Appl. Comput. Harmon. Anal. 49 (2020) 180–207

Contents lists available at ScienceDirect

Applied and Computational Harmonic Analysis

www.elsevier.com/locate/acha

Complex best r-term approximations almost always exist in finite dimensions

Yang Qi [a], Mateusz Michałek [b,c], Lek-Heng Lim [a,*]

[a] *Computational and Applied Mathematics Initiative, Department of Statistics, University of Chicago, Chicago, IL, 60637, USA*
[b] *Max Planck Institute for Mathematics in the Sciences, 04103 Leipzig, Germany*
[c] *Polish Academy of Sciences, 00656 Warsaw, Poland*

A R T I C L E   I N F O

Article history:
Received 27 December 2017
Received in revised form 22 December 2018
Accepted 31 December 2018
Available online 4 January 2019
Communicated by Gregory Beylkin

MSC:
15A69
41A46
41A52

Keywords:
Nonlinear approximations
Best rank-r approximations
Sparse-plus-low-rank approximations
Tensor ranks
Tensor completion
Tensor networks

A B S T R A C T

We show that in finite-dimensional nonlinear approximations, the best r-term approximant of a function f almost always exists over $\mathbb{C}$ but that the same is not true over $\mathbb{R}$, i.e., the infimum $\inf_{f_1,\dots,f_r\in D}\|f - f_1 - \cdots - f_r\|$ is almost always attainable by complex-valued functions $f_1, \dots, f_r$ in D, a set (dictionary) of functions (atoms) with some desired structures. Our result extends to functions that possess properties like symmetry or skew-symmetry under permutations of arguments. When D is the set of separable functions, this is the best rank-r tensor approximation problem. We show that over $\mathbb{C}$, any tensor almost always has a unique best rank-r approximation. This extends to other notions of ranks such as symmetric and alternating ranks, to best r-block-terms approximations, and to best approximations by tensor networks. Applied to sparse-plus-low-rank approximations, we obtain that for any given r and k, a general tensor has a unique best approximation by a sum of a rank-r tensor and a k-sparse tensor with a fixed sparsity pattern; a problem arising in covariance estimation of Gaussian model with k observed variables conditionally independent given r hidden variables. The existential (but not uniqueness) part of our result also applies to best approximations by a sum of a rank-r tensor and a k-sparse tensor with no fixed sparsity pattern, and to tensor completion problems.

1. Introduction

There are numerous problems in scientific and engineering applications that may ultimately be put in the following form: Given a real or complex-valued function $f : \Omega \to \mathbb{K}$ (with $\mathbb{K} = \mathbb{R}$ or $\mathbb{C}$ respectively), find a *best approximation* of f by a sum of functions $f_1, \dots, f_r$ with some special structure, i.e.,

$$\min_{f_1,\dots,f_r\in D}\|f - (f_1 + \cdots + f_r)\|, \tag{1}$$

* Corresponding author.
E-mail addresses: yangqi@galton.uchicago.edu (Y. Qi), Mateusz.Michalek@mis.mpg.de (M. Michałek), lekheng@galton.uchicago.edu (L.-H. Lim).

https://doi.org/10.1016/j.acha.2018.12.003

where D is a subset of functions possessing that special structure. In the applied harmonic analysis literature, the set D is often called a *dictionary*, a function $f_i \in D$ an *atom*, and the function to be approximated f the *target function*. The problem (1), called the *best r-term approximation problem*[1] or *best rank-r approximation problem* is ubiquitous by its generality and simplicity. A slight generalization involves

$$\min_{f_1,\dots,f_r \in D_1;\; g_1,\dots,g_s \in D_2} \|f - (f_1 + \cdots + f_r + g_1 + \cdots + g_s)\|, \tag{2}$$

where D_1 and D_2 denote two dictionaries/subsets of atoms/functions each with a different structure (e.g., this arises in so-called sparse plus low-rank approximations). More generally, the approximation could involve even more dictionaries/subsets $D_1, \dots, D_k$ of atoms/functions, each capturing a different structure in the target function f.

In practice, if the problem is not already discrete, then it has to be discretized for the purpose of computations. This usually involves discretizing Ω (e.g., sampling points, triangulation, mesh generation) or finding a finite-dimensional approximation of the function spaces (e.g., Galerkin method, quadrature) or both (e.g., collocation methods). The result of which is that we may in effect assume that Ω is a finite set or that the function space $L^2_{\mathbb{K}}(\Omega)$ is a finite-dimensional vector space (as our notation implies, we assume that the norm in (1) is an l^2-norm). In which case, up to a choice of basis, $L^2_{\mathbb{K}}(\Omega) \cong \mathbb{K}^n$, where $n = \dim L^2_{\mathbb{K}}(\Omega)$ (or, if Ω is finite, $n = \#\Omega$).

The main issue with (1) is that when D is not finite up to scaling, the approximation problem in (1) often does not have a solution: Let $f \in L^2_{\mathbb{K}}(\Omega)$ and $D \subseteq L^2_{\mathbb{K}}(\Omega)$ be a closed subset (under the metric topology induced by $\|\cdot\|$), the infimum

$$\inf_{f_1,\dots,f_r \in D} \|f - (f_1 + \cdots + f_r)\|$$

is often not attainable when $r > 1$. The reason being that the set of r-term approximants

$$\Sigma_r^\circ(D) := \{f \in L^2_{\mathbb{K}}(\Omega) \mid f = f_1 + \cdots + f_r \text{ for some } f_1, \dots, f_r \in D\}$$

is often not closed when $r > 1$. Let $\Sigma_r(D)$ denote the closure of $\Sigma_r^\circ(D)$ under the metric topology. Each element in $\Sigma_r(D)$ is a limit of a sequence of r-term approximants but may not itself be an r-term approximant. While

$$\varphi_* \in \operatorname{argmin}_{\varphi \in \Sigma_r(D)} \|f - \varphi\|$$

always exist, it would not in general be an r-term approximant, i.e., φ_* may fail to be of the required form $f_1 + \cdots + f_r$ with $f_i \in D$ — this occurs when $\varphi_* \in \Sigma_2(D) \setminus \Sigma_2^\circ(D)$. The main result of this article is to show that it makes a vast difference whether $\mathbb{K} = \mathbb{R}$ or $\mathbb{K} = \mathbb{C}$ — for the latter, this failure almost never happens under some mild conditions.

We will elaborate this phenomenon more concretely by studying the case of separable approximations. In this case, $\Omega = \Omega_1 \times \cdots \times \Omega_d$ and

$$D = \{\varphi \in L^2_{\mathbb{K}}(\Omega) \mid \varphi = \varphi_1 \otimes \cdots \otimes \varphi_d \text{ where } \varphi_j \in L^2_{\mathbb{K}}(\Omega_j)\} \tag{3}$$

is the set of separable functions, i.e., functions $\varphi : \Omega_1 \times \cdots \times \Omega_d \to \mathbb{K}$ of the form

$$\varphi(x_1, x_2, \dots, x_d) = \varphi_1(x_1)\varphi_2(x_2)\cdots\varphi_d(x_d), \qquad x_j \in \Omega_j,\ j = 1, \dots, d.$$

[1] Sometimes (1) is expressed as a minimization over linear combinations instead of sums, with a corresponding restriction on the atoms in D, e.g., to unit norm.

Take $d = 3$ for simplicity. When $\mathbb{K} = \mathbb{R}$, it has been shown [1, Theorem 8.4] that there is a nonempty open subset $\mathcal{O} \subseteq L^2_{\mathbb{R}}(\Omega)$ such that for any $f \in \mathcal{O}$, the infimum

$$\inf_{\varphi_i, \psi_i \in L^2_{\mathbb{R}}(\Omega_i)} \|f - \varphi_1 \otimes \varphi_2 \otimes \varphi_3 - \psi_1 \otimes \psi_2 \otimes \psi_3\| \tag{4}$$

fails to be attainable. In particular the set of such failures, given that it contains an open set, must have positive volume, i.e., such failures occur with positive probability and cannot be ignored in practice. Given any $f \in L^2_{\mathbb{K}}(\Omega)$, the *rank* of f is the integer r such that f is a sum of r separable functions, i.e.,

$$\operatorname{rank}(f) = \min\left\{ r \;\middle|\; f = \sum_{i=1}^{r} \varphi_{1,i} \otimes \cdots \otimes \varphi_{d,i} \right\}, \tag{5}$$

and so $\Sigma^\circ_r(D) = \{f \in L^2_{\mathbb{K}}(\Omega) \mid \operatorname{rank}(f) \le r\}$. Therefore the preceding discussion says that the set of f that fails to have a best rank-two approximation has positive volume. In fact, there are more extreme examples [1, Theorem 8.1] where *every* $f \in L^2_{\mathbb{R}}(\Omega)$ of rank > 2 fails to have a best rank-two approximation; we defer further discussions to Section 2.

Our main objective is to show that: (a) for most choices of dictionary D, the best r-term approximation problem (1) will fail to have a solution for *some* choices of target function f (see Theorem 2.2); (b) when $\mathbb{K} = \mathbb{C}$, for any choice of D, the set of f for which (1) fails to have a solution has zero volume (see Corollary 6.1). Note that the previous paragraph shows that (b) is false when $\mathbb{K} = \mathbb{R}$. Also, the results in this article are nontrivial only when the dictionary D contains *infinitely* many linearly independent atoms; if $D = \{\lambda\varphi_i \mid i = 1, \dots, m;\ \lambda \in \mathbb{K}\}$, then (1) trivially has a solution, irrespective of whether $\mathbb{K} = \mathbb{R}$ or $\mathbb{C}$ (see Proposition 2.1).

Up until this point, we have presented our discussions entirely in the language of function approximation in order to put the problems (1) and (2) within the context they most frequently arise, which traditionally goes under the heading of *nonlinear approximation*. Nevertheless, the subject has witnessed many recent breakthroughs and is now studied under a number of different names [2–11]. In this article we will undertake an approach via complex algebraic geometry and real analytic geometry. From our perspective, the problems in (1) and (2) are respectively about *secant varieties* and *join varieties*. We will henceforth drop the function approximation description in order to focus on the crux of the issue. In the rest of this article, we will let U denote a real vector space and V denote a complex vector space. Readers who prefer the function-theoretic description may continue to think of U as $L^2_{\mathbb{R}}(\Omega)$ and V as $L^2_{\mathbb{C}}(\Omega)$. The difference between the geometry of the best r-term approximation problem over $\mathbb{R}$ and over $\mathbb{C}$ may be summarized as follows:

- $\mathbb{K} = \mathbb{R}$: Let X be a closed semianalytic subset of an $\mathbb{R}$-vector space U, and $Z \subseteq X$ be a semianalytic subset of X with $\dim Z < \dim X < \dim U$. The set Z here represents the 'bad points' to be avoided — given $p \in U$, we seek a best approximation $x_* \in X$ that attains $\min_{x\in X} \|p - x\|$ but we also want $x_* \notin Z$. The example we mentioned above with D as defined in (3) has $X = \Sigma_2(D)$, $Z = \Sigma_2(D) \setminus \Sigma^\circ_2(D)$, and $f \in L^2_{\mathbb{R}}(\Omega)$ in the role of $p \in U$. Here we see that there can be a nonempty open subset $\mathcal{O} \subseteq U$ such that for $p \in \mathcal{O}$, any best approximation of p lies in Z, which in this example represents the functions in X that cannot be written as a sum of two separable functions.
- $\mathbb{K} = \mathbb{C}$: Let X be a closed irreducible complex analytic variety in a $\mathbb{C}$-vector space V, and $Z \subseteq X$ be a complex analytic subvariety of X with $\dim Z < \dim X < \dim V$. We will prove that any general $p \in V$ will have its best approximation in X but not in Z, and moreover this best approximation will be unique. In the context of the best r-term approximation problem (1), $X = \Sigma_r(D)$, $Z = \Sigma_r(D) \setminus \Sigma^\circ_r(D)$, and $f \in L^2_{\mathbb{C}}(\Omega)$ plays the role of $p \in V$. But this will apply more generally to the problem (2) as well, giving us existence and uniqueness results for various nonlinear approximations common in applications (see Sections 6–8).

Figs. 2 and 3 illustrate the gist of the difference between $\mathbb{R}$ and $\mathbb{C}$. Readers who just want an intuitive grasp without going into the details may like to consult these figures.

2. Nonexistence of best r-term approximations is common

We begin by first noting that when D is finite up to scaling, then a best r-term approximation always exists, i.e., (1) always has a solution.

Proposition 2.1. *If the dictionary $D \subseteq V$ is finite in the projective space $\mathbb{P}V$, i.e., finite up to scaling, then $\Sigma_r^\circ(D)$ is closed in V.*

Proof. D is the form $\{\lambda\varphi_i \mid i = 1, \dots, m;\ \lambda \in \mathbb{K}\}$ for some $m \in \mathbb{N}$. So $\Sigma_r^\circ(D)$ is a union of $\binom{m}{r}$ subspaces in V and therefore closed. $\square$

Proposition 2.1 explains why we are only interested in the case when D is *not* finite up to scaling. Examples include exponential dictionaries like $D = \{e^{isx} \mid s \in [0, a]\}$ for some $a > 0$ [12], $D = \{e^{-tx} \mid t \in \mathbb{R}\}$, or $D = \{e^{\tau x} \mid \tau \in \mathbb{C}\}$ [13], the dictionary of separable functions we encountered earlier in (3), with many more to come in Sections 6–8.

We will briefly review the example in [1, Theorem 8.1] to give readers a better idea as to why (1) can fail on a positive-volumed set of target functions f over $\mathbb{R}$. Let $\Omega_1 = \Omega_2 = \Omega_3 = \{0, 1\}$. One can show that real-valued functions $f \in L^2_{\mathbb{R}}(\Omega_1 \times \Omega_2 \times \Omega_3) = L^2_{\mathbb{R}}(\{0, 1\}^3)$ with $\operatorname{rank}_{\mathbb{R}}(f) > 2$ must take either one of the two following forms:

$$f = f_1 \otimes f_0 \otimes f_0 + f_0 \otimes f_1 \otimes f_0 + f_0 \otimes f_0 \otimes f_1, \tag{6}$$

$$f = (f_0 + f_1) \otimes f_1 \otimes f_1 + (f_0 - f_1) \otimes f_0 \otimes f_0 + f_1 \otimes (f_0 + f_1) \otimes (f_0 - f_1), \tag{7}$$

where $f_0, f_1 \in L^2_{\mathbb{R}}(\{0, 1\})$ are linearly independent. Let $\mathcal{D}$ and $\mathcal{O}$ be the subsets of functions in $L^2_{\mathbb{R}}(\{0, 1\}^3)$ that take the forms in (6) and (7) respectively. One can show that the infimum in (4) cannot be attained for any $f \in \mathcal{D} \cup \mathcal{O}$. In addition $\mathcal{O}$ is an open subset in $L^2_{\mathbb{R}}(\{0, 1\}^3)$ with positive volume. So the best two-term approximation problem fails to have a solution on a positive-volumed set. On the other hand, if one had worked over $\mathbb{C}$, then one can show that the only complex-valued functions $f \in L^2_{\mathbb{C}}(\Omega_1 \times \Omega_2 \times \Omega_3) = L^2_{\mathbb{C}}(\{0, 1\}^3)$ with $\operatorname{rank}_{\mathbb{C}}(f) > 2$ are those taking the form in (6) where $f_0, f_1 \in L^2_{\mathbb{C}}(\{0, 1\})$ are linearly independent. While the infimum in (4) also cannot be attained for such an f, the subset of all such functions has zero volume in $L^2_{\mathbb{C}}(\{0, 1\}^3)$. So the best two-term approximation problem fails to have a solution only on a zero-volumed set.

Let D be the dictionary of separable functions in (3). The reason the infimum in (4) is not attained in the above example is a consequence of the fact that the set of two-term approximants $\Sigma_2^\circ(D)$ is not closed. This phenomenon is common — we will next show that it can happen to any D with an accumulation point, modulo an embedding in some higher-dimensional space.

Theorem 2.2. *Let W be a vector space over $\mathbb{R}$ or $\mathbb{C}$ and $X \subseteq W$ be a set with an accumulation point. Then for any $r > 1$, there exists a vector space V and an embedding $\nu : X \hookrightarrow V$ such that if $D := \nu(X)$, then $\Sigma_r^\circ(D)$ is not closed in V.*

Proof. Let $V := \mathsf{S}^{r+1}(W)$ be the $(r + 1)$st symmetric power of W and let $\nu : W \to V$, $w \mapsto w^{\circ(r+1)}$ be the affine Veronese embedding (see Section 7.1). Let $x \in X$ be an accumulation point. Consider a tangent vector t to $\nu(X)$ at $\nu(x)$. Clearly, it is also a tangent vector to $\nu(W)$. Hence $t = x^{\circ r} \circ w$ for some $w \in W$. Furthermore $t \in \Sigma_k(\nu(X))$ for all $k \geq 2$. However, t has symmetric rank $r + 1$ as a symmetric tensor, i.e.,

$t \in \Sigma_{r+1}^{\circ}(\nu(W))$ and $t \notin \Sigma_r^{\circ}(\nu(W))$. As $\Sigma_r^{\circ}(\nu(X)) \subseteq \Sigma_r^{\circ}(\nu(W))$, we see that $\Sigma_k^{\circ}(\nu(X))$ is not closed for any $k = 2, \dots, r$. $\square$

The ambient space V of D constructed in our proof is of much higher dimension than D. But if desired, we can add an additional step to our construction — project from an arbitrary affine subspace disjoint from the join (see (11)) of t and $\Sigma_r(\nu(W))$, and choose V to be the image of this projection — then we are in a situation where $\Sigma_r^{\circ}(D)$ is only one dimension lower than V but everything in the statement of Theorem 2.2 still holds.

3. Subanalytic geometry

We will study our approximation problems using selected tools from subanalytic geometry [14], which we review in the following. A *semianalytic* subset of $\mathbb{R}^n$ is a set that is locally of the form

$$X = \bigcup_{i=1}^{m} \bigcap_{j=1}^{k} X_{ij},$$

where each X_{ij} takes the form

$$\{(x_1, \dots, x_n) \in \mathbb{R}^n \mid f_{ij}(x_1, \dots, x_n) = 0\} \qquad \text{or} \qquad \{(x_1, \dots, x_n) \in \mathbb{R}^n \mid f_{ij}(x_1, \dots, x_n) > 0\}$$

for some real analytic function f_{ij}. Let W be a linear subspace of $\mathbb{R}^n$, and $\pi \colon \mathbb{R}^n \to W$ be a linear projection. A subset $X \subseteq W$ is called *subanalytic* if X is locally of the form $\pi(D)$ for some relatively compact semianalytic subset $D \subseteq \mathbb{R}^n$. Given subanalytic subsets $X \subseteq \mathbb{R}^m$ and $D \subseteq \mathbb{R}^n$, a continuous map $f \colon X \to D$ is called *subanalytic* if the graph of f, i.e., $\{(x, y) \in X \times D \mid y = f(x)\}$, is subanalytic in $\mathbb{R}^m \times \mathbb{R}^n$.

We will also make significant use of Whitney stratifications, which we recall here. Let $\mathcal{I}$ be a partially ordered set with order relation $\preceq$, and X a closed subset of a real smooth manifold M. A *Whitney stratification* of X is a collection $\{S_i \mid i \in \mathcal{I}\}$ of disjoint locally closed smooth submanifolds $S_i \subseteq X$ called *strata* that satisfy:

(i) $X = \bigcup_{i \in \mathcal{I}} S_i$.
(ii) Every $p \in X$ has a neighborhood U such that $U \cap S_i \neq \varnothing$ for finitely many $i \in \mathcal{I}$.
(iii) $i \preceq j$ if and only if $S_i \subseteq \overline{S}_j$ if and only if $S_i \cap \overline{S}_j \neq \varnothing$, where $\overline{S}_j$ is the Euclidean closure of $\overline{S}_j$.
(iv) Suppose $S_\alpha \subseteq \overline{S}_\beta$. Let $\{x_i\}_{i=1}^{\infty} \subseteq S_\beta$ and $\{y_i\}_{i=1}^{\infty} \subseteq S_\alpha$ be two sequences converging to the same $y \in S_\alpha$. If the secant lines $\operatorname{span}\{x_i, y_i\}$ converge to some limiting line ℓ and the tangent spaces $\mathrm{T}_{x_i}(S_\beta)$ converge to some limiting space τ, then $\mathrm{T}_y(S_\alpha) \subseteq \tau$ and $\ell \subseteq \tau$.

For a proper subanalytic map $f \colon X \to D$, there are Whitney stratifications of X and D into analytic manifolds such that f is a *stratified map*, i.e., for each stratum $B \subseteq D$, $f^{-1}(B)$ is a union of connected components of strata of X. We refer the reader to [14] for further details, but the pictures of *Whitney's umbrella* $z^2 - yx^2 = 0$ and its Whitney stratification in Fig. 1 ought to give readers an intuitive idea.

The following result is well-known in the semialgebraic setting [15, Theorem 3.3] but we need it in a subanalytic setting; fortunately a similar proof yields the required analogous result.

Lemma 3.1. *Let U_1 and U_2 be real vector spaces. Let $X \subseteq U_1$ be a k-dimensional subanalytic subset and $f \colon X \to U_2$ be a subanalytic map. Then the set of points of X where f is not differentiable is contained in a subanalytic subset of dimension strictly smaller than k.*

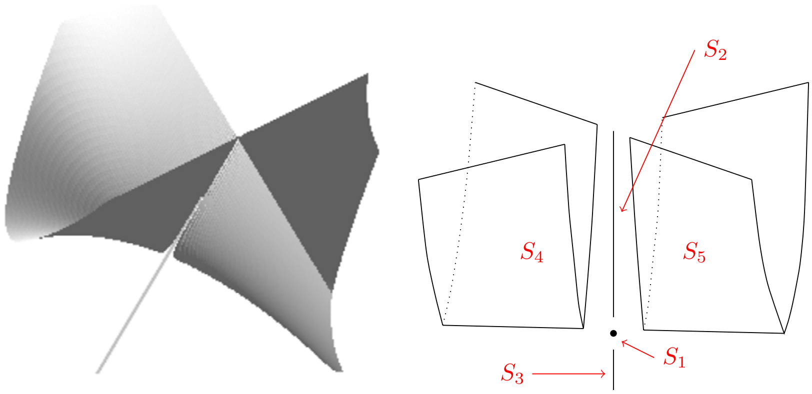

Fig. 1. *Left*: Whitney's umbrella $z^2 - yx^2 = 0$. *Right*: Stratification of Whitney's umbrella into five strata $S_1, \ldots, S_5$. Here S_1 is a point; S_2 and S_3 are the upper and lower half line in the singular locus; S_4 and S_5 are the left and right connected components of the smooth locus. In the partial order S_1 is the unique smallest element, i.e., $S_1 \preceq S_2, S_3, S_4, S_5$; we also have $S_2 \preceq S_4, S_5$; but all other strata are incomparable.

Proof. Let $\Gamma \subseteq U_1 \times U_2$ be the graph of f, and $\pi_1 \colon X \times U_2 \to X$ be the projection. Given a Whitney stratification $\Gamma = \bigcup_{i \in \mathcal{I}} \Gamma_i$ of Γ, since each $\pi_1(\Gamma_i)$ is subanalytic, there is a Whitney stratification $X = \bigcup_{j \in J} X_j$ such that $\pi_1(\Gamma_i)$ is a union of strata, namely $f|_{X_j} \colon X_j \to U_2$ is a subanalytic map whose graph is an analytic submanifold of $U_1 \times U_2$. For simplicity, we denote $f|_{X_j}$ by f_j and the graph of f_j by Δ_j. Then the set of points of X_j where f_j is not differentiable is contained in the critical values of $\pi_1 \colon \Delta_j \to X_j$. Hence the set of nondifferentiable points of f in X is contained in

$$\{x \in X_j \mid \dim X_j < k\} \cup \{x \in X_l \mid \dim X_l = k \text{ and } x \text{ is a critical value of } \pi_1 \colon \Delta_l \to X_l\}$$

which is a subanalytic subset whose dimension is less than k. $\square$

In this article we prove most of our results for *complex analytic varieties*, i.e., defined locally by the common zero loci of finitely many holomorphic functions; although the examples in Section 7 are all *complex algebraic varieties*, i.e., those holomorphic functions are polynomials.[2] Since a subset of $\mathbb{C}^n$ may be regarded as a subset of $\mathbb{R}^{2n}$, we have the following relations between (semi)algebraic and (semi)analytic sets (here '$\subseteq$' is interpreted to mean 'is a special case of'):

$$\begin{aligned} &\text{complex algebraic variety} \subseteq \text{real semialgebraic set} \subseteq \text{real semianalytic set}, \\ &\text{complex algebraic variety} \subseteq \text{complex analytic variety} \subseteq \text{real semianalytic set}. \end{aligned}$$

So for example any property of real semianalytic sets will automatically be satisfied by a complex analytic variety. In particular, the following important theorem [16] also holds true in subanalytic, semialgebraic, or complex algebraic contexts; but for our purpose, it will be stated for complex analytic sets.

Theorem 3.2. *Let M be a complex manifold and $Z \subseteq M$ be a closed complex analytic subset. Then there is a Whitney stratification $Z = \bigcup_{i \in \mathcal{I}} S_i$ of Z such that*

(a) *each S_i is a complex submanifold of M;*
(b) *if $S_i \subseteq \overline{S}_j$, there is a vector bundle $E_i \to S_i$, a neighborhood $U_i \subseteq E_i$ of the zero section S_i, and a homeomorphism $h : U_i \to \overline{S}_j$ of U_i to an open subset $h(U_i)$ in $\overline{S}_j$.*

[2] While every complex *projective* analytic variety is algebraic by Chow's theorem, there exist complex analytic varieties that are not algebraic.

3.1. Notations and terminologies

In the next sections, we will frame our discussion over an abstract real vector space U or an abstract complex vector space V. By a semialgebraic or subanalytic subset $X \subseteq U$ we mean that X can be identified with some semialgebraic or subanalytic subset in $\mathbb{R}^n$ when we identify $U \cong \mathbb{R}^n$ by fixing a basis of U. Likewise, by an algebraic or analytic variety $X \subseteq V$ we mean that X can be identified with some algebraic or analytic variety in $\mathbb{C}^n$ when we identify $V \cong \mathbb{C}^n$ by fixing a basis of V.

The main reason we prefer such coordinate-free descriptions instead of assuming at the outset that $U = \mathbb{R}^n$ or $V = \mathbb{C}^n$ is that we will ultimately be applying the results to cases where U and V have additional structures, e.g., $\mathbb{K}^{n_1} \otimes \cdots \otimes \mathbb{K}^{n_d}$, $\mathsf{S}^d(\mathbb{K}^n)$, $\mathsf{\Lambda}^k(\mathbb{K}^n)$, $\mathsf{S}^d(\mathsf{\Lambda}^k(\mathbb{K}^n))$, $\mathsf{S}^{d_1}(\mathbb{K}^n) \otimes \cdots \otimes \mathsf{S}^{d_n}(\mathbb{K}^n)$, $\mathsf{\Lambda}^{k_1}(\mathbb{K}^n) \otimes \cdots \otimes \mathsf{\Lambda}^{k_n}(\mathbb{K}^n)$, etc.

The results in this article will be proved for *general points*. A point in a semianalytic set X is said to be *general* with respect to a property if the subset of points that do not have that property is contained in a semianalytic subset whose dimension is strictly smaller than $\dim X$. In particular, a point in a real vector space is said to be *general* with respect to a property if the set of points that do not have that property is contained in a real analytic hypersurface. The same notion applies to a point in a complex vector space, regarded as a real vector space of real dimension twice its complex dimension.

Establishing that a result holds true for all general points is far stronger than the common practice in applied and computational mathematics of establishing it ‘with high probability’ or ‘almost surely’ or ‘almost everywhere’. If a result is valid for all general points, then not only do we know that it is valid almost surely/everywhere but also that the invalid points are all limited to a subset of strictly smaller dimension. In fact, when the result is about points in a vector space, which is the case in our article, this subset of invalid points can be described by a single equation that may in principle be determined, so that we know where the invalid points lie.

4. Best approximation by points in a closed subanalytic set

Let U be an n-dimensional real vector space with an inner product $\langle \cdot, \cdot \rangle$ and corresponding ℓ^2-norm $\|\cdot\|$. Let $X \subseteq U$ be a closed subanalytic set. We define the squared distance function d by

$$d\colon U \to \mathbb{R}, \quad p \mapsto \min_{q \in X} \|p - q\|^2.$$

For $p \notin X$, a *best X-approximation* of p is a minimizer $q \in X$ that attains $\min_{q \in X} \|p - q\|^2$, and it is customary to write $\operatorname{argmin}_{q \in X} \|p - q\|^2$ for the set of all such minimizers. Let

$$\mathcal{C}(X) := \{p \in U \setminus X \mid p \text{ does not have a unique best approximation in } X\}. \tag{8}$$

We have the following subanalytic analogue of [17, Theorem 3.7].

Lemma 4.1. *Let $m, n \in \mathbb{N}$. Let U be an n-dimensional real vector space. For any closed subanalytic set $X \subseteq U$, the set $\mathcal{C}(X)$ has the following two properties:*

(i) *$\mathcal{C}(X) \cap \overline{B}(0, m)$ is subanalytic, where $\overline{B}(0, m) := \{x \in U \mid \|x\| \le m\}$;*
(ii) *$\mathcal{C}(X)$ has dimension strictly less than n.*

Proof. We start by defining the two maps:

$$f\colon U \times \mathbb{R} \times X \to \mathbb{R}, \qquad (p, t, q) \mapsto \|p - q\|^2 - t,$$

$$g\colon U \times \mathbb{R} \times X \times (0,\infty) \to \mathbb{R}, \qquad (p,t,q,\varepsilon) \mapsto t + \varepsilon - \|p-q\|^2,$$

and the two projections:

$$\pi_1\colon U \times \mathbb{R} \times X \to U \times \mathbb{R}, \qquad \pi_2\colon U \times \mathbb{R} \times X \times (0,\infty) \to U \times \mathbb{R} \times X.$$

Let Γ be the graph of d. Then

$$\Gamma = \big[U \times [0,\infty)\big] \cap \big[U \times \mathbb{R} \setminus \pi_1\big(f^{-1}(-\infty,0)\big)\big] \cap \big[\pi_1\big(U \times \mathbb{R} \times X \setminus \pi_2(g^{-1}(-\infty,0])\big)\big]$$

because

$$\begin{aligned}
(p,t) \in U \times \mathbb{R} \setminus \pi_1\big(f^{-1}(-\infty,0)\big) &\iff \|p-q\|^2 \geq t \text{ for all } q \in X,\\
(p,t) \in \pi_1\big(U \times \mathbb{R} \times X \setminus \pi_2(g^{-1}(-\infty,0])\big) &\iff \text{for any } \varepsilon > 0,\ \text{there is a } q \in X \text{ with } \|p-q\|^2 < t + \varepsilon.
\end{aligned}$$

Therefore $d|_{\overline{B}(0,m)}$, the squared distance function restricted to $\overline{B}(0,m)$, is a subanalytic function for each $m \in \mathbb{N}$. By [17, Theorem 2.1], the set $\mathcal{C}(X)$ in (8) comprises precisely the nonsmooth points:

$$\mathcal{C}(X) = \{p \in U \setminus X \mid d \text{ is not differentiable at } p\}.$$

Hence, by Lemma 3.1, $\mathcal{C}(X) \cap \overline{B}(0,m)$ is a subanalytic subset with dimension less than n. Since

$$\mathcal{C}(X) = \bigcup_{m=1}^{\infty} \big(\mathcal{C}(X) \cap \overline{B}(0,m)\big),$$

$\mathcal{C}(X)$ must also have dimension less than n. $\square$

5. Best approximation by points in a closed complex variety

We will now switch our discussion from $\mathbb{R}$ to $\mathbb{C}$. Let V be an n-dimensional complex vector space with a Hermitian inner product $\langle\cdot,\cdot\rangle$ and corresponding ℓ^2-norm $\|\cdot\|$. For any closed irreducible complex analytic variety $X \subseteq V$, again we let $\mathcal{C}(X)$ denote the set of points which do not have a unique best approximation in X, i.e., with V in place of U in (8). Since V may be regarded as a real vector space of real dimension $2n$, X is naturally a real analytic variety. In particular, best X-approximations are as defined in Section 4.

By Lemma 4.1, $\mathcal{C}(X)$ has real dimension strictly less than $2n$ and we easily deduce the following.

Corollary 5.1. *Let X be a closed complex analytic variety in a complex vector space V. Then a general $p \in V$ has a unique best X-approximation.*

For any $p \in V \setminus (\mathcal{C}(X) \cup X)$, p has a unique best X-approximation, which we will denote by $\pi_X(p)$. Thus this gives us a map $\pi_X\colon V \setminus (\mathcal{C}(X) \cup X) \to X$ that sends p to its unique best X-approximation $\pi_X(p)$, i.e.,

$$\|p - \pi_X(p)\|^2 = \min_{q \in X} \|p-q\|^2.$$

The map π_X is also clearly subanalytic when restricted to $\overline{B}(0,m)$. With this observation, we will state our main result.

Theorem 5.2. *Let X be a closed irreducible complex analytic variety in V. Let $Z \subseteq X$ be any complex analytic subvariety with* $\dim Z < \dim X$. *For a general $p \in V$, its unique best X-approximation $\pi_X(p)$ does not lie in Z.*

Before proving Theorem 5.2, it will be instructive to study the simplest case where $\dim X = 1$ and $\dim Z = 0$, which will provide us with an intuitive idea as to why Theorem 5.2 holds; more importantly, it will illustrate the difference between working over $\mathbb{C}$ and working over $\mathbb{R}$.

Lemma 5.3. *Let $X \subseteq \mathbb{C}^n$ be an irreducible complex analytic curve and $Z \subseteq X$ be a zero-dimensional subvariety. Then for a general $p \in \mathbb{C}^n$, its unique best X-approximation $\pi_X(p)$ does not lie in Z.*

Proof. We proceed by contradiction. Suppose there is a nonempty open neighborhood $B(p, \varepsilon)$ of p such that $\pi_X(B(p, \varepsilon)) \subseteq Z$. Since Z is a collection of points of X, we may assume that $0 \in Z$ and $\pi_X(B(p, \varepsilon)) = 0$, i.e.,

$$|x_1|^2 + \cdots + |x_n|^2 \le |x_1 - y_1|^2 + \cdots + |x_n - y_n|^2$$

for all $x = (x_1, \ldots, x_n) \in B(p, \varepsilon)$ and all $y = (y_1, \ldots, y_n) \in X$. By applying elimination theory, in some nonempty open neighborhood $B(0, \eta) \cap X$ of 0 we may assume $(y_1(t), \ldots, y_n(t))$ is a local parametrization of X such that $y_1(0) = \cdots = y_n(0) = 0$, and each $y_j(t)$ is holomorphic in t around 0, i.e.,

$$y_j(t) = \sum_{k=m_j}^{\infty} a_{j,k} t^k,$$

with $m_j \ge 1$, $a_{j,m_j} \ne 0$, $j = 1, \ldots, n$. Let

$$f(x, t) \coloneqq \sum_{j=1}^{n} \bigl(|x_j|^2 - |x_j - y_j(t)|^2\bigr)$$

and let $m \coloneqq \min\{m_1, \ldots, m_n\}$. Then

$$f(x, t) = 2 \sum_{j=1}^{n} \bigl(\operatorname{Re}(a_{j,m_j} t^{m_j} \overline{x}_j) + O(t^{m_j+1})\bigr) = \operatorname{Re}\bigl(t^m g(x)\bigr) + O(t^{m+1}),$$

where $g(x) \coloneqq \sum_{j:m_j=m} 2a_{j,m_j} \overline{x}_j$, i.e., a sum over all j such that $m_j = m$, and $O(t^{m+1})$ denotes terms of degrees $\ge m + 1$. Note that $g(x) \ne 0$ for a general x. Let $t^m = \lambda \overline{g(x)}$ for some small $\lambda > 0$. Then $f(x, t) > 0$, which contradicts our assumption that $f(x, t) \le 0$ for all $x = (x_1, \ldots, x_n) \in B(p, \varepsilon)$ and all $y = (y_1, \ldots, y_n) \in X$. So $\pi_X(B(p, \varepsilon))$ cannot be just 0. $\square$

An important distinction between $\mathbb{R}$ and $\mathbb{C}$ is that Lemma 5.3 does not hold over $\mathbb{R}$. An example where

$$\{p = (x_1, x_2) \in \mathbb{R}^2 \mid q_* \in Z \text{ for all } q_* \in \operatorname{argmin}_{q \in X} \|p - q\|\}$$

contains a nonempty open subset is illustrated in Fig. 2. Here X is the real cuspidal curve $\{(x_1, x_2) \in \mathbb{R}^2 \mid x_2^2 = x_1^3\}$, and Z is the cusp $(0, 0)$ in X. We can see for any p_0 on the negative x_1-axis there is a small open neighborhood $B(p_0, \varepsilon)$ such that for any $p \in B(p_0, \varepsilon)$ the best approximation $\pi_X(p)$ of p in X is Z. This should be contrasted with the complex cuspidal curve $\{(x_1, x_2) \in \mathbb{C}^2 \mid x_2^2 = x_1^3\}$, which is a real surface in $\mathbb{C}^2 \simeq \mathbb{R}^4$. The projection to $\mathbb{R}^3$ of this curve is the surface[3] presented in Fig. 3. The real points of the curve form the black cuspidal curve. The crucial observation here is that the complex points of the curve prevent the cusp from being an exposed point, and thus generally preventing it from being a candidate for a best approximant.

[3] The apparent one-dimensional singular locus is an artifact of the projection; the complex curve has only one singular point.

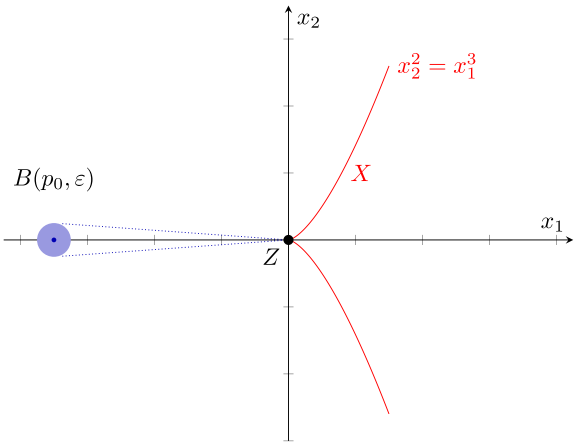

Fig. 2. Over $\mathbb{R}$, singular point is 'exposed.'

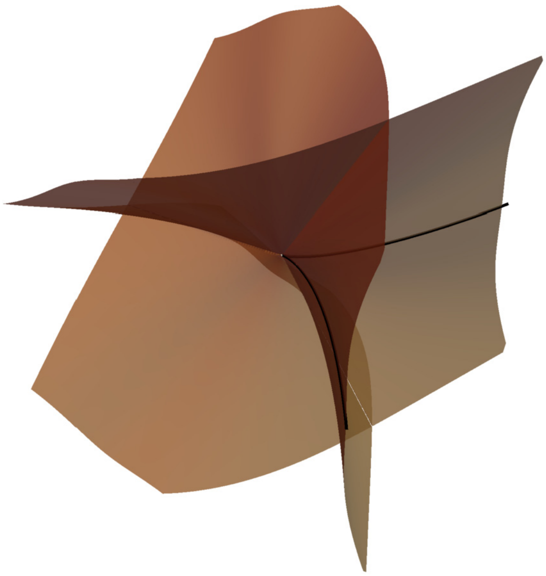

Fig. 3. Over $\mathbb{C}$, singular point is no longer 'exposed.' Figure shows $\mathbb{R}^3$ projection of complex cuspidal curve $x_2^2 = x_1^3$ in $\mathbb{C}^2 \simeq \mathbb{R}^4$.

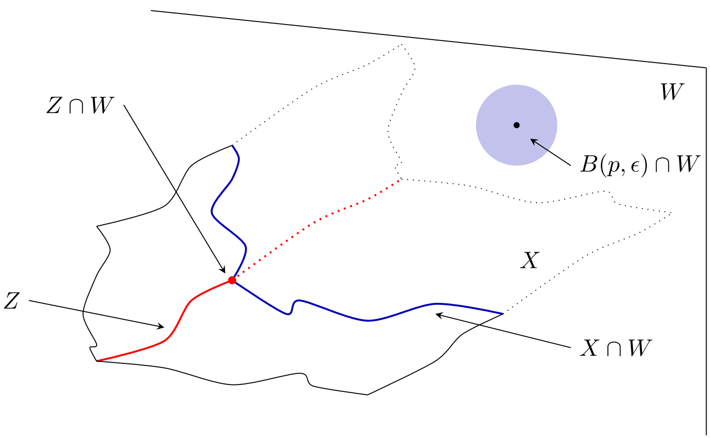

Fig. 4. Pictorial illustration of the proof of Theorem 5.2.

Now we are in a position to prove Theorem 5.2. A pictorial illustration of the proof is given in Fig. 4.

Proof of Theorem 5.2. We proceed by contradiction. Suppose there is a nonempty open subset $\mathcal{O} \subseteq V \setminus (\mathcal{C}(X) \cup X)$ such that for any $p \in \mathcal{O}$, the best X-approximation $\pi_X(p)$ lies in Z. Without loss of generality, we may assume that Z is a minimum subvariety of X containing $\pi_X(\mathcal{O})$ in the sense that any subvariety containing $\pi_X(\mathcal{O})$ has dimension at least $\dim_{\mathbb{C}} Z$. For simplicity, we choose a basis of V so that we may write $u = (u_1, \dots, u_n)$, $v = (v_1, \dots, v_n) \in V$, and the inner product is given by $\langle u, v\rangle = u_1\overline{v}_1 + \cdots + u_n\overline{v}_n$ under this basis. With this assumption, the norm is $\|v\| = (|v_1|^2 + \cdots + |v_n|^2)^{1/2}$.

First we assume $\dim Z > 0$. By taking a Whitney stratification, there is a nonsingular point $z \in Z$ and an open ball $B(z, \rho) \subseteq V$ around z, such that $B(z, \rho) \cap Z$ is a connected complex submanifold in V, and $\pi_X^{-1}(B(z, \rho) \cap Z)$ contains a nonempty open subset $\mathcal{O}' \subseteq \mathcal{O}$. For notational convenience, we will still denote $B(z, \rho) \cap Z$ by Z and $\mathcal{O}'$ by $\mathcal{O}$.

Given $p \in \mathcal{O}$, let $B(p, \varepsilon) \subseteq \mathcal{O}$ be a small open ball of p, and let $q = \pi_X(p)$. Then

$$\frac{d}{dt}\langle p - q(t), p - q(t)\rangle\Big|_{t=0} = 0$$

for any real analytic curve $q(t) \subseteq Z$ with $q(0) = q$, which implies that

$$\mathrm{Re}\langle p - q, v\rangle = 0 \quad \text{for any } v \in \mathrm{T}_q(Z),$$

where $\mathrm{T}_q(Z)$ is the tangent space of Z at q. Without loss of generality, let $q = 0 \in V$, and let

$$W := \{x \in V \mid \mathrm{Re}\langle x - q, v\rangle = 0 \text{ for any } v \in \mathrm{T}_q(Z)\}.$$

Since $\mathrm{T}_q(Z) = \mathrm{T}_0(Z)$ is a complex vector space, we in fact have

$$W = \{x \in V \mid \langle x, v\rangle = 0 \text{ for any } v \in \mathrm{T}_0(Z)\},$$

i.e., W is a complex vector subspace of V whose complex codimension is $\dim_{\mathbb{C}} Z$. Thus $W \cap Z \cap B(0, \delta) = \{0\}$ for some small open ball $B(0, \delta) \subseteq V$. Since X is irreducible and the set of nonsingular points of X is dense in X, $(X \setminus Z) \cap B(0, \delta) \neq \varnothing$. By Theorem 3.2, Z has a small tubular neighborhood in X. Thus $W \cap X \cap B(0, \delta)$ is of positive dimension. Hence it suffices to show that $\pi_{W\cap X\cap B(0,\delta)}(W \cap B(p, \varepsilon))$ cannot be just $\{0\}$.

We have reduced the problem to showing that in the complex vector space W, given a positive-dimensional complex analytic variety $W \cap X \cap B(0, \delta)$ and a point $0 \in W \cap X \cap B(0, \delta)$,

$$\pi_{W\cap X\cap B(0,\delta)}(W \cap B(p, \varepsilon)) \neq \{0\}$$

for a nonempty open ball $W \cap B(p, \varepsilon)$ in W. Again for notational simplicity, we will still use X to denote $W \cap X \cap B(0, \delta)$, and $B(p, \varepsilon)$ to denote $W \cap B(p, \varepsilon)$. One should note that this is in fact the case $\dim Z = 0$.

Pick an irreducible complex analytic curve $C \subseteq X$ passing through 0. By Lemma 5.3, we have that $\pi_C(B(p, \varepsilon)) \neq \{0\}$, and thus $\pi_X(B(p, \varepsilon)) \neq \{0\}$. $\square$

Since the singular locus of a complex analytic variety is a complex analytic subvariety of smaller dimension [18, Chapter 0], we obtain the following corollary of Theorem 5.2. Given its important implications in this article, we label it a theorem.

Theorem 5.4. *Let X be a closed irreducible complex analytic variety in V. For a general $p \in V$, its best X-approximation $\pi_X(p)$ is a nonsingular point of X.*

In fact, the requirement in Theorem 5.4 that X is irreducible can be relaxed, allowing for reducible varieties, and even constructible sets.

Corollary 5.5. *Let $X \subseteq V$ be a closed complex analytic variety, and $X = X_1 \cup \cdots \cup X_k$ be the decomposition of X into irreducible components. Let $Z \subseteq X$ be any complex analytic subvariety such that $\dim Z < \dim X$ and $Z \cap X_j$ is not Zariski dense in X_j for all $j = 1, \ldots, k$. Then for a general $p \in V$, its unique best X-approximation $\pi_X(p)$ does not lie in Z.*

Proof. Apply Theorem 5.2 to each X_j and $Z \cap X_j$. □

Corollary 5.6. *Let $f : X \to V$ be a morphism from a complex affine, possibly reducible, algebraic variety X to a complex vector space V. Then for a general $p \in V$, its unique best $\overline{f(X)}$-approximation $\pi_{\overline{f(X)}}(p)$ lies in $f(X)$.*

Proof. By Chevalley's theorem, $f(X)$ is a constructible set [19, Theorem 3.16]. We set $Z := \overline{\overline{f(X)} \setminus f(X)}$ and apply Corollary 5.5. □

6. Applications to nonlinear approximations

We will now apply the results in Section 5 back to the functional approximation problems described in Section 1. Since our results require that we work over $\mathbb{C}$ (in fact, they are false over $\mathbb{R}$), we will write $L^2(\Omega) = L^2_{\mathbb{C}}(\Omega)$ in the following.

We caution the reader that when applied to the best r-term approximation problem in (1), the set X in Theorem 5.4 is *not* the set of r-term approximant

$$\Sigma_r^\circ(D) = \{f \in L^2(\Omega) \mid f = f_1 + \cdots + f_r \text{ for some } f_1, \ldots, f_r \in D\},$$

but the closure of this set, which we denoted by $\Sigma_r(D)$. Since it is a closure, $X = \Sigma_r(D)$ is automatically closed; and requiring it be semianalytic is a very mild condition. In fact, we are unaware of any nonpathological example where this does not hold.

Corollary 6.1. *Let $f \in L^2(\Omega)$ and $D \subseteq L^2(\Omega)$. If $\Sigma_r(D)$ is semianalytic, then the best r-term approximation problem*

$$\inf_{f_1,\ldots,f_r \in D} \|f - (f_1 + \cdots + f_r)\| = \inf_{\varphi \in \Sigma_r^\circ(D)} \|f - \varphi\| \tag{9}$$

has a solution with probability one, i.e., the infimum is attained by a point in $\Sigma_r^\circ(D)$ with probability one.

Proof. The infimum in (9) may not be attained over $\Sigma_r^\circ(D)$ because this is in general not a closed set. However, the infimum must always be attained in its closure $\Sigma_r(D)$ even when this set is unbounded — the reason being that for a fixed f, the map $\Sigma_r(D) \to \mathbb{R}$, $\varphi \to \|f - \varphi\|$ is a *coercive* map [20] and has bounded sublevel sets. Now by Theorem 5.4, the minimizer $\varphi_* \in \Sigma_r(D)$ is a smooth point with probability one and therefore we have $\varphi_* \in \Sigma_r^\circ(D)$ with probability one. □

We will next apply our result to the join approximation problem

$$\min_{f_{i1},\ldots,f_{ir_i} \in D_i} \|f - (f_{11} + \cdots + f_{1r_1}) - (f_{21} + \cdots + f_{2r_2}) - \cdots - (f_{k1} + \cdots + f_{kr_k})\|. \tag{10}$$

Note that this is a generalization of (2) to $k \geq 2$.

Let $X_1, \dots, X_k \subseteq L^2(\Omega)$ and define their *join* as

$$J(X_1, \dots, X_k) := \{\varphi_1 + \cdots + \varphi_k \in L^2(\Omega) \mid \varphi_i \in X_i,\ i = 1, \dots, k\}. \tag{11}$$

We observe that (10) may be reexpressed in the form (12) below with $X_i = \Sigma^\circ_{r_i}(D_i)$, $i = 1, \dots, k$.

Corollary 6.2. *Let $f \in L^2(\Omega)$ and $D_1, \dots, D_k \subseteq L^2(\Omega)$. If the join $J\big(\Sigma_{r_1}(D_1), \dots, \Sigma_{r_k}(D_k)\big)$ is semianalytic, then the join approximation problem*

$$\inf_{\varphi_1 \in \Sigma^\circ_{r_1}(D_1), \dots, \varphi_k \in \Sigma^\circ_{r_k}(D_k)} \|f - (\varphi_1 + \cdots + \varphi_r)\| \tag{12}$$

has a solution with probability one, i.e., the infimum is attained by a point in $J\big(\Sigma^\circ_{r_1}(D_1), \dots, \Sigma^\circ_{r_k}(D_k)\big)$ with probability one.

Proof. Observe that $J(X_1, \dots, X_k)$ is an image of the map

$$J : L^2(\Omega) \times \cdots \times L^2(\Omega) \to L^2(\Omega), \quad (x_1, \dots, x_k) \mapsto x_1 + \cdots + x_k,$$

as $J(X_1, \dots, X_k) = J(X_1 \times \cdots \times X_k)$. Since J is a closed map, $J(X_1, \dots, X_k)$ is a closed set if $X_1, \dots, X_k$ are closed sets. Also, $\overline{J(X_1, \dots, X_k)} = J(\overline{X}_1, \dots, \overline{X}_k)$, and so

$$\overline{J\big(\Sigma^\circ_{r_1}(D_1), \dots, \Sigma^\circ_{r_k}(D_k)\big)} = J\big(\Sigma_{r_1}(D_1), \dots, \Sigma_{r_k}(D_k)\big).$$

Applying Theorem 5.4 to $X = J\big(\Sigma_{r_1}(D_1), \dots, \Sigma_{r_k}(D_k)\big)$ then yields the required result. □

Again we stress that it is crucial in Corollaries 6.1 and 6.2 that the functions be complex-valued, the results fail when we replace $\mathbb{C}$ by $\mathbb{R}$. We will discuss specific cases of Corollaries 6.1 and 6.2 in the next two sections.

7. Applications to separable approximations

We will now apply the results in Section 5 to more specific classes of functional approximation problems. As in Section 6, given that our results only hold over $\mathbb{C}$, we write $L^2(\Omega) = L^2_{\mathbb{C}}(\Omega)$ throughout this section. A point to recall is that

$$L^2(\Omega_1 \times \cdots \times \Omega_d) = L^2(\Omega_1) \otimes \cdots \otimes L^2(\Omega_d)$$

and so multivariate functions are in fact tensors. In many, if not most applications, the multivariate target function (i.e., the function to be approximated) possesses various forms of symmetries, the most common being *full* invariance or skew-invariance under all permutations of arguments:

$$f(x_{\tau(1)}, \dots, x_{\tau(d)}) = f(x_1, \dots, x_d) \quad \text{or} \quad g(x_{\tau(1)}, \dots, x_{\tau(d)}) = \operatorname{sgn}(\tau) g(x_1, \dots, x_d),$$

for all $\tau \in \mathfrak{S}_d$. In other words, f is a symmetric tensor and g an alternating tensor:

$$f \in \mathsf{S}^d(L^2(\Omega)) \qquad \text{or} \qquad g \in \mathsf{\Lambda}^d(L^2(\Omega)).$$

Note that in these cases, we need $\Omega_1 = \cdots = \Omega_d = \Omega$ in order to permute arguments. There are many types of *partial* symmetries as well. For example, we may have functions that satisfy

$$f(x_{\rho(1)}, \dots, x_{\rho(d)}, y_{\tau(1)}, \dots, y_{\tau(k)}) = f(x_1, \dots, x_d, y_1, \dots, y_k),$$
$$g(x_{\rho(1)}, \dots, x_{\rho(d)}, y_{\tau(1)}, \dots, y_{\tau(k)}) = \operatorname{sgn}(\rho) \operatorname{sgn}(\tau) g(x_1, \dots, x_d, y_1, \dots, y_k)$$

for all $\rho \in \mathfrak{S}_d$ and $\tau \in \mathfrak{S}_k$. These correspond respectively to functions

$$f \in \mathsf{S}^d(L^2(\Omega_1)) \otimes \mathsf{S}^k(L^2(\Omega_2)) \qquad \text{and} \qquad g \in \mathsf{\Lambda}^d(L^2(\Omega_1)) \otimes \mathsf{\Lambda}^k(L^2(\Omega_2)).$$

There are also more intricate symmetries. For example, a complex-valued function $f(X)$ of a matrix variable $X = (x_{ij})_{i,j=1}^{k,d}$ that satisfies

$$f(x_{\rho_1(1),\tau(1)}, \dots, x_{\rho_1(k),\tau(1)}, \dots, x_{\rho_d(1),\tau(d)}, \dots, x_{\rho_d(k),\tau(d)}) \\ = \operatorname{sgn}(\rho_1) \dots \operatorname{sgn}(\rho_d) f(x_{1,1}, \dots, x_{k,1}, \dots, x_{1,d}, \dots, x_{k,d})$$

for all $\rho_1, \dots, \rho_d \in \mathfrak{S}_k$ and $\tau \in \mathfrak{S}_d$. This corresponds to $f \in \mathsf{S}^d(\mathsf{\Lambda}^k(L^2(\Omega)))$. We will see more examples later.

Each of these classes of functions comes with its own nonlinear r-term approximation problem that preserves the respective symmetries or skew-symmetries, full or partial. For example, for $f \in L^2(\Omega \times \dots \times \Omega)$, we may just want

$$\inf \left\| f - \sum_{i=1}^{r} \varphi_{1,i} \otimes \dots \otimes \varphi_{d,i} \right\|, \tag{13}$$

as discussed in Section 1, but for a symmetric $f \in \mathsf{S}^d(L^2(\Omega))$ or an alternating $g \in \mathsf{\Lambda}^d(L^2(\Omega))$, the more natural approximation problems are respectively

$$\inf \left\| f - \sum_{i=1}^{r} \varphi_{1,i} \circ \dots \circ \varphi_{d,i} \right\| \qquad \text{or} \qquad \inf \left\| g - \sum_{i=1}^{r} \varphi_{1,i} \wedge \dots \wedge \varphi_{d,i} \right\|, \tag{14}$$

with $\varphi_{j,i} \in L^2(\Omega)$, $i = 1, \dots, r$, $j = 1, \dots, d$. In fact, there are other natural options for a symmetric f, where we may instead seek an approximation of the form

$$\inf \left\| f - \sum_{i=1}^{r} \varphi_i^{\otimes d} \right\|, \tag{15}$$

with $\varphi_i \in L^2(\Omega)$, $i = 1, \dots, r$. The symbols $\circ$ and $\wedge$ denote symmetric product and exterior product respectively; for those who have not encountered them, these and other notations will be defined in Section 7.1. Then in Section 7.2, we will describe other classes of r-term approximation problems for functions possessing more complicated symmetries than the ones in (13), (14), (15). All these approximation problems have two features in common:

(i) Each of them is associated with a complex algebraic variety; e.g., (13) with the Segre variety, (14) with the Chow variety and Grassmann variety respectively, (15) with the Veronese variety.
(ii) All of them may fail to have a solution; e.g., there exist target functions where the infima in (13), (14), (15) cannot be attained.

We will later see how our results in Section 5 combined with (i) allow us to rectify (ii) to the extent guaranteed by Theorem 5.2, i.e., for any general f.

We would like to mention a slightly different, perhaps more common, alternative for framing the function approximation problems above. As we described in Section 1, for computational purposes, $L^2(\Omega_i)$'s are all finite-dimensional (e.g., when Ω_i's are all finite sets) and in this case, $L^2(\Omega_i) \cong \mathbb{C}^{n_i}$ with a choice of basis, and thus

$$\begin{aligned}
L^2(\Omega_1 \times \Omega_2) &= L^2(\Omega_1) \otimes L^2(\Omega_2) \cong \mathbb{C}^{n_1} \otimes \mathbb{C}^{n_2} \cong \mathbb{C}^{n_1 \times n_2}, \\
L^2(\Omega_1 \times \Omega_2 \times \Omega_3) &= L^2(\Omega_1) \otimes L^2(\Omega_2) \otimes L^2(\Omega_3) \cong \mathbb{C}^{n_1} \otimes \mathbb{C}^{n_2} \otimes \mathbb{C}^{n_3} \cong \mathbb{C}^{n_1 \times n_2 \times n_3}, \\
&\vdots \\
L^2(\Omega_1 \times \cdots \times \Omega_d) &= L^2(\Omega_1) \otimes \cdots \otimes L^2(\Omega_d) \cong \mathbb{C}^{n_1} \otimes \cdots \otimes \mathbb{C}^{n_d} \cong \mathbb{C}^{n_1 \times \cdots \times n_d}.
\end{aligned}$$

The set $\mathbb{C}^{n_1 \times \cdots \times n_d}$ denotes the vector space of d-dimensional hypermatrices

$$(a_{i_1,\dots,i_d})_{i_1,\dots,i_d=1}^{n_1,\dots,n_d}$$

(when $d = 2$, this reduces to a usual matrix in linear algebra), which of course is nothing more than a convenient way to represent a function $f : \Omega_1 \times \cdots \times \Omega_d \to \mathbb{C}$ by storing its value

$$f(i_1, \dots, i_d) = a_{i_1,\dots,i_d}$$

at $(i_1, \dots, i_d) \in \Omega_1 \times \cdots \times \Omega_d$, assuming that Ω_i's are all finite sets. In fact, practitioners in computational mathematics overwhelmingly regard a tensor as such a hypermatrix. The symmetries described earlier carry verbatim to hypermatrices with the permutations acting on the indices.

We may either frame our results in this section in the form of function approximations or in the form of tensor approximations (or more accurately, matrix/hypermatrix approximations), but instead of favoring one over the other, we would simply resort to stating them for abstract vector spaces. So implicitly, $V = L^2(\Omega_1 \times \cdots \times \Omega_d)$ for function approximations and $V = \mathbb{C}^{n_1 \times \cdots \times n_d}$ for tensor approximations.

7.1. Segre/Veronese/Grassmann varieties and friends

We begin by reviewing some basic tensor constructions. The varieties that we define will be subsets of tensor spaces constructed via one or more of the following ways. We write $V_1 \otimes \cdots \otimes V_d$ for the tensor product of vector spaces $V_1, \dots, V_d$. A rank-one tensor is a nonzero decomposable tensor $v_1 \otimes \cdots \otimes v_d$, where $v_i \in V_i$. When $V_1 = \cdots = V_d = V$, we use the abbreviations $V^{\otimes d} = V \otimes \cdots \otimes V$ and $v^{\otimes d} = v \otimes \cdots \otimes v$. In this case, the symmetric group $\mathfrak{S}_d$ acts on rank-one tensors by

$$\tau(v_1 \otimes \cdots \otimes v_d) = v_{\tau(1)} \otimes \cdots \otimes v_{\tau(d)},$$

and this action extends linearly to an action on $V^{\otimes d}$. The subspaces

$$\begin{aligned}
\mathsf{S}^d(V) &= \{T \in V^{\otimes d} \mid \tau(T) = T \text{ for all } \tau \in \mathfrak{S}_d\}, \\
\mathsf{\Lambda}^d(V) &= \{T \in V^{\otimes d} \mid \tau(T) = \operatorname{sgn}(\tau) T \text{ for all } \tau \in \mathfrak{S}_d\},
\end{aligned}$$

are called the spaces of *symmetric d-tensors* and *alternating d-tensors* respectively. We may construct tensor spaces with more intricate symmetries and skew-symmetries by combining these, e.g., $\mathsf{S}^d(V) \otimes \mathsf{S}^k(V)$, $\mathsf{\Lambda}^d(V) \otimes \mathsf{\Lambda}^k(V)$, $\mathsf{S}^d(\mathsf{\Lambda}^k(V))$, $\mathsf{S}^d(\mathsf{S}^k(V))$, etc.

The *symmetric product* and *alternating product* of $v_1, \dots, v_d \in V$ are respectively defined by

$$v_1 \circ \cdots \circ v_d := \frac{1}{d!} \sum_{\tau \in \mathfrak{S}_d} v_{\tau(1)} \otimes \cdots \otimes v_{\tau(d)} \in \mathsf{S}^d(V),$$
$$v_1 \wedge \cdots \wedge v_d := \frac{1}{d!} \sum_{\tau \in \mathfrak{S}_d} \operatorname{sgn}(\tau) v_{\tau(1)} \otimes \cdots \otimes v_{\tau(d)} \in \mathsf{\Lambda}^d(V).$$

In finite dimensions, each of the approximation problems we have encountered earlier and yet others that we will see in Section 7.2 comes with an associated complex algebraic variety. These varieties are usually defined as complex projective varieties, i.e., subsets of projective spaces defined by the common zero loci of a finite collection of homogeneous polynomials, and we will not deviate from this standard practice in our definitions below:

Segre variety. This is the image $\sigma(\mathbb{P}V_1 \times \cdots \times \mathbb{P}V_d)$ of the Segre embedding:

$$\sigma \colon \mathbb{P}V_1 \times \cdots \times \mathbb{P}V_d \to \mathbb{P}(V_1 \otimes \cdots \otimes V_d), \quad ([v_1], \dots, [v_d]) \mapsto [v_1 \otimes \cdots \otimes v_d].$$

Veronese variety. This is the image $\nu(\mathbb{P}V)$ of the Veronese embedding:

$$\nu \colon \mathbb{P}V \to \mathbb{P}\mathsf{S}^d(V), \quad [v] \mapsto [v^{\otimes d}].$$

Chow variety. This is the image[4] $\mathrm{Ch}_d(V) := \kappa\big((\mathbb{P}V)^d\big)$ of the Chow map:

$$\kappa \colon \mathbb{P}V \times \cdots \times \mathbb{P}V \to \mathbb{P}\mathsf{S}^d(V), \quad ([v_1], \dots, [v_d]) \mapsto [v_1 \circ \cdots \circ v_d].$$

Grassmann variety. This is the image $\psi\big(\mathrm{Gr}_k(V)\big)$ of the *Grassmannian* $\mathrm{Gr}_k(V)$, the set of k-dimensional linear subspaces of V, under the *Plücker embedding*:

$$\psi \colon \mathrm{Gr}_k(V) \to \mathbb{P}\mathsf{\Lambda}^k(V), \quad \operatorname{span}\{v_1, \dots, v_k\} \mapsto [v_1 \wedge \cdots \wedge v_k].$$

Segre–Veronese variety. This is the image $\sigma\nu(\mathbb{P}V_1 \times \cdots \times \mathbb{P}V_m)$ of the Segre–Veronese embedding:

$$\sigma\nu \colon \mathbb{P}V_1 \times \cdots \times \mathbb{P}V_m \to \mathbb{P}\big(\mathsf{S}^{d_1}(V) \otimes \cdots \otimes \mathsf{S}^{d_m}(V)\big), \quad ([v_1], \dots, [v_m]) \mapsto [v_1^{\otimes d_1} \otimes \cdots \otimes v_m^{\otimes d_m}].$$

Segre–Chow variety. This is the image $\sigma\kappa\big(\mathrm{Ch}_{d_1}(V) \times \cdots \times \mathrm{Ch}_{d_m}(V)\big)$ of the Segre–Chow map:

$$\sigma\kappa \colon \mathrm{Ch}_{d_1}(V) \times \cdots \times \mathrm{Ch}_{d_m}(V) \to \mathbb{P}\big(\mathsf{S}^{d_1}(V) \otimes \cdots \otimes \mathsf{S}^{d_m}(V)\big),$$
$$([v_{1,1} \circ \cdots \circ v_{d_1,1}], \dots, [v_{1,m} \circ \cdots \circ v_{d_m,m}]) \mapsto [(v_{1,1} \circ \cdots \circ v_{d_1,1}) \otimes \cdots \otimes (v_{1,m} \circ \cdots \circ v_{d_m,m})].$$

Segre–Grassmann variety. This is the image $\sigma\psi\big(\mathrm{Gr}_{k_1}(V) \times \cdots \times \mathrm{Gr}_{k_m}(V)\big)$ of the Segre–Grassmann map:

$$\sigma\psi \colon \mathrm{Gr}_{k_1}(V) \times \cdots \times \mathrm{Gr}_{k_m}(V) \to \mathbb{P}\big(\mathsf{\Lambda}^{k_1}(V) \otimes \cdots \otimes \mathsf{\Lambda}^{k_m}(V)\big),$$
$$([v_{1,1} \wedge \cdots \wedge v_{k_1,1}], \dots, [v_{1,m} \wedge \cdots \wedge v_{k_m,m}]) \mapsto [(v_{1,1} \wedge \cdots \wedge v_{k_1,1}) \otimes \cdots \otimes (v_{1,m} \wedge \cdots \wedge v_{k_m,m})].$$

Veronese–Chow variety. This is the image $\nu\kappa\big(\mathrm{Ch}_d(V)\big)$ of the Veronese–Chow map:

[4] This is usually called the *Chow variety of zero cycles* in $\mathbb{P}V^*$. It is a subvariety of $\mathbb{P}\mathsf{S}^d(V)$ whose affine cone is the set of degree-d homogeneous polynomials on the dual space V^* that can be decomposed into a product of linear forms [21, Chapter 4, Proposition 2.1].

$$\nu\kappa\colon \operatorname{Ch}_d(V) \to \mathbb{P}\big(\mathsf{S}^k(\mathsf{S}^d(V))\big), \quad [v_1 \circ \cdots \circ v_d] \mapsto [(v_1 \circ \cdots \circ v_d)^{\otimes k}].$$

Veronese–Grassmann variety. This is the image $\nu\psi\big(\operatorname{Gr}_k(V)\big)$ of the Veronese–Grassmann map:

$$\nu\psi\colon \operatorname{Gr}_k(V) \to \mathbb{P}\big(\mathsf{S}^d(\mathsf{\Lambda}^k(V))\big), \quad [v_1 \wedge \cdots \wedge v_k] \mapsto [(v_1 \wedge \cdots \wedge v_k)^{\otimes d}].$$

Segre–Veronese–Chow variety. This is the image $\sigma\nu\kappa\big(\operatorname{Ch}_{d_1}(V)\times\cdots\times\operatorname{Ch}_{d_m}(V)\big)$ of the Segre–Veronese–Chow map:

$$\sigma\nu\kappa\colon \operatorname{Ch}_{d_1}(V) \times \cdots \times \operatorname{Ch}_{d_m}(V) \to \mathbb{P}\big(\mathsf{S}^{k_1}(\mathsf{S}^{d_1}(V)) \otimes \cdots \otimes \mathsf{S}^{k_m}(\mathsf{S}^{d_m}(V))\big),$$
$$([v_{1,1} \circ \cdots \circ v_{d_1,1}], \dots, [v_{1,m} \circ \cdots \circ v_{d_m,m}]) \mapsto [(v_{1,1} \circ \cdots \circ v_{d_1,1})^{\otimes k_1} \otimes \cdots \otimes (v_{1,m} \circ \cdots \circ v_{d_m,m})^{\otimes k_m}].$$

Segre–Veronese–Grassmann variety. This is $\sigma\nu\psi\big(\operatorname{Gr}_{k_1}(V) \times \cdots \times \operatorname{Gr}_{k_m}(V)\big)$, image of the Segre–Veronese–Grassmann map:

$$\sigma\nu\psi\colon \operatorname{Gr}_{k_1}(V) \times \cdots \times \operatorname{Gr}_{k_m}(V) \to \mathbb{P}\big(\mathsf{S}^{d_1}(\mathsf{\Lambda}^{k_1}(V)) \otimes \cdots \otimes \mathsf{S}^{d_m}(\mathsf{\Lambda}^{k_m}(V))\big),$$
$$([v_{1,1} \wedge \cdots \wedge v_{k_1,1}], \dots, [v_{1,m} \wedge \cdots \wedge v_{k_m,m}]) \mapsto [(v_{1,1} \wedge \cdots \wedge v_{k_1,1})^{\otimes d_1} \otimes \cdots \otimes (v_{1,m} \wedge \cdots \wedge v_{k_m,m})^{\otimes d_m}].$$

Note that $\sigma\nu, \sigma\kappa, \sigma\psi, \nu\kappa, \nu\psi, \sigma\nu\kappa, \sigma\nu\psi$ are all compositions of their respective constituent maps. In fact we may use successive compositions of the embeddings σ and ν to define more complicated algebraic varieties of the same nature; it is not possible to exhaust all such constructions. We refer the readers to [21,19,22] for basic properties of the Segre, Veronese, Grassmann, and Chow varieties from an algebraic geometric perspective, although none of which would be required for our subsequent discussions.

While we have defined all these varieties as projective varieties, it is important to note that for approximation problems, we will need to work in vector spaces rather than projective spaces. As such, when we discuss these varieties in the context of approximations, we would in fact be referring to their *affine cones*: For any $X \subseteq \mathbb{P}V$, its affine cone is $\widehat{X} \coloneqq \pi^{-1}(X) \cup \{0\}$, where $\pi\colon V \setminus \{0\} \to \mathbb{P}V$ is the quotient map taking a vector space V onto its projective space $\mathbb{P}V$. We will write $[v] \coloneqq \pi(v)$ for the projective equivalence class of $v \in V \setminus \{0\}$.

7.2. Best low-rank approximations of tensors

A variety $X \subseteq \mathbb{P}V$ is said to be *nondegenerate* if X is not contained in any hyperplane. An implication is that its affine cone $\widehat{X}$ would span V, i.e., $\operatorname{span}(\widehat{X}) = V$, and so any $p \in V$ can be expressed as a linear combination $p = \alpha_1 x_1 + \cdots + \alpha_r x_r$ with $x_1, \dots, x_r \in \widehat{X}$. Since, by the definition of affine cone, $x \in \widehat{X}$ iff $\alpha x \in \widehat{X}$ for any $\alpha \neq 0$, we may in fact replace the linear combination by a sum, i.e., every $p \in V$ may be expressed as $p = x_1 + \cdots + x_r$ for some $x_1, \dots, x_r \in \widehat{X}$.

Given an irreducible nondegenerate complex projective variety $X \subseteq \mathbb{P}V$, the X-rank of a nonzero $p \in V$ is defined as

$$\operatorname{rank}_X(p) \coloneqq \min\{r \in \mathbb{N} : p = x_1 + \cdots + x_r,\ x_i \in \widehat{X}\},$$

and $\operatorname{rank}_X(0) \coloneqq 0$. The *$X$-border rank* of p, denoted by $\overline{\operatorname{rank}}_X(p)$, is the minimum integer r such that p is a limit of a sequence of X-rank-r points. For irreducible nondegenerate complex projective varieties $X_1, \dots, X_r \subseteq \mathbb{P}V$, the *join map* is defined by

$$J : \widehat{X}_1 \times \cdots \times \widehat{X}_r \to V, \quad (x_1, \dots, x_r) \mapsto x_1 + \cdots + x_r.$$

The Euclidean closure of the image $J(\widehat{X}_1 \times \cdots \times \widehat{X}_r)$ in V is called the *join variety* of $X_1, \ldots, X_r$. In particular, when $X_1 = \cdots = X_r = X$, we denote the image of J by $\Sigma_r^\circ(X)$, and its Euclidean closure by $\Sigma_r(X)$, often called the *r-secant variety* of X. Note that

$$\Sigma_r^\circ(X) = \{p \in V : \operatorname{rank}_X(p) \le r\} \qquad \text{and} \qquad \Sigma_r(X) = \{p \in V : \overline{\operatorname{rank}}_X(p) \le r\}.$$

A notion closely related to secant varieties is that of a *tangent variety*, defined for a nonsingular projective variety, or more generally, for any projective variety as

$$\tau(X) := \bigcup_{x \in X} \widehat{\mathrm{T}}_x(X) \qquad \text{or} \qquad \tau(X) := \bigcup_{x \in X \setminus X_{\mathrm{sing}}} \widehat{\mathrm{T}}_x(X)$$

respectively. Here $\widehat{\mathrm{T}}_x(X)$ is the affine tangent space of X at x [22, Section 8.1.1] and X_{sing} is the *singular locus*, i.e., the subvariety of singular points in X, which has measure zero (in fact, positive codimension). When X is nonsingular, $\tau(X)$ is an algebraic variety. When X is singular, $\tau(X)$ is a quasiaffine variety. By abusing terminologies slightly, we call $\tau(X)$ a tangent variety in both cases. Among the varieties in Section 7.1 that we are interested in, the Chow variety $\mathrm{Ch}_d(V)$ is singular when $\dim V > 2$ and thus so are the Segre–Chow, Veronese–Chow, Segre–Veronese–Chow varieties; the rest are all nonsingular varieties.

In fact, for each of the varieties listed in Section 7.1, an element $A \in \tau(X)$ has a *normal form* that is either well-known or straightforward to determine, i.e., we may write down an expression for an arbitrary $A \in \tau(X)$.[5]

Proposition 7.1. *As each variety X in Section 7.1 is defined by a map, we will denote a tensor in the tangent space of that variety with a subscript given by the respective map. With this notation, the following is a list of normal forms for $\tau(X)$:*

$$\begin{aligned}
A_\sigma &= \sum_{i=1}^{d} v_1 \otimes \cdots \otimes v_{i-1} \otimes w_i \otimes v_{i+1} \otimes \cdots \otimes v_d, \\
A_\nu &= v^{\circ(d-1)} \circ w, \\
A_\kappa &= \sum_{i=1}^{d} v_1 \circ \cdots \circ v_{i-1} \circ w_i \circ v_{i+1} \circ \cdots \circ v_d, \\
A_\psi &= \sum_{i=1}^{k} v_1 \wedge \cdots \wedge v_{i-1} \wedge w_i \wedge v_{i+1} \wedge \cdots \wedge v_k, \\
A_{\sigma\nu} &= \sum_{i=1}^{m} v_1^{\otimes d_1} \otimes \cdots \otimes v_{i-1}^{\otimes d_{i-1}} \otimes \big(v_i^{\circ(d_i-1)} \circ w_i\big) \otimes v_{i+1}^{\otimes d_{i+1}} \otimes \cdots \otimes v_m^{\otimes d_m}, \\
A_{\sigma\kappa} &= \sum_{i=1}^{m} (v_{1,1} \circ \cdots \circ v_{d_1,1}) \otimes \cdots \otimes (v_{1,i-1} \circ \cdots \circ v_{d_{i-1},i-1}) \\
&\qquad \otimes \left(\sum_{j=1}^{d_i} v_{1,i} \circ \cdots \circ v_{j-1,i} \circ w_{j,i} \circ v_{j+1,i} \circ \cdots \circ v_{d_i,i} \right) \\
&\qquad\qquad \otimes (v_{1,i+1} \circ \cdots \circ v_{d_{i+1},i+1}) \otimes \cdots \otimes (v_{1,m} \circ \cdots \circ v_{d_m,m}),
\end{aligned}$$

[5] In general this is not possible for higher order tensors, i.e., we do not have an analogue of Jordan normal form for *all* d-tensors when $d > 2$; but in our case, we are restricting to a very small subset, namely, only the tensors in $\tau(X)$.

$$
\begin{aligned}
A_{\sigma\psi} = \sum_{i=1}^{m} &(v_{1,1} \wedge \cdots \wedge v_{k_1,1}) \otimes \cdots \otimes (v_{1,i-1} \wedge \cdots \wedge v_{k_{i-1},i-1}) \\
&\otimes \left(\sum_{j=1}^{k_i} v_{1,i} \wedge \cdots \wedge v_{j-1,i} \wedge w_{j,i} \wedge v_{j+1,i} \wedge \cdots \wedge v_{k_i,i} \right) \\
&\qquad \otimes (v_{1,i+1} \wedge \cdots \wedge v_{k_{i+1},i+1}) \otimes \cdots \otimes (v_{1,m} \wedge \cdots \wedge v_{k_m,m}), \\
A_{\nu\kappa} = (v_1 \circ &\cdots \circ v_d)^{\circ(k-1)} \circ \left(\sum_{i=1}^{d} v_1 \circ \cdots \circ v_{i-1} \circ w_i \circ v_{i+1} \circ \cdots \circ v_d \right), \\
A_{\nu\psi} = (v_1 \wedge &\cdots \wedge v_k)^{\circ(d-1)} \circ \left(\sum_{i=1}^{k} v_1 \wedge \cdots \wedge v_{i-1} \wedge w_i \wedge v_{i+1} \wedge \cdots \wedge v_k \right), \\
A_{\sigma\nu\kappa} = \sum_{i=1}^{m} &(v_{1,1} \circ \cdots \circ v_{d_1,1})^{\otimes k_1} \otimes \cdots \otimes (v_{1,i-1} \circ \cdots \circ v_{d_{i-1},i-1})^{\otimes k_{i-1}} \\
&\otimes \left[(v_{1,i} \circ \cdots \circ v_{d_i,i})^{\circ(k_i-1)} \circ \left(\sum_{j=1}^{d_i} v_{1,i} \circ \cdots \circ v_{j-1,i} \circ w_{j,i} \circ v_{j+1,i} \circ \cdots \circ v_{d_i,i} \right) \right] \\
&\qquad \otimes (v_{1,i+1} \circ \cdots \circ v_{d_{i+1},i+1})^{\otimes k_{i+1}} \otimes \cdots \otimes (v_{1,m} \circ \cdots \circ v_{d_m,m})^{\otimes k_m}, \\
A_{\sigma\nu\psi} = \sum_{i=1}^{m} &(v_{1,1} \wedge \cdots \wedge v_{k_1,1})^{\otimes d_1} \otimes \cdots \otimes (v_{1,i-1} \wedge \cdots \wedge v_{k_{i-1},i-1})^{\otimes d_{i-1}} \\
&\otimes \left[(v_{1,i} \wedge \cdots \wedge v_{k_i,i})^{\circ(d_i-1)} \circ \left(\sum_{j=1}^{k_i} v_{1,i} \wedge \cdots \wedge v_{j-1,i} \wedge w_{j,i} \wedge v_{j+1,i} \wedge \cdots \wedge v_{k_i,i} \right) \right] \\
&\qquad \otimes (v_{1,i+1} \wedge \cdots \wedge v_{k_{i+1},i+1})^{\otimes d_{i+1}} \otimes \cdots \otimes (v_{1,m} \wedge \cdots \wedge v_{k_m,m})^{\otimes d_m}.
\end{aligned}
$$

Proof. Let X be one of the varieties in Section 7.1. Then any $A \in \tau(X)$ is of the form

$$A = \lim_{t\to 0} \frac{\gamma(t) - \gamma(0)}{t},$$

where γ is any complex analytic curve $\gamma(t) \subseteq \widehat{X} \setminus \{0\}$ such that $\gamma(t)$ is nonconstant around 0 and $\gamma(0)$ is a nonsingular point in $\widehat{X}$. In fact, such a curve $\gamma(t)$ can be written down explicitly for any X in Section 7.1. Let $*$ denote either $\circ$ or $\wedge$. Then any point in $\widehat{X} \setminus \{0\}$ takes the form

$$(v_{1,1} * \cdots * v_{k_1,1})^{\otimes d_1} \otimes \cdots \otimes (v_{1,m} * \cdots * v_{k_m,m})^{\otimes d_m},$$

and any curve $\gamma(t) \subset \widehat{X} \setminus \{0\}$ is of the form

$$\gamma(t) = (v_{1,1}(t) * \cdots * v_{k_1,1}(t))^{\otimes d_1} \otimes \cdots \otimes (v_{1,m}(t) * \cdots * v_{k_m,m}(t))^{\otimes d_m},$$

where $v_{j,i}(t) \in V_j$ is a complex analytic curve with $v_{j,i}(0) = v_{j,i}$. Thus

$$
\begin{aligned}
A &= \lim_{t\to 0} \frac{\gamma(t) - \gamma(0)}{t} \\
&= \frac{d}{dt}(v_{1,1}(t) * \cdots * v_{k_1,1}(t))^{\otimes d_1}\Big|_{t=0} \otimes (v_{1,2} * \cdots * v_{k_2,2})^{\otimes d_2} \otimes \cdots \\
&\qquad \otimes (v_{1,m} * \cdots * v_{k_m,m})^{\otimes d_m} + \cdots + (v_{1,1} * \cdots * v_{k_1,1})^{\otimes d_1} \otimes \cdots
\end{aligned}
$$

$$
\begin{aligned}
&\otimes (v_{1,m-1} * \cdots * v_{k_{m-1},m-1})^{\otimes d_{m-1}} \otimes \frac{d}{dt}(v_{1,m}(t) * \cdots * v_{k_m,m}(t))^{\otimes d_m}\Big|_{t=0} \\
= &\sum_{i=1}^{m} (v_{1,1} * \cdots * v_{k_1,1})^{\otimes d_1} \otimes \cdots \otimes (v_{1,i-1} * \cdots * v_{k_{i-1},i-1})^{\otimes d_{i-1}} \\
&\otimes \left[(v_{1,i} * \cdots * v_{k_i,i})^{\circ(d_i-1)} \circ \left(\sum_{j=1}^{k_i} v_{1,i} * \cdots * v_{j-1,i} * w_{j,i} * v_{j+1,i} * \cdots * v_{k_i,i} \right) \right] \\
&\otimes (v_{1,i+1} * \cdots * v_{k_{i+1},i+1})^{\otimes d_{i+1}} \otimes \cdots \otimes (v_{1,m} * \cdots * v_{k_m,m})^{\otimes d_m}
\end{aligned}
$$

for some $w_{j,i} \in V_j$. $\square$

The simplest examples with $\overline{\operatorname{rank}}_X(A) \neq \operatorname{rank}_X(A)$ may be found in $\tau(X)$. Note that since every tangent is a limit of 2-secants, we clearly have $\tau(X) \subseteq \Sigma_2(X)$; but in general $\tau(X) \nsubseteq \Sigma_2^\circ(X)$. The next proposition gives sufficient conditions for a tensor A in Proposition 7.1 so that $A \in \tau(X) \setminus \Sigma_2^\circ(X)$, i.e.,

$$\overline{\operatorname{rank}}_X(A) = 2 < \operatorname{rank}_X(A).$$

As each variety X in Section 7.1 is defined by a map, we will denote the X-rank and border X-rank of a variety with a subscript given by the respective map.

Proposition 7.2. *Let $d, k, m \geq 3$. The tensors in $\tau(X)$ for any X that is one of the varieties in Section 7.1 have X-border-rank two and X-rank strictly greater than two given the respective sufficient conditions:*

(i) *If $\{v_i, w_i\}$ is linearly independent for $i = 1, \dots, d$, then $\overline{\operatorname{rank}}_\sigma(A_\sigma) = 2 < \operatorname{rank}_\sigma(A_\sigma)$.*
(ii) *If $\{v, w\}$ is linearly independent, then $\overline{\operatorname{rank}}_\nu(A_\nu) = 2 < \operatorname{rank}_\nu(A_\nu)$.*
(iii) *If $[v_1], \dots, [v_d], [w_1], \dots, [w_d]$ are distinct in $\mathbb{P}V$, then $\overline{\operatorname{rank}}_\kappa(A_\kappa) = 2 < \operatorname{rank}_\kappa(A_\kappa)$.*
(iv) *If $v_1 \wedge \cdots \wedge v_k \wedge w_1 \wedge \cdots \wedge w_k \neq 0$, then $\overline{\operatorname{rank}}_\psi(A_\psi) = 2 < \operatorname{rank}_\psi(A_\psi)$.*
(v) *If $[v_1], \dots, [v_m], [w_1], \dots, [w_m]$ are distinct in $\mathbb{P}V$, then $\overline{\operatorname{rank}}_{\sigma\nu}(A_{\sigma\nu}) = 2 < \operatorname{rank}_{\sigma\nu}(A_{\sigma\nu})$.*
(vi) *If $[v_{1,1}], \dots, [v_{d_m,m}], [w_{1,1}], \dots, [w_{d_m,m}]$ are distinct in $\mathbb{P}V$, then $\overline{\operatorname{rank}}_{\sigma\kappa}(A_{\sigma\kappa}) = 2 < \operatorname{rank}_{\sigma\kappa}(A_{\sigma\kappa})$.*
(vii) *If $v_{1,i} \wedge \cdots \wedge v_{k_i,i} \wedge w_{1,i} \wedge \cdots \wedge w_{k_i,i} \neq 0$ for $i = 1, \dots, m$, then $\overline{\operatorname{rank}}_{\sigma\psi}(A_{\sigma\psi}) = 2 < \operatorname{rank}_{\sigma\psi}(A_{\sigma\psi})$.*
(viii) *If $[v_1], \dots, [v_d], [w_1], \dots, [w_d]$ are distinct in $\mathbb{P}V$, then $\overline{\operatorname{rank}}_{\nu\kappa}(A_{\nu\kappa}) = 2 < \operatorname{rank}_{\nu\kappa}(A_{\nu\kappa})$.*
(ix) *If $v_1 \wedge \cdots \wedge v_k \wedge w_1 \wedge \cdots \wedge w_k \neq 0$, then $\overline{\operatorname{rank}}_{\nu\psi}(A_{\nu\psi}) = 2 < \operatorname{rank}_{\nu\psi}(A_{\nu\psi})$.*
(x) *If $[v_{1,1}], \dots, [v_{d_m,m}], [w_{1,1}], \dots, [w_{d_m,m}]$ are distinct in $\mathbb{P}V$, then $\overline{\operatorname{rank}}_{\sigma\nu\kappa}(A_{\sigma\nu\kappa}) = 2 < \operatorname{rank}_{\sigma\nu\kappa}(A_{\sigma\nu\kappa})$.*
(xi) *If $v_{1,i} \wedge \cdots \wedge v_{k_i,i} \wedge w_{1,i} \wedge \cdots \wedge w_{k_i,i} \neq 0$ for $i = 1, \dots, m$, then $\overline{\operatorname{rank}}_{\sigma\nu\psi}(A_{\sigma\nu\psi}) = 2 < \operatorname{rank}_{\sigma\nu\psi}(A_{\sigma\nu\psi})$.*

Proof. Since each of these tensors in Proposition 7.1 is a limit of the form

$$A = \lim_{t \to 0} \frac{\gamma(t) - \gamma(0)}{t},$$

we must have $\overline{\operatorname{rank}}_X(A) \leq 2$. On the other hand, because $\widehat{X}$ is complete as a metric space, $\operatorname{rank}_X(A) = 1$ if and only if $\overline{\operatorname{rank}}_X(A) = 1$. Hence requiring that $\operatorname{rank}_X(A) > 1$ ensures that $\overline{\operatorname{rank}}_X(A) > 1$. That the respective sufficient condition guarantees $\operatorname{rank}_X(A) > 2$ for each X in Section 7.1 follows from straightforward linear algebra. $\square$

Note that for each X in Section 7.1, the respective sufficient condition in Proposition 7.2 is satisfied by a general $A \in \tau(X)$.

The *generic X-rank*, denoted $r_g(X)$, is the minimum integer r such that $\Sigma_r(X) = V$. Let $r, s \in \mathbb{N}$ be such that $1 < r < s \leq r_g(X)$, and let $p \in V$ with $\operatorname{rank}_X(p) = s$. Then p does not necessarily have a best

X-rank-r approximation. While such failures are well-known for Segre variety [1] and Veronese variety [23], we deduce from Proposition 7.2 that they occur for all varieties in Section 7.1 when the orders of the tensors are higher than two.

Theorem 7.3. *Let X be any of the projective varieties listed in Section 7.1 and let $d, k, m \geq 3$. Then there exist a tensor $A \in \operatorname{span}(X)$, i.e., A has the required symmetry and/or skew-symmetry, and an $r < r_g(X)$ such that the best X-rank-r approximation of A does not exist, i.e.,*

$$\inf_{\operatorname{rank}_X(B) \leq r} \|A - B\|$$

is not attained by any tensor $B \in \operatorname{span}(X)$ with $\operatorname{rank}_X(B) \leq r$.

Proof. For each projective variety X in Section 7.1, we need to exhibit a tensor $A \in \operatorname{span}(X)$ with different X-rank and X-border rank; in which case A will not have a best X-rank-r approximation for $r = \overline{\operatorname{rank}}_X(A)$. By Proposition 7.2, we see that by setting $r = 2$ and picking a general point $A \in \tau(X)$ that satisfies the sufficient conditions in Proposition 7.2, we obtain a tensor with no best X-rank-two approximation for each of the varieties X in Section 7.1. $\square$

While Theorem 7.3 shows that the phenomenon where a tensor fails to have a best X-rank-r approximation can and does happen for every variety in Section 7.1, our next result is that such failures occur with probability zero, a consequence of Theorem 5.2. A high-level explanation goes as follows: When $r < r_g(X)$, the set of points in $\Sigma_r(X)$ whose X-rank is r contains a Zariski open subset, and so the set of "bad points" in $\Sigma_r(X)$ — those whose X-rank is not r — is contained in a subvariety; Theorem 5.2 then guarantees that these "bad points" are avoided almost always. A formal statement follows next.

Theorem 7.4. *Let $X \subseteq \mathbb{P}V$ be an irreducible nondegenerate complex projective variety. Then a general point $p \in V$ has a unique best X-rank-r approximation whenever $r < r_g(X)$.*

When applied to the varieties in Section 7.1, we obtain the following corollary.

Corollary 7.5. *Let V and $V_1, \dots, V_d$ be complex vector spaces.*

(i) *Any general tensor $A \in V_1 \otimes \cdots \otimes V_d$ has a unique best rank-r approximation when $r < r_g\big(\sigma(\mathbb{P}V_1 \times \cdots \times \mathbb{P}V_d)\big)$, i.e.,*

$$\inf_{v_{j,i} \in V_j} \left\| A - \sum_{i=1}^{r} v_{1,i} \otimes \cdots \otimes v_{d,i} \right\|$$

can be attained.

(ii) *Any general symmetric tensor $A \in \mathsf{S}^d(V)$ has a unique best symmetric rank-r approximation when $r < r_g\big(\nu(\mathbb{P}V)\big)$, and a unique best Chow rank-r approximation when $r < r_g\big(\operatorname{Ch}_d(V)\big)$, i.e.,*

$$\inf_{v_i \in V} \left\| A - \sum_{i=1}^{r} v_i^{\otimes d} \right\| \qquad \textit{and} \qquad \inf_{v_{j,i} \in V} \left\| A - \sum_{i=1}^{r} v_{1,i} \circ \cdots \circ v_{d,i} \right\|$$

can be attained.

(iii) *Any general alternating k-tensor $A \in \mathsf{\Lambda}^k(V)$ has a unique best alternating rank-r approximation when $r < r_g\big(\mathrm{Gr}_d(V)\big)$, i.e.,*

$$\inf_{v_{j,i} \in V} \left\| A - \sum_{i=1}^{r} v_{1,i} \wedge \cdots \wedge v_{k,i} \right\|$$

can be attained.

(iv) *Any general $A \in \mathsf{S}^{d_1}(V) \otimes \cdots \otimes \mathsf{S}^{d_m}(V)$ has a unique best Segre–Veronese rank-r approximation when $r < r_g\big(\sigma\nu(\mathbb{P}V_1 \times \cdots \times \mathbb{P}V_m)\big)$, and a unique best Segre–Chow rank-r approximation when $r < r_g\big(\sigma\kappa(\mathrm{Ch}_{d_1}(V) \times \cdots \times \mathrm{Ch}_{d_m}(V))\big)$, i.e.,*

$$\inf_{v_{j,i} \in V} \left\| A - \sum_{i=1}^{r} v_{1,i}^{\otimes d_1} \otimes \cdots \otimes v_{m,i}^{\otimes d_m} \right\|$$

and

$$\inf_{v_{j,l,i} \in V} \left\| A - \sum_{i=1}^{r} (v_{1,1,i} \circ \cdots \circ v_{d_1,1,i}) \otimes \cdots \otimes (v_{1,m,i} \circ \cdots \circ v_{d_m,m,i}) \right\|$$

can be attained.

(v) *Any general tensor $A \in \mathsf{\Lambda}^{k_1}(V) \otimes \cdots \otimes \mathsf{\Lambda}^{k_m}(V)$ has a unique best Segre–Grassmann rank-r approximation when $r < r_g\big(\sigma\psi(\mathrm{Gr}_{k_1}(V) \times \cdots \times \mathrm{Gr}_{k_m}(V))\big)$, i.e.,*

$$\inf_{v_{j,l,i} \in V} \left\| A - \sum_{i=1}^{r} \big(v_{1,1,i} \wedge \cdots \wedge v_{k_1,1,i}\big) \otimes \dots \big(v_{1,m,i} \wedge \cdots \wedge v_{k_m,m,i}\big) \right\|$$

can be attained.

(vi) *Any general tensor $A \in \mathsf{S}^k\big(\mathsf{S}^d(V)\big)$ has a unique best Veronese–Chow rank-r approximation when $r < r_g\big(\nu\kappa(\mathrm{Ch}_d(V))\big)$, i.e.,*

$$\inf_{v_{j,i} \in V} \left\| A - \sum_{i=1}^{r} (v_{1,i} \circ \cdots \circ v_{d,i})^{\otimes k} \right\|$$

can be attained.

(vii) *Any general tensor $A \in \mathsf{S}^d\big(\mathsf{\Lambda}^k(V)\big)$ has a unique best Veronese–Grassmann rank-r approximation when $r < r_g\big(\nu\psi(\mathrm{Gr}_k(V))\big)$, i.e.,*

$$\inf_{v_{j,i} \in V} \left\| A - \sum_{i=1}^{r} \big(v_{1,i} \wedge \cdots \wedge v_{k,i}\big)^{\otimes d} \right\|$$

can be attained.

(viii) *Any general tensor $A \in \mathsf{S}^{k_1}\big(\mathsf{S}^{d_1}(V)\big) \otimes \cdots \otimes \mathsf{S}^{k_m}\big(\mathsf{S}^{d_m}(V)\big)$ has a unique best Segre–Veronese–Chow rank-r approximation when $r < r_g\big(\sigma\nu\kappa(\mathrm{Ch}_{d_1}(V) \times \cdots \times \mathrm{Ch}_{d_m}(V))\big)$, i.e.,*

$$\inf_{v_{j,l,i} \in V} \left\| A - \sum_{i=1}^{r} (v_{1,1,i} \circ \cdots \circ v_{d_1,1,i})^{\otimes k_1} \otimes \cdots \otimes (v_{1,m,i} \circ \cdots \circ v_{d_m,m,i})^{\otimes k_m} \right\|$$

can be attained.

(ix) *Any general tensor* $A \in \mathsf{S}^{d_1}(\bigwedge^{k_1}(V)) \otimes \cdots \otimes \mathsf{S}^{d_m}(\bigwedge^{k_m}(V))$ *has a unique best Segre–Veronese–Grassmann rank-*r *approximation when* $r < r_g(\sigma\nu\psi(\mathrm{Gr}_{k_1}(V) \times \cdots \times \mathrm{Gr}_{k_m}(V)))$, *i.e.,*

$$\inf_{v_{j,l,i}\in V} \left\| A - \sum_{i=1}^{r} (v_{1,1,i} \wedge \cdots \wedge v_{k_1,1,i})^{\otimes d_1} \otimes \ldots (v_{1,m,i} \wedge \cdots \wedge v_{k_m,m,i})^{\otimes d_m} \right\|$$

can be attained.

8. Applications to other approximation problems

Our general existence and uniqueness result in Theorem 5.2 has implications beyond approximation by secant varieties in tensor spaces. In this section, we will look at some other approximation problems that arise in practical applications but do not fall naturally into the class of approximation problems considered in Section 7.2

8.1. Sparse-plus-low-rank approximations

Discussions of sparsity necessarily involve a choice of bases and so in this section we will assume that bases have been chosen on our vector spaces and that $V_1 = \mathbb{C}^{n_1}, \ldots, V_d = \mathbb{C}^{n_d}$. We start by considering the linear subspace of k-sparse hypermatrices with a *fixed* sparsity pattern in $\mathbb{C}^{n_1} \otimes \cdots \otimes \mathbb{C}^{n_d} \cong \mathbb{C}^{n_1\times\cdots\times n_d}$ and leave the case where the sparsity pattern is not fixed to later. More precisely, the sparsity pattern, $\Theta \subseteq \{1, \ldots, n_1\} \times \cdots \times \{1, \ldots, n_d\}$, is a subset of the index set and

$$S_\Theta := \{C \in \mathbb{C}^{n_1\times\cdots\times n_d} \mid c_{i_1\cdots i_d} = 0 \text{ if } (i_1, \ldots, i_d) \notin \Theta\}.$$

In the language of functions, this is just the vector space of complex-valued functions supported on Ω.

Denote the set of hypermatrices of rank not more than r by

$$R_r := \{B \in \mathbb{C}^{n_1\times\cdots\times n_d} \mid \mathrm{rank}(B) \le r\}, \tag{16}$$

and recall that its closure projectivizes to the r-secant variety of the Segre variety

$$\mathbb{P}\overline{R}_r = \Sigma_r(\sigma(\mathbb{P}\mathbb{C}^{n_1} \times \cdots \times \mathbb{P}\mathbb{C}^{n_d})).$$

Given a hypermatrix $A \in \mathbb{C}^{n_1\times\cdots\times n_d}$, a best k-sparse-plus-rank-r approximation of A is a solution to the problem:

$$\inf_{B\in R_r,\ C\in S_\Theta} \|A - (B + C)\|. \tag{17}$$

In other words, we would like study the best approximation problem whereby A is approximated by a hypermatrix in $J(\mathbb{P}\overline{R}_r, \mathbb{P}S)$. Recall that this is a join variety, as defined in Section 7.2. By Theorem 5.2, we may deduce the following existence and uniqueness result for such approximation problems.

Corollary 8.1. *For a general complex hypermatrix* $A \in \mathbb{C}^{n_1\times\cdots\times n_d}$*, the problem* (17) *has a unique solution* $B + C$ *such that* $\mathrm{rank}(B) = r$ *and* $C \in S_\Theta$ *is a* k*-sparse matrix with sparsity pattern* Θ*.*

Such scenarios where one requires a fixed sparsity arise in hypermatrix completion problems like the one in Section 8.2 and in statistical models like the one in Section 8.3. However, the better known k-sparse-plus-rank-r approximation problem as stated in [24,25] does not impose a fixed sparsity pattern on C but allows it be any k-sparse hypermatrix (or matrix, if $d = 2$). In other words, we consider the set of k-sparse hypermatrices (or matrices, if $d = 2$)

$$S_k := \bigcup_{\#\Theta = k} S_\Theta = \{C \in \mathbb{C}^{n_1 \times \cdots \times n_d} \mid \|C\|_0 \le k\},$$

where $\|C\|_0$ counts the number of nonzero entries in C. Note that such a union will no longer be a subspace of $\mathbb{C}^{n_1 \times \cdots \times n_d}$. In this case, we consider the *reducible* variety

$$\bigcup_{\#\Theta = k} J\bigl(\Sigma_r(\sigma(\mathbb{P}\overline{R}_r), \mathbb{P}S_\Theta\bigr)$$

and apply Corollary 5.6 to obtain our required existence and uniqueness result.

Corollary 8.2. *For a general complex hypermatrix $A \in \mathbb{C}^{n_1 \times \cdots \times n_d}$, the problem:*

$$\inf_{\operatorname{rank}(B) \le r,\ \|C\|_0 \le k} \|A - (B + C)\|$$

has a unique solution $B + C$ such that $\operatorname{rank}(B) = r$ and $\|C\|_0 = k$.

Note that in both Corollaries 8.1 and 8.2, the uniqueness applies to the solution $B + C$, i.e., not to B and C individually but to their sum.

8.2. Hypermatrix completion

The discussions in this section again depend on a choice of bases and we make the same assumptions as in Section 8.1. Suppose we are given the values of the entries

$$\{a_{i_1 \cdots i_d} \mid (i_1, \dots, i_d) \in \Theta\}$$

of a hypermatrix $A \in \mathbb{C}^{n_1 \times \cdots \times n_d}$ with $\Theta \subseteq \{1, \dots, n_1\} \times \cdots \times \{1, \dots, n_d\}$ a subset of the indices. Let

$$P_\Theta : \mathbb{C}^{n_1 \times \cdots \times n_d} \to \mathbb{C}^{n_1 \times \cdots \times n_d}, \quad A \mapsto \sum_{(i_1, \dots, i_d) \in \Theta} a_{i_1 \cdots i_d} e_{i_1} \otimes \cdots \otimes e_{i_d},$$

be the projection onto the subspace defined by coordinates of the known entries. The problem of finding $B \in \mathbb{C}^{n_1 \times \cdots \times n_d}$ of rank at most r that attains

$$\inf_{\operatorname{rank}(B) \le r} \|P_\Theta(A - B)\| \tag{18}$$

is called a best rank-r *hypermatrix completion problem.*[6] In this situation our result implies the following.

Corollary 8.3. *Let $\Theta \subseteq \{1, \dots, n_1\} \times \cdots \times \{1, \dots, n_d\}$. If the given entries $\{a_{i_1 \cdots i_d} \mid (i_1, \dots, i_d) \in \Theta\}$ of a hypermatrix $A \in \mathbb{C}^{n_1 \times \cdots \times n_d}$ are general, then there exists a hypermatrix $B \in \mathbb{C}^{n_1 \times \cdots \times n_d}$ of rank r that solves (18). Furthermore, although B may be not unique, its entries $\{b_{i_1 \cdots i_d} \mid (i_1, \dots, i_d) \in \Theta\}$ are unique.*

[6] Often called “tensor completion” but this is a misnomer since the problem clearly depends on coordinates.

Proof. Consider the image of rank-r hypermatrices $P_\Theta(R_r)$ with R_r as in (16). As $P_\Theta(A-B) = P_\Theta(A) - P_\Theta(B)$, a best rank-$r$ hypermatrix completion problem is equivalent to approximating $P_\Theta(A)$ by an element of $P_\Theta(R_r)$. Applying Theorem 5.2 with X as the closure of $P_\Theta(R_r)$ and $Z = X \setminus P_\Theta(R_r)$, we obtain a unique solution $P_\Theta(B) \in P_\Theta(R_r)$, i.e., there exists B of rank r that is a best rank-r completion of A. Note that B is not unique but two different B's must have the same projection $P_\Theta(B)$, i.e., the entries $\{b_{i_1\cdots i_d} \mid (i_1,\dots,i_d) \in \Theta\}$ are unique. $\square$

The similarity of the best rank-r hypermatrix completion problem (18) and the best k-sparse-plus-rank-r approximation problem (17) in Section 8.1 should not go unnoticed. For a given index set Θ, the solution B to (18) with Θ as the indices of known entries equals the B in (17) with Θ as the sparsity pattern and $C \in \ker(P_\Theta)$.

8.3. Gaussian r-factor analysis model with k observed variables

Consider a Gaussian hidden variable model with k observed variables $X_1,\dots,X_k$ and r hidden variables $Y_1,\dots,Y_r$ where $(X_1,\dots,X_k,Y_1,\dots,Y_r)$ follows a joint multivariate normal distribution with positive definite covariance matrix. If the observed variables are conditionally independent given the hidden variables, this model is called the *Gaussian r-factor analysis model with k observed variables* [26] and denoted by $\mathbf{F}_{k,r}$. In fact, by [26, Proposition 1], $\mathbf{F}_{k,r}$ is the family of multivariate normal distributions $\mathcal{N}(\mu,\Sigma)$ on $\mathbb{R}^k$ with $\mu \in \mathbb{R}^k$ and Σ belonging to

$$F_{k,r} := \{\Psi + LL^\intercal \in \mathbb{R}^{k\times k} \mid \Psi \succ 0 \text{ and diagonal}, L \in \mathbb{R}^{k\times r}\}.$$

The standard approach in *algebraic statistics* [27] and also that in [26] is to drop any semialgebraic conditions and complexify. In this case, it means to drop the condition $\Sigma \succ 0$ and regard all quantities as complex valued, i.e.,

$$F_{k,r}(\mathbb{C}) = \{\Psi + LL^\intercal \in \mathbb{C}^{k\times k} \mid \Psi \text{ diagonal}, L \in \mathbb{C}^{k\times r}\}.$$

This is the set of complexified covariance matrices for the model $\mathbf{F}_{k,r}$ and the algebraic approach undertaken in [26] effectively treats $F_{k,r}(\mathbb{C})$ as the parameter space of the Gaussian r-factor analysis model with k observed variables.

If we replace the space of matrices $V \otimes W$ in Section 8.1 by the space of symmetric matrices $\mathsf{S}^2(V)$ and let $S \subseteq \mathsf{S}^2(V)$ be the subspace of diagonal matrices, then we see that

$$F_{k,r}(\mathbb{C}) = J\big(\Sigma_r(\nu_2(\mathbb{P}V)), \mathbb{P}S\big).$$

It follows from Corollary 8.1 that every general complexified covariance matrix $\Sigma \in \mathsf{S}^2(V)$ has a unique best approximation by $\Psi + LL^\intercal$, a complexified covariance matrix for the model $\mathbf{F}_{k,r}$. To provide context, readers unfamiliar with factor analysis [28, Chapter 9] should note that its main parameter estimation problem is to determine the matrix of *loadings* L and the diagonal matrix of *specific variances* $\Psi = \operatorname{diag}(\psi_1,\dots,\psi_k)$ from a sample covariance matrix Σ.

8.4. Block-term tensor approximations

Our final example of finding a best approximation in a join variety brings us back to tensors. A class of tensor approximation problems with groundbreaking applications in signal processing is the so-called *block-term decompositions* [29]. Unlike the factor analysis application in Section 8.3, which is really a problem over $\mathbb{R}$ but is complexified to allow techniques of algebraic statistics, the use of block-term decompositions as a model in signal processing is naturally and necessarily over $\mathbb{C}$.

Any tensor $A \in V_1 \otimes \cdots \otimes V_d$ induces a linear map $A_i \colon V_i^* \to V_1 \otimes \cdots \otimes V_{i-1} \otimes V_{i+1} \otimes \cdots \otimes V_d$. The rank of this linear map is of course just the dimension of its image $\dim A_i(V_i^*)$. The *multilinear rank* of A is then defined to be the d-tuple

$$\mu\mathrm{rank}(A) \coloneqq \bigl(\dim A_1(V_1^*), \dots, \dim A_d(V_d^*)\bigr)$$

and the set

$$\mathrm{Sub}_{r_1,\dots,r_d}(V_1 \otimes \cdots \otimes V_d) \coloneqq \{A \in V_1 \otimes \cdots \otimes V_d \mid \mu\mathrm{rank}(A) \le (r_1, \dots, r_d)\}$$

is a complex algebraic variety called a *subspace variety*.

Given $(r_{1,1}, \dots, r_{1,d}), \dots, (r_{k,1}, \dots, r_{k,d})$, a tensor $B \in V_1 \otimes \cdots \otimes V_d$ is said to have a block-term decomposition if

$$B = B_1 + \cdots + B_k, \qquad \mu\mathrm{rank}(B_i) \le (r_{i,1}, \dots, r_{i,d}), \quad i = 1, \dots, k.$$

It is known [30,29] that the best block-term approximation problem

$$\inf\bigl\{\|A - (B_1 + \cdots + B_k)\| \bigm| \mu\mathrm{rank}(B_i) \le (r_{i,1}, \dots, r_{i,d}),\ i = 1, \dots, k\bigr\} \tag{19}$$

does not have a solution for certain choices of A, i.e., the infimum above cannot be attained, much like the tensor approximation problems we discussed in Section 7.2; an explicit example where the infimum in (19) is unattainable for a set of tensors with nonempty interior can be found in [30].

Theorem 5.2, when applied to the join variety of k subspace varieties, gives us the following.

Corollary 8.4. *Let $V_1, \dots, V_d$ be complex vector spaces. A general tensor $A \in V_1 \otimes \cdots \otimes V_d$ has a unique best approximation in the join variety*

$$J\bigl(\mathrm{Sub}_{r_{1,1},\dots,r_{1,d}}(V_1 \otimes \cdots \otimes V_d), \dots, \mathrm{Sub}_{r_{k,1},\dots,r_{k,d}}(V_1 \otimes \cdots \otimes V_d)\bigr),$$

i.e., the best block-term approximation problem in (19) *has a unique solution.*

8.5. Approximations by tensor networks

Tensor networks are tensors or functions that have separable decompositions indexed by an undirected graph. For example a *matrix product state* $f \in L^2(\Omega_1 \times \Omega_2 \times \Omega_3)$ corresponding to a triangle takes the form

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

where $\varphi_{ij} \in L^2(\Omega_1)$, $\psi_{jk} \in L^2(\Omega_2)$, $\theta_{ki} \in L^2(\Omega_3)$; note that the indices i, j, k form a triangle with edges $\{i, j\}$, $\{j, k\}$, $\{k, i\}$ (see [31] for more details). Other undirected graphs give other tensor networks, examples include: matrix product states (cycle graphs), tensor trains (path graphs), tree tensor network states (trees), projected entangled pair states (product of path graphs), etc. For a precise definition, see [31, Definition 2.1].

Tensor networks play a prominent role in quantum physics and quantum chemistry [32–37]. It is known that tensor networks X_G corresponding to cyclic graphs G are not closed [38,39] and thus may not have best X_G-approximation.

Corollary 8.5. *Let G be any undirected graph with d vertices and $V_1, \dots, V_d$ be complex vector spaces. If $X_G \subseteq V_1 \otimes \cdots \otimes V_d$ is a set of tensor network states corresponding to G. Then the best $\overline{X}_G$-approximation of a general tensor is unique and lies in X.*

Proof. By Theorem 5.2 it is enough to show that X_G contains a Zariski open subset of $\overline{X}_G$. For tensor networks this follows from [31, Proposition 2.6]. □

Corollary 8.5 may be extended to other variants of tensor networks such as uniform (or site-independent) matrix product states, which are symmetric versions of matrix product states [34]; the analogue of Corollary 8.5 in this case follows from [38, Proposition 4.2].

9. Conclusion

Our study in this article demonstrates a significant difference between approximation problems over $\mathbb{R}$ and over $\mathbb{C}$. In the case when the approximation involves tensors, this revelation is consistent with our knowledge of other aspects of tensors — tensor rank, tensor spectral norm, tensor nuclear norm are all known to be dependent on the base field [1,40]; even the fact that the Grothendieck constants have different values over $\mathbb{R}$ and over $\mathbb{C}$ is a manifestation of this phenomenon [41].

The knowledge that over $\mathbb{C}$, the best rank-r approximation problems for tensors and, more generally, any best r-term approximation problems, are well-posed except on a set of measure zero should provide a reasonable amount of justification for applications that depend on such approximations. In fact, such "almost everywhere" guarantees are about as much as one may hope for, since "everywhere" is certainly false — as we saw, for every single approximation problem that we considered in this article, there are indeed instances where best approximations do not exist.

Acknowledgments

We are grateful to J.M. Landsberg for posing the question about general existence of best low-rank approximations of tensors over $\mathbb{C}$ to us (when LHL visited College Station in 2010, and when he visited YQ in Grenoble in 2015). MM thanks W. Hackbusch for inspiring discussions on the topic. Special thanks go to the anonymous referees for their careful reading and numerous invaluable suggestions that significantly enriched the content of this article.

The work of YQ and LHL is supported by the Defense Advanced Research Projects Agency (DARPA) grant D15AP00109 and the National Science Foundation (NSF) grant IIS 1546413. The work of MM is supported by the Polish National Science Center grant 2013/08/A/ST1/00804. In addition LHL acknowledges support from a DARPA Director's Fellowship and the Eckhardt Faculty Fund.

References

[1] V. De Silva, L.-H. Lim, Tensor rank and the ill-posedness of the best low-rank approximation problem, SIAM J. Matrix Anal. Appl. 30 (3) (2008) 1084–1127, https://doi.org/10.1137/06066518X.
[2] E.J. Candès, Compressive sampling, in: International Congress of Mathematicians, vol. III, Eur. Math. Soc., Zürich, 2006, pp. 1433–1452.
[3] E.J. Candès, T. Tao, The power of convex relaxation: near-optimal matrix completion, IEEE Trans. Inform. Theory 56 (5) (2010) 2053–2080, https://doi.org/10.1109/TIT.2010.2044061.
[4] A. Cohen, W. Dahmen, R. DeVore, Compressed sensing and best k-term approximation, J. Amer. Math. Soc. 22 (1) (2009) 211–231, https://doi.org/10.1090/S0894-0347-08-00610-3.
[5] G. Cybenko, Approximation by superpositions of a sigmoidal function, Math. Control Signals Systems 2 (4) (1989) 303–314, https://doi.org/10.1007/BF02551274.
[6] D.L. Donoho, Compressed sensing, IEEE Trans. Inform. Theory 52 (4) (2006) 1289–1306, https://doi.org/10.1109/TIT.2006.871582.

[7] D.L. Donoho, M. Elad, Optimally sparse representation in general (nonorthogonal) dictionaries via l^1 minimization, Proc. Natl. Acad. Sci. USA 100 (5) (2003) 2197–2202, https://doi.org/10.1073/pnas.0437847100.

[8] A.C. Gilbert, S. Muthukrishnan, M.J. Strauss, Approximation of functions over redundant dictionaries using coherence, in: Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms, Baltimore, MD, 2003, ACM, New York, 2003, pp. 243–252.

[9] R. Gribonval, M. Nielsen, Sparse representations in unions of bases, IEEE Trans. Inform. Theory 49 (12) (2003) 3320–3325, https://doi.org/10.1109/TIT.2003.820031.

[10] D. Gross, Recovering low-rank matrices from few coefficients in any basis, IEEE Trans. Inform. Theory 57 (3) (2011) 1548–1566, https://doi.org/10.1109/TIT.2011.2104999.

[11] V. Temlyakov, Greedy Approximation, Cambridge Monographs on Applied and Computational Mathematics, vol. 20, Cambridge University Press, Cambridge, 2011.

[12] A. Poltoratski, A problem on completeness of exponentials, Ann. of Math. (2) 178 (3) (2013) 983–1016, https://doi.org/10.4007/annals.2013.178.3.4.

[13] G. Beylkin, L. Monzón, On approximation of functions by exponential sums, Appl. Comput. Harmon. Anal. 19 (1) (2005) 17–48, https://doi.org/10.1016/j.acha.2005.01.003.

[14] E. Bierstone, P.D. Milman, Semianalytic and subanalytic sets, Inst. Hautes Études Sci. Publ. Math. 67 (1988) 5–42, http://www.numdam.org/item?id=PMIHES_1988___67___5_0.

[15] A.H. Durfee, Neighborhoods of algebraic sets, Trans. Amer. Math. Soc. 276 (2) (1983) 517–530, https://doi.org/10.2307/1999065.

[16] M. Goresky, R. MacPherson, Stratified Morse theory, Ergebnisse der Mathematik und ihrer Grenzgebiete (3), vol. 14, Springer-Verlag, Berlin, 1988.

[17] S. Friedland, M. Stawiska, Some approximation problems in semi-algebraic geometry, in: Constructive Approximation of Functions, in: Banach Center Publ., vol. 107, Polish Acad. Sci. Inst. Math., Warsaw, 2015, pp. 133–147.

[18] P. Griffiths, J. Harris, Principles of Algebraic Geometry, Wiley Classics Library, John Wiley & Sons, Inc., New York, 1994, reprint of the 1978 original.

[19] J. Harris, Algebraic Geometry, Graduate Texts in Mathematics, vol. 133, Springer-Verlag, New York, 1992, a first course.

[20] L.-H. Lim, P. Comon, Nonnegative approximations of nonnegative tensors, J. Chemom. 23 (7–8) (2009) 432–441.

[21] I.M. Gel'fand, M.M. Kapranov, A.V. Zelevinsky, Discriminants, Resultants, and Multidimensional Determinants, Mathematics: Theory & Applications, Birkhäuser Boston, Inc., Boston, MA, 1994.

[22] J.M. Landsberg, Tensors: Geometry and Applications, Graduate Studies in Mathematics, vol. 128, American Mathematical Society, Providence, RI, 2012.

[23] P. Comon, G. Golub, L.-H. Lim, B. Mourrain, Symmetric tensors and symmetric tensor rank, SIAM J. Matrix Anal. Appl. 30 (3) (2008) 1254–1279, https://doi.org/10.1137/060661569.

[24] E.J. Candès, X. Li, Y. Ma, J. Wright, Robust principal component analysis?, J. ACM 58 (3) (2011) 11, https://doi.org/10.1145/1970392.1970395.

[25] V. Chandrasekaran, S. Sanghavi, P.A. Parrilo, A.S. Willsky, Rank-sparsity incoherence for matrix decomposition, SIAM J. Optim. 21 (2) (2011) 572–596, https://doi.org/10.1137/090761793.

[26] M. Drton, B. Sturmfels, S. Sullivant, Algebraic factor analysis: tetrads, pentads and beyond, Probab. Theory Related Fields 138 (3–4) (2007) 463–493, https://doi.org/10.1007/s00440-006-0033-2.

[27] M. Drton, B. Sturmfels, S. Sullivant, Lectures on Algebraic Statistics, Oberwolfach Seminars, vol. 39, Birkhäuser Verlag, Basel, 2009.

[28] R.A. Johnson, D.W. Wichern, Applied Multivariate Statistical Analysis, sixth edition, Pearson Prentice Hall, Upper Saddle River, NJ, 2007.

[29] L. De Lathauwer, Decompositions of a higher-order tensor in block terms. II. Definitions and uniqueness, SIAM J. Matrix Anal. Appl. 30 (3) (2008) 1033–1066, https://doi.org/10.1137/070690729.

[30] J.H. de Morais Goulart, P. Comon, On the minimal ranks of matrix pencils and the existence of a best approximate block-term tensor decomposition, preprint, https://arxiv.org/abs/1712.05742.

[31] K. Ye, L.-H. Lim, Tensor network ranks, preprint, https://arxiv.org/abs/1801.02662.

[32] W. Hackbusch, Tensor Spaces and Numerical Tensor Calculus, vol. 42, Springer Science & Business Media, 2012.

[33] D. Perez-Garcia, F. Verstraete, M. Wolf, J. Cirac, Matrix product state representations, Quantum Inf. Comput. 7 (5) (2007) 401–430.

[34] F. Verstraete, V. Murg, J.I. Cirac, Matrix product states, projected entangled pair states, and variational renormalization group methods for quantum spin systems, Adv. Phys. 57 (2) (2008) 143–224, https://doi.org/10.1080/14789940801912366.

[35] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, Ann. Physics 326 (1) (2011) 96–192, https://doi.org/10.1016/j.aop.2010.09.012.

[36] R. Orús, A practical introduction to tensor networks: matrix product states and projected entangled pair states, Ann. Physics 349 (2014) 117–158, https://doi.org/10.1016/j.aop.2014.06.013.

[37] S. Szalay, M. Pfeffer, V. Murg, G. Barcza, F. Verstraete, R. Schneider, Ö. Legeza, Tensor product methods and entanglement optimization for ab initio quantum chemistry, Int. J. Quant. Chem. 115 (19) (2015) 1342–1391, https://doi.org/10.1002/qua.24898.

[38] C. Harris, M. Michałek, E.C. Sertöz, Computing images of polynomial maps, preprint, https://arxiv.org/abs/1801.00827.

[39] J.M. Landsberg, Y. Qi, K. Ye, On the geometry of tensor network states, Quantum Inf. Comput. 12 (3–4) (2012) 346–354.

[40] S. Friedland, L.-H. Lim, Nuclear norm of higher-order tensors, Math. Comp. 87 (311) (2018) 1255–1281, https://doi.org/10.1090/mcom/3239.

[41] J. Zhang, S. Friedland, L.-H. Lim, Grothendieck constant is norm of Strassen matrix multiplication tensor, preprint, https://arxiv.org/abs/1711.04427.