

Statistical ranking and combinatorial Hodge theory

Xiaoye Jiang · Lek-Heng Lim · Yuan Yao ·
Yinyu Ye

Received: 31 October 2008 / Accepted: 10 June 2010 / Published online: 28 November 2010
© The Author(s) 2010. This article is published with open access at Springerlink.com

Abstract We propose a technique that we call *HodgeRank* for ranking data that may be incomplete and imbalanced, characteristics common in modern datasets coming from e-commerce and internet applications. We are primarily interested in cardinal data based on scores or ratings though our methods also give specific insights on ordinal data. From raw ranking data, we construct pairwise rankings, represented as edge flows on an appropriate graph. Our statistical ranking method exploits the graph

X. Jiang acknowledges support from ARO Grant W911NF-04-R-0005 BAA and the School of Engineering fellowship at Stanford. L.-H. Lim acknowledges support from the Gerald J. Liebermann fellowship at Stanford and the Charles B. Morrey assistant professorship at Berkeley. Y. Yao acknowledges supports from the National Basic Research Program of China (973 Program 2011CB809105), NSFC (61071157), Microsoft Research Asia, DARPA Grant HR0011-05-1-0007, NSF Grant DMS 0354543, and a professorship in the Hundred Talents Program at Peking University. Y. Ye acknowledges support from AFOSR Grant FA9550-09-1-0306 and DOE Grant DE-SC0002009.

X. Jiang
Institute for Computational and Mathematical Engineering,
Stanford University, Stanford, CA 94305, USA
e-mail: xiaoyej@stanford.edu

L.-H. Lim (✉)
Department of Mathematics, University of California, Berkeley, CA 94720, USA
e-mail: lekheng@galton.uchicago.edu; lekheng@math.berkeley.edu

Y. Yao
School of Mathematical Sciences, LMAM and Key Lab of Machine Perception (MOE),
Peking University, 100871 Beijing, People's Republic of China
e-mail: yuany@math.pku.edu.cn; yuany@math.stanford.edu

Y. Ye
Department of Management Science and Engineering, Stanford University,
Stanford, CA 94305, USA
e-mail: yinyu-ye@stanford.edu

Helmholtzian, which is the graph theoretic analogue of the Helmholtz operator or vector Laplacian, in much the same way the graph Laplacian is an analogue of the Laplace operator or scalar Laplacian. We shall study the graph Helmholtzian using combinatorial Hodge theory, which provides a way to unravel ranking information from edge flows. In particular, we show that every edge flow representing pairwise ranking can be resolved into two orthogonal components, a gradient flow that represents the l_2 -optimal global ranking and a divergence-free flow (cyclic) that measures the validity of the global ranking obtained—if this is large, then it indicates that the data does not have a good global ranking. This divergence-free flow can be further decomposed orthogonally into a curl flow (locally cyclic) and a harmonic flow (locally acyclic but globally cyclic); these provides information on whether inconsistency in the ranking data arises locally or globally. When applied to statistical ranking problems, Hodge decomposition sheds light on whether a given dataset may be globally ranked in a meaningful way or if the data is inherently inconsistent and thus could not have any reasonable global ranking; in the latter case it provides information on the nature of the inconsistencies. An obvious advantage over the NP-hardness of Kemeny optimization is that HodgeRank may be easily computed via a linear least squares regression. We also discuss connections with well-known ordinal ranking techniques such as Kemeny optimization and Borda count from social choice theory.

Keywords Statistical ranking · Rank aggregation · Combinatorial Hodge theory · Discrete exterior calculus · Combinatorial Laplacian · Hodge Laplacian · Graph Helmholtzian · HodgeRank · Kemeny optimization · Borda count

Mathematics Subject Classification (2000) 68T05 · 58A14 · 90C05 · 90C27 · 91B12 · 91B14

1 Introduction

The problem of ranking in various contexts has become increasingly important in machine learning. Many datasets¹ require some form of ranking to facilitate identification of important entries, extraction of principal attributes, and to perform efficient search and sort operations. Modern internet and e-commerce applications have spurred an enormous growth in such datasets: Google's search engine, CiteSeer's citation database, eBay's feedback-reputation mechanism, Netflix's movie recommendation system, all accumulate a large volume of data that needs to be ranked.

These modern datasets typically have one or more of the following features that render traditional ranking methods (such as those in social choice theory) inapplicable or ineffective: (1) unlike traditional ranking problems such as votings and tournaments, the data often contains *cardinal scores* instead of ordinal orderings; (2) the given data is largely *incomplete* with most entries missing a substantial amount of information; (3) the data will almost always be *imbalanced* where the amount of available informa-

¹ We use the word *data* in the singular (we have no use for the word *datum* in this paper). A collection of data would be called a *dataset*.

tion varies widely from entry to entry and/or from criterion to criterion; (4) the given data often lives on a large *complex network*, either explicitly or implicitly, and the structure of this underlying network is itself important in the ranking process. These new features have posed new challenges and call for new techniques. We will look at a method that addresses them to some extent.

A fundamental problem here is to globally rank a set of *alternatives* based on scores given by *voters*. Here the words ‘alternatives’ and ‘voters’ are used in a generalized sense that depends on the context. For example, the alternatives may be websites indexed by Google, scholarly articles indexed by CiteSeer, sellers on eBay, or movies on Netflix; the voters in the corresponding contexts may be other websites, other scholarly articles, buyers, or viewers. The ‘voters’ could also refer to groups of voters: e.g. websites, articles, buyers, or viewers grouped respectively by topics, authorship, buying patterns, or movie tastes. The ‘voters’ could even refer to something entirely abstract, such as a collection of different criteria used to judge the alternatives.

The features (1–4) can be observed in the aforementioned examples. In the eBay/Netflix context, a buyer/viewer would assign cardinal scores (1 through 5 stars) to sellers/movies instead of ranking them in an ordinal fashion; the eBay/Netflix datasets are highly incomplete since most buyers/viewers would have rated only a very small fraction of the sellers/movies, and also highly imbalanced since a handful of popular sellers/blockbuster movies would have received an overwhelming number of ratings while the vast majority would get only a moderate or small number of ratings. The datasets from Google and CiteSeer have obvious underlying network structures given by hyperlinks and citations, respectively. Somewhat less obvious are the network structures underlying the datasets from eBay and Netflix, which come from aggregating the pairwise comparisons of buyers/movies over all sellers/viewers. Indeed, we shall see that in all these ranking problems, graph structures naturally arise from *pairwise comparisons*, irrespective of whether there is an obvious underlying network (e.g. from citation, friendship, or hyperlink relations) or not, and this serves to place ranking problems of seemingly different nature on an equal graph-theoretic footing. The incompleteness and imbalance of the data are then manifested in the (edge) sparsity structure and (vertex) degree distribution of pairwise comparison graphs.

In collaborative filtering applications, one often encounters a personalized ranking problem, when one needs to find a global ranking of alternatives that generates the most consensus within a group of voters who share similar interests/tastes. While the statistical ranking problem investigated in this paper could play a fundamental role in such personalized ranking problems, there is also the equally important problem of clustering voters into interest groups, which we do not address in this paper.² We would like to stress that in this paper we only concern ourselves with the ranking problem but not the clustering problem. So while we have used the Netflix prize dataset in our studies, our paper should not be viewed as an attempt to solve the Netflix prize problem.

² Nevertheless the notions of local and global inconsistencies introduced in Sects. 2.4 and 5 could certainly be used as a criterion for such purposes, i.e. an interest group is one whose within-group inconsistency is small.

In a nutshell, our methods, collectively called *HodgeRank*, analyze pairwise rankings represented as edge flows on a graph using *discrete* or *combinatorial Hodge theory*. Among other things, combinatorial Hodge theory provides us with a mean to determine a global ranking that also comes with a ‘certificate of reliability’ for the validity of this global ranking. While Hodge theory is well-known to pure mathematicians as a corner stone of geometry and topology, and to applied mathematician as an important tool in computational electromagnetics and fluid dynamics, its application to statistical ranking problems has, to the best of our knowledge, never been studied.³

In HodgeRank, the graph in question has as its vertices the alternatives to be ranked, voters’ preferences are then quantified and aggregated (we will say how later) into an edge flow on this graph. Hodge theory then yields an orthogonal decomposition of the edge flow into three components: a *gradient flow* that is globally acyclic, a *harmonic flow* that is locally acyclic but globally cyclic, and a *curl flow* that is locally cyclic. This decomposition is known as the *Hodge* or *Helmholtz decomposition*. The usefulness of the decomposition lies in the fact that the gradient flow component induces a global ranking of the alternatives. Unlike the computationally intractable Kemeny optimal, this may be easily computed via a linear least squares problem. Furthermore, the l_2 -norm of the least squares residual, which represents the contribution from the sum of the remaining curl flow and harmonic flow components, quantifies the validity of the global ranking induced by the gradient flow component. If the residual is small, then the gradient flow accounts for most of the variation in the underlying data and therefore the global ranking obtained from it is expected to be a majority consensus. On the other hand, if the residual is large, then the underlying data is plagued with cyclic inconsistencies (i.e. intransitive preference relations of the form $a \succeq b \succeq c \succeq \dots \succeq z \succeq a$) and one may not assign any reasonable global ranking to it.

We would like to point out here that cyclic inconsistencies are not necessarily due to error or noise in the data but may very well be an inherent characteristic of the data. As the famous impossibility theorems from social choice theory [2, 39] have shown, inconsistency (or, rather, intransitivity) is inevitable in any societal preference aggregation that is sophisticated enough. Social scientists have, through empirical studies, observed that preference judgement of groups or individuals on a list of alternatives do in fact exhibit such irrational or inconsistent behavior. Indeed in any group decision making process, a lack of consensus is the norm rather than the exception in our everyday experience. This is the well-known *Condorcet paradox* [15]: the majority prefers $a-b$ and $b-c$, but may yet prefer $c-a$. Even a single individual making his own preference judgements could face such dilemma—if he uses multiple criteria to rank the alternatives. As such, the cyclic inconsistencies is intrinsic to any real world ranking data and should be thoroughly analyzed. Hodge theory again provides a mean to do so. The curl flow and harmonic flow components of an edge flow quantify respectively the local and global cyclic inconsistencies.

Loosely speaking, a dominant curl flow component suggests that the inconsistencies are of a local nature while a dominant harmonic flow component suggests that they are of a global nature. If most of the inconsistencies come from the curl (local) component

³ Hodge theory has recently also found other applications in statistical learning theory [5].

while the harmonic (global) component is small, then this roughly translates to mean that the ordering of closely ranked alternatives is unreliable but that of very differently ranked alternatives is reliable, i.e. we cannot say with confidence whether the ordering of the 27th, 28th, 29th ranked items makes sense but we can say with confidence that the 4th, 60th, 100th items should be ordered according to their rank. In other words, Condorcet paradox may well apply to items ranked closed together but not to items ranked far apart. For example, if a large number of gourmets (voters) are asked to state their preferences on an extensive range of food items (alternatives), there may not be a consensus for their preferences with regard to hamburgers, hot dogs, and pizzas and there may not be a consensus for their preferences with regard to caviar, foie gras, and truffles; but there may well be a near universal preference for the latter group of food items over the former group. In this case, the inconsistencies will be mostly local and we should expect a large curl flow component. If in addition the harmonic flow component is small, then most of the inconsistencies happen locally and we could interpret this to mean that the global ranking is valid on a coarse scale (ranking different groups of food) but not on a fine scale (ranking similar food items belonging to a particular group). We refer the reader to Sect. 7.1 for an explicit example based on the Netflix prize dataset.

1.1 What's new

The main contribution of this paper is in the application of Hodge decomposition to the analysis of ranking data. We show that this approach has several attractive features: (1) it generalizes the classical Borda Count method in voting theory to data that may have missing values; (2) it provides a way to analyze inherent inconsistencies or conflicts in the ranking data; (3) it is flexible enough to be combined with other techniques: these include other ways to form pairwise rankings reflecting prior knowledge. Although completely natural, Hodge theory has, as far as we know, never been applied to the study of ranking.

We emphasize two conceptual aspects underlying this work that are particularly unconventional: (1) We believe that obtaining a global ranking, which is the main if not the sole objective of all existing work on rank aggregation, gives only an incomplete picture of the ranking data—one also needs a ‘certificate of reliability’ for the global ranking. Our method provides this certificate by measuring also the local and global inconsistent components of the ranking data. (2) We believe that with the right mathematical model, rank aggregation need not be a computationally intractable task. The model that we proposed in this paper reduces rank aggregation to a linear least squares regression, avoiding usual NP-hard combinatorial optimization problems such as finding Kemeny optima or minimum feedback arc sets.

Hodge and Helmholtz decompositions are of course well-known in mathematics and physics, but usually in a continuous setting where the underlying spaces have the structure of a Riemannian manifold or an algebraic variety. The *combinatorial Hodge theory* that we presented here is arguably a trivial case with the simplest possible underlying space—a graph. Many of the difficulties in developing Hodge theory in differential and algebraic geometry simply do not surface in our case. However this

also makes combinatorial Hodge theory accessible—the way we developed and presented it essentially requires nothing more than some elementary matrix theory and multivariate calculus. We are unaware of similar treatments in the existing literature and would consider our elementary treatment a minor expository contribution that might help popularize the use of Hodge decomposition and the graph Helmholtzian, possibly to other areas in data analysis and machine learning.

1.2 Organization of this paper

In Sect. 2 we introduce the main problem and discuss how a pairwise comparison graph may be constructed from data comprising cardinal scores given by voters on alternatives and how a simple least squares regression may be used to compute the desired solution. We define the *combinatorial curl*, a measure of local (triangular) inconsistency for such data, and also the *combinatorial gradient* and *combinatorial divergence*. Section 3 describes a purely matrix-theoretic view of Hodge theory, but at the expense of some geometric insights. This deficiency is rectified in Sect. 4, where we introduce a graph-theoretic Hodge theory. We first remind the reader how one may construct a d -dimensional simplicial complex from any given graph (the pairwise comparison graph in our case) by simply filling-in all its k -cliques for $k \leq d$. Then we will introduce combinatorial Hodge theory for a general d -dimensional simplicial complex but focusing on the $d = 2$ case and its relevance to the ranking problem. In Sect. 5 we discuss the implications of Hodge decomposition applied to ranking, with a deeper analysis on the least squares method in Sect. 2. A discussion of the connections with Kemeny optimization and Borda count in social choice theory can be found in Sect. 6. Numerical experiments on three real datasets are given in Sect. 7 to illustrate the basic ideas.

1.3 Notations

Let V be a finite set. We will adopt the following notation from combinatorics:

$$\binom{V}{k} := \text{set of all } k\text{-element subset of } V.$$

In particular $\binom{V}{2}$ would be the set of all unordered pairs of elements of V and $\binom{V}{3}$ would be the set of all unordered triples of elements of V (the sets of ordered pairs and ordered triples will be denoted $V \times V$ and $V \times V \times V$ as usual). We will not distinguish between V and $\binom{V}{1}$. Ordered and unordered pairs will be delimited by parentheses (i, j) and braces $\{i, j\}$ respectively, and likewise for triples and n -tuples in general.

We will use positive integers to label alternatives and voters. Henceforth, V will always be the set $\{1, \dots, n\}$ and will denote a set of alternatives to be ranked. In our approach to statistical ranking, these alternatives are represented as vertices of a graph. $\Lambda = \{1, \dots, m\}$ will denote a set of voters. For $i, j \in V$, we write $i \succeq j$ to mean that

Graph theory	Linear algebra	Vec. calculus	Topology	Ranking
Function on vertices	Vector in \mathbb{R}^n	Potential function	0-cochain	Score function
Edge flow	Skew-symmetric matrix in $\mathbb{R}^{n \times n}$	Vector field	1-cochain	Pairwise ranking
Triangular flow	Skew-symmetric hyper-matrix in $\mathbb{R}^{n \times n \times n}$	Tensor field	2-cochain	Triplewise ranking

alternative i is preferred over alternative j . If we wish to emphasize the preference judgement of a particular voter $\alpha \in \Lambda$, we will write $i \succeq_\alpha j$.

Since in addition to ranking theoretic terms, we have borrowed terminologies from graph theory, vector calculus, linear algebra, and algebraic topology, we provide a table of correspondence for easy reference.

As the reader will see, the notions of gradient, divergence, curl, Laplace operator, and Helmholtz operator from vector calculus and topology will play important roles in statistical ranking. One novelty of our approach lies in extending these notions to the other three columns, where most of them have no well-known equivalent. For example, what we will call a *harmonic ranking* is central to the question of whether a global ranking is feasible. This notion is completely natural from the vector calculus or topology point-of-view, they correspond to solutions of the Helmholtz equation or homology classes. However, it will be hard to define harmonic ranking directly in social choice theory without this insight, and we suspect that it is the reason why the notion of harmonic ranking has never been discussed in existing studies of ranking in social choice theory and other fields.

2 Statistical ranking on graphs

The central problem discussed in this paper is that of determining a global ranking from a dataset comprising a number of alternatives ranked by a number of voters. This is a problem that has received attention in fields including decision science [22, 23, 37], financial economics [4, 30], machine learning [6, 12, 18, 21], social choice [2, 39, 36], statistics [16, 27, 28, 33, 34], among others. Our objective towards statistical ranking is two-fold: like everybody else, we want to deduce a global ranking from the data whenever possible; but in addition to that, we also want to detect when the data does not permit a statistically meaningful global ranking and in which case analyze the obstructions to ‘global rankability’.

Let $V = \{1, \dots, n\}$ be the set of alternatives to be ranked and $\Lambda = \{1, \dots, m\}$ be a set of voters. The implicit assumption is that each voter would have *rated*, i.e. assigned cardinal scores or given an ordinal ordering to, a small fraction of the alternatives. But no matter how incomplete the rated portion is, one may always convert such ratings into *pairwise rankings* that have no missing values with the following recipe. For each voter $\alpha \in \Lambda$, the *pairwise ranking matrix* of α is a skew-symmetric matrix $Y^\alpha \in \mathbb{R}^{n \times n}$, i.e. for each ordered pair $(i, j) \in V \times V$, we have

$$Y_{ij}^\alpha = -Y_{ji}^\alpha.$$

Informally, Y_{ij}^α measures the ‘degree of preference’ of the i th alternative over the j th alternative held by the α th voter. Studies of ranking problems in different disciplines have led to rather different ways of quantifying such ‘degree of preference’. In Sect. 2.3, we will see several ways of defining Y_{ij}^α (as score difference, score ratio, and score ordering) coming from decision science, machine learning, social choice theory, and statistics. If the voter α did not compare alternatives i and j , then Y_{ij}^α is considered a missing value and set to be 0 for convenience; this manner of handling missing values allows Y^α to be a skew-symmetric matrix for each $\alpha \in \Lambda$. Nevertheless we could have assigned any arbitrary value or a non-numerical symbol to represent missing values, and this would not have affected our algorithmic results because of our use of the following weight function.

Define the *weight function* $w : \Lambda \times V \times V \rightarrow [0, \infty)$ via

$$w_{ij}^\alpha = w(\alpha, i, j) = \begin{cases} 1 & \text{if } \alpha \text{ made a pairwise comparison for } \{i, j\}, \\ 0 & \text{otherwise.} \end{cases}$$

Therefore $w_{ij}^\alpha = 0$ iff Y_{ij}^α is a missing value. Note that $W^\alpha = [w_{ij}^\alpha]$ is a symmetric $\{0, 1\}$ -valued matrix; but more generally, w_{ij}^α may be chosen as the capacity (in the graph theoretic sense) if there are multiple comparisons of i and j by voter α . The pairs (i, j) for which $w(\alpha, i, j) = 1$ for some $\alpha \in \Lambda$ are known as *crucial pairs* in the machine learning literature.

Our approach towards statistical ranking is to minimize a weighted sum of pairwise loss of a global ranking on the given data over a model class \mathcal{M} of all global rankings. We begin with a simple sum-of-squares loss function,

$$\min_{X \in \mathcal{M}_G} \sum_{\alpha, i, j} w_{ij}^\alpha (X_{ij} - Y_{ij}^\alpha)^2, \tag{1}$$

where the *model class* \mathcal{M}_G is the set of rank-2 skew-symmetric matrices,

$$\mathcal{M}_G = \{X \in \mathbb{R}^{n \times n} \mid X_{ij} = s_j - s_i, s : V \rightarrow \mathbb{R}\}. \tag{2}$$

Any $X \in \mathcal{M}_G$ induces a global ranking on the alternatives $1, \dots, n$ via the rule $i \succeq j$ iff $s_i \geq s_j$. Note that ties, i.e. $i \succeq j$ and $j \succeq i$, are allowed and this happens precisely when $s_i = s_j$.

For ranking data given in terms of cardinal scores, this simple scheme preserves the magnitudes of the ratings, instead of merely the ordering, when we have globally consistent data (see Definition 3). Moreover, it may also be computed more easily than many other loss functions although the computational cost depends ultimately on the choice of \mathcal{M} . For example, Kemeny optimization in classical social choice theory, which is known to be NP-hard [17], may be realized as a special case where $Y_{ij}^\alpha \in \{\pm 1\}$ and \mathcal{M} is the *Kemeny* model class,

$$\mathcal{M}_K := \{X \in \mathbb{R}^{n \times n} \mid X_{ij} = \text{sign}(s_j - s_i), s : V \rightarrow \mathbb{R}\}. \tag{3}$$

The function $\text{sign} : \mathbb{R} \rightarrow \{\pm 1\}$ takes nonnegative numbers to 1 and negative numbers to -1 . A binary valued Y_{ij}^α is the standard scenario in binary pairwise comparisons [1, 2, 13, 21, 28]; in this context, a global ranking is usually taken to be synonymous as a Kemeny optimal. We will discuss Kemeny optimization in greater details in Sect. 6.

2.1 Pairwise comparison graphs and pairwise ranking flows

A graph structure arises naturally from ranking data as follows. Let $G = (V, E)$ be an undirected graph whose vertex set is V , the set of alternatives to be ranked, and whose edge set is

$$E = \left\{ \{i, j\} \in \binom{V}{2} \mid \sum_{\alpha} w_{ij}^\alpha > 0 \right\}, \tag{4}$$

i.e. the set of pairs $\{i, j\}$ where pairwise comparisons have been made. We call such G a *pairwise comparison graph*. One can further associate weights on the edges as capacity, e.g. $w_{ij} = \sum_{\alpha} w_{ij}^\alpha$.

A pairwise ranking can be viewed as *edge flows* on G , i.e. a function $X : V \times V \rightarrow \mathbb{R}$ that satisfies

$$\begin{cases} X(i, j) = -X(j, i) & \text{if } \{i, j\} \in E, \\ X(i, j) = 0 & \text{otherwise.} \end{cases} \tag{5}$$

It is clear that a skew-symmetric matrix $[X_{ij}]$ induces an edge flow and vice versa. So henceforth we will not distinguish between edge flows and skew-symmetric matrices and will often write X_{ij} in place of $X(i, j)$.

Borrowing terminologies from vector calculus, an edge flow of the form $X_{ij} = s_j - s_i$, i.e. $X \in \mathcal{M}_G$, may be regarded as the *gradient* of a *potential* function $s : V \rightarrow \mathbb{R}$ (or *negative potential*, depending on sign convention). In the language of ranking, a potential function is a *score function* or *utility function* on the set of alternatives, assigning a score $s(i) = s_i$ to alternative i . Note that any such function defines a *global ranking* as discussed after (2). To be precise, we will define gradient as follows.

Definition 1 The **combinatorial gradient** operator maps a potential function on the vertices $s : V \rightarrow \mathbb{R}$ to an edge flow $\text{grad } s : V \times V \rightarrow \mathbb{R}$ via

$$(\text{grad } s)(i, j) = s_j - s_i. \tag{6}$$

An edge flow that has this form will be called a **gradient flow**.

In other words, the combinatorial gradient takes global rankings to pairwise rankings. Pairwise rankings that arise in this manner will be called *globally consistent* (formally defined in Definition 3). Given a globally consistent pairwise ranking X , we can easily solve $\text{grad}(s) = X$ to determine a score function s (up to an additive constant), and from s we can obtain a global ranking of the alternatives in the manner

described after (2). Observe that the set of all globally consistent pairwise rankings in (2) may be written as $\mathcal{M}_G = \{\text{grad } s \mid s : V \rightarrow \mathbb{R}\} = \text{im}(\text{grad})$.

For convenience, we will drop the adjective ‘combinatorial’ from ‘combinatorial gradient’. We may sometimes also drop the adjective ‘pairwise’ in ‘globally consistent pairwise ranking’ when there is no risk of confusion.

The optimization problem (1) can be rewritten in the form of a weighted l_2 -minimization on a pairwise comparison graph

$$\min_{X \in \mathcal{M}_G} \|X - \bar{Y}\|_{2,w}^2 = \min_{X \in \mathcal{M}_G} \left[\sum_{\{i,j\} \in E} w_{ij} (X_{ij} - \bar{Y}_{ij})^2 \right] \quad (7)$$

where

$$w_{ij} := \sum_{\alpha} w_{ij}^{\alpha} \quad \text{and} \quad \bar{Y}_{ij} := \frac{\sum_{\alpha} w_{ij}^{\alpha} Y_{ij}^{\alpha}}{\sum_{\alpha} w_{ij}^{\alpha}}. \quad (8)$$

A minimizer thus corresponds to an l_2 -projection of a pairwise ranking edge flow \bar{Y} onto the space of gradient flows. We note that $W = [w_{ij}] = \sum_{\alpha} W^{\alpha}$ is a symmetric nonnegative-valued matrix. This choice of W is not intended to be rigid. One could for example define W to incorporate prior knowledge of the relative importance of paired comparisons as judged by voters.

Combinatorial Hodge theory provides a geometric interpretation of the minimizer and residual of (7). Before going further, we present several examples of pairwise ranking arising from applications.

2.2 Pairwise rankings

Humans are unable to make accurate preference judgement on even moderately large sets. In fact, it has been argued that most people can rank only between 5 and 9 alternatives at a time [31]. This is probably why many rating scales (e.g. the ones used by Amazon, eBay, Netflix, YouTube) are all based on a 5-star scale. Hence one expects large human-generated ranking data to be at best partially ordered (with chains of lengths mostly between 5 and 9, if [31] is right). For most people, it is a harder task to rank or rate 20 movies than to compare the movies a pair at a time. In certain settings such as tennis tournaments and wine tasting, only pairwise comparisons are possible. Pairwise comparison methods, which involve the smallest partial rankings, is thus natural for analyzing ranking data.

Pairwise comparisons also help reduce bias due to the arbitrariness of rating scale by adopting a relative measure. As we will see in Sect. 2.3, pairwise comparisons provide a way to handle missing values, which are expected because of the general lack of incentives or patience for a human to process a large dataset. For these reasons, pairwise comparison methods have been popular in psychology, management science, social choice theory, and statistics [2, 13, 23, 28, 37, 41]. Such methods are also getting increasing attention from the machine learning community as they may be

adapted for studying classification problems [18,20,21]. We will present two very different instances where pairwise rankings arise: recommendation systems and exchange economic systems.

2.2.1 Recommendation systems

The generic scenario in recommendation systems is that there are m voters rating n alternatives. For example, in the Netflix context, viewers rate movies on a scale of 5 stars [6]; in financial markets, analysts rate stocks or securities with five classes of recommendations [4]. In these cases, we let $A = [a_{\alpha i}] \in \mathbb{R}^{m \times n}$ represent the voter-alternative matrix. A typically has a large number of missing values; for example, the dataset that Netflix released for its prize competition contains a viewer-movie matrix with 99% of its values missing. The standard problem here is to predict these missing values from the given data but we caution the reader again that this is *not* the problem addressed in our paper. Instead of estimating the missing values of A , we want to learn a global ranking of the alternatives from A , without having to first estimate the missing values. We note here the striking difference that if one considers pairwise rankings instead, then only 0.22% of the pairwise comparison values are missing from the Netflix dataset. This partly motivates our discussion in Sect. 2.3. An actual numerical experiment can be found in Sect. 7.1.

2.2.2 Exchange economic systems

A purely exchange economic system may be described by a graph $G = (V, E)$ with vertex set $V = \{1, \dots, n\}$ representing the n goods and edge set $E \subseteq \binom{V}{2}$ representing feasible pairwise transactions. If the market is complete in the sense that every pair of goods is exchangeable, then G is a complete graph. Suppose the exchange rate between the i th and j th goods is given by

$$1 \text{ unit } i = a_{ij} \text{ unit } j, \quad a_{ij} > 0.$$

Then the exchange rate matrix $A = [a_{ij}]$ is a *reciprocal matrix* (possibly with missing values), i.e. $a_{ij} = 1/a_{ji}$ for all $i, j \in V$. The reciprocal matrix was first used for paired preference aggregation by Saaty [37] and later by Ma [30] for currency exchange analysis. The problem of *pricing* is to look for a *universal equivalent* that measures the values of goods (this is in fact an abstraction of the concept of money), i.e. $\pi : V \rightarrow \mathbb{R}$ such that

$$a_{ij} = \frac{\pi_j}{\pi_i}.$$

In complete markets where G is a complete graph, there exists a universal equivalent if and only if the market is *triangular arbitrage-free*, i.e. $a_{ij}a_{jk} = a_{ik}$ for all distinct $i, j, k \in V$; since in this case the transaction path $i \rightarrow j \rightarrow k$ provides no gain nor loss over a direct exchange $i \rightarrow k$.

Such a purely exchange economic system may be transformed into a pairwise ranking problem via the logarithmic map,

$$X_{ij} = \log a_{ij}.$$

The triangular arbitrage-free condition is then equivalent to the transitivity condition in (14), i.e.

$$X_{ij} + X_{jk} + X_{ki} = 0.$$

So asking if a universal equivalent exists is the same as asking if a global ranking $s : V \rightarrow \mathbb{R}$ exists so that $X_{ij} = s_j - s_i$ with $s_i = \log \pi_i$. This partly motivates our discussion in Sect. 2.4. An actual numerical experiment can be found in Sect. 7.2.

2.3 Filling in missing values: average pairwise ranking

While the available raw ranking data $A = [a_{\alpha i}] \in \mathbb{R}^{m \times n}$ may be highly incomplete (cf. Sect. 2.2.1), one may aggregate over all voters to get a pairwise ranking matrix Y that is usually much more complete. Each of the following four statistics may be regarded as a form of “average pairwise ranking” over all voters.

1. **Arithmetic mean of score differences:** The score difference refers to $Y_{ij}^\alpha = a_{\alpha j} - a_{\alpha i}$. The arithmetic mean over all customers who have rated both i and j is

$$\bar{Y}_{ij} = \frac{\sum_{\alpha} (a_{\alpha j} - a_{\alpha i})}{\#\{\alpha \mid a_{\alpha i}, a_{\alpha j} \text{ exist}\}}. \tag{9}$$

This is translation invariant.

2. **Geometric mean of score ratios:** The score ratio refers to $Y_{ij}^\alpha = a_{\alpha j} / a_{\alpha i}$. Assume that $a_{\alpha i} > 0$. The (log) geometric mean over all customers who have rated both i and j is

$$\bar{Y}_{ij} = \frac{\sum_{\alpha} (\log a_{\alpha j} - \log a_{\alpha i})}{\#\{\alpha \mid a_{\alpha i}, a_{\alpha j} \text{ exist}\}}. \tag{10}$$

This is scale invariant.

3. **Binary comparison:** Here $Y_{ij}^\alpha = \text{sign}(a_{\alpha j} - a_{\alpha i})$. Its average is the probability difference that the alternative j is preferred to i than the other way round,

$$\bar{Y}_{ij} = \Pr\{\alpha \mid a_{\alpha j} > a_{\alpha i}\} - \Pr\{\alpha \mid a_{\alpha j} < a_{\alpha i}\}. \tag{11}$$

This is invariant up to a monotone transformation.

4. **Logarithmic odds ratio:** As in the case of binary comparison, except that we adopt a logarithmic scale

$$\bar{Y}_{ij} = \log \frac{\Pr\{\alpha \mid a_{\alpha j} \geq a_{\alpha i}\}}{\Pr\{\alpha \mid a_{\alpha j} \leq a_{\alpha i}\}}. \tag{12}$$

This is also invariant up to a monotone transformation.

The first model leads to the concept of *position-rules* in social choice theory [36] and it has also been used in machine learning recently [12]. The second model has appeared in multi-criteria decision theory [37]. The third and fourth models are known as *linear model* [34] and *Bradley-Terry model* [7] respectively in the statistics and psychology literature. There are other plausible choices for defining \bar{Y}_{ij} , e.g. [33,41], but we will not discuss more of them here. It suffices to note that there is a rich variety of techniques to preprocess raw ranking data into the pairwise ranking edge flow \bar{Y}_{ij} that serves as input to our Hodge theoretic method. However, it should be noted that the l_2 -optimization on graphs in (7) may be applied with any of the four choices above since only the knowledge of \bar{Y}_{ij} is required but the sum-of-squares and Kemeny optimization in (1) and (3) require (respectively) the original score difference and original ordering be known for each voter.

2.4 Measuring local inconsistency: combinatorial curl

Upon constructing an average pairwise ranking from the raw data, we need a statistics to quantify its inconsistency. Again we will borrow a terminology from vector calculus and define a notion of *combinatorial curl* as a measure of triangular inconsistency.

Given a pairwise ranking represented as an edge flow X on a graph $G = (V, E)$, we expect a ‘consistency’ property: following a loop $i \rightarrow j \rightarrow \dots \rightarrow i$ where each edge is in E , the amount of scores raised should be equal to the amount of scores lowered; so after a loop of comparisons we should return to the same score on the same alternative. Since the simplest loop is a triangular loop $i \rightarrow j \rightarrow k \rightarrow i$, the ‘basic unit’ of inconsistency ought to be triangular in nature, which leads us to Definition 2.

We will first define a notion analogous to edge flows. The *triangular flow* on G is a function $\Phi : V \times V \times V \rightarrow \mathbb{R}$ that satisfies

$$\Phi(i, j, k) = \Phi(j, k, i) = \Phi(k, i, j) = -\Phi(j, i, k) = -\Phi(i, k, j) = -\Phi(k, j, i),$$

i.e. an odd permutation of the arguments of Φ changes its sign while an even permutation preserves its sign.⁴ A triangular flow describes *triplewise rankings* in the same way an edge flow describes pairwise rankings.

Definition 2 Let X be an edge flow on a graph $G = (V, E)$. Let

$$T(E) := \left\{ \{i, j, k\} \in \binom{V}{3} \mid \{i, j\}, \{j, k\}, \{k, i\} \in E \right\}$$

⁴ A triangular flow is an alternating 3-tensor and may be represented as a skew-symmetric hypermatrix $[\Phi_{ijk}] \in \mathbb{R}^{n \times n \times n}$, much like an edge flow is an alternating 2-tensor and may be represented by a skew-symmetric matrix $[X_{ij}] \in \mathbb{R}^{n \times n}$. We will often write Φ_{ijk} in place of $\Phi(i, j, k)$.

be the collection of triangles with every edge in E . We define the **combinatorial curl** operator that maps edge flows to triangular flows by

$$(\text{curl } X)(i, j, k) = \begin{cases} X_{ij} + X_{jk} + X_{ki} & \text{if } \{i, j, k\} \in T(E), \\ 0 & \text{otherwise.} \end{cases} \tag{13}$$

In other words, the combinatorial curl takes pairwise rankings to triplewise rankings. Again, we will drop the adjective ‘combinatorial’ when there is no risk of confusion. The skew-symmetry of X , i.e. $X_{ij} = -X_{ji}$, guarantees that $\text{curl } X$ is a triangular flow, i.e.

$$\begin{aligned} (\text{curl } X)(i, j, k) &= (\text{curl } X)(j, k, i) = (\text{curl } X)(k, i, j) \\ &= -(\text{curl } X)(j, i, k) = -(\text{curl } X)(i, k, j) \\ &= -(\text{curl } X)(k, j, i). \end{aligned}$$

The curl of a pairwise ranking measures its triangular inconsistency. This extends the *consistency index* of Kendall and Smith [28], which counts the number of circular triads, from ordinal settings to cardinal settings. Note that for binary pairwise ranking where $X_{ij} \in \{\pm 1\}$, the absolute value $|(\text{curl } X)(i, j, k)|$ may only take two values, 1 or 3. The triangle $\{i, j, k\} \in T(E)$ contains a cyclic ranking or *circular triad* if and only if $|(\text{curl } X)(i, j, k)| = 3$. If G is a complete graph, the number of circular triads is known [28] to be

$$N = \frac{n}{24}(n^2 - 1) - \frac{1}{8} \sum_i \left[\sum_j X_{ij} \right]^2.$$

For ranking data given in terms of cardinal scores and that is generally incomplete, curl plays an extended role in addition to just quantifying the triangular inconsistency. We now formally define some ranking theoretic notions in terms of the combinatorial gradient and combinatorial curl.

Definition 3 Let $X : V \times V \rightarrow \mathbb{R}$ be a pairwise ranking edge flow on a pairwise comparison graph $G = (V, E)$.

1. X is called **consistent** on $\{i, j, k\} \in T(E)$ if it is curl-free on $\{i, j, k\}$, i.e.

$$(\text{curl } X)(i, j, k) = X_{ij} + X_{jk} + X_{ki} = 0.$$


This implies that $\text{curl}(X)(\sigma(i), \sigma(j), \sigma(k)) = 0$ for every permutation σ .

2. X is called **globally consistent** if it is a gradient of a score function, i.e.

$$X = \text{grad } s \quad \text{for some } s : V \rightarrow \mathbb{R}.$$

3. X is called **locally consistent** or **triangularly consistent** if it is curl-free on every triangle in $T(E)$, i.e. every 3-clique of G .

Fig. 1 A harmonic pairwise ranking, which is locally consistent on every triangle but inconsistent along the loop $A \rightarrow B \rightarrow C \rightarrow D \rightarrow E \rightarrow F \rightarrow A$



4. X is called a **cyclic ranking** if it contains any **inconsistencies**, i.e. there exist $i, j, k, \dots, p, q \in V$ such that

$$X_{ij} + X_{jk} + \dots + X_{pq} + X_{qi} \neq 0.$$

Clearly any gradient flow must be curl-free everywhere, i.e. the well-known identity in vector calculus

$$\text{curl} \circ \text{grad} = 0$$

is also true for combinatorial curl and combinatorial gradient (a special case of Lemma 1). So global consistency implies local consistency. A qualified converse may be deduced from the Hodge decomposition theorem (see also Theorem 4): a curl-free flow on a complete graph must necessarily be a gradient flow, or putting it another way, a locally consistent pairwise ranking must necessarily be a globally consistent pairwise ranking when there are no missing values, i.e. if the pairwise comparison graph is a complete graph (every pair of alternatives has been compared).

When G is an incomplete graph, the condition that X is curl-free on every triangle in the graph will not be enough to guarantee that it is a gradient flow. The reason lies in that curl only takes into account triangular inconsistencies; but since there are missing edges in the pairwise comparison graph G , it is possible that non-triangular cyclic rankings of lengths greater than three can occur. For example, Fig. 1 shows a pairwise ranking that is locally consistent on every triangle but globally inconsistent, since it contains a cyclic ranking of length six. Fortunately, Hodge decomposition theorem will tell us that all such cyclic rankings lie in a subspace of *harmonic rankings*, which can be characterized as the kernel of some combinatorial Laplacians.

3 A matrix theoretic view of Hodge decomposition

We will see in this section that edge flows, gradient flows, harmonic flows, and curl flows can all be represented as specially structured skew-symmetric matrices. In this framework, the Hodge decomposition theorem may be viewed as an orthogonal direct sum decomposition of the space of skew-symmetric matrices into three subspaces. A formal geometric treatment of combinatorial Hodge theory will be given in Sect. 4.

Recall that a matrix $X \in \mathbb{R}^{n \times n}$ is said to be *skew-symmetric* if $X_{ij} = -X_{ji}$ for all $i, j \in V := \{1, \dots, n\}$. One knows from linear algebra that any square matrix A may be written uniquely as a sum of a symmetric and a skew-symmetric matrix,

$$A = \frac{1}{2}(A + A^\top) + \frac{1}{2}(A - A^\top).$$

We will denote⁵

$$\mathcal{A} := \{X \in \mathbb{R}^{n \times n} \mid X^\top = -X\}, \quad \text{and} \quad \mathcal{S} := \{X \in \mathbb{R}^{n \times n} \mid X^\top = X\}.$$

It is perhaps interesting to note that semidefinite programming takes place in the cone of symmetric positive definite matrices in \mathcal{S} but the optimization problems in this paper take place in the exterior space \mathcal{A} .

One simple way to construct a skew-symmetric matrix is to take a vector $s = [s_1, \dots, s_n]^\top \in \mathbb{R}^n$ and define X by

$$X_{ij} := s_i - s_j.$$

Note that if $X \neq 0$, then $\text{rank}(X) = 2$ since it can be expressed as $se^\top - es^\top$ with $e := [1, \dots, 1]^\top \in \mathbb{R}^n$. These are in a sense the simplest skew-symmetric matrices—they have the lowest possible rank among non-zero skew-symmetric matrices (recall that the rank of a skew-symmetric matrix is necessarily even). In this paper, we will call these *gradient* matrices and denote them collectively by \mathcal{M}_G ,

$$\mathcal{M}_G := \{X \in \mathcal{A} \mid X_{ij} = s_i - s_j \text{ for some } s \in \mathbb{R}^n\}.$$

For $T \subseteq \binom{V}{3}$, we define the set of *T-consistent* matrices as

$$\mathcal{M}_T := \{X \in \mathcal{A} \mid X_{ij} + X_{jk} + X_{ki} = 0 \text{ for all } \{i, j, k\} \in T\}. \tag{14}$$

We can immediately observe every $X \in \mathcal{M}_G$ is *T-consistent* for any $T \subseteq \binom{V}{3}$, i.e. $\mathcal{M}_G \subseteq \mathcal{M}_T$. Conversely, a matrix X that satisfies

$$X_{ij} + X_{jk} + X_{ki} = 0 \text{ for every triple } \{i, j, k\} \in \binom{V}{3}.$$

⁵ More common notations for \mathcal{A} are $\mathfrak{so}_n(\mathbb{R})$ (Lie algebra of $\text{SO}(n)$) and $\wedge^2(\mathbb{R}^n)$ (second exterior product of \mathbb{R}^n).

is necessarily a gradient matrix, i.e.

$$\mathcal{M}_G = \mathcal{M}_{\binom{Y}{3}}. \tag{15}$$

Given $T \subseteq \binom{Y}{3}$, it is straightforward to verify that both \mathcal{M}_G and \mathcal{M}_T are subspaces of $\mathbb{R}^{n \times n}$. The preceding discussions then imply

$$\mathcal{M}_G \subseteq \mathcal{M}_T \subseteq \mathcal{A}. \tag{16}$$

Since these are strict inclusions in general, several complementary subspaces arise naturally. With respect to the usual inner product $\langle X, Y \rangle = \text{tr}(X^\top Y) = \sum_{i,j} X_{ij} Y_{ij}$, we obtain orthogonal complements of \mathcal{M}_G and \mathcal{M}_T in \mathcal{A} as well as the orthogonal complement of \mathcal{M}_G in \mathcal{M}_T , which we denote by \mathcal{M}_H :

$$\mathcal{A} = \mathcal{M}_G \oplus \mathcal{M}_G^\perp, \quad \mathcal{A} = \mathcal{M}_T \oplus \mathcal{M}_T^\perp, \quad \mathcal{M}_T = \mathcal{M}_G \oplus \mathcal{M}_H.$$

Elements of \mathcal{M}_H will be called *harmonic* matrices as we will see that they are discrete analogues of solutions to the Laplace equation (or, more accurately, the Helmholtz equation). An alternative characterization of \mathcal{M}_H is

$$\mathcal{M}_H = \mathcal{M}_T \cap \mathcal{M}_G^\perp,$$

which may be viewed as a discrete analogue of the condition of being simultaneously curl-free and divergence-free. More generally, this discussion applies to any weighted inner product $\langle X, Y \rangle_w = \sum_{i,j} w_{ij} X_{ij} Y_{ij}$. The five subspaces $\mathcal{M}_G, \mathcal{M}_T, \mathcal{M}_H, \mathcal{M}_T^\perp, \mathcal{M}_G^\perp$ of \mathcal{A} play a central role in our techniques. As we shall see later, the Helmholtz decomposition in Theorem 2 may be viewed as the orthogonal direct sum decomposition

$$\mathcal{A} = \mathcal{M}_G \oplus \mathcal{M}_H \oplus \mathcal{M}_T^\perp.$$

4 Combinatorial Hodge theory

In this section we will give a brief introduction to combinatorial Hodge theory, paying special attention to its relevance in statistical ranking. One may wonder why we do not rely on our relatively simple matrix view in Sect. 3. The reasons are two fold: firstly, important geometric insights are lost when the actual motivations behind the matrix picture are disregarded; and secondly, the matrix approach applies only to the case of 2-dimensional simplicial complex but combinatorial Hodge theory extends to any k -dimensional simplicial complex. While so far we did not use any simplicial complex of dimension higher than two in our study of statistical ranking, it is conceivable that higher-dimensional simplicial complex could play a role in future studies.

4.1 Extension of pairwise comparison graph to simplicial complex

Let $G = (V, E)$ be a pairwise comparison graph. To characterize the triangular inconsistency or curl, one needs to study the triangles formed by the 3-cliques⁶, i.e. the set

$$T(E) := \left\{ \{i, j, k\} \in \binom{V}{3} \mid \{i, j\}, \{j, k\}, \{k, i\} \in E \right\}.$$

A combinatorial object of the form (V, E, T) where $E \subseteq \binom{V}{2}$, $T \subseteq \binom{V}{3}$, and $\{i, j\}, \{j, k\}, \{k, i\} \in E$ for all $\{i, j, k\} \in T$ is called a 2-dimensional simplicial complex. This is a generalization of a graph, which is a 1-dimensional simplicial complex. In particular, given a graph $G = (V, E)$, the 2-dimensional simplicial complex $(V, E, T(E))$ is often called the *3-clique complex* of G .

More generally, a *simplicial complex* (V, Σ) is a vertex set $V = \{1, \dots, n\}$ together with a collection Σ of subsets of V that is closed under inclusion, i.e. if $\tau \in \Sigma$ and $\sigma \subset \tau$, then $\sigma \in \Sigma$. The elements in Σ are called *simplices*. For example, a 0-simplex is just an element $i \in V$ (recall that we do not distinguish between $\binom{V}{1}$ and V), a 1-simplex is a pair $\{i, j\} \in \binom{V}{2}$, a 2-simplex is a triple $\{i, j, k\} \in \binom{V}{3}$, and so on. For $k \leq n$, a k -simplex is a $(k + 1)$ -element set in $\binom{V}{k+1}$ and $\Sigma_k \subset \binom{V}{k+1}$ will denote the set of all k -simplices in Σ . In the previous paragraph, $\Sigma_0 = V$, $\Sigma_1 = E$, $\Sigma_2 = T$, and $\Sigma = V \cup E \cup T$. In general, given any undirected graph $G = (V, E)$, one obtains a $(k - 1)$ -dimensional simplicial complex $K_G^k := (V, \Sigma)$ called the *k-clique complex*⁷ of G by ‘filling in’ all its j -cliques for $j = 1, \dots, k$, or more precisely, by setting $\Sigma = \{j\text{-cliques of } G \mid j = 1, \dots, k\}$. When k is maximal K_G^k is written K_G and called the *clique complex* of G .

In this paper, we will mainly concern ourselves with studying the 3-clique complex $K_G^3 = (V, E, T(E))$ where G is a pairwise comparison graph. Note that we could also look at the simplicial complex $(V, E, T_\gamma(E))$ where

$$T_\gamma(E) := \left\{ \{i, j, k\} \in T(E) \mid |X_{ij} + X_{jk} + X_{ki}| \leq \gamma \right\}$$

where $0 \leq \gamma \leq \infty$. For $\gamma = \infty$, we get K_G^3 but for general γ we get a subcomplex of K_G^3 . We have found this to be a useful multiscale characterization of inconsistencies but detailed discussion will be left to future work.

4.2 Cochains, coboundary maps, and combinatorial laplacians

We will now introduce some discrete exterior calculus on a simplicial complex where potential functions (scores or utility), edge flow (pairwise ranking), triangular flow (triplewise ranking), gradient (global ranking induced by scores), curl (local inconsistency) become just special cases of a much more general framework. We will now also define the notions of *combinatorial divergence* and *combinatorial Laplacians*.

⁶ Recall that a k -clique of G is just a complete subgraph of G with k vertices.

⁷ Note that a k -clique is a $(k - 1)$ -simplex.

The 0-dimensional combinatorial Laplacian is the usual graph Laplacian. The case of greatest interest here is the 1-dimensional combinatorial Laplacian, which we propose calling the *graph Helmholtzian*.

Definition 4 Let K be a simplicial complex and recall that Σ_k denotes its set of k -simplices. A k -dimensional **cochain** is a real-valued function on $(k + 1)$ -tuples of vertices that is alternating on each of the k -simplex and 0 otherwise, i.e. $f : V^{k+1} \rightarrow \mathbb{R}$ such that

$$f(i_{\sigma(0)}, \dots, i_{\sigma(k)}) = \text{sign}(\sigma) f(i_0, \dots, i_k),$$

for all $(i_0, \dots, i_k) \in V^{k+1}$ and all $\sigma \in \mathfrak{S}_{k+1}$, the permutation group on $k + 1$ elements, and that

$$f(i_0, \dots, i_k) = 0 \quad \text{if } \{i_0, \dots, i_k\} \notin \Sigma_k.$$

The set of all k -cochains on K is denoted $C^k(K, \mathbb{R})$.

For simplicity we often write C^k for $C^k(K, \mathbb{R})$. So C^0 is the space of potential functions, C^1 the space of edge flows, and C^2 the space of triangular flows.

The k -cochain space C^k can be given a choice of inner product. In view of the weighted l_2 -minimization for our statistical ranking problem (7), we will define the following inner product on C^1 ,

$$\langle X, Y \rangle_w = \sum_{\{i,j\} \in E} w_{ij} X_{ij} Y_{ij}, \tag{17}$$

for all edge flows $X, Y \in C^1$. In the context of a pairwise comparison graph G , it may not be immediately clear why this defines an inner product since we have noted after (8) that $W = [w_{ij}]$ is only a nonnegative matrix and it is possible that some entries are 0. However observe that by definition $w_{ij} = 0$ iff no voters have rated both alternatives i and j and therefore $\{i, j\} \notin E$ by (4) and so any edge flow X will automatically have $X_{ij} = 0$ by (5). Hence we indeed have that $\langle X, X \rangle_w = 0$ iff $X = 0$, as required for an inner product (the other properties are trivial to check).

The operators grad and curl are special instances of coboundary maps:

Definition 5 The k th **coboundary** operator $\delta_k : C^k(K, \mathbb{R}) \rightarrow C^{k+1}(K, \mathbb{R})$ is the linear map that takes a k -cochain $f \in C^k$ to a $(k + 1)$ -cochain $\delta_k f \in C^{k+1}$ defined by

$$(\delta_k f)(i_0, i_1, \dots, i_{k+1}) := \sum_{j=0}^{k+1} (-1)^j f(i_0, \dots, i_{j-1}, i_{j+1}, \dots, i_{k+1}).$$

Note that i_j is omitted from the j th summand, i.e. coboundary maps compute an alternating difference with one input left out. So $\delta_0 = \text{grad}$, i.e. $(\delta_0 s)(i, j) = s_j - s_i$, and $\delta_1 = \text{curl}$, i.e. $(\delta_1 X)(i, j, k) = X_{ij} + X_{jk} + X_{ki}$.

Given a choice of an inner product $\langle \cdot, \cdot \rangle_k$ on C^k , we may define the adjoint operator of the coboundary map, $\delta_k^* : C^{k+1} \rightarrow C^k$ in the usual manner, i.e. $\langle \delta_k f_k, g_{k+1} \rangle_{k+1} = \langle f_k, \delta_k^* g_{k+1} \rangle_k$.

Definition 6 The **combinatorial divergence** operator $\text{div} : C^1(K, \mathbb{R}) \rightarrow C^0(K, \mathbb{R})$ is the adjoint of $\delta_0 = \text{grad}$, i.e.

$$\text{div} := -\delta_0^*. \tag{18}$$

Divergence will appear in the minimum norm solution to (7) and can be used to characterize \mathcal{M}_G^\perp . As usual, we will drop the adjective ‘combinatorial’ when there is no cause for confusion.

For statistical ranking, it suffices to consider $k = 0, 1, 2$. Let G be a pairwise comparison graph and K_G its clique complex.⁸ The cochain maps,

$$C^0(K_G, \mathbb{R}) \xrightarrow{\delta_0} C^1(K_G, \mathbb{R}) \xrightarrow{\delta_1} C^2(K_G, \mathbb{R}) \tag{19}$$

and their adjoints,

$$C^0(K_G, \mathbb{R}) \xleftarrow{\delta_0^*} C^1(K_G, \mathbb{R}) \xleftarrow{\delta_1^*} C^2(K_G, \mathbb{R}), \tag{20}$$

have the following ranking theoretic interpretations: C^0, C^1, C^2 represent spaces of score functions, pairwise rankings, triplewise rankings, and

$$\begin{array}{ccccc} \text{scores} & \xrightarrow{\text{grad}} & \text{pairwise} & \xrightarrow{\text{curl}} & \text{triplewise}, \\ \text{scores} & \xleftarrow{-\text{div}=\text{grad}^*} & \text{pairwise} & \xleftarrow{\text{curl}^*} & \text{triplewise}. \end{array}$$

Recall that combinatorial gradient, curl, and divergence are defined by

$$\begin{aligned} (\text{grad } s)(i, j) &= (\delta_0 s)(i, j) = s_j - s_i, \\ (\text{curl } X)(i, j, k) &= (\delta_1 X)(i, j, k) = X_{ij} + X_{jk} + X_{ki}, \\ (\text{div } X)(i) &= -(\delta_0^* X)(i) = \sum_{j \text{ s.t. } \{i, j\} \in E} w_{ij} X_{ij} \end{aligned}$$

with respect to the inner product $\langle X, Y \rangle_w = \sum_{\{i, j\} \in E} w_{ij} X_{ij} Y_{ij}$ on C^1 .

As an aside, it is perhaps worth pointing out that there is no special name for the adjoint of curl coming from physics because in 3-space, C^1 may be identified with C^2 via a property called *Hodge duality* and in which case curl is a self-adjoint operator, i.e. $\text{curl}^* = \text{curl}$. This will not be true in our case.

⁸ It does not matter whether we consider K_G or K_G^3 or indeed any K_G^k where $k \geq 3$; the higher-dimensional k -simplices where $k \geq 3$ do not play a role in the coboundary maps $\delta_0, \delta_1, \delta_2$.

If we represent functions on vertices by n -vectors, edge flows by $n \times n$ skew-symmetric matrices, and triangular flows by $n \times n \times n$ skew-symmetric hypermatrices, i.e.

$$\begin{aligned}
 C^0 &= \mathbb{R}^n, \\
 C^1 &= \{[X_{ij}] \in \mathbb{R}^{n \times n} \mid X_{ij} = -X_{ji}\} = \mathcal{A}, \\
 C^2 &= \{[\Phi_{ijk}] \in \mathbb{R}^{n \times n \times n} \mid \Phi_{ijk} = \Phi_{jki} = \Phi_{kij} = -\Phi_{jik} = -\Phi_{ikj} = -\Phi_{kji}\},
 \end{aligned}$$

then in the language of linear algebra introduced in Sect. 3,

$$\begin{aligned}
 \text{im}(\delta_0) &= \text{im}(\text{grad}) = \mathcal{M}_G, & \ker(\delta_1) &= \ker(\text{curl}) = \mathcal{M}_T, \\
 \ker(\delta_0^*) &= \ker(\text{div}) = \mathcal{M}_G^\perp, & \text{im}(\delta_1^*) &= \text{im}(\text{curl}^*) = \mathcal{M}_T^\perp,
 \end{aligned}$$

where $T = T(E)$.

Coboundary maps have the following important property.

Lemma 1 (Closedness) $\delta_{k+1} \circ \delta_k = 0$.

For $k = 0$, this and its adjoint are well-known identities in vector calculus,

$$\text{curl} \circ \text{grad} = 0, \quad \text{div} \circ \text{curl}^* = 0. \tag{21}$$

Ranking theoretically, the first identity simply says that a global ranking must be consistent.

We will now define combinatorial Laplacians (also known as Hodge Laplacians), higher-dimensional analogues of the graph Laplacian.

Definition 7 Let K be a simplicial complex. The k -dimensional **combinatorial Laplacian** is the operator $\Delta_k : C^k(K, \mathbb{R}) \rightarrow C^k(K, \mathbb{R})$ defined by

$$\Delta_k = \delta_k^* \circ \delta_k + \delta_{k-1} \circ \delta_{k-1}^*. \tag{22}$$

In particular,

$$\Delta_0 = \delta_0^* \circ \delta_0 = -\text{div} \circ \text{grad}$$

is a discrete analogue of the scalar Laplacian or Laplace operator while

$$\Delta_1 = \delta_1^* \circ \delta_1 + \delta_0 \circ \delta_0^* = \text{curl}^* \circ \text{curl} - \text{grad} \circ \text{div}$$

is a discrete analogue of the vector Laplacian or Helmholtz operator. In the context of graph theory, if $K = K_G$, Δ_0 is called the graph Laplacian [11] while Δ_1 we will call the *graph Helmholtzian*.

The combinatorial Laplacian has some well-known, important properties.

Lemma 2 Δ_k is a positive semidefinite operator. Furthermore, the dimension of $\ker(\Delta_k)$ is equal to k th Betti number of K .

We will call a cochain $f \in \ker(\Delta_k)$ *harmonic* since they are solutions to higher-dimensional analogue of the Laplace equation

$$\Delta_k f = 0.$$

Strictly speaking, the Laplace equation refers to $\Delta_0 f = 0$. The equation $\Delta_1 X = 0$ is really the *Helmholtz equation*. But nonetheless, we will still call an edge flow $X \in \ker(\Delta_1)$ a harmonic flow.

4.3 Hodge decomposition theorem and HodgeRank

We now state the main theorem in combinatorial Hodge theory.

Theorem 1 (Hodge Decomposition Theorem) $C^k(K, \mathbb{R})$ admits an orthogonal decomposition

$$C^k(K, \mathbb{R}) = \text{im}(\delta_{k-1}) \oplus \ker(\Delta_k) \oplus \text{im}(\delta_k^*).$$

Furthermore,

$$\ker(\Delta_k) = \ker(\delta_k) \cap \ker(\delta_{k-1}^*).$$

An elementary proof targeted at a computer science readership may be found in [19]. For completeness we include a proof here.

Proof We will use Lemma 1. First, $C^k = \text{im}(\delta_{k-1}) \oplus \ker(\delta_{k-1}^*)$. Since $\delta_k \delta_{k-1} = 0$, taking adjoint yields $\delta_{k-1}^* \delta_k^* = 0$, which implies that $\text{im}(\delta_k^*) \subseteq \ker(\delta_{k-1}^*)$. Therefore $\ker(\delta_{k-1}^*) = [\text{im}(\delta_k^*) \oplus \ker(\delta_k)] \cap \ker(\delta_{k-1}^*) = [\text{im}(\delta_k^*) \cap \ker(\delta_{k-1}^*)] \oplus [\ker(\delta_k) \cap \ker(\delta_{k-1}^*)] = \text{im}(\delta_k^*) \oplus [\ker(\delta_k) \cap \ker(\delta_{k-1}^*)]$. It remains to show that $\ker(\delta_k) \cap \ker(\delta_{k-1}^*) = \ker(\Delta_k) = \ker(\delta_{k-1} \delta_{k-1}^* + \delta_k^* \delta_k)$. Clearly $\ker(\delta_k) \cap \ker(\delta_{k-1}^*) \subseteq \ker(\Delta_k)$. For any $X = \delta_k^* \Phi \in \text{im}(\delta_k^*)$ where $0 \neq \Phi \in C^{k+1}$, Lemma 1 again implies $\delta_{k-1} \delta_{k-1}^* X = \delta_{k-1} \delta_{k-1}^* \delta_k^* \Phi = 0$, but $\delta_k^* \delta_k X = \delta_k^* \delta_k \delta_k^* \Phi \neq 0$, which implies that $\Delta_k X \neq 0$. Similarly for $X \in \text{im}(\delta_{k-1})$. Hence $\ker(\Delta_k) = \ker(\delta_k) \cap \ker(\delta_{k-1}^*)$. \square

While Hodge decomposition holds in general for any simplicial complex of any dimension k , the case $k = 1$ is often called the *Helmholtz decomposition theorem*.⁹ We will state it here for the special case of a clique complex.

Theorem 2 (Helmholtz Decomposition Theorem) Let $G = (V, E)$ be an undirected, unweighted graph and K_G be its clique complex. The space of edge flows on G , i.e. $C^1(K_G, \mathbb{R})$, admits an orthogonal decomposition

$$\begin{aligned} C^1(K_G, \mathbb{R}) &= \text{im}(\delta_0) \oplus \ker(\Delta_1) \oplus \text{im}(\delta_1^*) \\ &= \text{im}(\text{grad}) \oplus \ker(\Delta_1) \oplus \text{im}(\text{curl}^*). \end{aligned} \tag{23}$$

⁹ On a simply connected manifold, the continuous version of the Helmholtz decomposition theorem is just the fundamental theorem of vector calculus.




Fig. 2 Hodge/Helmholtz decomposition of pairwise rankings

Furthermore,

$$\ker(\Delta_1) = \ker(\delta_1) \cap \ker(\delta_0^*) = \ker(\text{curl}) \cap \ker(\text{div}). \tag{24}$$

The clique complex K_G above may be substituted with any K_G^k with $k \geq 3$ (see Footnote 8). The Eq. 24 says that an edge flow is harmonic iff it is both curl-free and divergence-free. Figure 2 illustrates (23).

To understand the significance of this theorem, we need to discuss the ranking theoretic interpretations of each subspace in the theorem. These interpretations constitute a method for analyzing ranking data and will collectively be called **HodgeRank**.

1. $\text{im}(\delta_0) = \text{im}(\text{grad})$ denotes the subspace of pairwise rankings that are the gradient flows of score functions. Thus this subspace comprises the *globally consistent* or *acyclic* pairwise rankings. Given any pairwise ranking from this subspace, we may determine a score function on the alternatives that is unique up to an additive constant¹⁰ and then we may rank all alternatives globally in terms of their scores.
2. $\ker(\delta_0^*) = \ker(\text{div})$ denotes the subspace of *divergence-free* pairwise rankings, whose total in-flow equals total out-flow for each alternative $i \in V$. They may be regarded as *cyclic* rankings, i.e. rankings of the form $i \succeq j \succeq k \succeq \dots \succeq i$ and are clearly inconsistent. Since $\ker(\text{div}^*) = \text{im}(\text{grad})^\perp$, cyclic rankings have zero projection on global rankings.
3. $\ker(\delta_1) = \ker(\text{curl})$ denotes the subspace of *curl-free* pairwise rankings with zero flow-sum along any triangle in K_G . These correspond to *locally consistent* (i.e. triangularly consistent) pairwise rankings. By the Closedness Lemma $\text{curl} \circ \text{grad} = 0$ and so $\text{im}(\text{grad}) \subseteq \ker(\text{curl})$. In general, the globally consistent pairwise rankings induced by gradient flows of score functions only account for a subset of locally consistent rankings. The remaining ones are the locally

¹⁰ Note that $\ker(\delta_0) = \ker(\text{grad})$ is the set of constant functions on V and so $\text{grad}(s) = \text{grad}(s + \text{constant})$.

consistent rankings that are not globally consistent and they are precisely the harmonic rankings discussed below.

4. $\ker(\Delta_1) = \ker(\text{curl}) \cap \ker(\text{div})$ denotes the subspace of *harmonic* pairwise rankings, or just harmonic rankings in short. It is the space of solutions to the Helmholtz equation. Harmonic rankings are exactly those pairwise rankings that are both *curl-free* and *divergence-free*. These are locally consistent with zero curl on every triangle in $T(E)$ but not globally consistent. In other words, while there are no inconsistencies due to small loops of length 3, i.e. $i \geq j \geq k \geq i$, there are inconsistencies along larger loops of lengths > 3 , i.e. $a \geq b \geq c \geq \dots \geq z \geq a$. So these are also cyclic rankings. Rank aggregation on $\ker(\Delta_1)$ depends on edge paths traversed in the simplicial complex; along homotopy equivalent paths one obtains consistent rankings. Figure 1 has an example of harmonic ranking.
5. $\text{im}(\delta_1^*) = \text{im}(\text{curl}^*)$ denotes the subspace of *locally cyclic* pairwise rankings that have non-zero curls along triangles. By the Closedness Lemma, $\text{im}(\text{curl}^*) \subseteq \ker(\text{div})$ and so this subspace is in general a proper subspace of the divergence-free rankings; the orthogonal complement of $\text{im}(\text{curl}^*)$ in $\ker(\text{div})$ is precisely the space of harmonic rankings $\ker(\Delta_1)$.

5 Analysis of HodgeRank

We now state two immediate implications of the Helmholtz decomposition theorem when applied to statistical ranking. The first implication is that HodgeRank gives an interpretation of the solution and residual of the optimization problem (7); these are respectively the l_2 -projection on gradient flows and divergence-free flows. In the context of statistical ranking and in the l_2 -sense, the solution to (7) gives the nearest globally consistent pairwise ranking to the data while the residual gives the sum total of all inconsistent components (both local and harmonic) in the data. The second implication is the condition that local consistency guarantees global consistency whenever there is no harmonic component in the data (which happens iff the clique complex of the pairwise comparison graph is ‘loop-free’).

5.1 Structure theorem for global ranking and the residual of inconsistency

In order to cast our optimization problem (7) in the Hodge theoretic framework, we need to specify a choice of inner products on C^0, C^1, C^2 . As before, the inner product on the space of edge flows (pairwise rankings) C^1 will be a weighted Euclidean inner product:

$$\langle X, Y \rangle_w = \sum_{\{i,j\} \in E} w_{ij} X_{ij} Y_{ij}$$

for $X, Y \in C^1$. We use unweighted Euclidean inner products on C^0 and C^2 ,

$$\langle r, s \rangle = \sum_{i=1}^n r_i s_i, \quad \langle \Theta, \Phi \rangle = \sum_{\{i,j,k\} \in T(E)} \Theta_{ijk} \Phi_{ijk}$$

for $r, s \in C^0$ and $\Theta, \Phi \in C^2$. This is mainly to keep our notations uncluttered. Other choices could be made (e.g. inner products on C^0 and C^2 could have been weighted) with corresponding straightforward modification of (7) but this would not change the essential nature of our methods.

The optimization problem (7) is then equivalent to an l_2 -projection of an edge flow representing a pairwise ranking onto $\text{im}(\text{grad})$,

$$\min_{s \in C^0} \|\delta_0 s - \bar{Y}\|_{2,w} = \min_{s \in C^0} \|\text{grad } s - \bar{Y}\|_{2,w}.$$

The Helmholtz decomposition theorem then leads to the following result about the structure of solutions and residuals of (7). In Theorem 3 below, we assume that the pairwise ranking data \bar{Y} has been estimated from one of the methods in Sect. 2.3. The least squares solution s will be a score function that induces $\text{grad } s$, the l_2 -nearest global ranking to \bar{Y} . Since s is only unique up to a constant (see Footnote 10), we determine a unique minimum norm solution s^* for the sake of well-posedness; but nevertheless any s will yield the same global ordering of alternatives. The least squares residual R^* represents the inconsistent component of the ranking data \bar{Y} . The magnitude of R^* is a ‘certificate of reliability’ for s ; since if this is small, then the globally consistent component $\text{grad } s$ accounts for most of the variation in \bar{Y} and we may conclude that s gives a reasonably reliable ranking of the alternatives. But even when the magnitude of R^* is large, we will see that it may be further resolved into a global and a local component that determine when a comparison of alternatives with respect to s is valid.

Theorem 3 (i) *Solutions of (7) satisfy the following normal equation*

$$\Delta_0 s = -\text{div } \bar{Y}, \tag{25}$$

and thus the minimum norm solution is

$$s^* = -\Delta_0^\dagger \text{div } \bar{Y} \tag{26}$$

where \dagger indicates Moore-Penrose inverse. The divergence in (26) is

$$(\text{div } \bar{Y})(i) = \sum_{j \text{ s.t. } \{i,j\} \in E} w_{ij} \bar{Y}_{ij},$$

and the matrix representing the graph Laplacian is given by

$$[\Delta_0]_{ij} = \begin{cases} \sum_i w_{ii} & \text{if } j = i, \\ -w_{ij} & \text{if } j \text{ is such that } \{i, j\} \in E, \\ 0 & \text{otherwise.} \end{cases}$$

(ii) The residual $R^* = \bar{Y} - \delta_0 s^*$ is divergence-free, i.e. $\text{div } R^* = 0$. Moreover, it has a further orthogonal decomposition

$$R^* = \text{proj}_{\text{im}(\text{curl}^*)} \bar{Y} + \text{proj}_{\text{ker}(\Delta_1)} \bar{Y}, \tag{27}$$

where $\text{proj}_{\text{im}(\text{curl}^*)} \bar{Y}$ is a locally cyclic ranking accounting for local inconsistencies and $\text{proj}_{\text{ker}(\Delta_1)} \bar{Y}$ is a harmonic ranking accounting for global inconsistencies. In particular, the projections are given by

$$\text{proj}_{\text{im}(\text{curl}^*)} = \text{curl}^\dagger \text{curl} \quad \text{and} \quad \text{proj}_{\text{ker}(\Delta_1)} = I - \Delta_1^\dagger \Delta_1. \tag{28}$$

Proof The normal equation for $\min_{s \in C^0} \|\delta_0 s - \bar{Y}\|_{2,w}^2$ is

$$\delta_0^* \delta_0 s = \delta_0^* \bar{Y}.$$

(25, 26), and $\text{div } R^* = 0$ are obvious upon substituting $\Delta_0 = \delta_0^* \delta_0$ and $\text{div} = -\delta_0^*$. The expressions for divergence and graph Laplacian in (i) follow from their respective definitions. The Helmholtz decomposition theorem implies

$$\text{ker}(\Delta_1) \oplus \text{im}(\text{curl}^*) = \text{im}(\text{grad})^\perp.$$

Obviously $\text{proj}_{\text{im}(\text{grad})^\perp} \text{grad } s^* = 0$. Since $R^* = \bar{Y} - \text{grad } s^*$ is a least squares residual, we must have $\text{proj}_{\text{im}(\text{grad})} R^* = \text{proj}_{\text{im}(\text{grad})} \bar{Y} - \text{grad } s^* = 0$. These observations yield (27), as

$$R^* = \text{proj}_{\text{im}(\text{grad})} R^* + \text{proj}_{\text{im}(\text{grad})^\perp} R^* = 0 + \text{proj}_{\text{ker}(\Delta_1) \oplus \text{im}(\text{curl}^*)} \bar{Y}.$$

The expressions for the projections in (28) are standard. □

In the special case when the pairwise ranking matrix G is a complete graph and we have an unweighted Euclidean inner product on C^1 , the minimum norm solution s^* in (26) satisfies $\sum_i s_i^* = 0$ and is given by

$$s_i^* = -\frac{1}{n} \text{div}(\bar{Y})(i) = -\frac{1}{n} \sum_j \bar{Y}_{ij}. \tag{29}$$

In Sect. 6, we shall see that this is the well-known *Borda count* in social choice theory, a measure that is also widely used in psychology and statistics [13,28,33]. Since G is a complete graph only when the ranking data is complete, i.e. every voter has rated every alternative, this is an unrealistic scenario for the type of modern ranking data discussed in Sect. 1. Among other things, HodgeRank generalizes Borda count to scenarios where the ranking data is incomplete or even highly incomplete.

In (ii) the locally cyclic ranking component is obtained by solving

$$\min_{\Phi \in C^2} \|\text{curl}^* \Phi - R^*\|_{2,w} = \min_{\Phi \in C^2} \|\text{curl}^* \Phi - \bar{Y}\|_{2,w}.$$

The above equality implies that there is no need to first solve for R^* before we may obtain Φ ; one could get it directly from the pairwise ranking data \bar{Y} . Note that the solution is only determined up to an additive term of the form $\text{grad } s$ since by virtue of (21),

$$\text{curl}(\Phi + \text{grad } s) = \text{curl } \Phi. \tag{30}$$

For the sake of well-posedness, we seek the unique minimum norm solution

$$\Phi^* = (\delta_1 \circ \delta_1^*)^\dagger \delta_1 \bar{Y} = (\text{curl} \circ \text{curl}^*)^\dagger \text{curl } \bar{Y}$$

and the required component is given by $\text{proj}_{\text{im}(\text{curl}^*)} \bar{Y} = \text{curl}^* \Phi^*$. The reader may have noted a parallel between the two problems

$$\min_{s \in C^0} \|\text{grad } s - \bar{Y}\|_{2,w} \quad \text{and} \quad \min_{\Phi \in C^2} \|\text{curl}^* \Phi - \bar{Y}\|_{2,w}.$$

Indeed in many contexts, s is called the *scalar potential* while Φ is called the *vector potential*. As seen earlier in Definition 1, an edge flow of the form $\text{grad } s$ for some $s \in C^0$ is called a gradient flow; in analogy, we will call an edge flow of the form $\text{curl}^* \Phi$ for some $\Phi \in C^2$ a *curl flow*.

We note that the l_2 -residual R^* , being divergence-free, is a cyclic ranking. Much like (30), the divergence-free condition is satisfied by a whole family of edge flows that differ from R^* only by a term of the form $\text{curl}^* \Phi$ since

$$\text{div}(R^* + \text{curl}^* \Phi) = \text{div } R^*$$

because of (21). The subset of C^1 given by

$$\{R^* + \text{curl}^* \Phi \mid \Phi \in C^2\}$$

is called the *homology class* of R^* . The harmonic ranking $\text{proj}_{\text{ker}(\Delta_1)} \bar{Y}$ is just one element in this class.¹¹ In general, it will be dense in the sense that it will be nonzero on almost every edge in E . This is because in addition to the divergence-free condition, the harmonic ranking must also satisfy the curl-free condition by virtue of (24). So if parsimony or sparsity is the objective, e.g. if one wants to identify a small number of conflicting comparisons that give rise to the inconsistencies in the ranking data, then the harmonic ranking does not offer much information in this regard. To better understand ranking inconsistencies via the structure of R^* , it is often helpful to look for elements in the same homology class with the *sparsest support*, i.e.

$$\min_{\Phi \in C^2} \|\text{curl}^* \Phi - R^*\|_0 = \min_{\Phi \in C^2} \|\text{curl}^* \Phi - \text{proj}_{\text{ker}(\Delta_1)} \bar{Y}\|_0.$$

¹¹ Two elements of the same homology class are called *homologous*.

The widely used convex relaxation replacing the l_0 -‘norm’ by the l_1 -norm may be employed [40], i.e.

$$\min_{\Phi \in \mathcal{C}^2} \|\text{curl}^* \Phi - R^*\|_1 := \min_{\Phi \in \mathcal{C}^2} \sum_{i,j} |(\text{curl}^* \Phi)_{ij} - R^*_{ij}|.$$

A solution $\tilde{\Phi}$ of such an l_1 -minimization problem is expected to give a sparse element $R^* - \text{curl}^* \tilde{\Phi}$. The bottom line here is that we want to find the shortest cycles that represent the global inconsistencies and perhaps remove the corresponding edges in the pairwise comparison graph, in view of what we will discuss next in Sect. 5.2. One plausible strategy to get a globally consistent ranking is to remove a number of problematic ‘conflicting’ comparisons from the pairwise comparison graph. Since it is only reasonable to remove as few edges as possible, this translates to finding a homology class with the sparsest support. This is similar to the minimum feedback arc set approach discussed in Sect. 6.2.

We will end the discussion of this section with a note on computational costs of HodgeRank. Solving for a global ranking s^* in (26) only requires the solution of an $n \times n$ least squares problem, which comes with a modest cost of $O(n^3)$ flops ($n = |V|$). As we note later in Sect. 7.3, for web ranking analysis such a cost is no more than computing the PageRank. On the other hand, the analysis of inconsistency is generally harder. For example, evaluating curls requires $|T|$ flops and this is $\binom{n}{3} \sim O(n^3)$ in the worst case. Since an actual computation of Φ^* involves solving a least squares problem of size $|T| \times |T|$, the computation cost incurred is of order $O(n^9)$. Nevertheless, any sparsity in the data (when $|T| \ll n^3$) may be exploited by choosing the right least squares solver. For example, one may use the general sparse least squares solver LSQR [35] or the new MINRES-QLP [9, 10] that works specifically for symmetric matrices. We will leave discussions of actual computations and more extensive numerical experiments to a future article. It suffices to note here that it is in general harder to isolate the harmonic component of ranking data than the globally consistent component.

5.2 Local consistency versus global consistency

In this section, we discuss a useful result, that local consistency implies global consistency whenever the harmonic component is absent from the ranking data. Whether a harmonic component exists is dependent on the topology of the clique complex K_G^3 . We will invoke the recent work of Kahle [24] on topological properties of random graphs to argue that harmonic components are exceedingly unlikely to occur.

By Lemma 2, the dimension of $\ker(\Delta_1)$ is equal to the first Betti number $\beta_1(K)$ of the underlying simplicial complex K . Since $\ker(\Delta_1) = 0$ if $\beta_1(K) = 0$, the harmonic component of any edge flow on K is automatically absent when $\beta_1(K) = 0$ (roughly speaking, $\beta_1(K) = 0$ means that K does not have any 1-dimensional holes). This leads to the following result.

Theorem 4 *Let $K_G^3 = (V, E, T(E))$ be a 3-clique complex of a pairwise comparison graph $G = (V, E)$. If K_G^3 does not contain any 1-loops, i.e. $\beta_1(K_G^3) = 0$, then every locally consistent pairwise ranking is also globally consistent. In other words, if the*

edge flow $X \in C^1(K_G^3, \mathbb{R})$ is curl-free, i.e.

$$\text{curl}(X)(i, j, k) = 0$$

for all $\{i, j, k\} \in T(E)$, then it is a gradient flow, i.e. there exists $s \in C^0(K_G, \mathbb{R})$ such that

$$X = \text{grad } s.$$

Proof This follows from Helmholtz decomposition since $\dim(\ker \Delta_1) = \beta_1(K_G^3) = 0$ and so any X that is curl-free is automatically in $\text{im}(\text{grad})$. \square

When G is a complete graph, then we always have $\beta_1(K_G) = \beta_1(K_G^3) = 0$ and this justifies the discussion after Definition 3 about the equivalence of local and global consistencies for complete pairwise comparison graphs. In general, G will be incomplete due to missing data (not all voters have rated all alternatives) but as long as K_G^3 is loop-free, such a claim still holds. In finance, this theorem translates into the well-known result that “triangular arbitrage-free implies arbitrage-free.” The theorem enables us to infer global consistency from a local condition—whether the ranking data is curl-free. We note that being curl-free is a strong condition. If we instead have “triangular transitivity” in the ordinal sense, i.e. $a \succeq b \succeq c$ implies $a \succeq c$, then there is no result analogous to Theorem 4.

At least for Erdős-Rényi random graphs, the Betti number β_1 could only be non-zero when the edges are neither too sparse nor too dense. The following result by Kahle [24] quantifies this statement. He showed that β_1 undergoes two phase transitions from zero to nonzero and back to zero as the density of edges grows.

Theorem 5 (Kahle 2006) *For an Erdős-Rényi random graph $G(n, p)$ on n vertices where the edges are independently generated with probability p , its clique complex K_G almost always has $\beta_1(K_G) = 0$, except when*

$$\frac{1}{n^2} \ll p \ll \frac{1}{n}. \tag{31}$$

Without getting into a discussion as to whether an Erdős-Rényi random graph is a good model for pairwise comparison graphs of real-world ranking data, we note that the Netflix pairwise comparison graph has a high probability of having $\beta_1(K_G) = 0$ if Kahle’s result applies. Although the original customer-product rating matrix of the Netflix prize dataset is highly incomplete (more than 99% missing values), its pairwise comparison graph is very dense (less than 0.22% missing edges). So p (probability of an edge) and n (number of vertices) are both large and (31) is not satisfied.

6 Connections to social choice theory

Social choice theory is almost undoubtedly the discipline most closely associated with the study of ranking, having a long history dating back to Condorcet’s famous treatise in 1785 [15] and a large body of work that led to at least two Nobel prizes [3, 39].

The famous impossibility theorems of Arrow [2] and Sen [38] in social choice theory formalized the inherent difficulty of achieving a global ranking of alternatives by aggregating over the voters. However it is still possible to perform an approximate rank aggregation in reasonable, systematic ways. Among the various proposed methods, the best known ones are those by Condorcet [15], Borda [14], and Kemeny [25]. The Kemeny approach is often regarded as the best approximate rank aggregation method under some assumptions [43,44]. It is however NP-hard to compute and its sole reliance on ordinal information may be unnatural for score-based cardinal data.

We have described earlier how the minimization of (7) over the gradient flow model class

$$\mathcal{M}_G = \{X \in C^1 \mid X_{ij} = s_j - s_i, \quad s : V \rightarrow \mathbb{R}\}$$

leads to a Hodge theoretic generalization of Borda count but the minimization of (7) over the Kemeny model class

$$\mathcal{M}_K = \{X \in C^1 \mid X_{ij} = \text{sign}(s_j - s_i), \quad s : V \rightarrow \mathbb{R}\}$$

leads to Kemeny optimization. We will elaborate on these in Sect. 6.1.

The following are some desirable properties of ranking data that have been widely studied, used, and assumed in social choice theory. A ranking problem is called *complete* if each voter in Λ gives a total ordering or permutation of all alternatives in V ; this implies that $w_{ij}^\alpha > 0$ for all $\alpha \in \Lambda$ and all distinct $i, j \in V$, in the terminology of Sect. 2. It is *balanced* if the pairwise comparison graph $G = (V, E)$ is k -regular with equal weights $w_{ij} = c$ for all $\{i, j\} \in E$. A complete and balanced ranking induces a complete graph with equal weights on all edges. Moreover, it is *binary* if every pairwise comparison is allowed only two values, say, ± 1 without loss of generality. So $Y_{ij}^\alpha = 1$ if voter α prefers alternative j to alternative i , and $Y_{ij}^\alpha = -1$ otherwise. Ties are disallowed to keep the discussion simple.

Classical social choice theory often assumes complete, balanced, and binary rankings. However, these are all unrealistic assumptions for modern data coming from internet and e-commerce applications. Take the Netflix dataset for illustration, a typical user α of Netflix would have rated at most a very small fraction of the entire Netflix inventory. Indeed, as we have mentioned in Sect. 2.2.1, the viewer-movie rating matrix has 99% missing values. Moreover, while blockbuster movies would receive a disproportionately large number of ratings, since just about every viewer has watched them, the more obscure or special interest movies would receive very few ratings. In other words, the Netflix dataset is highly incomplete and highly imbalanced. Furthermore, the Netflix ratings are given in terms of scores, i.e. 1 through 5 stars. While it is possible to ignore the cardinal nature of the dataset and just use its ordinal information to construct a binary pairwise ranking, we would be losing valuable information—for example, a 5-star versus 1-star comparison is indistinguishable from a 3-star versus 2-star comparison when one only takes ordinal information into account.

Therefore, one is ill-advised to apply methods from classical social choice theory to modern ranking data directly. We will see next that our Hodge theoretic extension of Borda count adapts to these new features in modern datasets, i.e. incomplete, imbal-

anced, cardinal data, but restricts to the usual Borda count for data that is complete, balanced, and ordinal/binary.

6.1 Kemeny optimization and Borda count

The basic idea of Kemeny’s rule [25, 26] is to minimize the number of pairwise mismatches from a given ordering of alternatives to a voting profile, i.e. the collection of total orders on the alternatives by each voter. The minimizers are called *Kemeny optima* and are often regarded as the most reasonable candidates for a global ranking of the alternatives. To be precise, we define the binary pairwise ranking associated with a permutation $\sigma \in \mathfrak{S}_n$ (the permutation group on n elements) to be $Y_{ij}^\sigma = \text{sign}(\sigma(i) - \sigma(j))$. Given two total orders or permutations on the n alternatives, $\sigma, \tau \in \mathfrak{S}_n$, the *Kemeny distance* (also known as *Kemeny-Snell* or *Kendall τ distance*) is defined as

$$d_K(\sigma, \tau) := \frac{1}{2} \sum_{i < j} |Y_{ij}^\sigma - Y_{ij}^\tau| = \frac{1}{4} \sum_{i, j} |Y_{ij}^\sigma - Y_{ij}^\tau|,$$

i.e. the number of pairwise mismatches between σ and τ . Given a voting profile as a set of permutations on $V = \{1, \dots, n\}$ by m voters, $\{\tau_i \in \mathfrak{S}_n \mid i = 1, \dots, m\}$, the following combinatorial minimization problem

$$\min_{\sigma \in \mathfrak{S}_n} \sum_{i=1}^m d_K(\sigma, \tau_i) \tag{32}$$

is called *Kemeny optimization* and is known to be NP-hard [17] with respect to n when $m \geq 4$. For binary-valued rankings with $Y_{ij}^\alpha \in \{\pm 1\}$,

$$\min_{X \in \mathcal{M}_K} \sum_{\alpha, i, j} w_{ij}^\alpha (X_{ij} - Y_{ij}^\alpha)^2, \tag{33}$$

counts up to a constant the number of pairwise mismatches from a total order. Hence for a complete, balanced, and binary-valued ranking problem, our minimization problem (7) becomes Kemeny optimization if we replace the subspace \mathcal{M}_G by the discrete subset \mathcal{M}_K .

Another well-known method for rank aggregation is the *Borda count* [14], which assigns a voter’s top i th alternative a *position-based score* of $n - i$; the global ranking on V is then derived from the sum of its scores over all voters. This is equivalent to saying that the global ranking of the i th alternative is derived from the score

$$s_B(i) = - \sum_{\alpha, k=1}^{m, n} Y_{ik}^\alpha, \tag{34}$$

i.e. the alternative that has the most pairwise comparisons in favor of it from all voters will be ranked first, and so on. As we have found in (29), the minimum norm solution

of the l_2 -projection onto gradient flows is given by

$$s^*(i) = -\frac{1}{n} \sum_k \bar{Y}_{ik} = -c \sum_{\alpha,k=1}^{m,n} Y_{ik}^\alpha,$$

where c is a positive constant. Hence for a complete, balanced, and binary ranking problem, HodgeRank yields the Borda count (up to a positive multiplicative constant that has no effect on the ordering of alternatives by scores).

6.2 Comparative studies

The following theorem gives three equivalent characterizations of (33) when $Y_{ij}^\alpha \in \{\pm 1\}$. Note that we do not assume complete nor balanced data.

Theorem 6 *Suppose that $Y_{ij}^\alpha \in \{\pm 1\}$. The following optimization problems are all equivalent:*

- (i) *The weighted least squares problem,*

$$\min_{X \in \mathcal{M}_K} \sum_{\alpha,i,j} w_{ij}^\alpha (X_{ij} - Y_{ij}^\alpha)^2,$$

where $\mathcal{M}_K = \{X \in \mathcal{A} \mid X_{ij} = \text{sign}(s_j - s_i), \quad s : V \rightarrow \mathbb{R}\}$.

- (ii) *The linear programming problem,*

$$\max_{X \in \mathcal{K}_1} \langle X, \bar{Y} \rangle = \max_{X \in \mathcal{K}_1} \sum_{\{i,j\} \in E} w_{ij} X_{ij} \bar{Y}_{ij}, \tag{35}$$

where \mathcal{K}_1 is the set

$$\left\{ \sum_{\sigma \in \mathcal{S}_n} \mu_\sigma P^\sigma \mid \sum_{\sigma} \mu_\sigma = 1, \quad \mu_\sigma \geq 0, \quad P_{ij}^\sigma = \text{sign}(\sigma(j) - \sigma(i)) \right\}.$$

- (iii) *The weighted l_1 -minimization problem,*

$$\min_{X \in \mathcal{K}_2} \|X - \bar{Y}\|_{1,w} = \min_{X \in \mathcal{K}_2} \sum_{\{i,j\} \in E} w_{ij} |X_{ij} - \bar{Y}_{ij}|, \tag{36}$$

where \mathcal{K}_2 is the set

$$\{X \in \mathcal{A} \mid (s_j - s_i)X_{ij} \geq 0 \text{ for some } s : V \rightarrow \mathbb{R} \text{ and } \{i, j\} \in E\}.$$

- (iv) *The minimum feedback arc set of the weighted directed graph $G_{W \circ \bar{Y}} = (V, \mathbf{E}, W \circ \bar{Y})$, whose vertex set is V , directed edge $(i, j) \in \mathbf{E} \subseteq V \times V$ iff $\bar{Y}_{ij} > 0$ with weight $w_{ij} \bar{Y}_{ij}$.*

Proof Assuming (i). Since $X_{ij} \in \{\pm 1\}$, the problem in (i) is equivalent to

$$\max_{X \in \mathcal{M}_K} \sum_{\{i,j\} \in E} w_{ij} X_{ij} \bar{Y}_{ij}. \tag{37}$$

Since \mathcal{M}_K is a discrete set containing $n!$ points, a linear programming problem over \mathcal{M}_K is equivalent to searching over its convex hull, i.e. \mathcal{K}_1 , which gives (ii).

(iv) can also be derived from (37). Consider a weighted directed graph $G_{W \circ \bar{Y}}$ where an edge $(i, j) \in \mathbf{E}$ iff $\bar{Y}_{ij} > 0$, and in which case has weight $|w_{ij} \bar{Y}_{ij}|$. (37) is equivalent to finding a directed acyclic graph by reverting a set of edge directions whose weight sum is minimized. This is exactly the minimum feedback arc set problem.

Finally, we show that (iii) is also equivalent to the minimum feedback arc set problem. For any $X \in \mathcal{K}_2$, the transitive region, there is an associated weighted directed acyclic graph $G_{W \circ X}$ where an edge $(i, j) \in \mathbf{E}$ iff $X_{ij} > 0$, and in which case has weight $|w_{ij} X_{ij}|$. Note that a minimizer of (36) has either $X_{ij}^* = -X_{ji}^* = \bar{Y}_{ij}$ or $X_{ij}^* = -X_{ji}^* = 0$ on an edge $\{i, j\} \in E$, which is equivalent to the problem of finding a directed acyclic graph by deleting a set of edges from $G_{W \circ \bar{Y}}$ such that the sum of their weights is minimized. Again, this is exactly the minimum feedback arc set problem. □

The set \mathcal{K}_1 is the convex hull of *skew-symmetric permutation matrices* P^σ as defined in [44]. The set \mathcal{K}_2 is called the *transitive pairwise region* in [36], comprising $n!$ cones corresponding to each of the $n!$ permutations on V .

It is known that the minimum feedback arc set problem in (iv) is NP-hard, and therefore, so are the other three. Moreover, (iii) provides some geometric insights when viewed alongside (7), the l_2 -projection onto gradient flows $\mathcal{M}_G = \{X \in \mathcal{A} \mid X_{ij} = s_j - s_i, s : V \rightarrow \mathbb{R}\}$, which we have seen to be a Hodge theoretic extension of Borda count. We will illustrate their differences and similarities pictorially via an example from [36].

Consider the simplest case of three-item comparison with $V = \{i, j, k\}$. For simplicity, we will assume that $w_{ij} = w_{jk} = w_{ki} = 1$ and $\bar{Y}_{ij}, \bar{Y}_{jk}, \bar{Y}_{ki} \in [-1, 1]$. Figure 3 shows the unit cube in \mathbb{R}^3 . We will label the coordinates in \mathbb{R}^3 as $[X_{ij}, X_{jk}, X_{ki}]$ (instead of the usual $[x, y, z]$). The shaded plane corresponds to the set where $X_{ij} + X_{jk} + X_{ki} = 0$ in the unit cube. Note that this set is equal to the model class \mathcal{M}_G because of (15). On the other hand, the transitive pairwise region \mathcal{K}_2 consists of the six orthants within the cube with vertices $\{\pm 1, \pm 1, \pm 1\} - \{[1, 1, 1], [-1, -1, -1]\}$. The HodgeRank optimization (7) is the l_2 -projection onto the plane $X_{ij} + X_{jk} + X_{ki} = 0$, while by (iii), Kemeny optimization (33) is the l_1 -projection onto the aforementioned six orthants representing the transitive pairwise region \mathcal{K}_2 . The shaded plane in Fig. 3 is a special case of a *permutahedron* or *permutohedron*. We refer the reader to [32, 45] for excellent discussions of these. In particular, Zhang has shown [45] that the projection depicted in Fig. 3 for three-item comparison can be generalized to a comparison of an arbitrary number of items.

In the setting of social choice theory, the following theorem [36] connects order relations between Kemeny optimization and Borda count.




Fig. 3 The shaded region is the subspace $X_{ij} + X_{jk} + X_{ki} = 0$. The transitive region consists of six orthants whose corresponding vertices belong to $\{\pm 1, \pm 1, \pm 1\} - \{[1, 1, 1], [-1, -1, -1]\}$. The Borda count is the l_2 -projection onto the shaded plane while the Kemeny optimization is the l_1 -projection onto the transitive region

Theorem 7 (Saari-Merlin 2000) *The Kemeny winner (most preferred) is always strictly above the Kemeny loser (least preferred) under the Borda count; similarly the Borda winner is always strictly above the Borda loser under the Kemeny rule. Aside from these constraints, the two methods may generate arbitrarily different total orders.*

The Kemeny rule has several desirable properties in social choice theory which the Borda count lacks [44]. The Kemeny rule satisfies the Condorcet rule, in the sense that if an alternative in V wins all pairwise comparisons against other alternatives in V , then it must be the overall winner. A Condorcet winner is any alternative i such that $\sum_j \text{sign}(\sum_\alpha Y_{ij}^\alpha) = n$. Note that the Condorcet winner may not exist in general but Kemeny or Borda winners always exist. However, if a Condorcet winner exists, then it must be the Kemeny winner. On the other hand, Borda count can only ensure that the Condorcet winner is ranked strictly above the Condorcet loser (least preferred). Another major advantage of the Kemeny rule is its consistency in global rankings under the elimination of alternatives in V . The Borda count and many other position-based rules fail to meet this condition. In fact, the Kemeny rule is the unique rule [44] that meets all of following: (1) satisfies the Condorcet rule, (2) consistency under elimination, and (3) a natural property called neutral (which we will not discuss here).

Despite the many important features that the Kemeny rule has, its high computational cost (NP-hard) makes simpler rules like Borda count attractive in practice, especially when there is large number of alternatives to be ranked. Moreover, in cardinal rankings where it is desirable to preserve the magnitude of score differences [12] and not just the order relation, using the Hodge theoretic variant of Borda count with model class \mathcal{M}_G becomes more relevant than Kemeny optimization with model class \mathcal{M}_K .

7 Experimental studies

We present three examples of analysis of real data using HodgeRank with the hope that they would illustrate the basic ideas.

The first example ranks movies from the Netflix dataset. We show that (1) HodgeRank reduces temporal drift bias, and (2) the triangular curls provide a metric for characterizing inconsistencies in the ranking data. The second example illustrates the use of HodgeRank for finding a universal equivalent or price function (i.e. global ranking) in a currency exchange market, with data from Yahoo! Finance. The third example describes how the global ranking component in HodgeRank may be used to approximate PageRank via reversible Markov chains, with data from Webometrics.

7.1 Movie ranking on a subset of netflix data

The Netflix prize dataset¹² contains about 17,000 movies rated by 480,000 customers over 74 months from November 1999 to December 2005. Each customer rated 209 movies on average and around 99% of the ratings are absent from the customer-product matrix. We do not seek to address the Netflix prize problem of ratings prediction here. Instead we take advantage of this rare publicly available dataset and use it to test the rank aggregation capabilities of our method. We would like to aggregate viewers' ratings into a global ranking on movies, and to measure the reliability of such a global ranking. Note that such rank aggregation could be *personalized* if one also attempts to identify groups of viewers sharing similar tastes (say, via the suggestion in Footnote 2). This could then be used for rating prediction if desired.

For reasons that we will soon explain, we restrict our selections to movies that received ratings on all of the 74 months. There are not many such movies—only 25 in all. Several of these have monthly average scores that show substantial upward or downward drifts. In Fig. 4, we show the temporal variations in scores of six of these (numerical indices in the Netflix dataset are given in parentheses): *Dune* (17064), *Interview with the Vampire* (8079), *October Sky* (12473), *Shakespeare in Love* (17764), *The Waterboy* (14660), and *Witness* (15057). Such temporal variations make it dubious to rank movies by simply taking average score over all users, as ratings over different time periods may not be comparable under the same scale. It is perhaps worth noting that understanding the temporal dynamics in the Netflix dataset has been a key factor in the approach of Bell and Koren [6]. We will see below that the use of HodgeRank provides an effective method to globally rank these movies, detect any inherent inconsistency, and is robust under temporal variations.

Formation of pairwise ranking. Since pairwise rankings are relative measures, we expect them to reduce the effect of temporal drift. Using only ratings by the same customer in the same month, we employ three of the statistics described in Sect. 2.3 to form our pairwise rankings, i.e. the *arithmetic mean of score differences* (9), the *geometric mean of score ratios* (10), and the *binary comparisons* (11). Since there is

¹² Available from <http://www.netflixprize.com>.




Fig. 4 Average scores of six selected movies over 74 months. The three movies in the *top* row has a *decreasing* trend in monthly average scores, while in a contrast the other three movies in the *bottom* row exhibits an *increasing* trend

nothing to suggest that a logarithmic scale is relevant, the logarithmic odds ratio gives rather poor result as expected and we omitted it. For comparison, we compute the mean score of each movie over all customers, ignoring temporal information. A reference score is collected independently from MRQE (Movie Review Query Engine),¹³ the largest online directory of movie reviews on the internet.

Global ranking with HodgeRank. We then solve the regression problem in (7) to obtain a projection of pairwise ranking flows onto gradient flows, as given by Theorem 3(i). Note that in this example, the pairwise ranking graph is complete with $n = 6$ nodes. Table 1 collects the comparisons between different global rankings. The reference order of movies is again via the MRQE scores.

Inconsistencies and curls. Since the pairwise ranking graph is complete, its clique complex is a simplex with $n = 6$ vertices and so the harmonic term in the Hodge-Rank is always zero. Hence the residual in Theorem 3 is just the curl projection, i.e. $R^* = \text{proj}_{\text{im}(\text{curl}^*)} \bar{Y}$. We will define two indices of inconsistency to evaluate the results. The first, called *cyclicality ratio*, is a measure of global inconsistency given by $C_p = \|R^*\|_{2,w}^2 / \|\bar{Y}\|_{2,w}^2$; while the second, called *relative curl*, quantifies the local inconsistency, and is given by the following function of edges and triangles,

$$c_r(\{i, j\}, \{i, j, k\}) = \frac{(\text{curl } \bar{Y})(i, j, k)}{3(\text{grad } s^*)(i, j)} = \frac{\bar{Y}_{ij} + \bar{Y}_{jk} + \bar{Y}_{ki}}{3(s_j^* - s_i^*)}.$$

¹³ Available from <http://www.mrqe.com>.

Table 1 Global ranking of selected six movies via different methods: MRQE, mean score over customers, HodgeRank with arithmetic mean score difference, HodgeRank with geometric mean score ratio, and HodgeRank with binary comparisons

Movie	Global ranking (score)				
	MRQE	Mean	Hodge-difference	Hodge-ratio	Hodge-binary
Shakespeare in love	1 (85)	2 (3.87)	1 (0.247)	2 (0.0781)	1 (0.138)
Witness	2 (77)	3 (3.86)	2 (0.217)	1 (0.0883)	3 (0.107)
October sky	3 (76)	1 (3.93)	3 (0.213)	3 (0.0775)	2 (0.111)
The waterboy	4 (66)	6 (3.38)	6 (−0.464)	6 (−0.1624)	6 (−0.252)
Interview with the vampire	5 (65)	4 (3.71)	4 (−0.031)	4 (−0.0121)	4 (−0.012)
Dune	6 (44)	5 (3.49)	5 (−0.183)	5 (−0.0693)	5 (−0.092)
Cyclicity ratio (%)	–	–	0.77	1.15	0.30

It can be seen that the HodgeRank with binary comparisons has the smallest inconsistency in terms of the cyclicity ratio

Note that on every triangle $\{i, j, k\}$ the curl $\bar{Y}_{ij} + \bar{Y}_{jk} + \bar{Y}_{ki}$ measures the total sum of cyclic flow, therefore c_r measures the magnitude of its induced edge flow relative to the gradient edge flow of the global ranking s^* . If c_r has absolute value larger than 1, then the average cyclic flow has an effect larger than the global ranking s^* , which indicates that the global ranking s^* might be inconsistent on the pair of items.

Table 1 shows that in terms of cyclicity ratio, the best global ranking is obtained from HodgeRank with binary comparisons, which has the smallest cyclicity ratio, 0.30%. This global ranking is quite different from merely taking mean scores and is a better predictor of MRQE.

A closer look at relative curls allows us to identify the dubious scores. We will see that the placement of *Witness* and *October Sky* according to the global ranking contains significant inconsistency and should not be trusted. This inconsistency is largely due to the curls in the triangles $t_1 = \{\text{Witness}, \text{October Sky}, \text{The Waterboy}\}$, $t_2 = \{\text{Witness}, \text{October Sky}, \text{Interview with the Vampire}\}$. In fact, there are only two relative curls whose magnitudes exceed 1; both occurred on triangles that contain the edge $e = \{\text{Witness}, \text{October Sky}\}$. The relative curl of t_1 with respect to e is 3.6039 while that of t_2 with respect to e is 4.1338. As we can see from Table 1, the inconsistency (large curl) manifests itself as *instability* in the placement of *Witness* and *October Sky*—the results vary across different rank aggregation methods with no possibility of consensus. This illustrates the use of curl as a certificate of validity for global ranking.

7.2 Currency exchange market

This example illustrates a globally consistent pairwise ranking on a complete graph using currency exchange data taken from Yahoo! Finance.¹⁴ Consider a currency exchange market with V representing a collection of seven currencies, USD, JPY, EUR, CAD, GBP, AUD, and CHF. Here $G = (V, E)$ is a complete graph since every two currencies in V are exchangeable. Table 2 shows the exchange rates. By

¹⁴ Available from <http://www.finance.yahoo.com/currency-converter>.

Table 2 The last line is given by $\exp(-x^*)$ where x^* is the solution to (26)

	Currency exchange rate table						
	USD	JPY	EUR	CAD	GBP	AUD	CHF
1 USD	1.0000	114.6700	0.6869	0.9187	0.4790	1.0768	1.1439
1 JPY	0.0087	1.0000	0.0060	0.0080	0.0042	0.0094	0.0100
1 EUR	1.4558	166.9365	1.0000	1.3374	0.6974	1.5676	1.6653
1 CAD	1.0885	124.8177	0.7477	1.0000	0.5214	1.1721	1.2451
1 GBP	2.0875	239.3791	1.4340	1.9178	1.0000	2.2478	2.3879
1 AUD	0.9287	106.4940	0.6379	0.8532	0.4449	1.0000	1.0623
1 CHF	0.8742	100.2448	0.6005	0.8031	0.4188	0.9413	1.0000
Univ. equiv.	1.7097	0.0149	2.4890	1.8610	3.5691	1.5878	1.4946

The data was taken from the Currency Converter Yahoo! Finance on 6 November 2007

logarithmic transform the exchange rates is converted into pairwise rankings as in Example 2.2.2. The global ranking given by (26) (where $\delta_0^* = \delta_0^\top$) defines a *universal equivalent* that measures the ‘value’ of each currency. As the reader can check, the logarithmic transform of the data in Table 2 is curl-free (up to machine precision), which in this context means triangular arbitrage-free. In other words, there is no way one could profit from a cyclic exchange of any three currencies in V . Since G is a complete graph, the data has no harmonic components; so Hodge decomposition tells us that local consistency must imply global consistency, which in this context means arbitrage-free, i.e. there is no way one could profit from a cyclic exchange of any number of currencies in V either.

7.3 Comparisons with pageRank and HITS

We apply HodgeRank to the problem of web ranking, which we assumed here to mean any static linked objects, not necessarily the World Wide Web. As we shall see, HodgeRank provides an alternative to PageRank [8] and HITS [29]. In particular, it gives a new way to approximate PageRank and enables us to study the inconsistency or cyclicity in PageRank models.

Consider a link matrix L where L_{ij} is the number of links from site i to j . There are two well-known spectral approach to computing the global rankings of websites from L , HITS and PageRank. HITS computes the singular value decomposition $L = U\Sigma V^\top$, where the primary left-singular vector u_1 gives the *hub ranking* and the primary right-singular vector v_1 gives the *authority ranking* (both u_1 and v_1 are non-negative real-valued by the Perron-Frobenius theorem). PageRank constructs from L a Markov chain on the sites given by $P_{ij} = \frac{\alpha L_{ij}}{\sum_j L_{ij}} + \frac{1-\alpha}{n}$, where n is the number of sites and $\alpha = 0.85$ trades-off between Markovian link jumps and random surfing.

It is clear that we may define an edge flow via

$$Y_{ij} = \log \frac{P_{ij}}{P_{ji}}. \tag{38}$$

However what property does such a flow capture in PageRank? To answer this question we will need to recall the notion of a reversible Markov chain: An irreducible Markov chain with transition matrix P and stationary distribution π is *reversible* if $\pi_i P_{ij} = \pi_j P_{ji}$. Therefore a reversible Markov chain P has a pairwise ranking flow induced from a global ranking, $Y_{ij} = \log(P_{ij}/P_{ji}) = \log \pi_j - \log \pi_i$, where $\log \pi$ gives the global ranking. As we mentioned in Sect. 2.1, $\log \pi$ may be viewed as defining a *negative potential* on webpages if we regard ranking as being directed from a higher potential site to a lower potential site. This leads to the following interpretation.

Let P^* be the best reversible approximate of the PageRank Markov chain P , in the sense that

$$P^* = \operatorname{argmin}_{\tilde{P} \text{ reversible}} \left\| \log \frac{\tilde{P}_{ij}}{\tilde{P}_{ji}} - \log \frac{P_{ij}}{P_{ji}} \right\|_2. \tag{39}$$

Then the stationary distribution of P^* , denoted by π^* , is a Gibbs-Boltzmann distribution on webpages with potential $-s^*$, i.e. $\pi_i^* = e^{s_i^*} / \left(\sum_k e^{s_k^*}\right)$, where s^* is given by the Hodge projection of Y onto the space of gradient flows. Hence the Hodge decomposition of edge flow in (38) gives the stationary distribution of a best reversible approximate of the PageRank Markov chain.

We may further compute the Hodge decomposition of iterated flows, $Y_{ij}^k = \log(P_{ij}^k/P_{ji}^k)$. Clearly when $k \rightarrow \infty$, the global ranking given by HodgeRank converges to that given by PageRank. The benefit of the HodgeRank lies in that (1) it provides a way to approximate the PageRank stationary distribution; and (2) it enables us to study the inconsistency or cyclicity in PageRank Markov model. The cost of computing the global ranking in HodgeRank in Theorem 3(i) only involves a least squares problem of the graph Laplacian, which is less expensive than eigenvector computations in PageRank. For readers unfamiliar with numerical linear algebra, it might be worth pointing out that even the most basic algorithms for linear least squares problems guarantee global convergence in a finite number of steps whereas there are (a) no algorithms for eigenvalue problems that would terminate in a finite number of steps as soon as the matrix dimension exceeds 4; and (b) no algorithms with guaranteed global convergence for arbitrary input matrices.

To illustrate this discussion, we use the UK Universities Web Link Structure dataset.¹⁵ The dataset contains the number of web links between 111 UK universities in 2002. Independent of this link structure is a research score for each university, RAE 2001, performed during the 5-yearly Research Assessment Exercise.¹⁶ The RAE scores are widely used in UK for measuring the quality of research in universities. We retained 107 universities and eliminated four that are missing either RAE score, in-link, or out-link. The data has also been used in [42] recently but for a different purpose. Table 3 summarizes the comparisons among nine global rankings: RAE 2001, in-degree, out-degree, HITS authority, HITS hub, PageRank, HodgeRank with

¹⁵ Available from <http://www.cybermetrics.wlv.ac.uk/database/stats/data> from Webometrics. We used counts at the directory level.

¹⁶ Available from <http://www.rae.ac.uk>.

Table 3 Kendall τ -distance between different global rankings

	Kendall τ -distance								
	RAE '01	in-deg	out-deg	HITS auth	HITS hub	PageRank	Hodge-1	Hodge-2	Hodge-4
RAE '01	0	0.0994	0.1166	0.0961	0.1115	0.0969	0.1358	0.0975	0.0971
in-deg	0.0994	0	0.0652	0.0142	0.0627	0.0068	0.0711	0.0074	0.0065
out-deg	0.1166	0.0652	0	0.0672	0.0148	0.0647	0.1183	0.0639	0.0647
HITS auth	0.0961	0.0142	0.0672	0	0.0627	0.0119	0.0736	0.0133	0.0120
HITS hub	0.1115	0.0627	0.0148	0.0627	0	0.0615	0.1121	0.0607	0.0615
PageRank	0.0969	0.0068	0.0647	0.0119	0.0615	0	0.0710	0.0029	0.0005
Hodge-1	0.1358	0.0711	0.1183	0.0736	0.1121	0.0710	0	0.0692	0.0709
Hodge-2	0.0975	0.0074	0.0639	0.0133	0.0607	0.0029	0.0692	0	0.0025
Hodge-4	0.0971	0.0065	0.0647	0.0120	0.0615	0.0005	0.0709	0.0025	0

Note that HITS authority gives the nearest global ranking to the research score RAE '01, while HodgeRank for $k = 2, 4$ give closer results to PageRank which is the second closest to the RAE '01

$k = 1, 2$, and 4 , respectively. We then use Kendall τ -distance [27] to count the number of pairwise mismatches between global rankings, normalized by the total number of pairwise comparisons.

8 Summary and conclusion

We introduced combinatorial Hodge theory as a statistical ranking method based on minimizing pairwise ranking errors over a model space. In particular, we proposed HodgeRank, a Hodge theoretic approach for determining the global, local, and harmonic ranking components of a dataset of voters' scores on alternatives. The global ranking in HodgeRank is computed via an l_2 -projection of a pairwise ranking edge flow onto the space of gradient flows. We saw that among other connections to classical social choice theory, the score recovered from the global ranking in HodgeRank is a generalization of the well-known Borda count to ranking data that is cardinal, imbalanced, and incomplete. The remaining residual is the l_2 -projection onto the space of divergence-free flows. A subsequent l_2 -projection of this divergence-free residual onto the space of curl-free flows then yields a harmonic flow. The decomposition of pairwise ranking into a global ranking component, a locally cyclic ranking component, and a harmonic ranking component in HodgeRank is analogous to Helmholtz decomposition in fluid dynamics and electrodynamics.

Consistency of ranking data is governed to a large extent by the structure of its pairwise comparison graph; this is in turn revealed in the Helmholtz decomposition associated with the graph Helmholtzian, the combinatorial Laplacian of the 3-clique complex. The sparsity structure of a pairwise comparison graph imposes certain constraints on the topology and geometry of its clique complex, which in turn decides the properties of our statistical ranking algorithms.

Our results suggest that combinatorial Hodge theory could be a promising tool for the statistical analysis of ranking, especially for datasets with cardinal, incomplete, and imbalanced information.

Acknowledgments The authors thank the anonymous referees for their useful suggestions. The authors also gratefully acknowledge helpful discussions with Gunnar Carlsson, Leo Guibas, and Steve Smale.

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

1. Ailon, N., Charikar, M., Newman, A.: Aggregating inconsistent information: ranking and clustering. In: *Proceeding ACM Symposium Theory Computing (STOC '605)*, vol. 37, pp. 684–693 (2005)
2. Arrow, K.J.: A difficulty in the concept of social welfare. *J. Polit. Econ.* **58**(4), 328–346 (1950)
3. Arrow, K.J.: General economic equilibrium: purpose, analytic techniques, collective choice. Nobel Memorial Lecture, 12 December 1972. In: Lindbeck, A. (ed.) *Nobel Lectures: Economic Sciences 1969–1980*, pp. 109–131. World Scientific, Singapore (1992)
4. Barber, B.M., Lehavy, R., McNichols, M., Trueman, B.: Can investors profit from the prophets? Security analyst recommendations and stock returns. *J. Finance* **56**(2), 531–563 (2001)
5. Bartholdi, L., Schick, T., Smale, N., Smale, S., Baker, A.W.: Hodge theory on metric spaces. preprint. <http://www.arxiv.org/abs/0912.0284> (2009)
6. Bell, R.M., Koren, Y.: Scalable collaborative filtering with jointly derived neighborhood interpolation weights. In: *Proceedings IEEE International Conference Data Mining (ICDM '607)*, vol. 7, pp. 43–52. (2007)
7. Bradley, R., Terry, M.: The rank analysis of incomplete block designs: I. The method of paired comparisons. *Biometrika* **39**, 324–345 (1952)
8. Brin, S., Page, L.: The anatomy of a large-scale hypertextual web search engine. *Comput. Networks ISDN Syst.* **30**(1–7), 107–117 (1998)
9. Choi, S.-C.: Iterative methods for singular linear equations and least-squares problems. Ph.D. thesis, Stanford University, Stanford, CA. <http://www.stanford.edu/group/SOL/dissertations/sou-cheng-choi-thesis.pdf> (2006)
10. Choi, S.-C., Paige, C.C., Saunders, M.A.: MINRES-QLP: a Krylov subspace method for indefinite or singular symmetric systems, preprint. <http://www.arxiv.org/abs/1003.4042> (2010)
11. Chung, F.: Spectral Graph Theory. In: *CBMS Regional Conference Series in Mathematics*, vol. 92. AMS, Providence, RI (1997)
12. Cortes, C., Mohri, M., Rastogi, A.: Magnitude-preserving ranking algorithms. In: *Proceeding International Conference Machine Learning (ICML '607)*, vol. 24, pp. 169–176 (2007)
13. David, H.A.: *The Method of Paired Comparisons*, 2nd edn. Griffin's Statistical Monographs and Courses, vol. 41. Oxford University Press, New York, NY (1988)
14. de Borda, J.-C.: Mémoire sur les élections au scrutin. *Histoire de l'Académie Royale Des Sci.* **102**, 657–665 (1781)
15. de Condorcet A.-N.: *Éssai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Imprimerie Royale, Paris (1785)
16. Diaconis, P.: A generalization of spectral analysis with application to ranked data. *Ann. Statist.* **17**(3), 949–979 (1989)
17. Dwork, C., Kumar, R., Naor, M., Sivakumar, D.: Rank aggregation methods for the web. In: *Proceedings International Conference World Wide Web (WWW '01)*, vol. 10, pp. 613–622 (2001)
18. Freund, Y., Iyer, R., Shapire, R., Singer, Y.: An efficient boosting algorithm for combining preferences. *J. Mach. Learn. Res.* **4**(6), 933–969 (2004)
19. Friedman, J.: Computing Betti numbers via combinatorial Laplacians. *Algorithmica* **21**(4), 331–346 (1998)
20. Hastie, T., Tibshirani, R.: Classification by pairwise coupling. *Ann. Statist.* **26**(2), 451–471 (1998)

21. Herbrich, R., Graepel, T., Obermayer, K.: Large margin rank boundaries for ordinal regression. In: *Advances in Large Margin Classifiers*, pp. 115–132. MIT Press, Cambridge, MA (2000)
22. Hochbaum, D.: The separation, and separation-deviation methodology for group decision making and aggregate ranking. In: Hasenbein, J.J. (ed.) *Tutorials in Operations Research*. vol. 7, pp. 116–141. INFORMS, Hanover, MD (2010)
23. Hochbaum, D., Levin, A.: Methodologies and algorithms for group-rankings decision. *Manag. Sci.* **52**(9), 1394–1408 (2006)
24. Kahle, M.: Topology of random clique complexes. *Discrete Math.* (to appear, 2008)
25. Kemeny, J.G.: Mathematics without numbers. *Daedalus* **88**, 571–591 (1959)
26. Kemeny, J.G., Snell, L.J.: Preference ranking: an axiomatic approach. In: Kemeny, J.G., Snell, L.J. (eds.) *Mathematical Models in the Social Sciences*, pp. 9–23. MIT Press, Cambridge (1973)
27. Kendall, M.G., Gibbons, J.D.: *Rank Correlation Methods*. 5th ed. Oxford University Press, Oxford (1990)
28. Kendall, M.G., Smith, B.B.: On the method of paired comparisons. *Biometrika* **31**(3–4), 324–345 (1940)
29. Kleinberg, J.: Authoritative sources in a hyperlinked environment. *J. ACM* **46**(5), 604–632 (1999)
30. Ma, M.: A matrix approach to asset pricing in foreign exchange market. preprint. <http://www.ssrn.com/abstract=921755> (2008)
31. Miller, G.A.: The magical number seven, plus or minus two some limits on our capacity for processing information. *Psychol. Rev.* **63**(2), 81–97 (1956)
32. Morton, J., Pachter, L., Shiu, A., Sturmfels, B., Wienand, O.: Convex rank tests and semigraphoids. *SIAM J. Discrete Math.* **23**(3), 1117–1134 (2009)
33. Mosteller, F.: Remarks on the method of paired comparisons I, II, III. *Psychometrika* **16**(1), pp. 3–9, no. 2, pp. 203–218 (1951)
34. Noether, G.E.: Remarks about a paired comparison model. *Psychometrika* **25**, 357–367 (1960)
35. Paige, C.C., Saunders, M.A.: LSQR: an algorithm for sparse linear equations and sparse least squares. *ACM Trans. Math. Software* **8**(1), 43–71 (1982)
36. Saari, D.G., Merlin, V.R.: A geometric examination of Kemeny’s rule. *Soc. Choice Welf.* **17**(3), 403–438 (2000)
37. Saaty, T.L.: A scaling method for priorities in hierarchical structures. *J. Math. Psych.* **15**(3), 234–281 (1977)
38. Sen, A.K.: The impossibility of a paretian liberal. *J. Polit. Econ.* **78**(1), 152–157 (1970)
39. Sen, A.K.: The possibility of social choice, Nobel Lecture, 8 Dec 1998. In: Persson, T. (ed.) *Nobel Lectures: Economic Sciences 1996–2000*, pp. 178–215. World Scientific, Singapore (2003)
40. Tahbaz-Salehi, A., Jadbabaie, A.: Distributed coverage verification in sensor networks without location information. In: *Proceeding IEEE Conference Decision Control*, vol. 47. (2008), to appear
41. Thurstone, L.L.: The method of paired comparisons for social values. *J. Abnorm. Soc. Psychol.* **21**, 384–400 (1927)
42. Xu, Y., Dyer, J.S., Owen, A.B.: Empirical stationary correlations for semi-supervised learning on graphs. preprint (2009)
43. Young, H.P.: Condorcet’s theory of voting. *Am. Polit. Sci. Rev.* **82**, 1231–1244 (1988)
44. Young, H.P., Levenglick, A.: A consistent extension of Condorcet’s election principle. *SIAM J. Appl. Math.* **35**(2), 285–300 (1978)
45. Zhang, J.: Binary choice, subset choice, random utility, and ranking: a unified perspective using the permutahedron. *J. Math. Psych.* **48**(2), 107–134 (2004)