensor, Nijenhuis tensor, Weyl tensor all play a role. Unfortunately, the gap between these results and our brief treatment is too wide for us to convey anything beyond a glimpse of what is involved.

Example 4.33 (separation of variables: PDEs). The technique of separation of variables for solving PDEs evidently depends on the choice of coordinates. It works perfectly for the one-dimensional wave equation in the form

$$\frac{\partial^2 f}{\partial x^2} - \frac{\partial^2 f}{\partial t^2} = 0. \quad (4.100)$$

The standard recipe is to use the ansatz $f(x, t) = \varphi(x)\psi(t)$, obtain $\varphi''(x)/\varphi(x) = \psi''(t)/\psi(t)$, and argue that since the left-hand side is a function of x and the right-hand side of t , they must both be equal to some separation constant $-\omega^2$, and we obtain two ODEs $\varphi'' + \omega^2\varphi = 0$, $\psi'' + \omega^2\psi = 0$. We may deduce the same thing in one step from (4.97):

$$[\partial_x^2 \otimes I + I \otimes (-\partial_t^2)](\varphi \otimes \psi) = 0 \quad \longrightarrow \quad \begin{cases} \partial_x^2 \varphi = -\omega^2 \varphi, \\ -\partial_t^2 \psi = \omega^2 \psi. \end{cases}$$

The ODEs have solutions $\varphi(x) = a_1 e^{\omega x} + a_2 e^{-\omega x}$, $\psi(x) = a_3 e^{\omega t} + a_4 e^{-\omega t}$ if $\omega \neq 0$, or $\varphi(x) = a_1 + a_2 x$, $\psi(t) = a_3 + a_4 t$ if $\omega = 0$, and any finite linear combinations of $\varphi \otimes \psi$ give us solutions for (4.100). Nevertheless, a change of coordinates $\xi = x - t$, $\eta = x + t$ transforms (4.100) into

$$\frac{\partial^2 f}{\partial \xi \partial \eta} = 0. \quad (4.101)$$

Note that the operator here is $\partial_x \otimes \partial_y$, and does not have the required form in (4.96). Indeed, the solution of (4.101) is easily seen to take the form $f(\xi, \eta) = \varphi(\xi) + \psi(\eta)$, so the usual multiplicatively separable ansatz $f(\xi, \eta) = \varphi(\xi)\psi(\eta)$ will not work.

More generally, the same argument works in any dimension n to separate the spatial coordinates $x = (x_1, \dots, x_n)$ (which do not need to be Cartesian) from the temporal ones. Applying (4.97) to the n -dimensional wave equation,

$$\Delta f - \frac{\partial^2 f}{\partial t^2} = 0, \quad (4.102)$$

we get

$$[\Delta \otimes I + I \otimes (-\partial_t^2)](\varphi \otimes \psi) = 0 \quad \rightarrow \quad \begin{cases} \Delta \varphi = -\omega^2 \varphi, \\ -\partial_t^2 \psi = \omega^2 \psi, \end{cases}$$

with separation constant $-\omega^2$. In the remainder of this example, we will focus on the first equation, called the n -dimensional Helmholtz equation,

$$\Delta f + \omega^2 f = 0, \tag{4.103}$$

which may also be obtained by taking the Fourier transform of (4.102) in time. For $n = 2$, (4.103) in Cartesian and polar coordinates are given by

$$\frac{\partial^2 f}{\partial x^2} + \frac{\partial^2 f}{\partial y^2} + \omega^2 f = 0, \quad \frac{\partial^2 f}{\partial r^2} + \frac{1}{r} \frac{\partial f}{\partial r} + \frac{1}{r^2} \frac{\partial^2 f}{\partial \theta^2} + \omega^2 f = 0 \tag{4.104}$$

respectively. Separation of variables works in both cases but gives entirely different solutions. Applying (4.97) to $\partial_x^2 \otimes I + I \otimes (\partial_y^2 + \omega^2 I)$ gives us

$$\frac{d^2 \varphi}{dx^2} + k^2 \varphi = 0, \quad \frac{d^2 \psi}{dy^2} + (\omega^2 - k^2) \psi = 0,$$

with separation constant k^2 and therefore the solution

$$f_k(x, y) := a_1 e^{i[kx + (\omega^2 - k^2)^{1/2}y]} + a_2 e^{i[-kx + (\omega^2 - k^2)^{1/2}y]} \\ + a_3 e^{i[kx - (\omega^2 - k^2)^{1/2}y]} + a_4 e^{i[-kx - (\omega^2 - k^2)^{1/2}y]}.$$

Applying (4.97) to $[(r^2 \partial_r^2 + \omega^2 r^2 I) \otimes I + I \otimes \partial_\theta^2](\varphi \otimes \psi) = 0$ gives us

$$r^2 \frac{d^2 \varphi}{dr^2} + r \frac{d\varphi}{dr} + (\omega^2 r^2 - k^2) \varphi = 0, \quad \frac{d^2 \psi}{d\theta^2} + k^2 \psi = 0,$$

with separation constant k^2 and therefore the solution

$$f_k(r, \theta) := a_1 e^{ik\theta} J_k(\omega r) + a_2 e^{-ik\theta} J_k(\omega r) + a_3 e^{ik\theta} J_{-k}(\omega r) + a_4 e^{-ik\theta} J_{-k}(\omega r)$$

where J_k is a Bessel function. Any solution of the two-dimensional Helmholtz equation in Cartesian coordinates is a sum or integral of $f_k(x, y)$ over k and any solution in polar coordinates is one of $f_k(r, \theta)$ over k . Analytic solutions in different coordinate systems provide different insights. There are exactly two more such coordinate systems where separation of variables works; we call these *separable coordinates*. For $n = 2$, (4.103) has exactly four systems of separable coordinates: Cartesian, polar, parabolic and elliptic. For $n = 3$, there are exactly eleven (Eisenhart 1934).

The fundamental result that allows us to deduce these numbers is the Stäckel condition: the n -dimensional Helmholtz equation in coordinates x_1, \dots, x_n can be solved using the separation-of-variables technique if and only if (a) the Euclidean metric tensor g is a diagonal matrices in this coordinate system, and (b) if $g =$

$\text{diag}(g_{11}, \dots, g_{nn})$, then there exists an invertible matrix of the form

$$S = \begin{bmatrix} s_{11}(x_1) & s_{12}(x_1) & \cdots & s_{1n}(x_1) \\ s_{21}(x_2) & s_{22}(x_2) & \cdots & s_{2n}(x_2) \\ \vdots & \vdots & & \vdots \\ s_{n1}(x_n) & s_{n2}(x_n) & \cdots & s_{nn}(x_n) \end{bmatrix}$$

with

$$g_{jj}^{-1} = (S^{-1})_{1j}, \quad j = 1, \dots, n. \tag{4.105}$$

Here S is called a Stäckel matrix for g in coordinates x_1, \dots, x_n ; note that the i th row of S depends only on the i th coordinate. The Euclidean metric tensor g is a covariant 2-tensor field that parametrizes the Euclidean inner product at different points in a coordinate system. As we saw in Example 3.12, any 2-tensor field over \mathbb{R}^n may be represented as an $n \times n$ matrix of functions. Take $n = 3$; the Euclidean metric tensor on \mathbb{R}^3 in Cartesian, cylindrical, spherical and parabolic³¹ coordinates are

$$\begin{aligned} g(x, y, z) &= \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, & g(r, \theta, z) &= \begin{bmatrix} 1 & 0 & 0 \\ 0 & r^2 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \\ g(r, \theta, \phi) &= \begin{bmatrix} 1 & 0 & 0 \\ 0 & r^2 & 0 \\ 0 & 0 & r^2 \sin^2 \theta \end{bmatrix}, & g(\sigma, \tau, \phi) &= \begin{bmatrix} \sigma^2 + \tau^2 & 0 & 0 \\ 0 & \sigma^2 + \tau^2 & 0 \\ 0 & 0 & \sigma^2 \tau^2 \end{bmatrix}. \end{aligned}$$

The Stäckel condition is satisfied for these four coordinate systems using the following matrices and their inverses:

$$\begin{aligned} \text{Cartesian} \quad S &= \begin{bmatrix} 0 & -1 & -1 \\ 0 & 1 & 0 \\ 1 & 0 & 1 \end{bmatrix}, & S^{-1} &= \begin{bmatrix} 1 & 1 & 1 \\ 0 & 1 & 0 \\ -1 & -1 & 0 \end{bmatrix}, \\ \text{cylindrical} \quad S &= \begin{bmatrix} 0 & -\frac{1}{r^2} & -1 \\ 0 & 1 & 0 \\ 1 & 0 & 1 \end{bmatrix}, & S^{-1} &= \begin{bmatrix} 1 & \frac{1}{r^2} & 1 \\ 0 & 1 & 0 \\ -1 & -\frac{1}{r^2} & 0 \end{bmatrix}, \\ \text{spherical} \quad S &= \begin{bmatrix} 1 & -\frac{1}{r^2} & 0 \\ 0 & 1 & -\frac{1}{\sin^2 \theta} \\ 0 & 0 & 1 \end{bmatrix}, & S^{-1} &= \begin{bmatrix} 1 & \frac{1}{r^2} & \frac{1}{(r^2 \sin^2 \theta)} \\ 0 & 1 & \frac{1}{\sin^2 \theta} \\ 0 & 0 & 1 \end{bmatrix}, \\ \text{parabolic} \quad S &= \begin{bmatrix} \sigma^2 & -1 & -\frac{1}{\sigma^2} \\ \tau^2 & 1 & -\frac{1}{\tau^2} \\ 0 & 0 & 1 \end{bmatrix}, & S^{-1} &= \begin{bmatrix} \frac{1}{(\sigma^2 + \tau^2)} & \frac{1}{(\sigma^2 + \tau^2)} & \frac{1}{(\sigma^2 \tau^2)} \\ -\frac{\tau^2}{(\sigma^2 + \tau^2)} & \frac{\sigma^2}{(\sigma^2 + \tau^2)} & \frac{1}{\tau^2 - 1/\sigma^2} \\ 0 & 0 & 1 \end{bmatrix} \end{aligned}$$

³¹ Given by $x = \sigma\tau \cos \varphi$, $y = \sigma\tau \sin \varphi$, $z = (\sigma^2 - \tau^2)/2$.

The matrices on the left are Stäckel matrices for g in the respective coordinate system, that is, the entries in the first row of S^{-1} are exactly the reciprocal of the entries on the diagonal of g .

What we have ascertained is that the three-dimensional Helmholtz equation can be solved by separation of variables in these four coordinate systems, without writing down a single differential equation. In fact, with more effort, one can show that there are exactly eleven such separable coordinate systems:

- (i) Cartesian,
- (ii) cylindrical,
- (iii) spherical,
- (iv) parabolic,
- (v) paraboloidal,
- (vi) ellipsoidal,
- (vii) conical,
- (viii) prolate spheroidal,
- (ix) oblate spheroidal,
- (x) elliptic cylindrical,
- (xi) parabolic cylindrical.

These eleven coordinate systems have been thoroughly studied in engineering (Moon and Spencer 1988, 1961), where they are used for different tasks: (viii) for modelling radiation from a slender spheroidal antenna, (x) for heat flow in a bar of elliptic cross section, *etc.* More generally, the result can be extended to a Riemannian manifold M of arbitrary dimension: a system of local coordinates is separable if and only if both the Riemannian metric tensor g and the Ricci curvature tensor \bar{R} are diagonal matrices in those coordinates,³² and g satisfies the Stäckel condition (Eisenhart 1934). In particular, we now see why separation of variables does not apply to PDEs with mixed partials like (4.101): when the metric tensor is diagonal, $g = \text{diag}(g_{11}, \dots, g_{nn})$, the Laplacian takes the form

$$\Delta = \sum_{i=1}^n \frac{1}{\sqrt{\det(g)}} \frac{\partial}{\partial x_i} \frac{\sqrt{\det(g)}}{g_{ii}} \frac{\partial f}{\partial x_i} \quad (4.106)$$

and thus does not contain any mixed derivatives. We may use (4.106) to write down

³² We did not need to worry about the Ricci tensor because \mathbb{R}^n is a so-called Einstein manifold where g and \bar{R} differ by a scalar multiple. See Example 3.13 for a cursory discussion of g and \bar{R} .

the Helmholtz equation on \mathbb{R}^3 in cylindrical, spherical and parabolic coordinates:

$$\begin{aligned} & \frac{\partial^2 f}{\partial r^2} + \frac{1}{r} \frac{\partial f}{\partial r} + \frac{1}{r^2} \frac{\partial^2 f}{\partial \theta^2} + \frac{\partial^2 f}{\partial z^2} + \omega^2 f = 0, \\ & \frac{\partial^2 f}{\partial r^2} + \frac{2}{r} \frac{\partial f}{\partial r} + \frac{1}{r^2 \sin^2 \phi} \frac{\partial^2 f}{\partial \theta^2} + \frac{\cos \phi}{r^2 \sin^2 \phi} \frac{\partial f}{\partial \phi} + \frac{1}{r^2} \frac{\partial^2 f}{\partial \phi^2} + \omega^2 f = 0, \\ & \frac{1}{\sigma^2 + \tau^2} \left[\frac{\partial^2 f}{\partial \sigma^2} + \frac{1}{\sigma} \frac{\partial f}{\partial \sigma} + \frac{\partial^2 f}{\partial \tau^2} + \frac{1}{\tau} \frac{\partial f}{\partial \tau} \right] + \frac{1}{\sigma^2 \tau^2} \frac{\partial^2 f}{\partial \phi^2} + \omega^2 f = 0. \end{aligned}$$

It is not so obvious that these are amenable to separation of variables, speaking to the power of the Stäckel condition. While we have restricted to the Helmholtz equation for concreteness, the result has been generalized to higher-order semilinear PDEs (Koornwinder 1980, Theorem 3.8).

Separation of variables is manifestly about tensors in the form of definition ③ and particularly Definition 4.4, as it involves, in one way or another, a rank-one tensor $\varphi \otimes \psi \otimes \dots \otimes \theta$ or a sum of these. It may have come as a surprise that tensors in the sense of definition ① also play a major role in Example 4.33, but this should be expected as the tensor transformation rules originally came from a study of the relations between different coordinate systems.

By our discussion in Example 4.5, whether a function is of continuous variables or discrete ones, whether the variables are written as arguments or indices, should not make a substantive difference. Indeed, as we saw in Example 4.32, separation of variables ought to apply verbatim in cases where the variables are finite or discrete. Our next example, adapted from Thomas (1995, Example 3.2.2), is one where the rank-one tensors are of the form $u = a \otimes b$ with $a = (a_0, \dots, a_m) \in \mathbb{R}^{m+1}$ and $b = (b_n)_{n=0}^\infty \in c_0(\mathbb{N} \cup \{0\})$.

Example 4.34 (separation of variables: finite difference). Consider the recurrence relation

$$\begin{cases} u_{k,n+1} = ru_{k-1,n} + (1 - 2r)u_{k,n} + ru_{k+1,n} & k = 1, \dots, m - 1, \\ u_{0,n+1} = 0 = u_{m,n+1} \\ u_{k,0} = f\left(\frac{k}{m}\right) & k = 0, 1, \dots, m, \end{cases} \tag{4.107}$$

with $n = 0, 1, 2, \dots$, and $r > 0$ is some fixed constant. This comes from a forward-time centred space discretization scheme applied to an initial–boundary value problem for a one-dimensional heat equation on $[0, 1]$, although readers lose nothing by simply treating this example as one of solving recurrence relations. To be consistent with our notation, we have written $u_{k,n} = u(x_k, t_n)$ instead of the more common u_k^n with the time index in superscript. Applying (4.97) to (4.107) with

$$\Phi_k(u_{k,n}) := ru_{k-1,n} + (1 - 2r)u_{k,n} + ru_{k+1,n}, \quad \Psi_n(u_{k,n}) := u_{k,n+1},$$

we get

$$[\Phi_k \otimes I + I \otimes (-\Psi_n)](a \otimes b) = 0 \quad \longrightarrow \quad \begin{cases} \Phi_k(a_k) = \lambda a_k, \\ -\Psi_n(b_n) = -\lambda b_n, \end{cases}$$

with separation constant λ . We write these out in full:

$$\begin{aligned} r a_{k-1} + (1 - 2r) a_k + r a_{k+1} &= \lambda a_k, & k = 1, \dots, m - 1, \\ b_{n+1} &= \lambda b_n, & n = 0, 1, 2, \dots \end{aligned}$$

The second equation is trivial to solve: $b_n = \lambda^n b_0$. Noting that the boundary conditions $u_{0,n+1} = 0 = u_{m,n+1}$ give $a_0 = 0 = a_m$, we see that the first equation is a tridiagonal eigenproblem:

$$\begin{bmatrix} 1 - 2r & r & & & & \\ r & 1 - 2r & r & & & \\ & r & 1 - 2r & r & & \\ & & & \ddots & \ddots & \\ & & & & r & 1 - 2r \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \\ a_3 \\ \vdots \\ a_{m-1} \end{bmatrix} = \lambda \begin{bmatrix} a_1 \\ a_2 \\ a_3 \\ \vdots \\ a_{m-1} \end{bmatrix}.$$

The eigenvalues and eigenvectors of a tridiagonal Toeplitz matrices have well-known closed-form expressions:

$$\lambda_j = 1 - 4r \sin^2\left(\frac{j\pi}{2m}\right), \quad a_{jk} = \sin\left(\frac{jk\pi}{m}\right), \quad j, k = 1, \dots, m - 1,$$

where a_{jk} is the k th coordinate of the j th eigenvector. Hence we get

$$u_{k,n} = \sum_{j=1}^{m-1} c_j b_0 \left[1 - 4r \sin^2\left(\frac{j\pi}{2m}\right) \right]^n \sin\left(\frac{jk\pi}{m}\right).$$

The initial condition $u_{k,0} = f(k/m)$, $k = 0, 1, \dots, m$, may be used to determine the coefficients c_1, \dots, c_{m-1} and b_0 .

The von Neumann stability analysis is a special case of such a discrete separation of variables with $a_k = e^{ijk\pi/m}$ and $b_n = \lambda^n$, that is,

$$u_{k,n} = \lambda^n e^{ijk\pi/m}.$$

Substituting this into the recursion in (4.107) and simplifying gives us

$$\lambda = 1 - 4r \sin^2\left(\frac{j\pi}{2m}\right),$$

and since we require $|\lambda| \leq 1$ for stability, we get that $r \leq 1/2$, an analysis familiar to readers who have studied finite difference methods.

The next example is adapted from [Kostoglou \(2005\)](#), just to show the range of applicability of [Example 4.32](#).

Example 4.35 (separation of variables: integro-differential equations). Let us consider the following integro-differential equation arising from the study of heterogeneous heat transfer:

$$\frac{\partial f}{\partial t} = a \frac{\partial^2 f}{\partial x^2} + b \int_0^x f(y, t) dy - f, \quad (4.108)$$

with $a, b \geq 0$ and $f: [0, 1] \times \mathbb{R} \rightarrow \mathbb{R}$. Note that at this juncture, if we simply differentiate both sides to eliminate the integral, we will introduce mixed derivatives and thus prevent ourselves from using separation of variables. Nevertheless, our interpretation of separation of variables in [Example 4.32](#) allows for integrals. If we let

$$\Phi_x(f) := \frac{\partial^2 f}{\partial x^2} - f + b \int_0^x f(y, t) dy, \quad \Psi_t(f) := -\frac{\partial f}{\partial t},$$

then [\(4.97\)](#) gives us

$$[\Phi_x \otimes I + I \otimes \Psi_t](\varphi \otimes \psi) = 0 \quad \longrightarrow \quad \begin{cases} \Phi_x(\varphi) = \lambda\varphi, \\ \Psi_t(\psi) = -\lambda\psi, \end{cases}$$

with separation constant λ . Writing these out in full, we have

$$a \frac{d^2 \varphi}{dx^2} + (\lambda - 1)\varphi + b \int_0^x \varphi(y) dy = 0, \quad \frac{d\psi}{dt} + \lambda\psi = 0.$$

The second equation is easy: $\psi(t) = c e^{-\lambda t}$ for an arbitrary constant c that could be determined with an initial condition. [Kostoglou \(2005\)](#) solved the first equation in a convoluted manner involving Laplace transforms and partial fractions, but this is unnecessary; at this point it is harmless to simply differentiate and eliminate the integral. With this, we obtain a third-order homogeneous ODE with constant coefficients, $a\varphi''' + (\lambda - 1)\varphi' + b\varphi = 0$, whose solution is standard. Physical considerations show that its characteristic polynomial

$$r^3 + \frac{\lambda - 1}{a}r + \frac{b}{a} = 0$$

must have one real and two complex roots $r_1, r_2 \pm ir_3$, and thus the solution is given by $c_1 e^{r_1 x} + e^{r_2 x}(c_2 \cos r_3 x + c_3 \sin r_3 x)$ with arbitrary constants c_1, c_2, c_3 that could be determined with appropriate boundary conditions.

The last four examples are about exploiting separability in the structure of the solutions; the next few are about exploiting separability in the structure of the problems.

Example 4.36 (separable ODEs). The easiest ODEs to solve are probably the separable ones,

$$\frac{dy}{dx} = f(x)g(y), \tag{4.109}$$

with special cases $dy/dx = f(x)$ and $dy/dx = g(y)$ when one of the functions is constant. Solutions are, at least in principle, given by direct integration,

$$\int \frac{dy}{g(y)} = \int f(x) dx + c,$$

even though closed-form expressions still depend on having closed-form integrals. The effort is not so much in solving (4.109) but in seeking a transformation into (4.109), that is, given an inseparable ODE

$$\frac{dy}{dx} = K(x, y),$$

we would like to find a differentiable $\Phi: \mathbb{R}^2 \rightarrow \mathbb{R}^2, (x, y) \mapsto (u, v)$, so that

$$\frac{dv}{du} = f(u)g(v). \tag{4.110}$$

These transformations may be regarded as a special case of the transformation rules for tensor fields (two-dimensional vector fields in this case) discussed in Example 3.12. For instance,

$$\frac{dy}{dx} = \frac{y-x}{y+x} \quad \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} r \cos \theta \\ r \sin \theta \end{bmatrix} \quad \frac{dr}{d\theta} = -r \tag{4.111}$$

is the two-dimensional version of the tensor transformation rule in (3.55). There is a wide variety of scenarios where this is possible:

$$\begin{array}{lll} \frac{dy}{dx} = f(ax + by + c), & v = ax + by + c, & \frac{dv}{dx} = a + bf(v), \\ \frac{dy}{dx} = f\left(\frac{y}{x}\right), & v = \frac{y}{x}, & \frac{dv}{dx} = \frac{f(v) - v}{x}, \\ \frac{dy}{dx} = f(x)y + g(x), & v = y e^{-\int f(x) dx}, & \frac{dv}{dx} = g(x) e^{-\int f(x) dx}, \\ \frac{d^n y}{dx^n} = f\left(\frac{d^{n-1} y}{dx^{n-1}}\right)g(x), & v = \frac{d^{n-1} y}{dx^{n-1}}, & \frac{dv}{dx} = f(v)g(x), \\ \frac{dy}{dx} = f\left(\frac{ax + by + c}{a'x + b'y + c'}\right), & \begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} x - \frac{b'c + bc'}{a'b - ab'} \\ y - \frac{a'c + ac'}{a'b - ab'} \end{bmatrix}, & \frac{dv}{du} = f\left(\frac{a + bv/du}{a' + b'v/du}\right), \end{array}$$

noting that the last case reduces to the second, and if $u = x$, we do not introduce a new variable (Walter 1998, Chapter 1). Nevertheless, unlike the case of PDEs discussed in Example 4.33, there does not appear to be a systematic study of when

transformation to (4.110) is possible. Something along the lines of the Stäckel condition, but for ODEs, would be interesting and useful.

There is a close analogue of separable ODEs for integral equations with kernels that satisfy a finite version of (4.81). As we will see below, they have neat and simple solutions, but unfortunately such kernels are not too common in practice; in fact they are regarded as a degenerate case in the study of integral equations. The following discussion is adapted from Kanwal (1997, Chapter 2).

Example 4.37 (separable integral equations). Let us consider Fredholm integral equations of the first and second kind:

$$g(x) = \int_a^b K(x, y)f(y) dy, \quad f(x) = g(x) + \lambda \int_a^b K(x, y)f(y) dy, \quad (4.112)$$

with given constants $a < b$ and functions $g \in C[a, b]$ and $K \in C([a, b] \times [a, b])$. The goal is to solve for $f \in C[a, b]$ and, in the second case, also $\lambda \in \mathbb{R}$. The kernel K is said to be degenerate or separable if

$$K(x, y) = \sum_{i=1}^n \varphi_i(x)\psi_i(y) \quad (4.113)$$

for some $\varphi_1, \dots, \varphi_n, \psi_1, \dots, \psi_n \in C[a, b]$, assumed known and linearly independent. We set

$$v_i := \int_a^b \psi_i(y)f(y) dy, \quad b_i := \int_a^b \psi_i(x)g(x) dx, \quad a_{ij} := \int_a^b \varphi_j(x)\psi_i(x) dx$$

for $i, j = 1, \dots, n$. Plugging (4.113) into (4.112) and integrating with respect to y gives us

$$g(x) = \sum_{j=1}^n v_j \varphi_j(x), \quad f(x) = g(x) + \lambda \sum_{i=1}^n v_i \varphi_i(x). \quad (4.114)$$

Multiplying by $\psi_i(x)$ and integrating with respect to x gives us

$$b_i = \sum_{j=1}^n a_{ij} v_j, \quad v_i = b_i + \lambda \sum_{j=1}^n a_{ij} v_j$$

for $i = 1, \dots, n$, or, in matrix forms,

$$Av = b, \quad (I - \lambda A)v = b$$

respectively. For the latter, any λ^{-1} not an eigenvalue of A gives a unique solution for v and thus for f by (4.114); for the former, considerations similar to those for obtaining minimum-norm solutions in numerical linear algebra allow us to obtain f .

It is straightforward to extend the above to Fredholm integro-differential equations (Wazwaz 2011, Chapter 16) such as

$$\frac{d^k f}{dx^k} = g(x) + \lambda \int_a^b K(x, y) f(y) dy.$$

Note that if $k = 0$, this reduces to the second equation in (4.112), but it is of a different nature from the integro-differential considered in Example 4.35, which is of Volterra type and involves partial derivatives.

In principle, the discussion extends to the L^2 -kernels in Examples 4.19 and 4.20 with $n = \infty$, although in this case we would just be reformulating Hilbert–Schmidt operators in terms of infinite-dimensional matrices in $l^2(\mathbb{N} \times \mathbb{N})$; the resulting equations are no easier to analyse or solve (Kanwal 1997, Chapter 7).

While separable kernels are uncommon in integral equations, they are quite common in multidimensional integral transforms.

Example 4.38 (separable integral transforms). Whenever one iteratively applies a univariate integral transform to each variable of a multivariate function separately, one obtains an iterated integral

$$\int_{a_1}^{b_1} \cdots \int_{a_d}^{b_d} f(x_1, \dots, x_d) K_1(x_1, y_1) \cdots K_d(x_d, y_d) dx_1 \cdots dx_d$$

with respect to the separable kernel $K_1 \otimes \cdots \otimes K_d$ in Example 4.10(iv). We will just mention two of the most common. The multidimensional Fourier transform is given by

$$\mathcal{F}(f)(\xi_1, \dots, \xi_d) = \int_{-\infty}^{\infty} \cdots \int_{-\infty}^{\infty} f(t_1, \dots, t_d) e^{-i\xi_1 t_1 - \cdots - i\xi_d t_d} dt_1 \cdots dt_d$$

and the multidimensional Laplace transform is given by

$$\mathcal{L}(f)(s_1, \dots, s_d) = \int_0^{\infty} \cdots \int_0^{\infty} f(t_1, \dots, t_d) e^{-s_1 t_1 - \cdots - s_d t_d} dt_1 \cdots dt_d.$$

We point readers to Osgood (2019, Chapter 9) and Bracewell (1986, Chapter 13) for properties and applications of the former, and Debnath and Dahiya (1989) and Jaeger (1940) for the latter.

We briefly discuss another closely related notion of separability before our next examples, which will require it.

Example 4.39 (additive separability). Our interpretation of separability thus far is multiplicative separability of the form

$$f(x_1, x_2, \dots, x_d) = \varphi_1(x_1)\varphi_2(x_2) \cdots \varphi_d(x_d)$$

and not additive separability of the form

$$g(x_1, x_2, \dots, x_d) = \varphi_1(x_1) + \varphi_2(x_2) + \cdots + \varphi_d(x_d). \quad (4.115)$$

However, as we saw in the solution of the one-dimensional wave equation in Example 4.33, what is multiplicatively separable in one coordinate system is additively separable in another. The two notions are intimately related although the exact relation depends on the context. In PDEs, the relation is simple (Kalnins *et al.* 2018, Section 3.1.1), at least in principle. If

$$F(D^d f(x), D^{d-1} f(x), \dots, Df(x), f(x), x) = 0$$

has a multiplicatively separable solution, then the transformation $g(x) := e^{f(x)}$ gives us a PDE

$$G(D^d g(x), D^{d-1} g(x), \dots, Dg(x), g(x), x) = 0$$

with an additively separable solution of the form (4.115). The Stäckel condition in Example 4.33 was in fact originally proved for an additively separable solution of the form (4.115) for the Hamilton–Jacobi equation.

We describe a readily checkable condition due to Scheffé (1970) for deciding additive separability, which is important in statistics, particularly in generalized additive models (Hastie and Tibshirani 1990) and in the analysis of variance (Scheffé 1999, Section 4.1). Given any C^2 -function $f: \Omega \subseteq \mathbb{R}^d \rightarrow \mathbb{R}$, there exists a transformation $\varphi_0: \mathbb{R} \rightarrow \mathbb{R}$ so that $g = \varphi_0 \circ f$ has a decomposition of the form (4.115) if and only if its Hessian takes the form

$$\nabla^2 f(x) = h(f(x))\nabla f(x)\nabla f(x)^\top$$

for some function $h: \mathbb{R} \rightarrow \mathbb{R}$. In this case

$$\varphi_0(t) = c \int e^{-\int h(t) dt} dt + c_0, \quad \varphi_i(t) = c \int e^{-\int h(t) dt} \frac{\partial f}{\partial x_i}(t) dt + c_i,$$

for $i = 1, \dots, d$, with constants $c > 0$ and $c_0, c_1, \dots, c_d \in \mathbb{R}$ arising from the indefinite integrals. For instance, the condition may be used to verify that

$$\begin{aligned} f(x, y, z) &= x\sqrt{(1-y^2)(1-z^2)} + y\sqrt{(1-x^2)(1-z^2)} \\ &\quad + z\sqrt{(1-x^2)(1-y^2)} - xyz, \\ f(a, b, c, d) &= \frac{a + b + c + d + abc + abd + acd + bcd}{1 + ab + ac + ad + bc + bd + cd + abcd} \end{aligned}$$

are both additively separable with $\varphi_i(t) = \sin^{-1}(t)$, $i = 0, 1, 2, 3$, in the first case and $\varphi_i(t) = \log(1+t) - \log(1-t)$, $i = 0, 1, 2, 3, 4$, in the second.

The next two examples show that it is sometimes more natural to formulate separability in the additive form.

Example 4.40 (separable convex optimization). Let $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$ and $c \in \mathbb{R}^n$. We consider the linear programming problem

$$\begin{array}{ll} \text{LP} & \text{minimize } c^\top x \\ & \text{subject to } Ax \leq b, x \geq 0 \end{array}$$

and integer linear programming problem

$$\begin{array}{ll} \text{ILP} & \text{minimize } c^\top x \\ & \text{subject to } Ax \leq b, x \geq 0, x \in \mathbb{Z}^n \end{array}$$

where, as is customary in optimization, inequalities between vectors are interpreted in a coordinate-wise sense. As we will be discussing complexity and will assume finite bit length inputs with A, b, c having rational entries, by clearing denominators, there is no loss of generality in assuming that they have integer entries, that is, $A \in \mathbb{Z}^{m \times n}, b \in \mathbb{Z}^m, c \in \mathbb{Z}^n$.

LP is famously solvable in polynomial time to arbitrary ε -relative accuracy (Khachiyan 1979), with its time complexity improved over the years, notably in Karmarkar (1984) and Vaidya (1990), to its current bound in Cohen, Lee and Song (2019) that is essentially in terms of ω , the exponent of matrix multiplication we saw in Example 3.9. The natural question is: Polynomial in what? The aforementioned complexity bounds invariably involve the bit length of the input,

$$L := \sum_{i=1}^m \sum_{j=1}^n \log_2(|a_{ij}| + 1) + \sum_{i=1}^m \log_2(|b_i| + 1) + \log_2(mn) + 1,$$

but ideally we want an algorithm that runs in time polynomial in the size of the structure, rather than the size of the actual numbers involved. This is called a strongly polynomial-time algorithm. Whether it exists for LP is still famously unresolved (Smale 1998, Problem 9), although a result of Tardos (1986) shows that it is plausible: there is a polynomial-time algorithm for LP whose time complexity is independent of the vectors b and c and depends only on m, n and the largest subdeterminant³³ of A ,

$$\Delta := \max\{|\det A_{\sigma \times \tau}| : \sigma \subseteq [m], \tau \subseteq [n]\}.$$

Note that L depends on A, b, c and Δ only on A .

We will now assume that we have a box constraint $0 \leq x \leq 1$ in our LP and ILP. In this case ILP becomes a zero-one ILP with $x \in \{0, 1\}^n$; such an ILP is also known to have time complexity that does not depend on b and c although it generally still depends on the entries of A and not just Δ (Frank and Tardos 1987). We will let the time complexity of the box-constrained LP be $\text{LP}(m, n, \Delta)$ and let that of the zero-one ILP be $\text{ILP}(m, n, A)$. The former has time complexity polynomial in m, n, Δ by Tardos (1986); the latter is polynomial-time if A is *totally unimodular*, i.e. when $\Delta = 1$ (Schrijver 1986, Chapter 19).

Observe that the linear objective

$$c^\top x = c_1 x_1 + \cdots + c_n x_n$$

³³ Recall from Example 4.5 that a matrix is a function $A: [m] \times [n] \rightarrow \mathbb{R}$ and here $A_{\sigma, \tau}$ is the function restricted to the subset $\sigma \times \tau \subseteq [m] \times [n]$. Recall from page 568 that a non-square determinant is identically zero.

is additively separable. It turns out that this separability, not so much that it is linear, is the key to polynomial-time solvability. Consider an objective function $f: \mathbb{R}^n \rightarrow \mathbb{R}$ of the form

$$f(x) = f_1(x_1) + \dots + f_n(x_n),$$

where $f_1, \dots, f_n: \mathbb{R} \rightarrow \mathbb{R}$ are all convex. Note that if we set $f_i(x_i) = c_i x_i$, which is linear and therefore convex, we recover the linear objective. A surprising result of Hochbaum and Shanthikumar (1990) shows that the separable convex programming problem

$$\begin{array}{ll} \text{SP} & \text{minimize} \quad f_1(x_1) + \dots + f_n(x_n) \\ & \text{subject to} \quad Ax \leq b, 0 \leq x \leq 1 \end{array}$$

and its zero-one variant

$$\begin{array}{ll} \text{ISP} & \text{minimize} \quad f_1(x_1) + \dots + f_n(x_n) \\ & \text{subject to} \quad Ax \leq b, x \in \{0, 1\}^n \end{array}$$

are solvable to ε -accuracy in time complexities

$$\begin{aligned} \text{SP}(m, n, \Delta) &= \log_2 \left(\frac{\beta}{2\varepsilon} \right) \text{LP}(m, 8n^2\Delta, \Delta), \\ \text{ISP}(m, n, A) &= \log_2 \left(\frac{\beta}{2n\Delta} \right) \text{LP}(m, 8n^2\Delta, \Delta) + \text{ILP}(m, 4n^2\Delta, A \otimes \mathbb{1}_{4n\Delta}^\top) \end{aligned}$$

respectively. Here $\mathbb{1}_n \in \mathbb{R}^n$ denotes the vector of all ones, \otimes the Kronecker product, and $\beta > 0$ is some constant. A consequence is that SP is always polynomial-time solvable and ISP is polynomial-time solvable for totally unimodular A . The latter is a particularly delicate result, and even a slight deviation yields an NP-hard problem (Baldick 1995). While the objective $f = f_1 + \dots + f_n$ is convex, we remind readers of Example 3.16: it is a mistake to think that all convex optimization problems have polynomial-time algorithms.

Note that we could have stated the results above in terms of multiplicatively separable functions: if $g_1, \dots, g_n: \mathbb{R} \rightarrow \mathbb{R}_{++}$ are log convex, and $g: \mathbb{R}^n \rightarrow \mathbb{R}_{++}$ is defined by $g(x_1, \dots, x_n) = g_1(x_1) \cdots g_n(x_n)$, then SP and ISP are equivalent to

$$\begin{array}{ll} \text{minimize} & \log g(x_1, \dots, x_n) & \text{minimize} & \log g(x_1, \dots, x_n) \\ \text{subject to} & Ax \leq b, 0 \leq x \leq 1, & \text{subject to} & Ax \leq b, x \in \{0, 1\}^n, \end{array}$$

although it would be unnatural to discuss LP and ILP in these forms.

We will discuss another situation where additive separability arises.

Example 4.41 (separable Hamiltonians). The Hamilton equations

$$\frac{dp_i}{dt} = -\frac{\partial H}{\partial x_i}(x, p), \quad \frac{dx_i}{dt} = \frac{\partial H}{\partial p_i}(x, p), \quad i = 1, \dots, n, \tag{4.116}$$

are said to be separable if the Hamiltonian $H: \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ is additively separable:

$$H(x, p) = V(x) + T(p), \tag{4.117}$$

that is, the kinetic energy $T: \mathbb{R}^n \rightarrow \mathbb{R}$ depends only on momentum $p = (p_1, \dots, p_n)$ and the potential energy $V: \mathbb{R}^n \rightarrow \mathbb{R}$ depends only on position $x = (x_1, \dots, x_n)$, a common scenario. In this case the system (4.116) famously admits a finite difference scheme that is both *explicit*, that is, iteration depends only on quantities already computed in a previous step, and *symplectic*, that is, it conforms to the tensor transformation rules with change-of-coordinates matrices from $\text{Sp}(2n, \mathbb{R})$ on page 576. Taking $n = 1$ for illustration, the equations in (4.116) are

$$\frac{dp}{dt} = -V'(x), \quad \frac{dx}{dt} = -T'(p),$$

and with backward Euler on the first, forward Euler on the second, we obtain the finite difference scheme

$$p^{(k+1)} = p^{(k)} + V'(x^{(k+1)})\Delta t, \quad x^{(k+1)} = x^{(k)} - T'(x^{(k)})\Delta t,$$

which is easily seen to be explicit as $x^{(k+1)}$ may be computed before $p^{(k+1)}$; to show that it is a symplectic integrator requires a bit more work and the reader may consult [Stuart and Humphries \(1996, pp. 580–582\)](#). Note that we could have written (4.117) in an entirely equivalent multiplicatively separable form with $e^{H(x,p)} = e^{H(x)} e^{T(p)}$, although our subsequent discussion would be somewhat awkward.

We now present an example with separability structures in both problem and solution and where both additive and multiplicative separability play a role.

Example 4.42 (one-electron and Hartree–Fock approximations). The time-dependent Schrödinger equation for a system of d particles in \mathbb{R}^3 is

$$i\hbar \frac{\partial}{\partial t} f(x, t) = \left[-\frac{\hbar^2}{2m} \Delta + V(x) \right] f(x, t).$$

Here $x = (x_1, \dots, x_d) \in \mathbb{R}^{3n}$ represents the positions of the d particles, V a real-valued function representing potential, and

$$\Delta = \Delta_1 + \Delta_2 + \dots + \Delta_d$$

with each $\Delta_i: L^2(\mathbb{R}^3) \rightarrow L^2(\mathbb{R}^3)$ a copy of the Laplacian on \mathbb{R}^3 corresponding to the i th particle. Note that these do not need to be in Cartesian coordinates. For instance, we might have

$$\Delta_i = \frac{1}{r_i^2} \frac{\partial}{\partial r_i} \left(r_i^2 \frac{\partial}{\partial r_i} \right) + \frac{1}{r_i^2 \sin \theta_i} \frac{\partial}{\partial \theta_i} \left(\sin \theta_i \frac{\partial}{\partial \theta_i} \right) + \frac{1}{r_i^2 \sin^2 \theta_i} \frac{\partial^2}{\partial \phi_i^2}$$

with $x_i = (r_i, \theta_i, \phi_i)$, and $i = 1, \dots, d$.

We will drop the constants, which are unimportant to our discussions, and just keep their signs:

$$(-\Delta + V)f - i\partial_t f = 0. \quad (4.118)$$

Separation of variables (4.97) applies to give us

$$(-\Delta + V) \otimes I + I \otimes (-i\partial_t) \longrightarrow \begin{cases} (-\Delta + V)\varphi = E\varphi, \\ -i\partial_t\psi = -E\psi, \end{cases}$$

where we have written our separation constant as $-E$. The second equation is trivial to solve, $\psi(t) = e^{-iEt}$, and as (4.118) is linear, the solution f is given by a linear combination of $\psi \otimes \varphi$ over all possible values of E and our main task is to determine φ and E from the first equation, called the time-independent Schrödinger equation for d particles. Everything up to this stage is similar to our discussion for the wave equation in Example 4.33. The difference is that we now have an extra V term: (E, φ) are an eigenpair of $-\Delta + V$.

The motivation behind the one-electron, Hartree–Fock and other approximations is as follows. If the potential V is additively separable,

$$V(x) = V_1(x_1) + V_2(x_2) + \cdots + V_d(x_d), \quad (4.119)$$

then the eigenfunction φ is multiplicatively separable,

$$\varphi(x) = \varphi_1(x_1)\varphi_2(x_2) \cdots \varphi_d(x_d),$$

and the eigenvalue E is additively separable,

$$E = E_1 + E_2 + \cdots + E_d.$$

The reason is that with (4.119) we have

$$\sum_{i=1}^d (-\Delta_i + V_i)\varphi - E\varphi = 0,$$

and we may apply (4.97) to get

$$\begin{aligned} & [(-\Delta_1 + V_1) \otimes I \otimes \cdots \otimes I + I \otimes (-\Delta_2 + V_2) \otimes \cdots \otimes I \\ & + I \otimes \cdots \otimes I \otimes (-\Delta_d + V_d - E)](\varphi_1 \otimes \varphi_2 \otimes \cdots \otimes \varphi_d) = 0 \\ & \longrightarrow \begin{cases} (-\Delta_1 + V_1)\varphi_1 = E_1\varphi_1, \\ (-\Delta_2 + V_2)\varphi_2 = E_2\varphi_2, \\ \vdots \\ (-\Delta_d + V_d)\varphi_d = (E - E_1 - \cdots - E_{d-1})\varphi_d. \end{cases} \end{aligned} \quad (4.120)$$

Note that we may put the $-E\varphi$ term with any $-\Delta_i + V_i$; we just chose $i = d$ for convenience. If we write $E_d := E - E_1 - \cdots - E_{d-1}$ and $\varphi = \varphi_1 \otimes \cdots \otimes \varphi_d$, we obtain the required expressions. So (4.120) transforms a d -particle Schrödinger equation into d one-particle Schrödinger equations.

If one could solve the d -particle Schrödinger equation, then one could in principle determine all chemical and physical properties of atoms and molecules, so the above discussion seems way too simple, and indeed it is: the additivity in (4.119) can only happen if the particles do not interact. In this context an operator of the form in

(4.96) is called a non-interacting Hamiltonian, conceptually useful but unrealistic. It is, however, the starting point from which various approximation schemes are developed. To see the necessity of approximation, consider what appears to be a slight deviation from (4.119). Let V include only pairwise interactions:

$$V(x) = \sum_{i=1}^d V_i(x_i) + \sum_{i<j} V_{ij}(x_i, x_j),$$

that is, no higher-order terms of the form $V_{ijk}(x_i, x_j, x_k)$, and we may even fix $V_{ij}(x_i, x_j) = 1/\|x_i - x_j\|$. But the equation $(-\Delta + V)\varphi = E\varphi$ then becomes computationally intractable in multiple ways (Whitfield, Love and Aspuru-Guzik 2013).

The one-electron approximation and Hartree–Fock approximation are based on the belief that

$$V(x) \approx V_1(x_1) + \dots + V_d(x_d) \quad \Rightarrow \quad \begin{cases} \varphi(x) \approx \varphi_1(x_1) \cdots \varphi_d(x_d), \\ E \approx E_1 + \dots + E_d, \end{cases}$$

with ‘ \approx ’ interpreted differently and with different tools: the former uses perturbation theory and the latter calculus of variations. While these approximation methods are only tangentially related to the topic of this section, they do tell us how additive and multiplicative separability can be approximated and thus we will briefly describe the key ideas. Our discussion below is based on Fischer (1977), Faddeev and Yakubovskii (2009, Chapters 33, 34, 50, 51) and Hannabuss (1997, Chapter 12).

In most scenarios, $V(x)$ will have an additively separable component $V_1(x_1) + \dots + V_d(x_d)$ and an additively inseparable component comprising the higher-order interactions $V_{ij}(x_i, x_j), V_{ijk}(x_i, x_j, x_k), \dots$. We will write

$$H_0 := \sum_{i=1}^d (-\Delta_i + V_i), \quad H_1 := -\Delta + V = H_0 + W,$$

with W accounting for the inseparable terms. The one-electron approximation method uses perturbation theory to recursively add correction terms to the eigenpairs of H_0 to obtain those of H_1 . Ignoring rigour, the method is easy to describe. Set $H_\varepsilon := (1 - \varepsilon)H_0 + \varepsilon W$ and plug the power series

$$\lambda_\varepsilon = \lambda^{(0)} + \varepsilon\lambda^{(1)} + \varepsilon^2\lambda^{(2)} + \dots, \quad \psi_\varepsilon = \psi^{(0)} + \varepsilon\psi^{(1)} + \varepsilon^2\psi^{(2)} + \dots$$

into $H_\varepsilon\psi_\varepsilon = \lambda_\varepsilon\psi_\varepsilon$ to see that

$$H_0\psi^{(0)} = \lambda^{(0)}\psi^{(0)}, \tag{4.121}$$

$$H_0\psi^{(1)} + W\psi^{(0)} = \lambda^{(0)}\psi^{(1)} + \lambda^{(1)}\psi^{(0)},$$

⋮

$$H_0\psi^{(k)} + W\psi^{(k-1)} = \lambda^{(0)}\psi^{(k)} + \lambda^{(1)}\psi^{(k-1)} + \dots + \lambda^{(k)}\psi^{(0)}. \tag{4.122}$$

By (4.121), $\lambda^{(0)}$ and $\psi^{(0)}$ form an eigenpair of H_0 . The subsequent terms $\lambda^{(k)}$ and $\psi^{(k)}$ are interpreted as the k th-order corrections to the eigenpair of H_0 , and as $\lambda^{(k)}$ and $\psi^{(k)}$ may be recursively calculated from (4.122), the hope is that

$$\lambda^{(0)} + \lambda^{(1)} + \lambda^{(2)} + \dots, \quad \psi^{(0)} + \psi^{(1)} + \psi^{(2)} + \dots$$

would converge to the desired eigenpair of H_1 . For instance, if H_0 has an orthonormal basis of eigenfunctions $\{\psi_i : i \in \mathbb{N}\}$ with eigenvalues $\{\lambda_i : i \in \mathbb{N}\}$ and λ_j is a simple eigenvalue, then (Hannabuss 1997, Theorem 12.4.3)

$$\lambda_j^{(1)} = \langle W\psi_j, \psi_j \rangle, \quad \psi_j^{(1)} = \sum_{i \neq j} \frac{\langle W\psi_j, \psi_i \rangle}{\lambda_j - \lambda_i} \psi_i, \quad \lambda_j^{(2)} = \sum_{i \neq j} \frac{|\langle W\psi_j, \psi_i \rangle|^2}{\lambda_j - \lambda_i}.$$

Whereas one-electron approximation considers perturbation of additive separability in the problem (the Hamiltonian), Hartree–Fock approximation considers variation of multiplicative separability in the solution (the wave function). Observe that the non-linear functional given by the Rayleigh quotient $\mathcal{E}(\varphi) = \langle (-\Delta + V)\varphi, \varphi \rangle / \|\varphi\|^2$ is stationary, *i.e.* $\delta\mathcal{E} = 0$, if and only if $(-\Delta + V)\varphi = E\varphi$, that is, its Euler–Lagrange equation gives the Schrödinger equation and its Lagrange multiplier gives the eigenvalue E . The Hartree–Fock approximation seeks stationarity under the additional condition of multiplicative separability, *i.e.* $\varphi = \varphi_1 \otimes \dots \otimes \varphi_d$. If the non-linear functional

$$\mathcal{E}(\varphi_1, \dots, \varphi_d) = \langle (-\Delta + V)\varphi_1 \otimes \dots \otimes \varphi_d, \varphi_1 \otimes \dots \otimes \varphi_d \rangle$$

with constraints $\|\varphi_1\|^2 = \dots = \|\varphi_d\|^2 = 1$ is stationary with respect to variations in $\varphi_1, \dots, \varphi_d$, *i.e.* $\delta\mathcal{L} = 0$, where

$$\mathcal{L}(\varphi_1, \dots, \varphi_d, \lambda_1, \dots, \lambda_d) = \mathcal{E}(\varphi_1, \dots, \varphi_d) - \lambda_1 \|\varphi_1\|^2 - \dots - \lambda_d \|\varphi_d\|^2$$

is the Lagrangian with Lagrange multipliers $\lambda_1, \dots, \lambda_d$, then the Euler–Lagrange equations give us

$$\left[-\Delta_i + \sum_{j \neq i} \int_{\mathbb{R}^3} |\varphi_j(y)|^2 V(x, y) dy \right] \varphi_i = \lambda_i \varphi_i, \quad i = 1, \dots, d. \quad (4.123)$$

These equations make physical sense. For a fixed i , (4.123) is the Schrödinger equation for particle i in a potential field due to the charge of particle j ; this charge is spread over space with density $|\varphi_j|^2$, and we have to sum over the potential fields created by all particles $j = 1, \dots, i-1, i+1, \dots, d$. While the system of coupled integro-differential equations (4.123) generally would not have an analytic solution like the one in Example 4.35, it readily lends itself to numerical solution via a combination of quadrature and finite difference (Fischer 1977, Chapter 6, Section 7).

We have ignored the spin variables in the last and the next examples, as spin has already been addressed in Examples 4.7 and would only be an unnecessary distraction here. So, strictly speaking, these examples are about Hartree approximation

(Hartree 1928), *i.e.* no spin, and not Hartree–Fock approximation (Fock 1930), *i.e.* with spin.

Example 4.43 (multiconfiguration Hartree–Fock). We will continue our discussion in the last example but with some simplifications. We let $d = 2$ and $V(x, y) = V(y, x)$. We emulate the Hartree–Fock approximation for the time-independent Schrödinger equation $(-\Delta + V)f = Ef$ but now with an ansatz of the form

$$f = a_1\varphi_1 \otimes \varphi_1 + a_2\varphi_2 \otimes \varphi_2, \quad (4.124)$$

where $\varphi_1, \varphi_2 \in L^2(\mathbb{R}^3)$ are orthonormal and $a = (a_1, a_2) \in \mathbb{R}^2$ is a unit vector, that is, we consider the non-linear functional

$$\begin{aligned} \mathcal{E}(\varphi_1, \varphi_2, a_1, a_2) &= a_1^2 \langle (-\Delta + V)\varphi_1 \otimes \varphi_1, \varphi_1 \otimes \varphi_1 \rangle \\ &\quad + 2a_1a_2 \langle (-\Delta + V)\varphi_1 \otimes \varphi_1, \varphi_2 \otimes \varphi_2 \rangle \\ &\quad + a_2^2 \langle (-\Delta + V)\varphi_2 \otimes \varphi_2, \varphi_2 \otimes \varphi_2 \rangle \end{aligned}$$

with constraints $\|\varphi_1\|^2 = \|\varphi_2\|^2 = 1$, $\langle \varphi_1, \varphi_2 \rangle = 0$, $\|a\|^2 = 1$. Introducing Lagrange multipliers $\lambda_{11}, \lambda_{12}, \lambda_{22}$ and λ for the constraints, we obtain the Lagrangian

$$\begin{aligned} \mathcal{L}(\varphi_1, \varphi_2, a_1, a_2, \lambda_{11}, \lambda_{12}, \lambda_{22}, \lambda) \\ = \mathcal{E}(\varphi_1, \varphi_2, a_1, a_2) + \lambda_{11}\|\varphi_1\|^2 + \lambda_{12}\langle \varphi_1, \varphi_2 \rangle + \lambda_{22}\|\varphi_2\|^2 - \lambda\|a\|^2. \end{aligned}$$

For $i, j \in \{1, 2\}$ we write

$$b_{ij} = \langle (-\Delta + V)\varphi_i \otimes \varphi_i, \varphi_j \otimes \varphi_j \rangle, \quad c_{ij}(x) = \int_{\mathbb{R}^3} \varphi_i(x)\varphi_j(y)V(x, y) dy.$$

The stationarity conditions $\nabla_a \mathcal{L} = 0$ and $\delta \mathcal{L} = 0$ give us

$$\begin{bmatrix} b_{11} - \lambda & b_{12} \\ b_{12} & b_{22} - \lambda \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \end{bmatrix} = 0$$

and

$$\begin{bmatrix} -\Delta_1 \varphi_1 \\ -\Delta_2 \varphi_2 \end{bmatrix} = \begin{bmatrix} c_{11}(x) - \lambda_{11} & (a_2/a_1)(c_{12}(x) - \lambda_{12}) \\ (a_1/a_2)(c_{12}(x) - \lambda_{12}) & c_{22}(x) - \lambda_{22} \end{bmatrix} \begin{bmatrix} \varphi_1 \\ \varphi_2 \end{bmatrix}$$

respectively. What we have described here is a simplified version of the multiconfiguration Hartree–Fock approximation (Fischer 1977, Chapters 3 and 4), and it may also be viewed as the starting point for a variety of other heuristics, one of which will be our next example.

We will describe an extension of the ansatz in (4.124) to more summands and to higher order.

Example 4.44 (tensor networks). Instead of restricting ourselves to $L^2(\mathbb{R}^3)$ as in the last two examples, we will now assume arbitrary separable Hilbert spaces $\mathbb{H}_1, \dots, \mathbb{H}_d$ to allow for spin and other properties. We will seek a solution $f \in$

$\mathbb{H}_1 \otimes \cdots \otimes \mathbb{H}_d$ to the Schrödinger equation for d particles. Readers are reminded of the discussion at the end of Example 4.16. Note that by definition of \otimes , f has finite rank regardless of whether $\mathbb{H}_1, \dots, \mathbb{H}_d$ are finite- or infinite-dimensional. Let $\mu\text{rank}(f) = (r_1, \dots, r_d)$. Then there is a decomposition

$$f = \sum_{i=1}^{r_1} \sum_{j=1}^{r_2} \cdots \sum_{k=1}^{r_d} c_{ij\dots k} \varphi_i \otimes \psi_j \otimes \cdots \otimes \theta_k \tag{4.125}$$

with orthogonal factors (De Lathauwer, De Moor and Vandewalle 2000, equation 13)

$$\langle \varphi_i, \varphi_j \rangle = \langle \psi_i, \psi_j \rangle = \cdots = \langle \theta_i, \theta_j \rangle = \begin{cases} 0 & i \neq j, \\ 1 & i = j. \end{cases}$$

Note that when our spaces have inner products, any multilinear rank decomposition (4.63) may have its factors orthogonalized (De Lathauwer *et al.* 2000, Theorem 2), that is, the orthogonality constraints do not limit the range of possibilities in (4.125). By definition of multilinear rank, there exist subspaces $\mathbb{H}'_1, \dots, \mathbb{H}'_d$ of dimensions r_1, \dots, r_d with $f \in \mathbb{H}'_1 \otimes \cdots \otimes \mathbb{H}'_d$ that attain the minimum in (4.62). As such, we may replace \mathbb{H}_i with \mathbb{H}'_i at the outset, and to simplify our discussion we may as well assume that $\mathbb{H}_1, \dots, \mathbb{H}_d$ are of dimensions r_1, \dots, r_d .

The issue with (4.125) is that there is an exponential number of rank-one terms as d increases. Suppose $r_1 = \cdots = r_d = r$; then there are r^d summands in (4.125). This is not unexpected because (4.125) is the most general form a finite-rank tensor can take. Here the ansatz in (4.125) describes the whole space $\mathbb{H}_1 \otimes \cdots \otimes \mathbb{H}_d$ and does not quite serve its purpose; an ansatz is supposed to be an educated guess, typically based on physical insights, that captures a small region of the space where the solution likely lies. The goal of *tensor networks* is to provide such an ansatz by limiting the coefficients $[c_{ij\dots k}] \in \mathbb{R}^{r_1 \times \cdots \times r_d}$ to a much smaller set. The first and best-known example is the *matrix product states* tensor network (Anderson 1959, White 1992, White and Huse 1993), which imposes on the coefficients the structure

$$c_{ij\dots k} = \text{tr}(A_i B_j \cdots C_k), \quad A_i \in \mathbb{R}^{n_1 \times n_2}, B_j \in \mathbb{R}^{n_2 \times n_3}, \dots, C_k \in \mathbb{R}^{n_d \times n_1}$$

for $i = 1, \dots, r_1, j = 1, \dots, r_2, \dots, k = 1, \dots, r_d$. An ansatz of the form

$$f = \sum_{i=1}^{r_1} \sum_{j=1}^{r_2} \cdots \sum_{k=1}^{r_d} \text{tr}(A_i B_j \cdots C_k) \varphi_i \otimes \psi_j \otimes \cdots \otimes \theta_k \tag{4.126}$$

is called a matrix product state or MPS (Affleck, Kennedy, Lieb and Tasaki 1987). Note that the coefficients are now parametrized by $r_1 + r_2 + \cdots + r_d$ matrices of various sizes. For easy comparison, if $r_1 = \cdots = r_d = r$ and $n_1 = \cdots = n_d = n$, then the coefficients in (4.125) have r^d degrees of freedom whereas those in (4.126) only have rdn^2 . When $n_1 = 1$, the first and last matrices in (4.126) are a row and a column vector respectively; as the trace of a 1×1 matrix is itself, we may drop the ‘tr’ in (4.126). This special case with $n_1 = 1$ is sometimes called MPS with open

boundary conditions (Anderson 1959) and the more general case is called MPS with periodic conditions.

The above discussion of matrix product states conceals an important structure. Take $d = 3$ and denote the entries of the matrices as $A_i = [a_{\alpha\beta}^{(i)}]$, $B_j = [b_{\beta\gamma}^{(j)}]$ and $C_k = [c_{\gamma\alpha}^{(k)}]$. Then

$$\begin{aligned} f &= \sum_{i,j,k=1}^{r_1,r_2,r_3} \text{tr}(A_i B_j C_k) \varphi_i \otimes \psi_j \otimes \theta_k \\ &= \sum_{i,j,k=1}^{r_1,r_2,r_3} \left[\sum_{\alpha,\beta,\gamma=1}^{n_1,n_2,n_3} a_{\alpha\beta}^{(i)} b_{\beta\gamma}^{(j)} c_{\gamma\alpha}^{(k)} \varphi_i \otimes \psi_j \otimes \theta_k \right] \\ &= \sum_{\alpha,\beta,\gamma=1}^{n_1,n_2,n_3} \left[\sum_{i=1}^{r_1} a_{\alpha\beta}^{(i)} \varphi_i \right] \otimes \left[\sum_{j=1}^{r_2} b_{\beta\gamma}^{(j)} \psi_j \right] \otimes \left[\sum_{k=1}^{r_3} c_{\gamma\alpha}^{(k)} \theta_k \right] \\ &= \sum_{\alpha,\beta,\gamma=1}^{n_1,n_2,n_3} \varphi_{\alpha\beta} \otimes \psi_{\beta\gamma} \otimes \theta_{\gamma\alpha}, \end{aligned}$$

where

$$\varphi_{\alpha\beta} := \sum_{i=1}^{r_1} a_{\alpha\beta}^{(i)} \varphi_i, \quad \psi_{\beta\gamma} := \sum_{j=1}^{r_2} b_{\beta\gamma}^{(j)} \psi_j \quad \text{and} \quad \theta_{\gamma\alpha} := \sum_{k=1}^{r_3} c_{\gamma\alpha}^{(k)} \theta_k.$$

In other words, the indices have the incidence structure of an undirected graph, in this case a triangle. This was first observed in Landsberg, Qi and Ye (2012) and later generalized in Ye and Lim (2018b), the bottom line being that any tensor network state is a sum of separable functions indexed by a graph. In the following, we show some of the most common tensor network states, written in this simplified form, together with the graphs they correspond to in Figure 4.4.

Periodic matrix product states:

$$f(x, y, z) = \sum_{i,j,k=1}^{n_1,n_2,n_3} \varphi_{ij}(x) \psi_{jk}(y) \theta_{ki}(z).$$

Tree tensor network states:

$$f(x, y, z, w) = \sum_{i,j,k=1}^{n_1,n_2,n_3} \varphi_{ijk}(x) \psi_i(y) \theta_j(z) \pi_k(w).$$

Open matrix product states:

$$f(x, y, z, u, v) = \sum_{i,j,k,l=1}^{n_1,n_2,n_3,n_4} \varphi_i(x) \psi_{ij}(y) \theta_{jk}(z) \pi_{kl}(u) \rho_l(v).$$

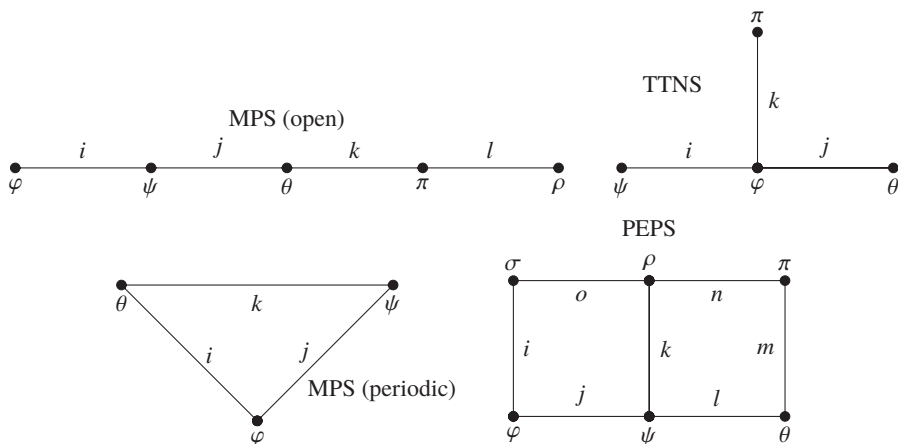


Figure 4.4. Graphs associated with common tensor networks.

Projected entangled pair states:

$$f(x, y, z, u, v, w) = \sum_{i,j,k,l,m,n,o=1}^{n_1, n_2, n_3, n_4, n_5, n_6, n_7} \varphi_{ij}(x)\psi_{jkl}(y)\theta_{lm}(z)\pi_{mn}(u)\rho_{nko}(v)\sigma_{oi}(w).$$

The second and the last, often abbreviated to TTNS and PEPS, were proposed by Shi, Duan and Vidal (2006) and Verstraete and Cirac (2004) respectively. More generally, any periodic MPS corresponds to a cycle graph and any open MPS corresponds to a line graph. We have deliberately written them without the \otimes symbol to emphasize that all these ansätze are just sums of separable functions, differing only in terms of how their factors are indexed.

One difference between the Hartree–Fock approximations in Examples 4.42 and 4.43 and the tensor network methods here is that the factors $\varphi_i, \psi_j, \dots, \theta_k$ in (4.126) are often fixed in advance as some standard bases of $\mathbb{H}_1, \dots, \mathbb{H}_d$, called a *local basis* in this context. The computational effort is solely in determining the coefficients $c_{ij\dots k}$; in the case of MPS this can be done via several singular value decompositions (Orús 2014). Indeed, because of the way it is computed, the coefficients of the MPS ansatz are sometimes represented as

$$\text{tr}(Q_1 \Sigma_1 Q_2 \Sigma_2 \dots \Sigma_d Q_{d+1}), \quad Q_i \in O(n_i), \Sigma_i \in \mathbb{R}^{n_i \times n_{i+1}}$$

for $i = 1, \dots, d$, with Σ_i a diagonal matrix. That any $\text{tr}(A_1 A_2 \dots A_d)$ with $A_i \in \mathbb{R}^{n_i \times n_{i+1}}$ may be written in this form follows from the singular value decompositions $A_i = U_i \Sigma_i V_i^T$ and setting $Q_{i+1} = V_i^T U_{i+1}$ for $i = 1, \dots, d - 1$, and $Q_1 = U_1, Q_{d+1} = V_d^T$.

We end with a brief word about our initial assumption that $f \in \mathbb{H}_1 \otimes \dots \otimes \mathbb{H}_d$. For a solution to a PDE, we expect that $f \in \mathbb{H}_1 \widehat{\otimes}_F \dots \widehat{\otimes}_F \mathbb{H}_d$, where $\widehat{\otimes}_F$ is the Hilbert–Schmidt tensor product discussed in Examples 4.18 and 4.19. Since any

tensor product basis (4.43) formed from orthonormal bases of $\mathbb{H}_1, \dots, \mathbb{H}_d$ is an orthonormal basis of $\mathbb{H}_1 \widehat{\otimes} \cdots \widehat{\otimes} \mathbb{H}_d$ with respect to the inner product in (4.64), it is always possible to express $f \in \mathbb{H}_1 \widehat{\otimes} \cdots \widehat{\otimes} \mathbb{H}_d$ in the form (4.125) as long as we allow r_1, \dots, r_d to be infinite. The ones with finite r_1, \dots, r_d are precisely the finite-rank tensors in $\mathbb{H}_1 \otimes \cdots \otimes \mathbb{H}_d$, but any other $f \in \mathbb{H}_1 \widehat{\otimes}_F \cdots \widehat{\otimes}_F \mathbb{H}_d$ can be approximated to arbitrary accuracy by elements of $\mathbb{H}_1 \otimes \cdots \otimes \mathbb{H}_d$, that is, f is a limit of finite-rank tensors with respect to the Hilbert–Schmidt norm.

We will now embark on a different line of applications of definition ③ tied to our discussions of integral kernels in Examples 3.7, 4.2 and 4.20. As we saw in Example 3.7, some of the most important kernels take the form $K(v, v') = f(\|v - v'\|)$ for some real-valued function f on a norm space \mathbb{V} . In this case we usually write $\kappa(v - v') = f(\|v - v'\|)$ and call κ a radial basis kernel (Buhmann 2003). These are usually not separable like those in Examples 4.37 or 4.38. The fast multipole method (Greengard and Rokhlin 1987) may be viewed as an algorithm for approximating the convolution of certain radial basis kernels with certain functions. We will discuss how tensors arise in this algorithm.

Example 4.45 (multipole tensors). Let \mathbb{V} be an inner product space. As usual we will denote the inner product and its induced norm by $\langle \cdot, \cdot \rangle$ and $\|\cdot\|$ respectively. As we saw in Examples 4.17 and 4.22, this induces an inner product and norm on $\mathbb{V}^{\otimes d}$ that we may denote with the same notation. For any $s \geq 0$, the function

$$\kappa_s : \mathbb{V} \setminus \{0\} \rightarrow \mathbb{R}, \quad \kappa_s(v) = \begin{cases} \log \|v\| & s = 0, \\ 1/\|v\|^s & s > 0, \end{cases}$$

is called a Newtonian kernel. Henceforth we will drop the subscript s from κ_s since this is normally fixed throughout. If \mathbb{V} is equipped with some measure, then the convolution

$$\kappa * f(v) = \int_{w \in \mathbb{V}} \kappa(v - w) f(w) dw$$

is called the Newtonian potential³⁴ of f . We will assume $s > 0$ in the following; the $s = 0$ case is similar. Our interpretation of higher-order derivatives as multilinear maps in Example 3.2 makes their calculation an effortless undertaking:

$$\begin{aligned} [D\kappa(v)](h) &= -\frac{s}{\|v\|^{s+2}} \langle v, h \rangle, \\ [D^2\kappa(v)](h_1, h_2) &= -\frac{s}{\|v\|^{s+2}} \langle h_2, h_1 \rangle + \frac{s(s+2)}{\|v\|^{s+4}} \langle v, h_1 \rangle \langle v, h_2 \rangle, \end{aligned}$$

³⁴ Strictly speaking, the names are used for the case when $s = \dim \mathbb{V} - 2$, and for more general s they are called the Riesz kernel and Riesz potentials. Because of the singularity in κ , the integral is interpreted in a principal value sense as in (3.32) but we ignore such details (Stein 1993, Chapter 1, Section 8.18). We have also dropped some multiplicative constants.

$$\begin{aligned}
[D^3 \kappa(v)](h_1, h_2, h_3) &= \frac{s(s+2)}{\|v\|^{s+4}} [\langle v, h_1 \rangle \langle h_2, h_3 \rangle + \langle v, h_2 \rangle \langle h_1, h_3 \rangle + \langle v, h_3 \rangle \langle h_1, h_2 \rangle] \\
&\quad - \frac{s(s+2)(s+4)}{\|v\|^{s+6}} \langle v, h_1 \rangle \langle v, h_2 \rangle \langle v, h_3 \rangle, \\
[D^4 \kappa(v)](h_1, h_2, h_3, h_4) &= \frac{s(s+2)}{\|v\|^{s+4}} [\langle h_1, h_2 \rangle \langle h_3, h_4 \rangle + \langle h_1, h_3 \rangle \langle h_2, h_4 \rangle + \langle h_1, h_4 \rangle \langle h_2, h_3 \rangle] \\
&\quad - \frac{s(s+2)(s+4)}{\|v\|^{s+6}} [\langle v, h_1 \rangle \langle v, h_2 \rangle \langle h_3, h_4 \rangle + \langle v, h_1 \rangle \langle v, h_3 \rangle \langle h_2, h_4 \rangle \\
&\quad + \langle v, h_1 \rangle \langle v, h_4 \rangle \langle h_2, h_3 \rangle + \langle v, h_2 \rangle \langle v, h_3 \rangle \langle h_1, h_4 \rangle \\
&\quad + \langle v, h_2 \rangle \langle v, h_4 \rangle \langle h_1, h_3 \rangle + \langle v, h_3 \rangle \langle v, h_4 \rangle \langle h_1, h_2 \rangle] \\
&\quad + \frac{s(s+2)(s+4)(s+6)}{\|v\|^{s+8}} \langle v, h_1 \rangle \langle v, h_2 \rangle \langle v, h_3 \rangle \langle v, h_4 \rangle.
\end{aligned}$$

To obtain these, all we need is the expression for $[D\kappa(v)](h)$, which follows from binomial expanding to linear terms in h , that is,

$$\kappa(v+h) = \frac{1}{\|v\|^s} \left[1 + \frac{2\langle v, h \rangle}{\|v\|^2} + \frac{\|h\|^2}{\|v\|^2} \right]^{-s/2} = \kappa(v) - \frac{s}{\|v\|^{s+2}} \langle v, h \rangle + O(\|h\|^2),$$

along with the observation that $[D\langle v, h' \rangle](h) = \langle h, h' \rangle$ and the product rule $Df \cdot g(v) = f(v) \cdot Dg(v) + g(v) \cdot Df(v)$. Applying these repeatedly gives us $D^k \kappa(v)$ in any coordinate system without having to calculate a single partial derivative.

As we saw in Example 4.29, these derivatives may be linearized via the universal factorization property. With an inner product, this is particularly simple. We will take $s = 1$, as this gives us the Coulomb potential, the most common case. We may rewrite the above expressions as

$$\begin{aligned}
[D\kappa(v)](h) &= \left\langle -\frac{v}{\|v\|^3}, h \right\rangle, \\
[D^2 \kappa(v)](h_1, h_2) &= \left\langle -\frac{I}{\|v\|^3} + \frac{3v \otimes v}{\|v\|^5}, h_1 \otimes h_2 \right\rangle, \\
[D^3 \kappa(v)](h_1, h_2, h_3) &= \left\langle \frac{3}{\|v\|^5} \left(v \otimes I + \frac{I}{v} + I \otimes v \right) - \frac{15v \otimes v \otimes v}{\|v\|^7}, h_1 \otimes h_2 \otimes h_3 \right\rangle.
\end{aligned}$$

The symbol where I and v appear above and below \otimes is intended to mean the following. Take any orthonormal basis e_1, \dots, e_n of \mathbb{V} ; then $I = \sum_{i=1}^n e_i \otimes e_i$, and we have

$$v \otimes I = \sum_{i=1}^n v \otimes e_i \otimes e_i, \quad \frac{I}{v} = \sum_{i=1}^n e_i \otimes v \otimes e_i, \quad I \otimes v = \sum_{i=1}^n e_i \otimes e_i \otimes v.$$

The tensors appearing in the first argument of the inner products are precisely the linearized derivatives and they carry important physical meanings:

$$\begin{aligned}
 \text{monopole} \quad \kappa(v) &= \frac{1}{\|v\|}, \\
 \text{dipole} \quad \partial\kappa(v) &= \frac{-1}{\|v\|^2} \hat{v}, \\
 \text{quadrupole} \quad \partial^2\kappa(v) &= \frac{1}{\|v\|^3} (3 \hat{v} \otimes \hat{v} - I), \\
 \text{octupole} \quad \partial^3\kappa(v) &= \frac{-1}{\|v\|^4} \left[15 \hat{v} \otimes \hat{v} \otimes \hat{v} - 3 \left(\hat{v} \otimes I + \frac{I}{\hat{v}} + I \otimes \hat{v} \right) \right],
 \end{aligned}$$

where $\hat{v} := v/\|v\|$. The norm-scaled k th linearized derivative

$$M^k \kappa(v) := \|v\|^{k+1} \partial^k \kappa(v) \in \mathbb{V}^{\otimes k}$$

is called a *multipole tensor* or k -pole tensor. The Taylor expansion (4.92) may then be written in terms of the multipole tensors,

$$\kappa(v) = \kappa(v_0) \sum_{k=0}^d \frac{1}{k!} \left\langle M^k \kappa(v_0), \frac{(v - v_0)^{\otimes k}}{\|v_0\|^k} \right\rangle + R(v - v_0),$$

and this is called a *multipole expansion* of $\kappa(v) = 1/\|v\|$. We may rewrite this in terms of a tensor Taylor series and a tensor geometric series as in (4.93),

$$\partial\kappa(v_0) = \sum_{k=0}^{\infty} \frac{1}{k!} M^k \kappa(v_0) \in \widehat{\mathbb{T}}(\mathbb{V}), \quad S\left(\frac{v - v_0}{\|v_0\|}\right) = \sum_{k=0}^{\infty} \frac{(v - v_0)^{\otimes k}}{\|v_0\|^k} \in \widehat{\mathbb{T}}(\mathbb{V}).$$

Note that

$$\partial\kappa(v) = 1 - \frac{1}{2!} \hat{v} + \frac{1}{3!} (3 \hat{v} \otimes \hat{v} - I) - \frac{1}{4!} \left[15 \hat{v} \otimes \hat{v} \otimes \hat{v} - 3 \left(\hat{v} \otimes I + \frac{I}{\hat{v}} + I \otimes \hat{v} \right) \right] + \dots$$

is an expansion in the tensor algebra $\widehat{\mathbb{T}}(\mathbb{V})$.

Let $f: \mathbb{V} \rightarrow \mathbb{R}$ be a compactly supported function that usually represents a charge distribution confined to a region within \mathbb{V} . The tensor-valued integral

$$(M^k \kappa) * f(v) = \int_{v \in \mathbb{V}} M^k \kappa(v - w) f(w) dw$$

is called a *multipole moment* or k -pole moment, and the Taylor expansion

$$\kappa * f(v) = \kappa(v_0) \sum_{k=0}^d \frac{1}{k!} \left\langle (M^k \kappa) * f(v_0), \frac{(v - v_0)^{\otimes k}}{\|v_0\|^k} \right\rangle + R(v - v_0),$$

is called a *multipole expansion* of f . One of the most common scenarios is when we have a finite collection of n point charges at positions $v_1, \dots, v_n \in \mathbb{V}$ with

charges $q_1, \dots, q_n \in \mathbb{R}$. Then f is given by

$$f(v) = q_1 \delta(v - v_1) + \dots + q_n \delta(v - v_n). \quad (4.127)$$

Our goal is to estimate the potential at a point v some distance away from v_1, \dots, v_n . Now consider a distinct point v_0 that is much nearer to v_1, \dots, v_n than to v in the sense that $\|v_i - v_0\| \leq c \|v - v_0\|$, $i = 1, \dots, n$, for some $c < 1$. Suppose there is just a single point charge with $f(v) = q_i \delta(v - v_i)$; then $(M^k \kappa) * f(v - v_0) = q_i M^k \kappa(v_i - v_0)$ and the multipole expansion of f at the point $v - v_i = (v - v_0) + (v_0 - v_i)$ about the point $v - v_0$ is

$$\begin{aligned} \kappa * f(v - v_i) &= \kappa(v - v_0) \sum_{k=0}^d \frac{q_i}{k!} \left\langle M^k \kappa(v_i - v_0), \frac{(v_0 - v_i)^{\otimes k}}{\|v - v_0\|^k} \right\rangle + R(v_0 - v_i) \\ &= \sum_{k=0}^d \varphi_k(v_i - v_0) \psi_k(v - v_0) + O(c^{d+1}), \end{aligned}$$

where we have written

$$\varphi_k(v) := \frac{(-1)^k q_i}{k!} \langle M^k \kappa(v), v^{\otimes k} \rangle, \quad \psi_k(v) := \frac{\kappa(v)}{\|v\|^k},$$

and since $R(v_0 - v_i) / \|v_0 - v_i\|^d \rightarrow 0$, we assume $R(v_0 - v_i) = O(c^{d+1})$ for simplicity. For the general case in (4.127) with n point charges, the potential at v is then approximated by summing over i :

$$\kappa * f(v) \approx \sum_{k=0}^d \left[\sum_{i=1}^n \varphi_k(v_i - v_0) \right] \psi_k(v - v_0). \quad (4.128)$$

This sum can be computed in $O(nd)$ complexity or $O(n \log(1/\varepsilon))$ for an ε -accurate algorithm. While the high level idea in (4.128) is still one of approximating a function by a sum of d separable functions, the fast multipole method involves a host of other clever ideas, not least among which are the techniques for performing such sums in more general situations (Demaine *et al.* 2005), for subdividing the region containing v_1, \dots, v_n into cubic cells (when $\mathbb{V} = \mathbb{R}^3$) and thereby organizing these computations into a tree-like multilevel algorithm (Barnes and Hut 1986). Clearly the approximation is good only when v is far from v_1, \dots, v_n relative to v_0 but the algorithm allows one to circumvent this requirement. We refer to Greengard and Rokhlin (1987) for further information.

Multipole moments, multipole tensors and multipole expansions are usually discussed in terms of coordinates (Jackson 1999, Chapter 4). The coordinate-free approach, the multipole expansion as an element of the tensor algebra, *etc.*, are results of our working with definitions ② and ③.

Although non-separable kernels make for more interesting examples, separable kernels arise as some of the most common multidimensional integral transforms,

as we saw in Example 4.38. Also, they warrant a mention if only to illustrate why separability in a kernel is computationally desirable.

Example 4.46 (discrete multidimensional transforms). Three of the best-known discrete transforms are the discrete Fourier transform we encountered in Example 3.14, the discrete Z-transform and the discrete cosine transform:

$$F(x_1, \dots, x_d) = \sum_{k_1=-\infty}^{\infty} \cdots \sum_{k_d=-\infty}^{\infty} f(k_1, \dots, k_d) e^{-ik_1x_1 - \cdots - ik_dx_d},$$

$$F(z_1, \dots, z_d) = \sum_{k_1=-\infty}^{\infty} \cdots \sum_{k_d=-\infty}^{\infty} f(k_1, \dots, k_d) z_1^{-k_1} \cdots z_d^{-k_d},$$

$$F(j_1, \dots, j_d) = \sum_{k_1=0}^{n_1-1} \cdots \sum_{k_d=0}^{n_d-1} f(k_1, \dots, k_d) \cos\left(\frac{\pi(2j_1+1)k_1}{2n_1}\right) \cdots \cos\left(\frac{\pi(2j_d+1)k_d}{2n_d}\right),$$

where the x_i are real, the z_i are complex, and j_i and k_i are integer variables. We refer the reader to Dudgeon and Mersereau (1983, Sections 2.2 and 4.2) for the first two and Rao and Yip (1990, Chapter 5) for the last. While we have stated them for general d , in practice $d = 2, 3$ are the most useful. The separability of these kernels is exploited in the *row-column decomposition* for their evaluation (Dudgeon and Mersereau 1983, Section 2.3.2). Assuming, for simplicity, that we have $d = 2$, a kernel $K(j, k) = \varphi(j)\psi(k)$ and all integer variables, then

$$F(x, y) = \sum_{j=-\infty}^{\infty} \sum_{k=-\infty}^{\infty} \varphi(j)\psi(k)f(x-j, y-k) = \sum_{j=-\infty}^{\infty} \varphi(j) \left[\sum_{k=-\infty}^{\infty} \psi(k)f(x-j, y-k) \right].$$

We store the sum in the bracket, which we then re-use when evaluating F at other points (x', y) where only the first argument $x' = x + \delta$ is changed:

$$\begin{aligned} F(x + \delta, y) &= \sum_{j=-\infty}^{\infty} \varphi(j) \left[\sum_{k=-\infty}^{\infty} \psi(k)f(x-j+\delta, y-k) \right] \\ &= \sum_{j=-\infty}^{\infty} \varphi(j+\delta) \left[\sum_{k=-\infty}^{\infty} \psi(k)f(x-j, y-k) \right]. \end{aligned}$$

We have assumed that the indices run over all integers to avoid having to deal with boundary complications. In reality, when we have a finite sum as in the discrete cosine transform, evaluating F in the direct manner would have taken $n_1^2 n_2^2 \cdots n_d^2$ additions and multiplications, whereas the row-column decomposition would just require

$$n_1 n_2 \cdots n_d (n_1 + n_2 + \cdots + n_d)$$

additions and multiplications. In cases where there are fast algorithms available

for the one-dimensional transform, say, if we employ one-dimensional FFT in an evaluation of the d -dimensional DFT via row–column decomposition, the number of additions and multiplications could be further reduced to

$$n_1 n_2 \cdots n_d \log_2(n_1 + n_2 + \cdots + n_d) \quad \text{and} \quad \frac{1}{2} n_1 n_2 \cdots n_d \log_2(n_1 + n_2 + \cdots + n_d)$$

respectively. Supposing $d = 2$ and $n_1 = n_2 = 2^{10}$, the approximate number of multiplications required to evaluate a two-dimensional DFT using the direct method, the row–column decomposition method and the row–column decomposition with FFT method are 10^{12} , 2×10^9 , and 10^7 respectively (Dudgeon and Mersereau 1983, Section 2.3.2).

The fact that we may construct a tensor product basis $\mathcal{B}_1 \otimes \cdots \otimes \mathcal{B}_d$ out of given bases $\mathcal{B}_1, \dots, \mathcal{B}_d$ allows us to build multivariate bases out of univariate ones. While this is not necessarily the best way to construct multivariate bases, it is the simplest and often serves as a basic standard that more sophisticated multivariate bases are compared against. This discussion is most interesting with Hilbert spaces, and this will be the backdrop for our next example, where we will see the construction for various notions of bases and beyond.

Example 4.47 (tensor product wavelets and splines). We let \mathbb{H} be any separable Banach space with norm $\|\cdot\|$. Then $\mathcal{B} = \{\varphi_i \in \mathbb{H} : i \in \mathbb{N}\}$ is said to be a *Schauder basis* of \mathbb{H} if, for every $f \in \mathbb{H}$, there is a unique sequence $(a_i)_{i=1}^\infty$ such that

$$f = \sum_{i=1}^{\infty} a_i \varphi_i,$$

where, as usual, this means the series on the right converges to f in $\|\cdot\|$. A Banach space may not have a Schauder basis but a Hilbert space always does. If \mathbb{H} is a Hilbert space with inner product $\langle \cdot, \cdot \rangle$ and $\|\cdot\|$ its induced norm, then its Schauder basis has specific names when it satisfies additional conditions; two of the best-known ones are the orthonormal basis and the *Riesz basis*. Obviously the elements of a Schauder basis must be linearly independent, but this can be unnecessarily restrictive since overcompleteness can be a desirable feature (Mallat 2009), leading us to the notion of *frames*. For easy reference, we define them as follows (Heil 2011):

$$\begin{array}{ll} \text{orthonormal basis} & \|f\|^2 = \sum_{i=1}^{\infty} |\langle f, \varphi_i \rangle|^2, \quad f \in \mathbb{H}, \\ \text{Riesz basis} & \alpha \sum_{i=1}^{\infty} |a_i|^2 \leq \left\| \sum_{i=1}^{\infty} a_i \varphi_i \right\|^2 \leq \beta \sum_{i=1}^{\infty} |a_i|^2, \quad (a_i)_{i=1}^\infty \in l^2(\mathbb{N}), \\ \text{frame} & \alpha \|f\|^2 \leq \sum_{i=1}^{\infty} |\langle f, \varphi_i \rangle|^2 \leq \beta \|f\|^2, \quad f \in \mathbb{H}, \end{array}$$

where the constants $0 < \alpha < \beta$ are called *frame constants* and if $\alpha = \beta$, then the frame is *tight*. Clearly every orthonormal basis is a Riesz basis and every Riesz basis is a frame.

Let $\mathbb{H}_1, \dots, \mathbb{H}_d$ be separable Hilbert spaces and let $\mathcal{B}_1, \dots, \mathcal{B}_d$ be countable dense spanning sets. Let $\mathbb{H}_1 \widehat{\otimes} \dots \widehat{\otimes} \mathbb{H}_d$ be their Hilbert–Schmidt tensor product as discussed in Examples 4.18 and 4.19 and let $\mathcal{B}_1 \otimes \dots \otimes \mathcal{B}_d$ be as defined in (4.43). Then

$$\mathcal{B}_1, \dots, \mathcal{B}_d \text{ are } \begin{cases} \text{orthonormal bases,} \\ \text{Riesz bases,} \\ \text{frames,} \end{cases} \Leftrightarrow \mathcal{B}_1 \otimes \dots \otimes \mathcal{B}_d \text{ is } \begin{cases} \text{an orthonormal basis,} \\ \text{a Riesz basis,} \\ \text{a frame.} \end{cases}$$

The forward implication is straightforward, and if the frame constants of \mathcal{B}_i are α_i and β_i , $i = 1, \dots, d$, then the frame constants of $\mathcal{B}_1 \otimes \dots \otimes \mathcal{B}_d$ are $\prod_{i=1}^d \alpha_i$ and $\prod_{i=1}^d \beta_i$ (Feichtinger and Gröchenig 1994, Lemma 8.18). The converse is, however, more surprising (Bourouhiya 2008).

If $\psi \in L^2(\mathbb{R})$ is a wavelet, that is, $\mathcal{B}_\psi := \{\psi_{m,n} : (m, n) \in \mathbb{Z} \times \mathbb{Z}\}$ with $\psi_{m,n}(x) := 2^{m/2}\psi(2^m x - n)$ is an orthonormal basis, Riesz basis or frame of $L^2(\mathbb{R})$, then

$$\mathcal{B}_\psi \otimes \mathcal{B}_\psi = \{\psi_{m,n} \otimes \psi_{p,q} : (m, n, p, q) \in \mathbb{Z} \times \mathbb{Z} \times \mathbb{Z} \times \mathbb{Z}\} \tag{4.129}$$

is an orthonormal basis, Riesz basis or frame of $L^2(\mathbb{R}^2)$ by our general discussion above. More generally, $\mathcal{B}_\psi^{\otimes d}$ gives a tensor product wavelet basis/frame for $L^2(\mathbb{R}^d)$ (Mallat 2009, Section 7.7.4). While this is certainly the most straightforward way to obtain a wavelet basis for multivariate functions out of univariate wavelets, there are often better options.

Some orthonormal wavelets have a *multiresolution analysis*, that is, a sequence of nested subspaces whose intersection is $\{0\}$ and union is dense in $L^2(\mathbb{R})$,

$$\{0\} \subseteq \dots \subseteq \mathbb{V}_1 \subseteq \mathbb{V}_0 \subseteq \mathbb{V}_{-1} \subseteq \mathbb{V}_{-2} \subseteq \dots \subseteq L^2(\mathbb{R}),$$

defined by a *scaling function* $\varphi \in L^2(\mathbb{R})$ so that $\{\varphi_{m,n} : n \in \mathbb{Z}\}$ is an orthonormal basis of \mathbb{V}_m for any $m \in \mathbb{Z}$. If we write the orthogonal complement of \mathbb{V}_m in \mathbb{V}_{m-1} as \mathbb{W}_m , then $\{\psi_{m,n} : n \in \mathbb{Z}\}$ is an orthonormal basis of \mathbb{W}_m (Meyer 1992, Chapter 2). Multiresolution analysis interacts with tensor products via the rules in Example 4.31:

$$\mathbb{V}_{m-1} \otimes \mathbb{V}_{m-1} = (\mathbb{V}_m \otimes \mathbb{V}_m) \oplus (\mathbb{V}_m \otimes \mathbb{W}_m) \oplus (\mathbb{W}_m \otimes \mathbb{V}_m) \oplus (\mathbb{W}_m \otimes \mathbb{W}_m).$$

Here \oplus denotes orthogonal direct sum. In image processing lingo, the subspaces

$$\begin{aligned} \mathbb{V}_m \otimes \mathbb{V}_m &= \overline{\text{span}}\{\varphi_{m,n} \otimes \varphi_{p,q} : (n, q) \in \mathbb{Z} \times \mathbb{Z}\}, \\ \mathbb{V}_m \otimes \mathbb{W}_m &= \overline{\text{span}}\{\varphi_{m,n} \otimes \psi_{p,q} : (n, q) \in \mathbb{Z} \times \mathbb{Z}\}, \\ \mathbb{W}_m \otimes \mathbb{V}_m &= \overline{\text{span}}\{\psi_{m,n} \otimes \varphi_{p,q} : (n, q) \in \mathbb{Z} \times \mathbb{Z}\}, \\ \mathbb{W}_m \otimes \mathbb{W}_m &= \overline{\text{span}}\{\psi_{m,n} \otimes \psi_{p,q} : (n, q) \in \mathbb{Z} \times \mathbb{Z}\} \end{aligned} \tag{4.130}$$

are called the LL-, LH-, HL- and HH-subbands respectively, with L for low-pass and

H for high-pass. For an image $f \in \mathbb{V}_{m-1} \otimes \mathbb{V}_{m-1}$, its component in the LL-subband $\mathbb{V}_m \otimes \mathbb{V}_m$ represents a low-resolution approximation to f ; the high-resolution details are contained in its components in the LH-, HL- and HH-subbands, whose direct sum is sometimes called the detailed space.

The orthonormal basis obtained via (4.130) can be a better option than simply taking the tensor product (4.129). For instance, in evaluating an $n \times n$ dense matrix–vector product representing an integral transform like the ones in Example 4.45, the *Beylkin et al. (1991)* wavelet-variant of the fast multipole algorithm mentioned in Example 2.14 runs in time $O(n \log n)$ using the basis in (4.129) and in time $O(n)$ using that in (4.130), both impressive compared to the usual $O(n^2)$, but the latter is clearly superior when n is very large.

The main advantage of (4.129) is its generality; with other types of bases or frames we also have similar constructions. For instance, the simplest types of multivariate B-splines (*Höllig and Hörner 2013*) are constructed out of tensor products of univariate B-splines; expressed in a multilinear rank decomposition (4.63), we have

$$B_{l,m,n}(x, y, z) = \sum_{i=1}^p \sum_{j=1}^q \sum_{k=1}^r a_{ijk} B_{i,l}(x) B_{j,m}(y) B_{k,n}(z), \quad (4.131)$$

where $B_{i,l}, B_{j,m}, B_{k,n}$ are univariate B-splines of degrees l, m, n . Nevertheless, the main drawback of straightforward tensor product constructions is that they attach undue importance to the directions of the coordinate axes (*Cohen and Daubechies 1993*). There are often better alternatives such as box splines or beamlets, curvelets, ridgelets, shearlets, wedgelets, *etc.*, that exploit the geometry of \mathbb{R}^2 or \mathbb{R}^3 .

We next discuss a covariant counterpart to the contravariant example above: a univariate quadrature on $[-1, 1]$ is a covariant 1-tensor, and a multivariate quadrature on $[-1, 1]^d$ is a covariant d -tensor.

Example 4.48 (quadrature). Let $f: [-1, 1] \rightarrow \mathbb{R}$ be a univariate polynomial function of degree not more than n . Without knowing anything else about f , we know that there exist n distinct points $x_0, x_1, \dots, x_n \in [-1, 1]$, called *nodes*, and coefficients $w_0, w_1, \dots, w_n \in \mathbb{R}$, called *weights*, so that

$$\int_{-1}^1 f(x) dx = w_0 f(x_0) + w_1 f(x_1) + \dots + w_n f(x_n).$$

In fact, since f is arbitrary, the formula holds with the same nodes and weights for all polynomials of degree n or less. This is called a *quadrature* formula and its existence is simply a consequence of the following observations. A definite integral is a linear functional,

$$I: C([-1, 1]) \rightarrow \mathbb{R}, \quad I(f) = \int_{-1}^1 f(x) dx,$$

as is point evaluation, introduced in Example 4.6,

$$\varepsilon_x : C([-1, 1]) \rightarrow \mathbb{R}, \quad \varepsilon_x(f) = f(x).$$

Since $\mathbb{V} = \{f \in C([-1, 1]): f(x) = c_0 + c_1x + \cdots + c_nx^n\}$ is a vector space of dimension $n + 1$, its dual space \mathbb{V}^* also has dimension $n + 1$, and the $n + 1$ linear functionals $\varepsilon_{x_0}, \varepsilon_{x_1}, \dots, \varepsilon_{x_n}$, obviously linearly independent, form a basis of \mathbb{V}^* . Thus it must be possible to write I as a unique linear combination

$$I = w_0\varepsilon_{x_0} + w_1\varepsilon_{x_1} + \cdots + w_n\varepsilon_{x_n}. \quad (4.132)$$

This observation is presented in Lax (2007, Chapter 2, Theorem 7) as an instance in linear algebra where ‘even trivial-looking material has interesting consequences’.

More generally, a quadrature formula is simply a linear functional in \mathbb{V}^* that is a linear combination of point evaluation functionals, as on the right-hand side of (4.132). If $\mathbb{V} \subseteq C(X)$ is any $(n + 1)$ -dimensional subspace and X is compact Hausdorff with a Borel measure, then the same argument in the previous paragraph shows that there is a quadrature for $I = \int_X$ that gives the exact value of the integral for all $f \in \mathbb{V}$ in terms of the values of f at $n + 1$ nodes $x_0, \dots, x_n \in X$. However, this abstract existential argument does not make it any easier to find $x_0, \dots, x_n \in X$ for a specific type of function space \mathbb{V} and a specific domain X , nor does it tell us about the error if we apply the formula to some $f \notin \mathbb{V}$. In fact, for specific \mathbb{V} and X we can often do better. For instance, the three-node Gauss quadrature on $X = [-1, 1]$ has $x_0 = 0$, $x_1 = \sqrt{3/5} = -x_2$ and $w_0 = 8/9$, $w_1 = 5/9 = w_2$, with

$$\int_{-1}^1 f(x) dx = \frac{8}{9}f(0) + \frac{5}{9}f\left(\sqrt{\frac{3}{5}}\right) + \frac{5}{9}f\left(-\sqrt{\frac{3}{5}}\right)$$

giving exact answers for all *quintic* polynomials f . So $\dim \mathbb{V} = 6$ but we only need three point evaluations to determine I . There is no contradiction here: we are not saying that every linear functional in \mathbb{V}^* can be determined with these three point evaluations, just this one linear functional I . This extra efficacy comes from exploiting the structure of \mathbb{V} as a space of polynomials; while the moment functionals $f \mapsto \int_{-1}^1 x^n f(x) dx$ are linearly independent, they satisfy non-linear relations (Schmüdgen 2017) that can be exploited to obtain various n -node quadratures with exact answers for all polynomials up to degree $2n - 1$ (Golub and Welsch 1969). Generally speaking, the more nodes there are in a quadrature, the more accurate it is but also the more expensive it is. As such we will take the number of nodes as a crude measure of accuracy and computational cost.

As univariate quadratures are well studied in theory (Brass and Petras 2011) and practice (Davis and Rabinowitz 2007), it is natural to build multivariate quadratures from univariate ones. We assume without too much loss of generality that $X = [-1, 1]^d$, noting that any $[a_1, b_1] \times \cdots \times [a_d, b_d]$ can be transformed into $[-1, 1]^d$ with an affine change of variables. Since a univariate quadrature is a linear functional $\varphi : C([-1, 1]) \rightarrow \mathbb{R}$, it seems natural to define a multivariate quadrature as a d -linear

functional

$$\Phi: C([-1, 1]) \times \cdots \times C([-1, 1]) \rightarrow \mathbb{R}, \tag{4.133}$$

but this does not appear to work as we want multivariate quadratures to be defined on $C([-1, 1]^d)$. Fortunately the universal factorization property (4.88) gives us a unique linear map

$$F_\Phi: C([-1, 1]) \otimes \cdots \otimes C([-1, 1]) \rightarrow \mathbb{R}$$

with $\Phi = F_\Phi \circ \sigma_\otimes$, and as $C([-1, 1]^d) = C([-1, 1])^{\otimes d}$ by (4.73), we get a linear functional

$$F_\Phi: C([-1, 1]^d) \rightarrow \mathbb{R}. \tag{4.134}$$

A multivariate quadrature is a linear functional $F: C([-1, 1]^d) \rightarrow \mathbb{R}$ that is a linear combination of point evaluation functionals on $[-1, 1]^d$. Note that given F we may also work backwards to obtain a multilinear functional $\Phi_F = F \circ \sigma_\otimes$ as in (4.133). So (4.133) and (4.134) are equivalent.

The simplest multivariate quadratures are simply separable products of univariate quadratures. Let

$$\varphi_1 = \sum_{i=1}^{n_1} a_i \varepsilon_{x_i}, \quad \varphi_2 = \sum_{j=1}^{n_2} b_j \varepsilon_{y_j}, \dots, \varphi_d = \sum_{k=1}^{n_d} c_k \varepsilon_{z_k}$$

be d univariate quadratures with nodes $N_i \subseteq [-1, 1]$ of size $n_i, i = 1, \dots, d$. Then the covariant d -tensor

$$\varphi_1 \otimes \varphi_2 \otimes \cdots \otimes \varphi_d = \sum_{i=1}^{n_1} \sum_{j=1}^{n_2} \cdots \sum_{k=1}^{n_d} a_i b_j \cdots c_k \varepsilon_{x_i} \otimes \varepsilon_{y_j} \otimes \cdots \otimes \varepsilon_{z_k}$$

is a multivariate quadrature on $[-1, 1]^d$ usually called the tensor product quadrature. The downside of such a quadrature is immediate: for an $f \in C([-1, 1]^d)$, evaluating $\varphi_1 \otimes \cdots \otimes \varphi_d(f)$ requires f to be evaluated on $n_1 n_2 \cdots n_d$ nodes, but for $d = 2, 3$ this is still viable. For instance, the nine-node tensor product Gauss quadrature on $[-1, 1] \times [-1, 1]$ is the covariant 2-tensor

$$\left(\frac{5}{9} \varepsilon_{-\sqrt{3/5}} + \frac{8}{9} \varepsilon_0 + \frac{5}{9} \varepsilon_{\sqrt{3/5}} \right)^{\otimes 2},$$

with nodes

$$\sqrt{\frac{3}{5}} \left\{ \begin{bmatrix} -1 \\ -1 \end{bmatrix}, \begin{bmatrix} 0 \\ -1 \end{bmatrix}, \begin{bmatrix} 1 \\ -1 \end{bmatrix}, \begin{bmatrix} -1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -1 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right\},$$

and corresponding weights

$$\left\{ \frac{25}{81}, \frac{40}{81}, \frac{25}{81}, \frac{40}{81}, \frac{64}{81}, \frac{40}{81}, \frac{25}{81}, \frac{40}{81}, \frac{25}{81} \right\}.$$

One observation is that the weights in the tensor product quadrature are always positive whenever those in its univariate factors $\varphi_1, \dots, \varphi_d$ are. So tensor product quadrature is essentially just a form of weighted average. Taking a leaf from Strassen’s algorithm in Example 3.9, allowing for negative coefficients in a tensor decomposition can be a starting point for superior algorithms; we will examine this in the context of quadrature.

The Smolyak quadrature (Smolyak 1963) is a more sophisticated multivariate quadrature that is also based on tensor products. Our description is adapted from Kaarnioja (2013) and Keshavarzzadeh, Kirby and Narayan (2018). As before, we are given univariate quadratures φ_i with nodes $N_i \subseteq [-1, 1]$ of cardinality n_i , $i = 1, \dots, d$. An important difference is that we now order $\varphi_1, \varphi_2, \dots, \varphi_d$ so that $n_1 \leq n_2 \leq \dots \leq n_d$, that is, these univariate quadratures get increasingly accurate but also increasingly expensive as i increases. A common scenario would have $n_i = 2^{i-1}$. The goal is to define a tensor product quadrature that uses the cheaper quadratures on the left end of the list $\varphi_1, \varphi_2, \dots$ liberally and the expensive ones on the right end $\dots, \varphi_{d-1}, \varphi_d$ sparingly. We set $\varphi_0 := 0$,

$$\Delta_i := \varphi_i - \varphi_{i-1}, \quad i = 1, \dots, d,$$

and observe that

$$\varphi_i = \Delta_1 + \Delta_2 + \dots + \Delta_i, \quad i = 1, \dots, d.$$

The separable multivariate quadrature we defined earlier is

$$\varphi_1 \otimes \varphi_2 \otimes \dots \otimes \varphi_d = \Delta_1 \otimes (\Delta_1 + \Delta_2) \otimes \dots \otimes (\Delta_1 + \Delta_2 + \dots + \Delta_d);$$

it could be made less expensive and accurate by replacing the φ_i with φ_1 ,

$$\varphi_1 \otimes \varphi_1 \otimes \dots \otimes \varphi_1 = \Delta_1 \otimes \Delta_1 \otimes \dots \otimes \Delta_1,$$

or more expensive and accurate by replacing the φ_i with φ_d ,

$$\varphi_d \otimes \varphi_d \otimes \dots \otimes \varphi_d = (\Delta_1 + \Delta_2 + \dots + \Delta_d) \otimes \dots \otimes (\Delta_1 + \Delta_2 + \dots + \Delta_d).$$

The Smolyak quadrature allows one to adjust the level of accuracy and computational cost between these two extremes. For any $r \in \mathbb{N}$, the level- r Smolyak quadrature is the covariant d -tensor

$$\begin{aligned} S_{r,d} &= \sum_{k=0}^{r-1} \sum_{i_1+\dots+i_d=d+k} \Delta_{i_1} \otimes \dots \otimes \Delta_{i_d} \\ &= \sum_{k=r-d}^{r-1} (-1)^{r-1-k} \binom{d-1}{r-1-k} \sum_{i_1+\dots+i_d=d+k} \varphi_{i_1} \otimes \dots \otimes \varphi_{i_d}. \end{aligned}$$

Note that when $r = 1$ and $(d - 1)d$, we get

$$S_{1,d} = \Delta_1^{\otimes d} = \varphi_1^{\otimes d}, \quad S_{(d-1)d,d} = (\Delta_1 + \Delta_2 + \dots + \Delta_d)^{\otimes d} = \varphi_d^{\otimes d}$$

respectively. Take the smallest non-trivial case with $d = r = 2$; we have $\Delta_1 = \varphi_1$,

$\Delta_2 = \varphi_2 - \varphi_1$ and

$$\begin{aligned} S_{2,2} &= \Delta_1 \otimes \Delta_1 + \Delta_1 \otimes \Delta_2 + \Delta_2 \otimes \Delta_1 \\ &= \varphi_1 \otimes \varphi_2 + \varphi_2 \otimes \varphi_1 - \varphi_1 \otimes \varphi_1. \end{aligned}$$

Observe that the weights can now be negative. For the one-point and two-point Gauss quadratures $\varphi_1 = 2\varepsilon_0$ and $\varphi_2 = \varepsilon_{-1/\sqrt{3}} + \varepsilon_{1/\sqrt{3}}$, we have

$$S_{2,2}(f) = 2f\left(0, \frac{-1}{\sqrt{3}}\right) + 2f\left(0, \frac{1}{\sqrt{3}}\right) + 2f\left(\frac{-1}{\sqrt{3}}, 0\right) + 2f\left(\frac{1}{\sqrt{3}}, 0\right) - 4f(0, 0) \quad (4.135)$$

for any $f \in C([-1, 1]^2)$. We will mention just one feature of Smolyak quadrature: if the univariate quadrature φ_i is exact for univariate polynomials of degree $2i - 1$, $i = 1, \dots, d$, then $S_{r,d}$ is exact for d -variate polynomials of degree $2r - 1$ (Heiss and Winschel 2008). Thus (4.135) gives the exact answers when integrating bivariate cubic polynomials f .

For any $f \in C([-1, 1]^d)$, $S_{r,d}(f)$ has to be evaluated at the nodes in

$$\bigcup_{r \leq i_1 + \dots + i_d \leq r+d-1} N_{i_1} \times \dots \times N_{i_d},$$

which is called a *sparse grid* (Gerstner and Griebel 1998). Smolyak quadrature has led to further developments under the heading of sparse grids; we refer interested readers to Garcke (2013) for more information. Nevertheless, many multivariate quadratures tend to be ‘grid-free’ and involve elements of randomness or pseudorandomness, essentially (4.132) with the nodes x_0, \dots, x_n chosen (pseudo)randomly in $[-1, 1]^d$ to have low discrepancy, as defined in Niederreiter (1992). As a result, modern study of multivariate quadratures tends to be quite different from the classical approach described above.

We revisit a point that we made in Example 4.15: it is generally not a good idea to identify tensors in $\mathbb{V}_1 \otimes \dots \otimes \mathbb{V}_d$ with hypermatrices in $\mathbb{R}^{n_1 \times \dots \times n_d}$, as the bases often carry crucial information, that is, in a decomposition such as

$$T = \sum_{i=1}^p \sum_{j=1}^q \dots \sum_{k=1}^r a_{ij\dots k} e_i \otimes f_j \otimes \dots \otimes g_k,$$

the basis elements e_i, f_j, \dots, g_k can be far more important than the coefficients $a_{ij\dots k}$. Quadrature provides a fitting example. The basis elements are the point evaluation functionals at the nodes and these are key to the quadrature: once they are chosen, the coefficients, *i.e.* the weights, can be determined almost as an afterthought. Furthermore, while these basis elements are by definition *linearly* independent, in the context of quadrature there are *non-linear* relations between them that will be lost if one just looks at the coefficients.

All our examples in this article have focused on computations in the classical sense; we will end with a quantum computing example, adapted from Nakahara

and Ohmi (2008, Chapter 7), Nielsen and Chuang (2000, Chapter 6) and Wallach (2008, Section 2.3). While tensors appear at the beginning, ultimately Grover's quantum search algorithm (Grover 1996) reduces to the good old power method.

Example 4.49 (Grover's quantum search). Let \mathbb{C}^2 be equipped with its usual Hermitian inner product $\langle x, y \rangle = x^*y$ and let e_0, e_1 be any pair of orthonormal vectors in \mathbb{C}^2 . Let $d \in \mathbb{N}$ and $n = 2^d$. Suppose we have a function

$$f: \{0, 1, \dots, n-1\} \rightarrow \{-1, +1\}, \quad f(i) = \begin{cases} -1 & \text{if } i = j, \\ +1 & \text{if } i \neq j, \end{cases} \quad (4.136)$$

that is, f takes the value -1 for exactly one $j \in \{0, 1, \dots, n-1\}$ but we do not know which j . Also, f is only accessible as a black box; we may evaluate $f(i)$ for any given i and so to find j in this manner would require $O(n)$ evaluations in the worst case. Grover's algorithm finds j in $O(\sqrt{n})$ complexity with a quantum computer.

We begin by observing that the d -tensor u below may be expanded as

$$u := \left[\frac{1}{\sqrt{2}}(e_0 + e_1) \right]^{\otimes d} = \frac{1}{\sqrt{n}} \sum_{i=0}^{n-1} u_i \in (\mathbb{C}^2)^{\otimes d}, \quad (4.137)$$

with

$$u_i := e_{i_1} \otimes e_{i_2} \otimes \cdots \otimes e_{i_d},$$

where $i_1, \dots, i_d \in \{0, 1\}$ are given by expressing the integer i in binary:

$$[i]_2 = i_d i_{d-1} \cdots i_2 i_1.$$

Furthermore, with respect to the inner product on $(\mathbb{C}^2)^{\otimes d}$ given by (4.64) and its induced norm, u is of unit norm and $\{u_0, \dots, u_{n-1}\}$ is an orthonormal basis. Recall the notion of a Householder reflector in Example 2.12. We define two Householder reflectors H_f and $H_u: (\mathbb{C}^2)^{\otimes d} \rightarrow (\mathbb{C}^2)^{\otimes d}$,

$$H_f(v) = v - 2\langle u_j, v \rangle u_j, \quad H_u(v) = v - 2\langle u, v \rangle u,$$

reflecting v about the hyperplane orthogonal to u_j and u respectively. Like the function f , the Householder reflectors H_f and H_u are only accessible as black boxes. Given any $v \in (\mathbb{C}^2)^{\otimes d}$, a quantum computer allows us to evaluate $H_f(v)$ and $H_u(v)$ in logarithmic complexity $O(\log n) = O(d)$ (Nielsen and Chuang 2000, p. 251).

Grover's algorithm is essentially the power method with $H_u H_f$ applied to u as the initial vector. Note that $H_u H_f$ is unitary and thus norm-preserving and we may skip the normalization step. If we write

$$(H_u H_f)^k u = a_k u_j + b_k \sum_{i \neq j} u_i,$$

then we may show (Nakahara and Ohmi 2008, Propositions 7.2–7.4) that

$$a_0 = b_0 = \frac{1}{\sqrt{n}}, \quad \begin{cases} a_k = \frac{2-n}{n}a_{k-1} + \frac{2(1-n)}{n}b_{k-1}, \\ b_k = \frac{2}{n}a_{k-1} + \frac{2-n}{n}b_{k-1}, \end{cases}$$

for $k \in \mathbb{N}$, with the solution given by

$$a_k = -\sin(2k+1)\theta, \quad b_k = \frac{-1}{\sqrt{n-1}}\cos(2k+1)\theta, \quad \theta := \sin^{-1}(1/\sqrt{n}).$$

The probability of a quantum measurement giving the correct answer j is given by

$$|a_k|^2 = \sin^2(2k+1)\theta,$$

and for $n \gg 1$ this is greater than $1 - 1/n$ when $k \approx \pi\sqrt{n}/4$. This algorithm can be extended to search an f in (4.136) with m copies of -1 s (Nakahara and Ohmi 2008, Section 7.2).

Observe that after setting up the stage with tensors, Grover's algorithm deals only with *linear* operators on tensor spaces. It is the same with Shor's quantum algorithm for integer factorization and discrete log (Shor 1994), which also relies on the tensor u in (4.137) but uses DFT instead of Householder reflectors, although the details (Nakahara and Ohmi 2008, Chapter 8) are too involved to go into here. In either case there is no computation, classical or quantum, involving higher-order tensors.

5. Odds and ends

We conclude our article by recording some parting thoughts. Section 5.1 is about the major topics that had to be omitted. Sections 5.2 and 5.3 may be viewed respectively as explanations as to why our article is not written like a higher-order analogue of Golub and Van Loan (2013) or Horn and Johnson (2013).

5.1. Omissions

There are three major topics in our original plans for this article that we did not manage to cover: (i) symmetric and alternating tensors, (ii) algorithms based on tensor contractions, and (iii) provably correct algorithms for higher-order tensor problems.

In (i) we would have included a discussion of the following: polynomials and differential forms as coordinate representations of symmetric and alternating tensors; polynomial kernels as Veronese embedding of symmetric tensors, Slater determinants as Plücker embedding of alternating tensors; moments and sums-of-squares theory for symmetric tensors; matrix–matrix multiplication as a special case of the Cauchy–Binet formula for alternating tensors. We also planned to discuss the symmetric and alternating analogues of the three definitions of a tensor, the three

constructions of tensor products and the tensor algebra. The last in particular gives us the symmetric and alternating Fock spaces with their connections to bosons and fermions. We would also have discussed various notions of symmetric and alternating tensor ranks.

In (ii) we had intended to demonstrate how the Cooley–Tukey FFT, the multi-dimensional DFT, the Walsh transform, wavelet packet transform, Yates’ method in factorial designs, even FFT on finite non-Abelian groups, *etc.*, are all tensor contractions. We would also show how the Strassen tensor in Example 3.9, the tensors corresponding to the multiplication operations in Lie, Jordan and Clifford algebras, and any tensor network in Example 4.44 may each be realized as the self-contraction of a rank-one tensor. There would also be an explanation of matchgate tensors and holographic algorithms tailored for a numerical linear algebra readership.

In (iii) we had planned to provide a reasonably complete overview of the handful of provably correct algorithms for various NP-hard problems involving higher-order tensors from three different communities: polynomial optimization, symbolic computing/computer algebra and theoretical computer science. We will have more to say about these below.

5.2. Numerical multilinear algebra?

This article is mostly about how tensors arise in computations and not so much about how one might compute tensorial quantities. Could one perhaps develop algorithms analogous to those in numerical linear algebra (Demmel 1997, Golub and Van Loan 2013, Higham 2002, Trefethen and Bau 1997) but for higher-order tensors, or more accurately, for hypermatrices since most models of computations require that we work in terms of coordinates? To be more precise, given a hypermatrix $A \in \mathbb{R}^{n_1 \times \dots \times n_d}$ with $d > 2$, can we compute its ranks, norms, determinants, eigenvalues and eigenvectors, various decompositions and approximations, solutions to multilinear systems of equations, multilinear programming and multilinear regression with coefficients given by A , *etc.*? The answer we think is no. One reason is that most of these problems are NP-hard (Hillar and Lim 2013). Another related reason – but more of an observation – is that it seems difficult to construct *provably correct* algorithms when $d > 2$.

If we sort a list of numbers in Excel or compute the QR factorization of a matrix in MATLAB, we count on the program to give us the correct answer: a sorted list or a QR factorization up to some rounding error. Such is the nature of algorithms: producing the correct answer or an approximation with a controllable error to the problem under consideration is an integral part of its definition (Skiena 2020, Chapter 1). An algorithm needs to have a proof of correctness or, better, a certificate of correctness that the user may employ to check the output. In numerical computing, an example of the former would be the Eckart–Young theorem, which guarantees that SVD produces a best rank- r approximation, or the Lax equivalence theorem, which guarantees convergence of consistent finite difference schemes; an

example of the latter would be relative backward error for a matrix decomposition or duality gap for a convex optimization problem having strong duality.

Note that we are not asking for efficient algorithms: exponential running time or numerical instability are acceptable. The basic requirement is just that for some non-trivial set of inputs, the output in exact arithmetic can be guaranteed to be the answer we seek up to some quantifiable error. While such algorithms for tensor problems do exist in the literature – notably the line of work in Brachat, Comon, Mourrain and Tsingaridas (2010) and Nie (2017) that extended Sylvester (1886) and Reznick (1992) to give a provably correct algorithm for symmetric tensor decomposition – they tend to be the exception rather than the rule. It is more common to find ‘algorithms’ for tensor problems that are demonstrably wrong.

The scarcity of provably correct algorithms goes hand in hand with the NP-hardness of tensor problems. For illustration, we consider the best rank-one approximation problem for a hypermatrix $A \in \mathbb{R}^{n \times n \times n}$:

$$\min_{x, y, z \in \mathbb{R}^n} \|A - x \otimes y \otimes z\|_F. \quad (5.1)$$

This is particularly pertinent as every tensor approximation problem contains (5.1) as a special case; all notions of rank – tensor rank, multilinear rank, tensor network rank, slice rank, *etc.* – agree on the non-zero lowest-ranked tensor and (5.1) represents the lowest order beyond $d = 2$. If we are given a purported solution (x, y, z) , there is no straightforward way to ascertain that it is indeed a solution unless A happens to be rank-one and equal to $x \otimes y \otimes z$. One way to put this is that (5.1) is NP-hard but not NP-complete: given a candidate solution to an NP-complete problem, one can check whether it is indeed a solution in polynomial time; an NP-hard problem does not have this feature. In fact, any graph G can be encoded as a hypermatrix $A_G \in \mathbb{R}^{n \times n \times n}$ in such a way that finding a best rank-one approximation to A_G gives us the chromatic number of the graph (Hillar and Lim 2013), which is NP-hard even to approximate (Khanna, Linial and Safra 2000).

In case it is not clear, applying general-purpose methods such as gradient descent or alternating least-squares to an NP-hard optimization problem does not count as an algorithm. With few exceptions, an algorithm for a problem, say, the Golub–Kahan bidiagonalization for SVD, must exploit the special structure of the problem. Even though one would not expect gradient descent to yield SVD, it is not uncommon to find ‘algorithms’ for tensor problems that do no more than apply off-the-shelf non-linear optimization methods (*e.g.* steepest descent, Newton’s method) or coordinate cycling methods (*e.g.* non-linear Jacobi or Gauss–Seidel).

Furthermore, it is a misconception to think that it is only NP-hard to find global optimizers, and that if one just needs a local optimizer then NP-hardness is no obstacle. In fact, finding local optimizers (Vavasis 1991, Theorem 5.1) or even just stationary points (Khachiyan 1996) may well be NP-hard problems too. It is also a fallacy to assume that since NP-hardness is an asymptotic notion, one should be fine with ‘real-world’ instances of moderate dimensions. As David S. Johnson,

co-author of the standard compendium (Garey and Johnson 1979), puts it (Johnson 2012):

What does it mean for practitioners? I believe that the years have taught us to take the warnings of NP-completeness seriously. If an optimization problem is NP-hard, it is rare that we find algorithms that, even when restricted to ‘real-world’ instances, always seem to find optimal solutions, and do so in empirical polynomial time.

A pertinent example for us is the Strassen tensor $M_{3,3,3}$ in Example 3.9 with $m = n = p = 3$. With respect to the standard basis on $\mathbb{R}^{3 \times 3}$, it is a simple $9 \times 9 \times 9$ hypermatrix with mostly zeros and a few ones, but its tensor rank is still unknown today.

5.3. Hypermatrix analysis?

Readers conversant with matrix analysis (Bellman 1997, Bhatia 1997, Horn and Johnson 2013, 1994) or applied linear algebra (Boyd and Vandenberghe 2018, Dym 2013, Lax 2007, Strang 1980) would certainly have noticed that *abstract linear algebra* – vector spaces, dual spaces, bases, dual bases, change of basis, *etc.* – can be largely or, with some effort, even completely avoided. Essentially, one just needs to know how to work with matrices. In fact, pared down to an absolute minimum, there is just one prerequisite: how to add and multiply matrices.

There is a highbrow partial explanation for this. The category of matrices Mat and the category of vector spaces Vec are equivalent: abstract linear algebra and concrete linear algebra are by and large the same thing (Riehl 2016, Corollary 1.5.11). The analogous statement is false for hypermatrices and tensor spaces except when interpreted in ways that reduce it to variations of the Mat – Vec equivalence as in Example 4.13. A main reason is that it is impossible to multiply hypermatrices in the manner one multiplies matrices, as we have discussed in Example 2.7. The reader may notice that if we remove everything that involves matrix multiplication from Bellman (1997), Bhatia (1997), Horn and Johnson (2013, 1994) or Boyd and Vandenberghe (2018), Dym (2013), Lax (2007), Strang (1980), there would hardly be any content left.

Another big difference in going to order $d \geq 3$ is the lack of canonical forms that we mentioned in Section 2. The 2-tensor transformation rules, *i.e.* equivalence, similarity and congruence, allow us to transform matrices to various canonical forms: rational, Jordan, Weyr, Smith, Turnbull–Aitken, Hodge–Pedoe, *etc.*, or Schur and singular value decompositions with orthogonal/unitary similarity/congruence and equivalence. These play an outsize role in matrix analysis and applied linear algebra. For $d \geq 3$, while there is the Kronecker–Weierstrass form for $m \times n \times 2$ hypermatrices, this is the only exception; there is no canonical form for $m \times n \times p$ hypermatrices, even if we fix $p = 3$, and therefore none for $n_1 \times \cdots \times n_d$ hypermatrices since these include $m \times n \times 3$ as a special case by setting $n_4 = \cdots = n_d = 1$. Note that we did not say ‘no *known*’: the ‘no’ above is in the sense of ‘mathematically proven to be non-existent’.

Tensor rank appears to give us a handle – after all, writing a matrix A as a sum of $\text{rank}(A)$ matrices, each of rank one, is essentially the Smith normal form – but this is an illusion. The $M_{3,3,3}$ example mentioned at the end of Section 5.2 is one among the hundreds of things we do not know about tensor rank. In fact, $4 \times 4 \times 4$ hypermatrices appear to be more or less the boundary of our current knowledge. Two of the biggest breakthroughs in recent times – finding the border rank of $M_{2,2,2} \in \mathbb{C}^{4 \times 4 \times 4}$ (Landsberg 2006) and finding the equations that define border rank-4 tensors $\{A \in \mathbb{C}^{4 \times 4 \times 4} : \text{rank}(A) \leq 4\}$ (Friedland 2013) – were both about $4 \times 4 \times 4$ hypermatrices and both required Herculean efforts, notwithstanding the fact that border rank is a simpler notion than rank.

Acknowledgement

Most of what I learned about tensors in the last ten years I learned from my two former postdocs Yang Qi and Ke Ye; I gratefully acknowledge the education they have given me. Aside from them, I would like to thank Keith Conrad, Zhen Dai, Shmuel Friedland, Edinah Gngang, Shenglong Hu, Risi Kondor, and J. M. Landsberg, Visu Makam, Peter McCullagh, Emily Riehl and Thomas Schultz for answering my questions and/or helpful discussions. I would like to express heartfelt gratitude to Arieh Iserles for his kind encouragement and Glennis Starling for her excellent copy-editing. Figure 4.1 is adapted from Walmes M. Zeviani's gallery of TikZ art and Figure 4.3 is reproduced from Yifan Peng's blog. This work is partially supported by DARPA HR00112190040, NSF DMS-11854831, and the Eckhardt Faculty Fund.

This article was written while taking refuge from Covid-19 at my parents' home in Singapore. After two decades in the US, it is good to be reminded of how fortunate I am to have such wonderful parents. I dedicate this article to them: my father Lim Pok Beng and my mother Ong Aik Kuan.

References

- R. Abraham, J. E. Marsden and T. Ratiu (1988), *Manifolds, Tensor Analysis, and Applications*, Vol. 75 of Applied Mathematical Sciences, second edition, Springer.
- D. Aerts and I. Daubechies (1978), Physical justification for using the tensor product to describe two quantum systems as one joint system, *Helv. Phys. Acta* **51**, 661–675.
- D. Aerts and I. Daubechies (1979a), A characterization of subsystems in physics, *Lett. Math. Phys.* **3**, 11–17.
- D. Aerts and I. Daubechies (1979b), A mathematical condition for a sublattice of a propositional system to represent a physical subsystem, with a physical interpretation, *Lett. Math. Phys.* **3**, 19–27.
- I. Affleck, T. Kennedy, E. H. Lieb and H. Tasaki (1987), Rigorous results on valence-bond ground states in antiferromagnets, *Phys. Rev. Lett.* **59**, 799–802.
- S. S. Akbarov (2003), Pontryagin duality in the theory of topological vector spaces and in topological algebra, *J. Math. Sci.* **113**, 179–349.

- J. Alman and V. V. Williams (2021), A refined laser method and faster matrix multiplication, in *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)* (D. Marx, ed.), Society for Industrial and Applied Mathematics (SIAM), pp. 522–539.
- P. W. Anderson (1959), New approach to the theory of superexchange interactions, *Phys. Rev. (2)* **115**, 2–13.
- A. Arias and J. D. Farmer (1996), On the structure of tensor products of l_p -spaces, *Pacific J. Math.* **175**, 13–37.
- R. Baldick (1995), A unified approach to polynomially solvable cases of integer ‘non-separable’ quadratic optimization, *Discrete Appl. Math.* **61**, 195–212.
- J. Barnes and P. Hut (1986), A hierarchical $O(N \log N)$ force-calculation algorithm, *Nature* **324**, 446–449.
- R. Bellman (1997), *Introduction to Matrix Analysis*, Vol. 19 of Classics in Applied Mathematics, Society for Industrial and Applied Mathematics (SIAM).
- A. Bényi and R. H. Torres (2013), Compact bilinear operators and commutators, *Proc. Amer. Math. Soc.* **141**, 3609–3621.
- S. K. Berberian (2014), *Linear Algebra*, Dover Publications.
- G. Beylkin (1993), Wavelets and fast numerical algorithms, in *Different Perspectives on Wavelets (San Antonio, TX, 1993)*, Vol. 47 of Proceedings of Symposia in Applied Mathematics, American Mathematical Society, pp. 89–117.
- G. Beylkin, R. Coifman and V. Rokhlin (1991), Fast wavelet transforms and numerical algorithms I, *Comm. Pure Appl. Math.* **44**, 141–183.
- G. Beylkin, R. Coifman and V. Rokhlin (1992), Wavelets in numerical analysis, in *Wavelets and their Applications* (M. B. Ruskai et al., eds), Jones & Bartlett, pp. 181–210.
- R. Bhatia (1997), *Matrix Analysis*, Vol. 169 of Graduate Texts in Mathematics, Springer.
- D. Bini (1980), Relations between exact and approximate bilinear algorithms: Applications, *Calcolo* **17**, 87–97.
- D. Bini, M. Capovani, F. Romani and G. Lotti (1979), $O(n^{2.7799})$ complexity for $n \times n$ approximate matrix multiplication, *Inform. Process. Lett.* **8**, 234–235.
- D. Bini, G. Lotti and F. Romani (1980), Approximate solutions for the bilinear form computational problem, *SIAM J. Comput.* **9**, 692–697.
- C. M. Bishop (2006), *Pattern Recognition and Machine Learning*, Information Science and Statistics, Springer.
- A. Blass (1984), Existence of bases implies the axiom of choice, in *Axiomatic Set Theory (Boulder, CO, 1983)*, Vol. 31 of Contemporary Mathematics, American Mathematical Society, pp. 31–33.
- D. Bleeker (1981), *Gauge Theory and Variational Principles*, Vol. 1 of Global Analysis Pure and Applied Series A, Addison-Wesley.
- K. Blum (1996), *Density Matrix Theory and Applications*, Physics of Atoms and Molecules, second edition, Plenum Press.
- J. Board and L. Schulten (2000), The fast multipole algorithm, *Comput. Sci. Eng.* **2**, 76–79.
- A. Bogatskiy, B. Anderson, J. Offermann, M. Roussi, D. Miller and R. Kondor (2020), Lorentz group equivariant neural network for particle physics, in *Proceedings of the 37th International Conference on Machine Learning (ICML 2020)* (H. Daumé III and A. Singh, eds), Vol. 119 of Proceedings of Machine Learning Research, PMLR, pp. 992–1002.
- D. Boneh and A. Silverberg (2003), Applications of multilinear forms to cryptography, in *Topics in Algebraic and Noncommutative Geometry (Luminy/Annapolis, MD, 2001)*, Vol. 324 of Contemporary Mathematics, American Mathematical Society, pp. 71–90.

- W. M. Boothby (1986), *An Introduction to Differentiable Manifolds and Riemannian Geometry*, Vol. 120 of Pure and Applied Mathematics, second edition, Academic Press.
- S. F. Borg (1990), *Matrix–Tensor Methods in Continuum Mechanics*, second edition, World Scientific.
- A. I. Borisenko and I. E. Tarapov (1979), *Vector and Tensor Analysis with Applications*, Dover Publications.
- N. Bourbaki (1998), *Algebra I, Chapters 1–3*, Elements of Mathematics (Berlin), Springer.
- A. Bourouhiya (2008), The tensor product of frames, *Sampl. Theory Signal Image Process.* **7**, 65–76.
- S. Boyd and L. Vandenberghe (2004), *Convex Optimization*, Cambridge University Press.
- S. Boyd and L. Vandenberghe (2018), *Introduction to Applied Linear Algebra*, Cambridge University Press.
- R. N. Bracewell (1986), *The Fourier Transform and its Applications*, McGraw-Hill Series in Electrical Engineering, third edition, McGraw-Hill.
- J. Brachat, P. Comon, B. Mourrain and E. Tsigaridas (2010), Symmetric tensor decomposition, *Linear Algebra Appl.* **433**, 1851–1872.
- L. Brand (1947), *Vector and Tensor Analysis*, Wiley.
- H. Brass and K. Petras (2011), *Quadrature Theory: The Theory of Numerical Integration on a Compact Interval*, Vol. 178 of Mathematical Surveys and Monographs, American Mathematical Society.
- R. P. Brent (1970), Algorithms for matrix multiplication. Report Stan-CS-70-157, Stanford University.
- R. P. Brent and P. Zimmermann (2011), *Modern Computer Arithmetic*, Vol. 18 of Cambridge Monographs on Applied and Computational Mathematics, Cambridge University Press.
- T. Bröcker and T. tom Dieck (1995), *Representations of Compact Lie Groups*, Vol. 98 of Graduate Texts in Mathematics, Springer.
- R. A. Brualdi (1992), The symbiotic relationship of combinatorics and matrix theory, *Linear Algebra Appl.* **162/164**, 65–105.
- M. D. Buhmann (2003), *Radial Basis Functions: Theory and Implementations*, Vol. 12 of Cambridge Monographs on Applied and Computational Mathematics, Cambridge University Press.
- P. Bürgisser and F. Cucker (2013), *Condition: The Geometry of Numerical Algorithms*, Vol. 349 of Grundlehren der Mathematischen Wissenschaften, Springer.
- P. Bürgisser, M. Clausen and M. A. Shokrollahi (1997), *Algebraic Complexity Theory*, Vol. 315 of Grundlehren der Mathematischen Wissenschaften, Springer.
- E. Callaway (2020), ‘It will change everything’: Deepmind’s AI makes gigantic leap in solving protein structures, *Nature* **588** (7837), 203–204.
- A. Cayley (1845), On the theory of linear transformations, *Cambridge Math. J.* **4**, 193–209.
- S. S. Chern, W. H. Chen and K. S. Lam (1999), *Lectures on Differential Geometry*, Vol. 1 of Series on University Mathematics, World Scientific.
- Y. Choquet-Bruhat, C. DeWitt-Morette and M. Dillard-Bleick (1982), *Analysis, Manifolds and Physics*, second edition, North-Holland.
- P. C. Chou and N. J. Pagano (1992), *Elasticity: Tensor, Dyadic and Engineering Approaches*, Dover Publications.
- F. Cobos, T. Kühn and J. Peetre (1992), Schatten–von Neumann classes of multilinear forms, *Duke Math. J.* **65**, 121–156.

- F. Cobos, T. Kühn and J. Peetre (1999), On \mathfrak{G}_p -classes of trilinear forms, *J. London Math. Soc.* (2) **59**, 1003–1022.
- A. Cohen and I. Daubechies (1993), Nonseparable bidimensional wavelet bases, *Rev. Mat. Iberoamericana* **9**, 51–137.
- M. B. Cohen, Y. T. Lee and Z. Song (2019), Solving linear programs in the current matrix multiplication time, in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing (STOC 2019)*, ACM, pp. 938–942.
- T. Cohen and M. Welling (2016), Group equivariant convolutional networks, in *Proceedings of the 33rd International Conference on Machine Learning (ICML 2016)* (M. F. Balcan and K. Q. Weinberger, eds), Vol. 48 of Proceedings of Machine Learning Research, PMLR, pp. 2990–2999.
- C. Cohen-Tannoudji, B. Diu and F. Laloë (2020a), *Quantum Mechanics 1: Basic Concepts, Tools, and Applications*, second edition, Wiley-VCH.
- C. Cohen-Tannoudji, B. Diu and F. Laloë (2020b), *Quantum Mechanics 2: Angular Momentum, Spin, and Approximation Methods*, second edition, Wiley-VCH.
- K. Conrad (2018), Tensor products. Available at <https://kconrad.math.uconn.edu/blurbs/linmultialg/tensorprod.pdf>.
- J. B. Conway (1990), *A Course in Functional Analysis*, Vol. 96 of Graduate Texts in Mathematics, second edition, Springer.
- S. A. Cook and S. O. Aanderaa (1969), On the minimum computation time of functions, *Trans. Amer. Math. Soc.* **142**, 291–314.
- F. Cucker and S. Smale (2002), On the mathematical foundations of learning, *Bull. Amer. Math. Soc. (N.S.)* **39**, 1–49.
- E. R. Davidson (1976), *Reduced Density Matrices in Quantum Chemistry*, Vol. 6 of Theoretical Chemistry, Academic Press.
- A. M. Davie (1985), Matrix norms related to Grothendieck’s inequality, in *Banach Spaces (Columbia, MO, 1984)*, Vol. 1166 of Lecture Notes in Mathematics, Springer, pp. 22–26.
- P. J. Davis and P. Rabinowitz (2007), *Methods of Numerical Integration*, Dover Publications.
- C. De Concini and C. Procesi (2017), *The Invariant Theory of Matrices*, Vol. 69 of University Lecture Series, American Mathematical Society.
- R. de la Madrid (2005), The role of the rigged Hilbert space in quantum mechanics, *European J. Phys.* **26**, 287–312.
- L. De Lathauwer, B. De Moor and J. Vandewalle (2000), A multilinear singular value decomposition, *SIAM J. Matrix Anal. Appl.* **21**, 1253–1278.
- V. De Silva and L.-H. Lim (2008), Tensor rank and the ill-posedness of the best low-rank approximation problem, *SIAM J. Matrix Anal. Appl.* **30**, 1084–1127.
- F. De Terán (2016), Canonical forms for congruence of matrices and T -palindromic matrix pencils: A tribute to H. W. Turnbull and A. C. Aitken, *SeMA J.* **73**, 7–16.
- J. Debnath and R. S. Dahiya (1989), Theorems on multidimensional Laplace transform for solution of boundary value problems, *Comput. Math. Appl.* **18**, 1033–1056.
- A. Defant and K. Floret (1993), *Tensor Norms and Operator Ideals*, Vol. 176 of North-Holland Mathematics Studies, North-Holland.
- E. D. Demaine, M. L. Demaine, A. Edelman, C. E. Leiserson and P.-O. Persson (2005), Building blocks and excluded sums, *SIAM News* **38**, 1, 4, 6.
- J. W. Demmel (1997), *Applied Numerical Linear Algebra*, Society for Industrial and Applied Mathematics (SIAM).

- L. Deng (2012), The MNIST database of handwritten digit images for machine learning research, *IEEE Signal Process. Mag.* **29**, 141–142.
- P. J. C. Dickinson and L. Gijben (2014), On the computational complexity of membership problems for the completely positive cone and its dual, *Comput. Optim. Appl.* **57**, 403–415.
- J. Diestel, J. H. Fourie and J. Swart (2008), *The Metric Theory of Tensor Products: Grothendieck's Résumé Revisited*, American Mathematical Society.
- C. T. J. Dodson and T. Poston (1991), *Tensor Geometry: The Geometric Viewpoint and its Uses*, Vol. 130 of Graduate Texts in Mathematics, second edition, Springer.
- J. Dongarra and F. Sullivan (2000), Guest editors' introduction to the top 10 algorithms, *Comput. Sci. Eng.* **2**, 22–23.
- D. E. Dudgeon and R. M. Mersereau (1983), *Multidimensional Digital Signal Processing*, Prentice Hall.
- D. S. Dummit and R. M. Foote (2004), *Abstract Algebra*, third edition, Wiley.
- N. Dunford and J. T. Schwartz (1988), *Linear Operators, Part I*, Wiley Classics Library, Wiley.
- G. Dunn (1988), Tensor product of operads and iterated loop spaces, *J. Pure Appl. Algebra* **50**, 237–258.
- H. Dym (2013), *Linear Algebra in Action*, Vol. 78 of Graduate Studies in Mathematics, second edition, American Mathematical Society.
- J. Earman and C. Glymour (1978), Lost in the tensors: Einstein's struggles with covariance principles 1912–1916, *Stud. Hist. Philos. Sci. A* **9**, 251–278.
- A. Einstein (2002), Fundamental ideas and methods of the theory of relativity, presented in its development, in *The Collected Papers of Albert Einstein* (M. Janssen *et al.*, eds), Vol. 7: *The Berlin Years, 1918–1921*, Princeton University Press, pp. 113–150.
- L. P. Eisenhart (1934), Separable systems of Stackel, *Ann. of Math. (2)* **35**, 284–305.
- P. Etingof, S. Gelaki, D. Nikshych and V. Ostrik (2015), *Tensor Categories*, Vol. 205 of Mathematical Surveys and Monographs, American Mathematical Society.
- L. D. Faddeev and O. A. Yakubovskii (2009), *Lectures on Quantum Mechanics for Mathematics Students*, Vol. 47 of Student Mathematical Library, American Mathematical Society.
- C. L. Fefferman (2006), Existence and smoothness of the Navier–Stokes equation, in *The Millennium Prize Problems* (J. Carlson, A. Jaffe and A. Wiles, eds), Clay Mathematics Institute, pp. 57–67.
- H. G. Feichtinger and K. Gröchenig (1994), Theory and practice of irregular sampling, in *Wavelets: Mathematics and Applications*, CRC, pp. 305–363.
- R. P. Feynman, R. B. Leighton and M. Sands (1963), *The Feynman Lectures on Physics*, Vol. 1: *Mainly Mechanics, Radiation, and Heat*, Addison-Wesley.
- C. F. Fischer (1977), *The Hartree–Fock Method for Atoms*, Wiley.
- V. Fock (1930), Näherungsmethode zur Lösung des quantenmechanischen Mehrkörperproblems, *Z. Physik* **61**, 126–148.
- L. Fortnow (2013), *The Golden Ticket: P, NP, and the Search for the Impossible*, Princeton University Press.
- A. Frank and E. Tardos (1987), An application of simultaneous Diophantine approximation in combinatorial optimization, *Combinatorica* **7**, 49–65.
- M. Frazier and B. Jawerth (1990), A discrete transform and decompositions of distribution spaces, *J. Funct. Anal.* **93**, 34–170.

- S. H. Friedberg, A. J. Insel and L. E. Spence (2003), *Linear Algebra*, fourth edition, Prentice Hall.
- S. Friedland (2013), On tensors of border rank l in $\mathbb{C}^{m \times n \times l}$, *Linear Algebra Appl.* **438**, 713–737.
- S. Friedland and E. Gross (2012), A proof of the set-theoretic version of the salmon conjecture, *J. Algebra* **356**, 374–379.
- S. Friedland and L.-H. Lim (2018), Nuclear norm of higher-order tensors, *Math. Comp.* **87**, 1255–1281.
- S. Friedland, L.-H. Lim and J. Zhang (2019), Grothendieck constant is norm of Strassen matrix multiplication tensor, *Numer. Math.* **143**, 905–922.
- J. Friedman (1991), The spectra of infinite hypertrees, *SIAM J. Comput.* **20**, 951–961.
- J. Friedman and A. Wigderson (1995), On the second eigenvalue of hypergraphs, *Combinatorica* **15**, 43–65.
- F. Fuchs, D. E. Worrall, V. Fischer and M. Welling (2020), SE(3)-transformers: 3D rotation-equivariant attention networks, in *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)* (H. Larochelle *et al.*, eds), Curran Associates, pp. 1970–1981.
- W. Fulton and J. Harris (1991), *Representation Theory: A First Course*, Vol. 129 of Graduate Texts in Mathematics, Springer.
- M. Fürer (2009), Faster integer multiplication, *SIAM J. Comput.* **39**, 979–1005.
- A. A. García, F. W. Hehl, C. Heinicke and A. Macías (2004), The Cotton tensor in Riemannian spacetimes, *Classical Quantum Gravity* **21**, 1099–1118.
- J. Garcke (2013), Sparse grids in a nutshell, in *Sparse Grids and Applications*, Vol. 88 of Lecture Notes in Computational Science and Engineering, Springer, pp. 57–80.
- M. R. Garey and D. S. Johnson (1979), *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W. H. Freeman.
- S. Garg, C. Gentry and S. Halevi (2013), Candidate multilinear maps from ideal lattices, in *Advances in Cryptology (EUROCRYPT 2013)*, Vol. 7881 of Lecture Notes in Computer Science, Springer, pp. 1–17.
- P. Garrett (2010), Non-existence of tensor products of Hilbert spaces. Available at http://www-users.math.umn.edu/~garrett/m/v/nonexistence_tensors.pdf.
- I. M. Gel'fand, M. M. Kapranov and A. V. Zelevinsky (1992), Hyperdeterminants, *Adv. Math.* **96**, 226–263.
- I. M. Gel'fand, M. M. Kapranov and A. V. Zelevinsky (1994), *Discriminants, Resultants, and Multidimensional Determinants*, Mathematics: Theory & Applications, Birkhäuser.
- C. Gentry, S. Gorbunov and S. Halevi (2015), Graph-induced multilinear maps from lattices, in *Theory of Cryptography (TCC 2015)*, part II, Vol. 9015 of Lecture Notes in Computer Science, Springer, pp. 498–527.
- R. Geroch (1985), *Mathematical Physics*, Chicago Lectures in Physics, University of Chicago Press.
- T. Gerstner and M. Griebel (1998), Numerical integration using sparse grids, *Numer. Algorithms* **18**, 209–232.
- G. H. Golub and C. F. Van Loan (2013), *Matrix Computations*, Johns Hopkins Studies in the Mathematical Sciences, fourth edition, Johns Hopkins University Press.
- G. H. Golub and J. H. Welsch (1969), Calculation of Gauss quadrature rules, *Math. Comp.* **23**, 221–230, A1–A10.

- G. H. Golub and J. H. Wilkinson (1976), Ill-conditioned eigensystems and the computation of the Jordan canonical form, *SIAM Rev.* **18**, 578–619.
- R. Goodman and N. R. Wallach (2009), *Symmetry, Representations, and Invariants*, Vol. 255 of Graduate Texts in Mathematics, Springer.
- L. Grafakos (2014), *Classical Fourier Analysis*, Vol. 249 of Graduate Texts in Mathematics, third edition, Springer.
- L. Grafakos and R. H. Torres (2002a), Discrete decompositions for bilinear operators and almost diagonal conditions, *Trans. Amer. Math. Soc.* **354**, 1153–1176.
- L. Grafakos and R. H. Torres (2002b), Multilinear Calderón–Zygmund theory, *Adv. Math.* **165**, 124–164.
- L. Greengard and V. Rokhlin (1987), A fast algorithm for particle simulations, *J. Comput. Phys.* **73**, 325–348.
- W. Greub (1978), *Multilinear Algebra*, Universitext, second edition, Springer.
- D. Griffiths (2008), *Introduction to Elementary Particles*, second edition, Wiley-VCH.
- A. Grothendieck (1953), Résumé de la théorie métrique des produits tensoriels topologiques, *Bol. Soc. Mat. São Paulo* **8**, 1–79.
- A. Grothendieck (1955), *Produits Tensoriels Topologiques et Espaces Nucléaires*, Vol. 16 of Memoirs of the American Mathematical Society, American Mathematical Society.
- L. K. Grover (1996), A fast quantum mechanical algorithm for database search, in *Proceedings of the 28th Annual ACM Symposium on the Theory of Computing (STOC 1996)*, ACM, pp. 212–219.
- K. Hannabuss (1997), *An Introduction to Quantum Theory*, Vol. 1 of Oxford Graduate Texts in Mathematics, The Clarendon Press, Oxford University Press.
- J. Harris (1995), *Algebraic Geometry: A First Course*, Vol. 133 of Graduate Texts in Mathematics, Springer.
- E. Hartmann (1984), *An Introduction to Crystal Physics*, Vol. 18 of Commission on Crystallographic Teaching: Second series pamphlets, International Union of Crystallography, University College Cardiff Press.
- D. R. Hartree (1928), The wave mechanics of an atom with a non-Coulomb central field, I: Theory and methods, *Proc. Cambridge Philos. Soc.* **24**, 89–132.
- R. Hartshorne (1977), *Algebraic Geometry*, Vol. 52 of Graduate Texts in Mathematics, Springer.
- D. Harvey and J. van der Hoeven (2021), Integer multiplication in time $O(n \log n)$, *Ann. of Math. (2)* **193**, 563–617.
- S. Hassani (1999), *Mathematical Physics: A Modern Introduction to its Foundations*, Springer.
- T. J. Hastie and R. J. Tibshirani (1990), *Generalized Additive Models*, Vol. 43 of Monographs on Statistics and Applied Probability, Chapman & Hall.
- A. Hatcher (2002), *Algebraic Topology*, Cambridge University Press.
- R. A. Hauser and Y. Lim (2002), Self-scaled barriers for irreducible symmetric cones, *SIAM J. Optim.* **12**, 715–723.
- G. E. Hay (1954), *Vector and Tensor Analysis*, Dover Publications.
- C. Heil (2011), *A Basis Theory Primer*, Applied and Numerical Harmonic Analysis, expanded edition, Birkhäuser / Springer.
- F. Heiss and V. Winschel (2008), Likelihood approximation by numerical integration on sparse grids, *J. Econometrics* **144**, 62–80.

- S. Helgason (1978), *Differential Geometry, Lie Groups, and Symmetric Spaces*, Vol. 80 of Pure and Applied Mathematics, Academic Press.
- J. W. Helton and M. Putinar (2007), Positive polynomials in scalar and matrix variables, the spectral theorem, and optimization, in *Operator Theory, Structured Matrices, and Dilations*, Vol. 7 of Theta Ser. Adv. Math., Theta, Bucharest, pp. 229–306.
- N. J. Higham (1992), Stability of a method for multiplying complex matrices with three real matrix multiplications, *SIAM J. Matrix Anal. Appl.* **13**, 681–687.
- N. J. Higham (2002), *Accuracy and Stability of Numerical Algorithms*, second edition, Society for Industrial and Applied Mathematics (SIAM).
- C. J. Hillar and L.-H. Lim (2013), Most tensor problems are NP-hard, *J. Assoc. Comput. Mach.* **60**, 45.
- F. L. Hitchcock (1927), The expression of a tensor or a polyadic as a sum of products, *J. Math. Phys. Mass. Inst. Tech.* **6**, 164–189.
- D. S. Hochbaum and J. G. Shanthikumar (1990), Convex separable optimization is not much harder than linear optimization, *J. Assoc. Comput. Mach.* **37**, 843–862.
- T. Hofmann, B. Schölkopf and A. J. Smola (2008), Kernel methods in machine learning, *Ann. Statist.* **36**, 1171–1220.
- K. Höllig and J. Hörner (2013), *Approximation and Modeling with B-Splines*, Society for Industrial and Applied Mathematics (SIAM).
- J. R. Holub (1970), Tensor product bases and tensor diagonals, *Trans. Amer. Math. Soc.* **151**, 563–579.
- R. A. Horn and C. R. Johnson (1994), *Topics in Matrix Analysis*, Cambridge University Press.
- R. A. Horn and C. R. Johnson (2013), *Matrix Analysis*, second edition, Cambridge University Press.
- F. Irgens (2019), *Tensor Analysis*, Springer.
- J. D. Jackson (1999), *Classical Electrodynamics*, third edition, Wiley.
- N. Jacobson (1975), *Lectures in Abstract Algebra*, Vol II: *Linear Algebra*, Vol. 31 of Graduate Texts in Mathematics, Springer.
- J. C. Jaeger (1940), The solution of boundary value problems by a double Laplace transformation, *Bull. Amer. Math. Soc.* **46**, 687–693.
- D. S. Johnson (2012), A brief history of NP-completeness, 1954–2012, *Doc. Math.* pp. 359–376.
- T. F. Jordan (2005), *Quantum Mechanics in Simple Matrix Form*, Dover Publications.
- A. Joux (2004), A one round protocol for tripartite Diffie–Hellman, *J. Cryptology* **17**, 263–276.
- V. Kaarnioja (2013), Smolyak quadrature. Master’s thesis, University of Helsinki, Finland.
- E. G. Kalnins, J. M. Kress and W. Miller, Jr. (2018), *Separation of Variables and Superintegrability: The Symmetry of Solvable Systems*, IOP Expanding Physics, IOP Publishing.
- R. P. Kanwal (1997), *Linear Integral Equations*, second edition, Birkhäuser.
- A. A. Karatsuba and Y. Ofman (1962), Multiplication of many-digit numbers by automatic computers, *Dokl. Akad. Nauk SSSR* **145**, 293–294.
- N. Karmarkar (1984), A new polynomial-time algorithm for linear programming, *Combinatorica* **4**, 373–395.
- E. A. Kearsley and J. T. Fong (1975), Linearly independent sets of isotropic Cartesian tensors of ranks up to eight, *J. Res. Nat. Bur. Standards B* **79**, 49–58.

- V. Keshavarzzadeh, R. M. Kirby and A. Narayan (2018), Numerical integration in multiple dimensions with designed quadrature, *SIAM J. Sci. Comput.* **40**, A2033–A2061.
- L. Khachiyan (1979), A polynomial algorithm in linear programming, *Soviet Math. Dokl.* **20**, 191–194.
- L. Khachiyan (1996), Diagonal matrix scaling is NP-hard, *Linear Algebra Appl.* **234**, 173–179.
- S. Khanna, N. Linial and S. Safra (2000), On the hardness of approximating the chromatic number, *Combinatorica* **20**, 393–415.
- T. Kleinjung and B. Wesolowski (2021), Discrete logarithms in quasi-polynomial time in finite fields of fixed characteristic, *J. Amer. Math. Soc.* doi:10.1090/jams/985.
- V. V. Kljuev and N. I. Kokovkin-Ščerbak (1965), On the minimization of the number of arithmetic operations for solving linear algebraic systems of equations, *Ž. Vyčisl. Mat i Mat. Fiz.* **5**, 21–33.
- A. A. Klyachko (1998), Stable bundles, representation theory and Hermitian operators, *Selecta Math. (N.S.)* **4**, 419–445.
- D. E. Knuth (1998), *The Art of Computer Programming*, Vol. 2: *Seminumerical Algorithms*, third edition, Addison-Wesley.
- A. Knutson and T. Tao (1999), The honeycomb model of $GL_n(\mathbb{C})$ tensor products I: Proof of the saturation conjecture, *J. Amer. Math. Soc.* **12**, 1055–1090.
- R. Kondor and S. Trivedi (2018), On the generalization of equivariance and convolution in neural networks to the action of compact groups, in *Proceedings of the 35th International Conference on Machine Learning (ICML 2018)* (J. Dy and A. Krause, eds), Vol. 80 of *Proceedings of Machine Learning Research*, PMLR, pp. 2747–2755.
- R. Kondor, Z. Lin and S. Trivedi (2018), Clebsch–Gordan nets: A fully Fourier space spherical convolutional neural network, in *Advances in Neural Information Processing Systems 31 (NeurIPS 2018)* (S. Bengio *et al.*, eds), Curran Associates, pp. 10138–10147.
- T. H. Koornwinder (1980), A precise definition of separation of variables, in *Geometrical Approaches to Differential Equations (Proc. Fourth Scheveningen Conference, 1979)*, Vol. 810 of *Lecture Notes in Mathematics*, Springer, pp. 240–263.
- M. Kostoglou (2005), On the analytical separation of variables solution for a class of partial integro-differential equations, *Appl. Math. Lett.* **18**, 707–712.
- A. I. Kostrikin and Y. I. Manin (1997), *Linear Algebra and Geometry*, Vol. 1 of *Algebra, Logic and Applications*, Gordon & Breach Science Publishers.
- S. G. Krantz (2009), *Explorations in Harmonic Analysis: With Applications to Complex Function Theory and the Heisenberg Group*, Applied and Numerical Harmonic Analysis, Birkhäuser.
- S. Krishna and V. Makam (2018), On the tensor rank of 3×3 permanent and determinant. Available at [arXiv:1801.00496](https://arxiv.org/abs/1801.00496).
- J.-L. Krivine (1979), Constantes de Grothendieck et fonctions de type positif sur les sphères, *Adv. Math.* **31**, 16–30.
- A. Krizhevsky, I. Sutskever and G. E. Hinton (2012), Imagenet classification with deep convolutional neural networks, in *Advances in Neural Information Processing Systems 25 (NIPS 2012)* (F. Pereira *et al.*, eds), Curran Associates, pp. 1097–1105.
- M. Lacey and C. Thiele (1997), L^p estimates on the bilinear Hilbert transform for $2 < p < \infty$, *Ann. of Math. (2)* **146**, 693–724.
- M. Lacey and C. Thiele (1999), On Calderón’s conjecture, *Ann. of Math. (2)* **149**, 475–496.

- J. M. Landsberg (2006), The border rank of the multiplication of 2×2 matrices is seven, *J. Amer. Math. Soc.* **19**, 447–459.
- J. M. Landsberg (2012), *Tensors: Geometry and Applications*, Vol. 128 of Graduate Studies in Mathematics, American Mathematical Society.
- J. M. Landsberg (2017), *Geometry and Complexity Theory*, Vol. 169 of Cambridge Studies in Advanced Mathematics, Cambridge University Press.
- J. M. Landsberg (2019), *Tensors: Asymptotic Geometry and Developments 2016–2018*, Vol. 132 of CBMS Regional Conference Series in Mathematics, American Mathematical Society.
- J. M. Landsberg, Y. Qi and K. Ye (2012), On the geometry of tensor network states, *Quantum Inform. Comput.* **12**, 346–354.
- K. Landsman (2017), *Foundations of Quantum Theory: From Classical Concepts to Operator Algebras*, Vol. 188 of Fundamental Theories of Physics, Springer.
- S. Lang (1993), *Real and Functional Analysis*, Vol. 142 of Graduate Texts in Mathematics, third edition, Springer.
- S. Lang (2002), *Algebra*, Vol. 211 of Graduate Texts in Mathematics, third edition, Springer.
- P. D. Lax (2007), *Linear Algebra and its Applications*, Pure and Applied Mathematics (Hoboken), second edition, Wiley-Interscience.
- J. M. Lee (2013), *Introduction to Smooth Manifolds*, Vol. 218 of Graduate Texts in Mathematics, second edition, Springer.
- W. A. Light and E. W. Cheney (1985), *Approximation Theory in Tensor Product Spaces*, Vol. 1169 of Lecture Notes in Mathematics, Springer.
- D. Lovelock and H. Rund (1975), *Tensor, Differential Forms, and Variational Principles*, Wiley-Interscience.
- S. Lovett (2019), The analytic rank of tensors and its applications, *Discrete Anal.* **2019**, 7.
- G. W. Mackey (1978), *Unitary Group Representations in Physics, Probability, and Number Theory*, Vol. 55 of Mathematics Lecture Note Series, Benjamin / Cummings.
- S. Mallat (2009), *A Wavelet Tour of Signal Processing*, third edition, Elsevier / Academic Press.
- D. Martin (1991), *Manifold Theory: An Introduction for Mathematical Physicists*, Ellis Horwood Series in Mathematics and its Applications, Ellis Horwood.
- A. J. McConnell (1957), *Application of Tensor Analysis*, Dover Publications.
- P. McCullagh (1987), *Tensor Methods in Statistics*, Monographs on Statistics and Applied Probability, Chapman & Hall.
- W. S. McCulloch and W. Pitts (1943), A logical calculus of the ideas immanent in nervous activity, *Bull. Math. Biophys.* **5**, 115–133.
- J. Mercer (1909), Functions of positive and negative type and their connection with the theory of integral equations, *Philos. Trans. Roy. Soc. London Ser. A* **209**, 415–446.
- Y. Meyer (1992), *Wavelets and Operators*, Vol. 37 of Cambridge Studies in Advanced Mathematics, Cambridge University Press.
- Y. Meyer and R. Coifman (1997), *Wavelets: Calderón–Zygmund and Multilinear Operators*, Vol. 48 of Cambridge Studies in Advanced Mathematics, Cambridge University Press.
- A. D. Michal (1947), *Matrix and Tensor Calculus with Applications to Mechanics, Elasticity, and Aeronautics*, Wiley / Chapman & Hall.
- W. Miller, Jr (1977), *Symmetry and Separation of Variables*, Vol. 4 of Encyclopedia of Mathematics and its Applications, Addison-Wesley.

- J. W. Milnor and J. D. Stasheff (1974), *Characteristic Classes*, Vol. 76 of Annals of Mathematics Studies, Princeton University Press / University of Tokyo Press.
- C. W. Misner, K. S. Thorne and J. A. Wheeler (1973), *Gravitation*, W. H. Freeman.
- P. Moon and D. E. Spencer (1961), *Field Theory for Engineers*, The Van Nostrand Series in Electronics and Communications, Van Nostrand.
- P. Moon and D. E. Spencer (1988), *Field Theory Handbook: Including Coordinate Systems, Differential Equations and Their Solutions*, second edition, Springer.
- P. M. Morse and H. Feshbach (1953), *Methods of Theoretical Physics*, McGraw-Hill.
- K. G. Murty and S. N. Kabadi (1987), Some NP-complete problems in quadratic and nonlinear programming, *Math. Program.* **39**, 117–129.
- C. Muscalu and W. Schlag (2013), *Classical and Multilinear Harmonic Analysis II*, Vol. 138 of Cambridge Studies in Advanced Mathematics, Cambridge University Press.
- M. Nakahara and T. Ohmi (2008), *Quantum Computing: From Linear Algebra to Physical Realizations*, CRC Press.
- H. Narayanan (2006), On the complexity of computing Kostka numbers and Littlewood–Richardson coefficients, *J. Algebraic Combin.* **24**, 347–354.
- Y. Nesterov and A. Nemirovskii (1994), *Interior-Point Polynomial Algorithms in Convex Programming*, Vol. 13 of SIAM Studies in Applied Mathematics, Society for Industrial and Applied Mathematics (SIAM).
- J. Nie (2017), Generating polynomials and symmetric tensor decompositions, *Found. Comput. Math.* **17**, 423–465.
- H. Niederreiter (1992), *Random Number Generation and Quasi-Monte Carlo Methods*, Vol. 63 of CBMS-NSF Regional Conference Series in Applied Mathematics, Society for Industrial and Applied Mathematics (SIAM).
- M. A. Nielsen and I. L. Chuang (2000), *Quantum Computation and Quantum Information*, Cambridge University Press.
- P. Niyogi, S. Smale and S. Weinberger (2008), Finding the homology of submanifolds with high confidence from random samples, *Discrete Comput. Geom.* **39**, 419–441.
- K. C. O’Meara, J. Clark and C. I. Vinsonhaler (2011), *Advanced Topics in Linear Algebra: Weaving Matrix Problems Through the Weyr Form*, Oxford University Press.
- R. Orús (2014), A practical introduction to tensor networks: Matrix product states and projected entangled pair states, *Ann. Phys.* **349**, 117–158.
- B. G. Osgood (2019), *Lectures on the Fourier Transform and its Applications*, Vol. 33 of Pure and Applied Undergraduate Texts, American Mathematical Society.
- B. N. Parlett (2000), The QR algorithm, *Comput. Sci. Eng.* **2**, 38–42.
- D. S. Passman (1991), *A Course in Ring Theory*, The Wadsworth & Brooks/Cole Mathematics Series, Wadsworth & Brooks/Cole Advanced Books & Software.
- W. Pauli (1980), *General Principles of Quantum Mechanics*, Springer.
- G. Pisier (2012), Grothendieck’s theorem, past and present, *Bull. Amer. Math. Soc. (N.S.)* **49**, 237–323.
- J. Plebański and A. Krasinski (2006), *An Introduction to General Relativity and Cosmology*, Cambridge University Press.
- K. R. Rao and P. Yip (1990), *Discrete Cosine Transform: Algorithms, Advantages, Applications*, Academic Press.
- M. Reed and B. Simon (1980), *Functional Analysis*, Vol. I of Methods of Modern Mathematical Physics, second edition, Academic Press.

- J. Renegar (2001), *A Mathematical View of Interior-Point Methods in Convex Optimization*, MPS/SIAM Series on Optimization, Society for Industrial and Applied Mathematics (SIAM) / Mathematical Programming Society (MPS).
- B. Reznick (1992), *Sums of Even Powers of Real Linear Forms*, Vol. 96 of *Memoirs of the American Mathematical Society*, American Mathematical Society.
- M. M. G. Ricci and T. Levi-Civita (1900), Méthodes de calcul différentiel absolu et leurs applications, *Math. Ann.* **5**, 125–201. Translation by R. Hermann: *Ricci and Levi-Civita's Tensor Analysis Paper* (1975), Vol. II of *Lie Groups: History, Frontiers and Applications*, Math Sci Press.
- E. Riehl (2016), *Category Theory in Context*, Dover Publications.
- W. Rindler (2006), *Relativity: Special, General, and Cosmological*, second edition, Oxford University Press.
- M. D. Roberts (1995), The physical interpretation of the Lanczos tensor, *Nuovo Cimento B (11)* **110**, 1165–1176.
- D. N. Rockmore (2000), The FFT: An algorithm the whole family can use, *Comput. Sci. Eng.* **2**, 30–34.
- F. Rosenblatt (1958), The perceptron: A probabilistic model for information storage and organization in the brain, *Psychol. Rev.* **65**, 386–408.
- R. A. Ryan (2002), *Introduction to Tensor Products of Banach Spaces*, Springer Monographs in Mathematics, Springer.
- H. H. Schaefer and M. P. Wolff (1999), *Topological Vector Spaces*, Vol. 3 of *Graduate Texts in Mathematics*, second edition, Springer.
- R. Schatten (1950), *A Theory of Cross-Spaces*, Vol. 26 of *Annals of Mathematics Studies*, Princeton University Press.
- H. Scheffé (1970), A note on separation of variables, *Technometrics* **12**, 388–393.
- H. Scheffé (1999), *The Analysis of Variance*, Wiley Classics Library, Wiley.
- K. Schmüdgen (2017), *The Moment Problem*, Vol. 277 of *Graduate Texts in Mathematics*, Springer.
- K. Schöbel (2015), *An Algebraic Geometric Approach to Separation of Variables*, Springer Spektrum.
- B. Scholköpfung and A. J. Smola (2002), *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*, MIT Press.
- A. Schönhage (1972/73), Unitäre Transformationen grosser Matrizen, *Numer. Math.* **20**, 409–417.
- A. Schönhage and V. Strassen (1971), Schnelle Multiplikation grosser Zahlen, *Computing (Arch. Elektron. Rechnen)* **7**, 281–292.
- J. A. Schouten (1951), *Tensor Analysis for Physicists*, Clarendon Press.
- A. Schrijver (1986), *Theory of Linear and Integer Programming*, Wiley-Interscience Series in Discrete Mathematics, Wiley-Interscience.
- R. F. Service (2020), ‘The game has changed’: AI triumphs at protein folding, *Science* **370** (6521), 1144–1145.
- Y. Y. Shi, L. M. Duan and G. Vidal (2006), Classical simulation of quantum many-body systems with a tree tensor network, *Phys. Rev. A* **74**, 022320.
- P. W. Shor (1994), Algorithms for quantum computation: Discrete logarithms and factoring, in *35th Annual Symposium on Foundations of Computer Science (1994)*, IEEE Computer Society Press, pp. 124–134.

- W. Sickel and T. Ullrich (2009), Tensor products of Sobolev–Besov spaces and applications to approximation from the hyperbolic cross, *J. Approx. Theory* **161**, 748–786.
- J. G. Simmonds (1994), *A Brief on Tensor Analysis*, Undergraduate Texts in Mathematics, second edition, Springer.
- S. S. Skiena (2020), *The Algorithm Design Manual*, Texts in Computer Science, third edition, Springer.
- S. Smale (1998), Mathematical problems for the next century, *Math. Intelligencer* **20**, 7–15.
- S. A. Smolyak (1963), Quadrature and interpolation formulas for tensor products of certain classes of functions, *Dokl. Akad. Nauk SSSR* **148**, 1042–1045.
- N. P. Sokolov (1960), *Spatial Matrices and their Applications*, Gosudarstv. Izdat. Fiz.-Mat. Lit., Moscow. In Russian.
- N. P. Sokolov (1972), *Introduction to the Theory of Multidimensional Matrices*, Izdat. ‘Naukova Dumka’, Kiev. In Russian.
- B. Spain (1960), *Tensor Calculus: A Concise Course*, third edition, Oliver and Boyd / Interscience.
- E. M. Stein (1993), *Harmonic Analysis: Real-Variable Methods, Orthogonality, and Oscillatory Integrals*, Vol. 43 of Princeton Mathematical Series, Princeton University Press.
- I. Steinwart and A. Christmann (2008), *Support Vector Machines*, Information Science and Statistics, Springer.
- G. W. Stewart (2000), The decompositional approach to matrix computation, *Comput. Sci. Eng.* **2**, 50–59.
- G. Strang (1980), *Linear Algebra and its Applications*, second edition, Academic Press.
- V. Strassen (1969), Gaussian elimination is not optimal, *Numer. Math.* **13**, 354–356.
- V. Strassen (1973), Vermeidung von Divisionen, *J. Reine Angew. Math.* **264**, 184–202.
- V. Strassen (1987), Relative bilinear complexity and matrix multiplication, *J. Reine Angew. Math.* **375/376**, 406–443.
- V. Strassen (1990), Algebraic complexity theory, in *Handbook of Theoretical Computer Science*, Vol. A: *Algorithms and Complexity*, Elsevier, pp. 633–672.
- A. M. Stuart and A. R. Humphries (1996), *Dynamical Systems and Numerical Analysis*, Vol. 2 of Cambridge Monographs on Applied and Computational Mathematics, Cambridge University Press.
- J. J. Sylvester (1886), Sur une extension d’un théorème de Clebsch relatif aux courbes du quatrième degré, *C. R. Math. Acad. Sci. Paris* **102**, 1532–1534.
- J. L. Synge and A. Schild (1978), *Tensor Calculus*, Dover Publications.
- C.-T. Tai (1997), *Generalized Vector and Dyadic Analysis*, second edition, IEEE Press.
- M. Takesaki (2002), *Theory of Operator Algebras I*, Vol. 124 of Encyclopaedia of Mathematical Sciences, Springer.
- L. A. Takhtajan (2008), *Quantum Mechanics for Mathematicians*, Vol. 95 of Graduate Studies in Mathematics, American Mathematical Society.
- E. Tardos (1986), A strongly polynomial algorithm to solve combinatorial linear programs, *Oper. Res.* **34**, 250–256.
- G. Temple (1960), *Cartesian Tensors: An Introduction*, Methuen’s Monographs on Physical Subjects, Methuen / Wiley.
- G. Teschl (2014), *Mathematical Methods in Quantum Mechanics: With Applications to Schrödinger Operators*, Vol. 157 of Graduate Studies in Mathematics, second edition, American Mathematical Society.

- J. W. Thomas (1995), *Numerical Partial Differential Equations: Finite Difference Methods*, Vol. 22 of Texts in Applied Mathematics, Springer.
- K. S. Thorne and R. D. Blandford (2017), *Modern Classical Physics: Optics, Fluids, Plasmas, Elasticity, Relativity, and Statistical Physics*, Princeton University Press.
- A. L. Toom (1963), The complexity of a scheme of functional elements simulating the multiplication of integers, *Dokl. Akad. Nauk SSSR* **150**, 496–498.
- L. N. Trefethen and D. Bau, III (1997), *Numerical Linear Algebra*, Society for Industrial and Applied Mathematics (SIAM).
- F. Trèves (2006), *Topological Vector Spaces, Distributions and Kernels*, Dover Publications.
- P. M. Vaidya (1990), An algorithm for linear programming which requires $O((m+n)n^2 + (m+n)^{1.5}nL)$ arithmetic operations, *Math. Program.* **47**, 175–201.
- L. G. Valiant (1979), The complexity of computing the permanent, *Theoret. Comput. Sci.* **8**, 189–201.
- H. A. van der Vorst (2000), Krylov subspace iteration, *Comput. Sci. Eng.* **2**, 32–37.
- N. T. Varopoulos (1965), Sur les ensembles parfaits et les séries trigonométriques, *C. R. Acad. Sci. Paris* **260**, 4668–4670, 5165–5168, 5997–6000.
- N. T. Varopoulos (1967), Tensor algebras and harmonic analysis, *Acta Math.* **119**, 51–112.
- S. A. Vavasis (1991), *Nonlinear Optimization: Complexity Issues*, Vol. 8 of International Series of Monographs on Computer Science, The Clarendon Press, Oxford University Press.
- F. Verstraete and J. I. Cirac (2004), Renormalization algorithms for quantum-many body systems in two and higher dimensions. Available at [arXiv:cond-mat/0407066](https://arxiv.org/abs/cond-mat/0407066).
- E. B. Vinberg (2003), *A Course in Algebra*, Vol. 56 of Graduate Studies in Mathematics, American Mathematical Society.
- W. Voigt (1898), *Die fundamentalen physikalischen Eigenschaften der Krystalle in elementarer Darstellung*, Von Veit.
- R. M. Wald (1984), *General Relativity*, University of Chicago Press.
- N. R. Wallach (2008), Quantum computing and entanglement for mathematicians, in *Representation Theory and Complex Analysis*, Vol. 1931 of Lecture Notes in Mathematics, Springer, pp. 345–376.
- W. Walter (1998), *Ordinary Differential Equations*, Vol. 182 of Graduate Texts in Mathematics, Springer.
- A.-M. Wazwaz (2011), *Linear and Nonlinear Integral Equations: Methods and Applications*, Higher Education Press (Beijing) / Springer.
- G. Weinreich (1998), *Geometrical Vectors*, Chicago Lectures in Physics, University of Chicago Press.
- H. Weyl (1997), *The Classical Groups: Their Invariants and Representations*, Princeton Landmarks in Mathematics, Princeton University Press.
- S. R. White (1992), Density matrix formulation for quantum renormalization groups, *Phys. Rev. Lett.* **69**, 2863–2866.
- S. R. White and D. A. Huse (1993), Numerical renormalization-group study of low-lying eigenstates of the antiferromagnetic $S = 1$ Heisenberg chain, *Phys. Rev. B* **48**, 3844–3853.
- J. D. Whitfield, P. J. Love and A. Aspuru-Guzik (2013), Computational complexity in electronic structure, *Phys. Chem. Chem. Phys.* **15**, 397–411.

- N. M. J. Woodhouse (2003), *Special Relativity*, Springer Undergraduate Mathematics Series, Springer.
- R. C. Wrede (1963), *Introduction to Vector and Tensor Analysis*, Wiley.
- S.-T. Yau (2020), Shiing-Shen Chern: A great geometer of 20th century.
<https://cmsa.fas.harvard.edu/wp-content/uploads/2020/05/2020-04-22-essay-on-Chern-english-v1.pdf>.
- K. Ye and L.-H. Lim (2018a), Fast structured matrix computations: Tensor rank and Cohn–Umans method, *Found. Comput. Math.* **18**, 45–95.
- K. Ye and L.-H. Lim (2018b), Tensor network ranks. Available at [arXiv:1801.02662](https://arxiv.org/abs/1801.02662).