UC Berkeley
UC Berkeley Electronic Theses and Dissertations

Title
Topics in Tropical Linear Algebra and Applied Probability

Permalink
https://escholarship.org/uc/item/0gc3m4p1

Author
Tran, Ngoc Mai

Publication Date
2013

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

Topics in Tropical Linear Algebra and Applied Probability

by

Ngoc Mai Tran

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Statistics

in the

Graduate Division

of the

University of California, BERKELEY

Committee in charge:

Professor Bernd Sturmfels, Chair
Professor Jim Pitman
Professor Lior Pachter

Fall 2013

Topics in Tropical Linear Algebra and Applied Probability

Copyright 2013
by
Ngoc Mai Tran

Abstract

Topics in Tropical Linear Algebra and Applied Probability

by

Ngoc Mai Tran

Doctor of Philosophy in Statistics

University of California, BERKELEY

Professor Bernd Sturmfels, Chair

Tropical linear algebra is the study of classical linear algebra problems with arithmetic done over the tropical semiring, namely with addition replaced by max, and multiplication replaced by addition. It allows one to reformulate nonlinear problems into piecewise-linear ones. This approach has successfully been employed to solve and characterize solutions to many problems in combinatorial optimization, control theory and game theory [5]. Tropical spectral theory, the study of tropical eigenvalues and eigenspaces, often plays a central role in these applications. We derive the basics of this theory in Chapter 1.

In Chapter 2 we give a combinatorial description of the cones of linearity of the tropical eigenvector map. In Chapter 3 we extend this work to cones of linearity of the tropical eigenspace and polytrope map. Our results contribute to a better understanding of the polyhedral foundations of tropical linear algebra.

Chapter 4 illustrates the above results in the context of pairwise ranking. Here one assigns to each pair of candidates a comparison score, and the algorithm produces a cardinal (numerically quantified) ranking of candidates. This setup is natural in sport competitions, business and decision making. The difficulty lies in the existence of inconsistencies of the form $A > B > C > A$, since pairwise comparisons are performed independently. TropicalRank is an algorithm pioneered by Elsner and van den Driessche. Solution sets of this ranking method are precisely the polytropes studied in Chapter 3. For generic input pairwise comparison matrices, this set contains one unique point that is the tropical eigenvector, which is then interpreted as the comparison score. In particular, the results in Chapter 3 provide a complete classification of all possible solution sets to the optimization problem that TropicalRank solves. This answers open questions from several papers [22, 32] in the area.

In Chapter 4 we also show that TropicalRank belongs to the same parametrized family of ranking methods as two other commonly used algorithms, PerronRank and HodgeRank. Furthermore, we show that HodgeRank and PerronRank asymptotically give the same score under certain random ranking models. Despite their mathematical connections, we can construct instances in which these three methods produce arbitrarily different rank order.

The last two chapters are topics in applied probability. Chapter 5 studies the exact and asymptotic distribution of size-biased permutations of finite sequences with independent and identically distributed (i.i.d) terms. The size-biased permutation of a positive summable sequence $(x_1, x_2, \ldots)$ is the same sequence presented in a random order $(x[1], x[2], \ldots)$, where $x[1]$ is defined to be x_i with probability proportional to its 'size' x_i; given that $x[1]$ is x_i, the next term $x[2]$ is defined to be x_j for $j \neq i$ with probability again proportional to its 'size' x_j, and so on. Size-biased permutations model the successive sampling method in ecology and oil discovery, where species (or oil reserves) are discovered in a random order proportional to their abundances. In the ranking literature it is known as the Plackett-Luce model, a parametrized family modeling ranking distributions. Size-biased permutation is one of the building blocks of combinatorial stochastic processes, an area of probability with applications to computer science [78]. Finite i.i.d sequence setup serves as a simple model for successive sampling, or ranking with increasing number of items. We study the size-biased permutation of such a sequence using two separate methods: Markov chains and induced order statistics. By going back and forth between the two approaches, we arrive at more general results with simplified proofs, and provide a Poisson coupling argument which leads to an explicit formula for the asymptotic distribution of the last few terms in the size-biased permutation.

Chapter 6 is about the binary Little-Hopfield network. This is an established computational model of neural memory storage and retrieval which can have exponential capacity relative to the number of neurons. However, known algorithms have produced networks with linear capacity, and it has been a long-standing open problem whether robust exponential storage is possible. For a network with n neurons, the problem involves a linear program in n^2 variables and exponentially many constraints. We utilized the action of the symmetric group on the neuron labels and successfully reduced the problem to a linear program in three variables and three constraints. Thus we explicitly constructed simple networks that answer the question affirmatively, with the best possible asymptotic robustness index. This work calls for further research into Little-Hopfield networks and their applications to theoretical neuroscience and computer science.

To my family and teachers

For making me who I am.

Contents

List of Figures **iv**

1 Tropical Linear Algebra **1**
1.1 Introduction . . . 1
1.2 The tropical eigenvalue . . . 2
1.3 An application: longest paths . . . 6
1.4 Tropical eigenspace and polytropes . . . 7
1.5 Tropical and inverse tropical limit of classical eigenvector . . . 9
1.6 Applications to pairwise ranking . . . 10

2 Tropical Eigenvector: Cones of Linearity **12**
2.1 Introduction . . . 12
2.2 The partition $\mathcal{P}_n$ and its full-dimensional cones . . . 14
2.3 Eigenpair cones and failure of the fan property . . . 16
2.4 Skew-symmetric matrices . . . 19
2.5 Summary . . . 21

3 Polytropes and Tropical Eigenspaces: Cones of Linearity **22**
3.1 Introduction . . . 22
3.2 Construction of $\mathcal{F}_n$ . . . 24
3.3 Complete and compatible sets of connected relations . . . 27
3.4 The lattice isomorphism . . . 30
3.5 Codimension, facet-defining equations and symmetries . . . 37
3.6 Examples . . . 38

4 TropicalRank **42**
4.1 Introduction . . . 42
4.2 Pairwise ranking . . . 43
4.3 Definitions and basic properties . . . 44
4.4 Geometry of TropicalRank . . . 47
4.5 Proof of the pairwise inconsistency theorem . . . 53

4.6 HodgeRank is the limit of PerronRank 56
4.7 Random ranking models 60
4.8 Summary and open problems 61

5 Finite Size Biased Permutation **64**
5.1 Introduction 64
5.2 Markov property and stick-breaking 66
5.3 Size-biased permutation as induced order statistics 69
5.4 Limit in distributions of finite size-biased permutations 74
5.5 Asymptotics of the last u fraction of the size-biased permutation 81
5.6 Poisson coupling of size-biased permutation and order statistics 87
5.7 Summary 93

6 Exponential Binary Storage in Little-Hopfield Networks **94**
6.1 Introduction 95
6.2 Little-Hopfield networks and robust storage 96
6.3 Main theorem 98
6.4 Properties of the two-parameter family 104
6.5 Summary 109

Bibliography **110**

List of Figures

2.1 The simplicial complex Σ_3 of connected functions $\phi : [3] \to [3]$. Fixed-point free ϕ are on the left and functions with $\phi(3) = 3$ on the right. 18

3.1 Cones in the partition $\mathcal{N}_2 = \mathcal{P}_2 = \mathcal{F}_2$ of $\mathbb{R}^{2\times 2}$, their defining equations and inequalities, and their associated maximal paths. Figure is drawn modulo lineality space, and thus represents the fan $\partial\mathcal{J}_2 = \partial\mathcal{K}_2$ of $\mathbb{R}^2$. Edges in the critical cycles (cycles with maximal mean) are shown in red. 27
3.2 An incomplete set of connected functions $\mathcal{G} = (D, E)$. 29
3.3 All three possible connected relations that can complete the set $\mathcal{G}$ of Figure 3.2. 29
3.4 Some connected relations incompatible to $\mathcal{G}$ of Figure 3.2. 30
3.5 Table of $N(k, \lambda, p, c)$, the number of cones in $\mathcal{F}_3$ of a given dimension k whose indexing complete set of connected relations G satisfy $\lambda(G) = \lambda$, with p the number of equalities coming from paths, and c the number of equalities coming from cycles. 39
3.6 The five equivalence classes and their sizes, $\lambda = (1, 1, 1), p = c = 0$ 39
3.7 Facets of a non-simplicial full-dimensional cone in $\mathcal{F}_3$ labeled by $\mathcal{G}$, the complete connected relation shown in the middle box. 40
3.8 The face lattice of $\Psi(G)$. 41

4.1 The zonotope $Z_\perp$ and its subdivision to cones of linearity of the tropical eigenvector map. 50

5.1 A point scatter on the plane. Here $J_1 = 5, J_2 = 3, J_3 = 1, J_4 = 4, J_5 = 2$, and $K_1 = 3, K_2 = 5, K_3 = 2, K_4 = 4, K_5 = 1$. The permutations J and K are inverses. Conditioned on $x^\uparrow(2) = a$ and $y^*(2) = b$, the number of points lying in the shaded region determines $J_2 - 1$. 88

6.1 Illustration of the energy landscape of a Little-Hopfield network depicting the robust storage of all 4-cliques in graphs on 8 vertices. The network dynamics sends a graph that is almost a clique to a graph with smaller energy, until finally converging to the underlying 4-clique as a stable point. 99

6.2 Robustness of finite networks with $v = 2k = 50, 75, 100, 125$ and 150 vertices with parameters given in Corollary 6.9. The cut-off for p, the probability of an edge being corrupted, occurs around $p = 0.11$. The asymptotic cut-off for networks with these parameters as $v \to \infty$ is $p = 1/6 \approx 0.167$. These two cut-offs are indicated by vertical dotted lines. 100
6.3 Feasible region for network parameters giving exponential storage. The shaded region is the feasible polytope $\bigcap_{k=5}^{15} \mathbf{I}_k$ for network parameters giving clique storage for the range $5 \leq k \leq 15$. Black points are its vertices, the red R_k and blue B_k lines are linear constraints. 108
6.4 Feasible region for larger range of k. Lines R_k (red) and B_k (blue) for $1000 \leq k \leq 5500$. 108

Acknowledgments

I would like to thank my wonderful advisor Bernd Sturmfels for being a live definition of the word 'advisor'. Writing this thesis has been a lot of fun. Looking back, I realized how much I have grown up, not only as a researcher but also as a person. And I am sure I could not have done this without Bernd. He has a relaxed yet very effective advising style. Initially I was plagued with worries. He gave me 'too much' freedom in selecting topics to work on. Topics that attracted me were different from the 'popular' topics in the Statistics department. We had fewer 'technical' discussions than I had imagined. In retrospect, our discussions, both technical and others, direct and indirect, have been tremendously valuable. Bernd is enthusiastic, easy to approach and always encouraging. He has a profound influence on the way I think. I am very honored to be his student.

I am in debt to Lek-Heng Lim for acting as my second advisor during the early stages of my graduate program. Lek-Heng is enthusiastic, has great insights to problems, and would have also made a wonderful advisor. Without his generous help in that critical stage, this dissertation would not be here today. I am very grateful for our interactions.

I would like to thank my collaborators and teachers, Jim Pitman, Bernd Sturmfels, and Chris Hillar. Special thanks to Jim for showing me the joy of striving to understand a problem back to front, and to explain oneself clearly and succinctly. He helped me to realize the saying 'knowing is not enough'. Working with Jim, Bernd and Chris has been a very valuable experience.

I would like to thank all of my professors, especially Steve Evans, Bin Yu, Yuval Peres, Elchanan Mossel, David Aldous, Nourredine El Karoui, and Haiyan Huang. I sincerely thank everyone at the Berkeley Statistics department for creating a truly wonderful environment to live, work and play. I would not have discovered Berkeley nor Statistics without the encouragement of my undergraduate professors and mentors in Australia, Jacqui Ramagge and Iain Raeburn. Many thanks to my friends, officemate, teammate, housemate and almost housemates: Linda, Luis, Vince, Miki, Wayne, Jasmine and Rachel. A bunch of thanks to Po-Ling and Christine for the Saturday misery parties, an essential part of graduate school.

I would like to thank my family and in-laws for their encouragement and patience with the stack of dishes in the sink, the occasional brief visits and the solemn dinners with the 'I'm stuck' face. Last but not least, special thanks to Joe Neeman, my life companion.

Chapter 1

Tropical Linear Algebra

This chapter describes the basics of tropical linear algebra and in particular, tropical eigenvalues, eigenvectors and their applications. These are the main objects of study in the next three chapters of this dissertation. The main references are [5, 16]. Chapter 5 and 6 represent self-contained work on other unrelated topics, and hence have their own background introduction.

1.1 Introduction

The *max-plus* algebra $(\mathbb{R}, \oplus, \odot)$ is an algebra over the real numbers with addition replaced by max, and multiplication replaced by addition, that is, $u \oplus v = \max\{u, v\}, u \odot v = u + v$. Some authors attach the additive identity $-\infty$ to obtain the *max-plus semiring* $(\mathbb{R} \cup \{-\infty\}, \oplus, \odot)$. Replacing max by min yields the *min-plus* algebra. Replacing addition by multiplication yields the *max-times* algebra. All of these are known as (different flavors of) *tropical algebra* (and likewise, *tropical semiring*). For the majority of this dissertation we shall use the *max-plus* algebra unless otherwise indicated, although most of our results directly carry over to the other flavors listed. We quote [66] for an explanation of the word 'tropical'.

> The adjective 'tropical' was coined by French mathematicians, notably Jean-Eric Pin, to honor their Brazilian colleague Imre Simon, who pioneered the use of min-plus algebra in optimization theory. There is no deeper meaning in the adjective 'tropical'. It simply stands for the French view of Brazil.

Tropical mathematics is (roughly) the study of classical problems with arithmetic done over the tropical algebra. As hinted in the quote above, combinatorial optimization provides one of the many reasons to go 'tropical'. Many problems in this area become *linear* over the tropical algebra. This has at least two advantages. First of all, it simplifies notation and connects different areas of mathematics. Secondly, we can try solving these problems by searching for tropical analogues of classical tools. A staple example of such success stories in

tropical linear algebra is the tropical spectral theory, discussed in detail in Section 1.2 to 1.4 below. Conversely, tropical mathematics can enrich its classical counterparts. In many cases, tropical objects naturally arise as limits of classical objects; we shall see a concrete example in Section 1.5. This gives rise to the verb 'tropicalize'. One may hope to study classical objects by understanding analogous properties of their tropical versions. This approach has been very successful in areas such as tropical algebraic geometry, establishing deep connections with the classical theory as well as other fields.

Most areas of tropical mathematics existed as a field in their own right before the appearance of the term 'tropical'. For the case of tropical linear algebra, many developments came from certain areas of operations research, combinatorial optimization, signal processing and automata theory. These have been generalized to a discrete, combinatorial version of dynamical systems theory called *discrete event dynamic systems* [5]. One central problem in this area is finding fixed points of dynamical systems whose evolution equations are linear over $(\mathbb{R}, \oplus, \odot)$. This led to the development of tropical spectral theory, in other words, the study of tropical eigenvalues and eigenvectors. Other areas are under active research. Compared to their classical cousins, tropical matrices are less understood. The characterization of commuting tropical matrices, for example, is still an open problem.

Chapter 2 and 3 of this dissertation contribute to the intersection of tropical spectral theory and polyhedral geometry. Our work leads to new combinatorial classifications of tropical eigenspaces and tropical simplices, also known as full-dimensional polytropes. In Chapter 4, we derive an 'inverse' tropical limit of the Perron eigenvector of positive matrices. In the context of pairwise ranking, our results reveal novel mathematical connections between the three most popular pairwise ranking methods. We further compare these three methods and construct instances in which they can produce arbitrarily different rank orders.

In this chapter, we review tropical (max-plus) spectral theory with applications from discrete event dynamic systems in mind. In Section 1.2, we derive the combinatorial characterization of the unique tropical eigenvalue. Section 1.3 serves as the prelude to our study of tropical eigenspaces and polytropes in Section 1.4. In Section 1.5, we state our inverse tropical limit theorem. We briefly discuss the implications of these results in pairwise ranking in Section 1.6.

1.2 The tropical eigenvalue

1.2.1 A basic classical dynamical system

The starting point of many first courses on discrete dynamical systems is the difference equation

$$x(t+1) = Ax(t), \quad t = 0, 1, 2, \ldots \tag{1.1}$$

where $x(t)$ is a vector in $\mathbb{R}^n$ describing the state of the system at time t, $A \in \mathbb{R}^{n \times n}$ is a given matrix. One general goal of this field is to understand the long-run behavior of this system.

In this case, the evolution of states over time is completely governed by the eigenvalues and eigenvectors of A. For simplicity, let us assume A is a symmetric matrix, hence all of its eigenvalues $\lambda_1, \ldots, \lambda_n$ are real. Equip $\mathbb{R}^n$ with the eigenbasis of A, one can reduce to the case where A is a diagonal matrix. Thus (1.1) becomes

$$x_i(t+1) = \lambda_i x_i(t), \quad t = 0, 1, 2, \ldots, i = 1, \ldots, n.$$

For the i-th coordinate, if $|\lambda_i| > 1$, the system at this coordinate diverges exponentially fast. If $|\lambda_i| < 1$, the system converges exponentially fast to 0. If $\lambda_i = -1$, the system oscillates between the original input $x_i(0)$ and its negative $-x_i(0)$, and if $\lambda_i = 1$, then any input to this coordinate is a fixed point. In control theory, the log of the maximal eigenvalue of A is the maximal *Lyapunov exponent* of (1.1). It indicates the stability of a dynamical system, with positive values generally signify chaotic behavior (in this case, divergence). For general matrices $A \in \mathbb{R}^{n\times n}$, its maximal Lyapunov exponent is the log of its operator norm

$$\|A\|_{op} = \min\{c \in \mathbb{R} : \|Ax\| \leq c\|x\| \text{ for all } x \in \mathbb{R}^n\},$$

where $\|\cdot\|$ denotes the ℓ_2-norm in $\mathbb{R}^n$.

1.2.2 The analogous tropical dynamical system

The tropical (max-plus) analogue of (1.1) written out for each coordinate is

$$x_i(t+1) = (A \odot x(t))_i = \max_{j=1,\ldots n} (A_{ij} + x_j(t)). \tag{1.2}$$

Tropical analogues of dynamical systems arise in operations research, automata theory, combinatorial optimization, and other areas of discrete event dynamic systems [5]. In scheduling, for example, each node i may represent a machine producing a certain part, which cannot be started before receiving all the parts from the other machines. Let A_{ij} be the time it takes to ship parts from i to j and $x(0)$ be the time when the machines start production. Then $x(t)$ given by (1.2) is the time the t-th iteration of the factory produces the next set of products. (Admittedly our 'application' does not sound realistic, but it is not much worse than attempts to motivate (1.1) as a model for growth of economies due to trading. See [5] and references therein for various examples and applications of tropical methods in discrete event dynamic systems.)

To understand the behavior of (1.1), let us look for an analogue of the maximal Lyapunov exponent. The straight-forward max-plus analogue of the ℓ_2-norm is

$$\frac{1}{2} \max_{i=1,\ldots,n} 2x_i = \max_{i=1,\ldots,n} x_i.$$

The analogue of the operator norm is the solution to the following linear program

$$\text{Minimize } \lambda \text{ subject to } \quad A_{ij} + x_j \leq \lambda + x_i \text{ for all } 1 \leq i, j \leq n. \tag{1.3}$$

The linear program dual to (1.3) takes the form

$$\begin{array}{c} \text{Maximize } \sum_{i,j=1}^n A_{ij} p_{ij} \text{ subject to } p_{ij} \geq 0 \text{ for } 1 \leq i, j \leq n, \\ \sum_{i,j=1}^n p_{ij} = 1 \text{ and } \sum_{j=1}^n p_{ij} = \sum_{k=1}^n p_{ki} \text{ for all } 1 \leq i \leq n. \end{array} \tag{1.4}$$

The p_{ij}'s are the decision variables. One says that (p_{ij}) defines an *edge flow* (or simply a flow), where p_{ij} is the flow from i to j. The first two constraints require (p_{ij}) to be a probability distribution on the edges of the complete directed graph on n vertices. The last constraint, $\sum_{j=1}^n p_{ij} = \sum_{k=1}^n p_{ki}$, states that (p_{ij}) is a *total flow conservation*. (That is, at every node i, the total in-flow is equal to the total out-flow).

Let Γ_n be the set of simple directed cycles on the complete directed graph on n nodes, including self-loops. For $C \in \Gamma_n$, let $\chi_C \in \{0,1\}^{n^2}$ be its incidence vector, $|C|$ the number of edges in the cycle. For $i, j \in \{1, \ldots, n\}$, write χ_{ij} for the incidence vector of the edge $i \to j$. Altogether in (1.4) there are n independent equality constraints, thus the set $\mathcal{C}_n$ of all feasible solutions to (1.4) in $\mathbb{R}^{n \times n}$ is an $(n^2 - n)$-dimensional polytope. One can check that the vertices of $\mathcal{C}_n$ are precisely the vectors of the form $\frac{1}{|C|}\chi_C$ for some $C \in \Gamma_n$. Thus it is the *normalized cycle polytope*

$$\mathcal{C}_n = \operatorname{conv}\left\{ \frac{1}{|C|}\chi_C : C \in \Gamma_n \right\}. \tag{1.5}$$

By strong duality, the primal (1.3) and the dual (1.4) have the same optimal value. This implies that the max-plus operator norm function $A \mapsto \lambda(A)$ is the support function of the convex polytope $\mathcal{C}_n$. Since each vertex of $\mathcal{C}_n$ is the uniform probability distribution on a directed cycle, the max-plus operator norm $\lambda(A)$ is the maximum cycle mean of the directed graph with edge weights A.

Proposition 1.1. *The max-plus operator norm $\lambda(A)$ of a matrix $A \in \mathbb{R}^{n \times n}$ is the maximum cycle mean of the directed graph with edge weights A.*

Definition 1.2. Cycles which achieve the maximum cycle mean are the *critical cycles*, their graph union is the *critical graph*, and their nodes are the *critical nodes*.

Let us check that $\lambda(A)$ plays the role of the tropical Lyapunov exponent. Starting at $x(0) \in \mathbb{R}^n$, after k iterations of (1.2), we obtain

$$x(k) = A^k \odot x(0),$$

where matrix multiplication takes place in the max-plus algebra. Here A^k_{ij} is the maximal path of length k from i to j. Thus for all $i, j \in \{1, \ldots, n\}$, if $\lambda(A) \neq 0$, A^k_{ij} diverges to either

plus or minus infinity as $k \to \infty$ and so do the entries of $x(k)$. When $\lambda(A) = 0$, the A_{ij}^k's and hence $x(k)$ iterate through a finite set of values depending on the input and the length of the maximal cycle. It follows from Theorem 1.3 below that the fixed points of (1.2) are tropical eigenvectors of A. We give their characterization in Section 1.4, and a complete combinatorial description in Chapter 3.

1.2.3 The tropical eigenvalue

The classical operator norm of a matrix is also its largest singular value. In the case of symmetric matrices, this is precisely the largest absolute eigenvalue. Given this connections between operator norm and eigenvalue, let us proceed to characterize all tropical eigenvalues of a matrix $A \in \mathbb{R}^{n\times n}$. Let $(\lambda, x) \in \mathbb{R} \times \mathbb{R}^n$ be a max-plus eigenvalue-vector pair, that is,

$$\max_{j=1,\ldots,n} A_{ij} + x_j - x_i = \lambda \quad \text{for each } i = 1, \ldots, n. \tag{1.6}$$

By definition, $\lambda \geq \lambda(A)$. For each $i = 1, \ldots, n$, let $\phi(i)$ be one of the possibly many j's that attain the maximum in the left-hand side of (1.6). Then ϕ is a function from $[n] = \{1, \ldots, n\}$ to itself, thus its graph necessarily contains a cycle $C \in \Gamma_n$. Summing the two sides of (1.6) over edges in C gives

$$A \cdot \chi_C = \lambda |C|, \ \Rightarrow \lambda = A \cdot \frac{1}{|C|}\chi_C.$$

By Proposition 1.1, $\lambda \leq \lambda(A)$. Hence $\lambda = \lambda(A)$. In other words, a tropical matrix has precisely one tropical eigenvalue given by $\lambda(A)$. Combined with Proposition 1.1, we obtain the classical combinatorial characterization of tropical eigenvalues.

Theorem 1.3 (Cunninghame-Green [21])**.** *Each matrix $A \in \mathbb{R}^{n\times n}$ has a unique max-plus eigenvalue $\lambda(A)$, which is also its max-plus operator norm. Furthermore, $\lambda(A)$ is equal to the maximum cycle mean of the directed graph with edge weights A.*

For the benefit of future sections, we now collect some definitions and observations. One can check that the orthogonal complement to the affine span of the cycle polytope $\mathcal{C}_n$ is the space spanned by

$$W_n = \text{span}\left\{\sum_{i,j=1}^{n} \chi_{ij}\right\} + V_n \tag{1.7}$$

where $+$ denotes direct sum of subspaces, and V_n is the space of total flow conservation

$$V_n = \text{span}\left\{\sum_{j=1}^{n} (\chi_{ij} - \chi_{ji}) : i = 1, \ldots, n\right\}. \tag{1.8}$$

Each matrix $A \in V_n$ can be written as $A_{ij} = a_i - a_j$ for some vector $a \in \mathbb{R}^n$ unique up to additive constants [55]. The vector a can be interpreted as a potential vector, and the

flow from i to j is simply the potential difference between two nodes. For this reason, V_n is sometimes called the space of *gradient flows* [55]. In view of (1.3) where the term $x_i - x_j$ appeared, the appearance of V_n is not surprising. We shall revisit V_n in Section 1.6 from the viewpoint of pairwise ranking.

1.3 An application: longest paths

Define the graph of a matrix $A \in \mathbb{R}^{n\times n}$ to be the weighted graph with edge weights A_{ij}. Fix a target node $i^* \in [n] = \{1, \ldots, n\}$. The *single-target longest path* problem involves finding the maximal path from i to i^* in the graph of A for all $i \neq i^*$. Variations include *single-source longest path*, where one fixes the source instead of the target node, and *all-pairs longest path*, where one simultaneously finds the longest path over all pairs $(i, j) \in [n]^2$.

Often formulated in terms of shortest rather than longest paths, this is a basic and important combinatorial optimization problem. Recall our scheduling example in the previous section. If i^* outputs the final product, then the longest path to i^* gives a lower bound on the total production time. A less obvious example is problem of finding the most likely hidden states in a hidden Markov model [5, §1.2.2], [74]

Better known solutions of the longest path problem use dynamic programming, such as the Floyd-Warshall algorithm for all-pairs longest path, or the Viterbi algorithm for inferring hidden states previously mentioned [1]. Here we present a tropical solution. While it does not yield faster algorithms, the tropical approach reveals interesting mathematical properties of the solution set, in particular, its connections with tropical eigenvectors.

Definition 1.4. For a matrix $A \in \mathbb{R}^{n\times n}$ with $\lambda(A) \leq 0$, define

$$A^+ := A \oplus A^2 \oplus \ldots, \text{ and } \quad A^* := I \oplus A^+,$$

where I is the matrix with 0 on the diagonal and $-\infty$ elsewhere. The map $A \mapsto A^*$ is the *Kleene star* map, and A^* is the *Kleene star* of A.

This terminology originated from automata theory [91]. The condition $\lambda(A) \leq 0$ ensures that the Kleene star map is elementwise finite. By a direct computation, we obtain the following.

Lemma 1.5. *A^*_{ij} is the maximal path value amongst paths from i to j in the graph of A.*

Thus A^* is precisely the solution to the all-pairs longest path problem with edge weights A. This problem is known to have an elementwise finite solution if and only if A has non-positive maximal cycle, equivalently, $\lambda(A) \leq 0$. In particular, for a fixed column i^*, the map $A \mapsto A^*_{i^*}$ is the solution to the following classical linear program of single-target longest path [1, §4].

$$\text{Maximize } \sum_{i,j=1}^{n} A_{ij}x_{ij} \tag{1.9}$$
$$\text{subject to } x_{ij} \geq 0, \sum_{j=1}^{n} x_{ji} - \sum_{j=1}^{n} x_{ij} = \begin{cases} n-1 & \text{for} \quad i = i^* \\ -1 & \text{for} \quad i \neq i^* \end{cases}.$$

The constraint set of this linear program has a nice interpretation. Each node $i \neq i^*$ sends out a net flow of one unit, and the node i^* receives a net flow of $n-1$ units. One wants to find such a flow (x_{ij}) that maximizes the total weighted flow $\sum_{i,j=1}^{n} A_{ij}x_{ij}$. Thus flows which use larger edges are better, and in particular, the optimal flow picks out the best path to i^*.

This linear program is bounded if and only if $\lambda(A) \leq 0$. Its constraint set is a polyhedron in $\mathbb{R}^{n\times n}$ with the $(n-1)$ dimensional lineality space

$$\{A \in \mathbb{R}^{n\times n} : A \cdot \chi_C = 0 \text{ for all } C \in \Gamma_n\},$$

which, after a little computation, turns out to be the space of gradient flows V_n defined in (1.8). The vertices of the polyhedron are in bijection with in-directed spanning trees on n nodes with root i^* [1]. These are the *trees of longest path* to i^*, since it specifies the longest path to i^* from any other node $i \neq i^*$. From such an in-directed spanning tree, one recovers the flow by defining x_{ij} to be the number of parents of j if the edge $i \to j$ is in the tree, and 0 otherwise. This gives a combinatorial interpretation to the columns of A^*.

Proposition 1.6. *Each column $i^* \in [n]$ of A^* can be associated with a collection $\tilde{T}_{i^*}$ of in-directed spanning trees on n nodes with root i^*, such that for any tree $T \in \tilde{T}_{i^*}$, A^*_{ji} is the sum of the edge values along the path from j to i^* in T.*

1.4 Tropical eigenspace and polytropes

In this section we define the two terms in the title and study their properties. We shall see that they are tropically convex sets, and their tropically extreme vectors can be computed using Kleene stars. Proposition 1.6 then gives a combinatorial description of the set of tropical eigenspaces, and is the starting point for the discussion in Chapter 2.

We first need a notion of tropical convex hull. Since the additive identity is $-\infty$ in the max-plus algebra, every real number is tropically positive. Thus, the *tropical line segment between* two points $x, y \in \mathbb{R}^n$ is the set

$$\{(a \odot x) \oplus (b \odot y) : a, b \in \mathbb{R}\} = \{\max\{a + x, b + y\} : a, b \in \mathbb{R}\}.$$

We say that a subset $S \subset \mathbb{R}^n$ is *tropically convex* if it contains the tropical line segment between any two of its points. Since tropically convex sets are closed under tropical scalar

multiplication, one often views them as subsets of the *tropical projective torus*

$$\mathbb{TP}^{n-1} := \mathbb{R}^n/\mathbb{R} \odot (0,\dots,0)^T = \mathbb{R}^n/\mathbb{R} \cdot (1,\dots,1)^T.$$

The tropical convex hull of finitely many points in $\mathbb{TP}^{n-1}$ is a *tropical polytope*.

Definition 1.7. For $A \in \mathbb{R}^{n\times n}$, the *tropical eigenspace* $Eig(A)$ is the set of vectors $x \in \mathbb{TP}^{n-1}$ which satisfy

$$\max_{1\le j\le n}(A_{ij} + x_j) = \lambda(A) + x_i, \qquad \text{for each } i = 1,\dots,n. \tag{1.10}$$

The set of subeigenvectors, or the *polytrope of* A, is denoted $Pol(A)$. This is the set of vectors $x \in \mathbb{TP}^{n-1}$ which satisfy

$$\max_{1\le i,j\le n}(A_{ij} + x_j - x_i) = \lambda(A). \tag{1.11}$$

Joswig and Kulas [57] coined the name *polytrope* to describe sets which are convex both in the tropical and ordinary sense. A direct computation shows that $Pol(A)$ is indeed a polytrope in this sense, while $Eig(A) \subseteq Pol(A)$ is a tropical polytope. It follows from [29, Proposition 18] that all polytropes arise as $Pol(A)$ for some matrix A. Sometimes we shall abuse terminologies and identify $Eig(A)$ and $Pol(A)$ with their tropical extreme points.

The following proposition neatly characterizes the tropical extreme points of $Eig(A)$ and $Pol(A)$. The proof is a direct computation, see [91].

Proposition 1.8 ([91])**.** *Define* $\bar{A} := A \odot (-\lambda(A))$*, that is,* $\bar{A}_{ij} = A_{ij} - \lambda(A)$ *for all* $i, j \in [n]$*. Then up to tropical scaling, the extreme tropical vertices of* $Pol(A)$ *are the columns of* $\bar{A}^*$*. The extreme tropical vertices of* $Eig(A)$ *are precisely the common columns of* $\bar{A}^+$ *and* $\bar{A}^*$*. The indices of such columns are precisely the critical nodes of* A*.*

The following lemma tells us exactly which columns of $Pol(A)$ are tropically dependent, and hence, which are precisely the tropical extreme points of $Eig(A)$ and $Pol(A)$. Recall the definition of critical graph in Definition 1.2.

Lemma 1.9 ([91])**.** *For* $A \in \mathbb{R}^{n\times n}$*, any subset of columns of* $(\bar{A})^*$ *are either independent, or at least two are tropical scalar multiples of each other. Two columns* i, j *of* $(\bar{A})^*$ *are equal up to tropical scaling if and only if* i *and* j *are connected in the critical graph of* A*.*

Corollary 1.10. *The maximal number of extreme tropical vertices a polytrope in* $\mathbb{TP}^{n-1}$ *can have is* n*. These polytropes are tropical simplices.*

Combined with Proposition 1.6, we have the following combinatorial interpretation of tropical eigenvectors as trees of longest paths rooted at critical nodes.

Theorem 1.11. *For a matrix* $A \in \mathbb{R}^{n\times n}$*, the number of extreme tropical vertices of its tropical eigenspace* $Eig(A)$ *is the number of maximal strongly connected components in its*

critical graph. In particular, A has a unique tropical eigenvector $x(A)$ if and only if its critical graph is strongly connected. In such cases, the eigenvector $x(A)$ can be calculated by first fixing a critical vertex ℓ, and then setting $x(A)_i$ to be the maximal length among paths from i to ℓ in the graph of $\bar{A} := A \odot (-\lambda(A))$. Equivalently, $x(A)$ is the ℓ-th column of $\bar{A}^*$.

Theorem 1.11 falls slightly short of being the analogue of Theorem 1.3 for tropical eigenspaces. While each extreme tropical eigenvector can be associated with a collection of trees of longest path, it is not clear which sets of trees can appear together as the extreme tropical vertices of $Eig(A)$. Furthermore, each column of $\bar{A}^*$ is the solution of a linear program. Can $Pol(A)$ or $Eig(A)$ also be computed as the solution of a linear program? In that case, what would be the analogue of $\mathcal{C}_n$? We address all of these questions in Chapter 2 and Chapter 3, first for the case of unique tropical eigenvectors, followed by the general case.

In this dissertation we will only consider matrices with finite entries. We mention that over the max-plus semiring $(\mathbb{R} \cup \{-\infty\}, \oplus, \odot)$, if $A_{ij} = -\infty$, then the graph of A does not contain the edge $i \to j$. One says that A is irreducible if its graph is strongly connected, otherwise it is reducible, and each strongly connected component is an irreducible component. The spectral theory we have discussed so far applies to irreducible matrices. The general theory can be built by considering the spectra of the irreducible components, see [16, §4]. The spectral theory of reducible tropical matrices largely parallels the classical Perron-Frobenius theory for nonnegative matrices. For instance, each matrix admits a unique lower-diagonal representation called the Frobenius normal form, which partitions the node set into classes correspond to strongly connected components. Accessibility between classes play a key role in the behavior of eigenvalues. The maximum cycle mean corresponds to the Perron root for classical irreducible nonnegative matrices, and finite eigenvectors correspond to classical positive eigenvectors. There is a good reason for this parallel, namely the fact that tropical (max-times) eigenvectors appear as the limit of the Perron vector. We make this statement explicit in Section 1.5 below.

1.5 Tropical and inverse tropical limit of classical eigenvector

Consider the tropical max-times algebra defined over the positive real numbers. Since the max-times and max-plus algebra are in bijection via the log/exponential map, all of the spectral theory carries through mostly unchanged. Note that the positive projective space $\mathbb{PR}_+^{n-1}$, which is isomorphic to the simplex Δ_{n-1}, plays the role of $\mathbb{TP}^{n-1}$. The additive identity is 0, and the maximum cycle mean is now the maximum geometric cycle mean.

Let $X = [X_{ij}]$ be a $n \times n$ matrix, $\mathcal{K} := \mathbb{R}_+^{n \times n}$ be the open cone of elementwise positive matrices. (We suppress the dependent on n of $\mathcal{K}$). Write $X^{(k)} := [X_{ij}^k]$ for its kth Hadamard power. For $X \in \mathcal{K}$, let $v(X), \rho(X)$ denote its unique classical principal eigenvector-value pair in $\Delta_{n-1} \times \mathbb{R}_+$, also known as the Perron eigenvector and the Perron root of X. We

reserve $\lambda(X)$ for the tropical max-times eigenvalue of X, which is its maximum geometric cycle mean. The following theorem connects tropical max-times spectral theory with its classical version, known as the Perron-Frobenius theory for positive matrices.

Theorem 1.12 ([2]). *For $X \in \mathcal{K}$,*

$$\lim_{k\to\infty} \rho(X^{(k)})^{1/k} = \lambda(X).$$

Furthermore, if the tropical max-times eigenvector $m(X)$ of X is unique in Δ_{n-1}, then it is the coordinate-wise limit of the sequence $v(X^{(k)})^{1/k}$ as $k \to \infty$

$$\lim_{k\to\infty} v(X^{(k)})^{1/k} = m(X)$$

See [2] for generalizations of this statement to sequence of nonnegative matrices $\{X(k)\}$ such that $(X(k))^{1/k}$ converges elementwise as $k \to \infty$. Interestingly, there is also an 'inverse tropical limit', that is, the sequence $(X(k))^{1/k}$ also converges elementwise as $k \to 0$. We shall prove this fact in Chapter 3.

Theorem 1.13. *For $X \in \mathcal{K}$, define $h(X)_i = (\prod_{j=1}^{n} X_{ij})^{1/n}$. Then*

$$\lim_{k\to 0} v(X^{(k)})^{1/k} = h(X)$$

1.6 Applications to pairwise ranking

Theorems 1.12 and 1.13 play an important role in the context of pairwise ranking. Here one assigns to each pair of candidates a comparison score, forming a pairwise comparison matrix X. A ranking algorithm produces a cardinal (numerically quantified) ranking of candidates. This setup is natural in sport competitions, business and decision making. Recent applications include crowd-sourcing in large e-commerce databases such as Netflix, where each viewer only views and rates a tiny fraction of all movies, yet 99% of the movie pairs have comparisons by some viewers. Thus the 'opinion of the crowd', such as the fraction of viewers who prefer one movie to another, can serve as pairwise comparison scores, forming an almost complete data set from an otherwise sparse input. The difficulty of this setup lies in the existence of inconsistencies of the form $A > B > C > A$, since pairwise comparisons are performed independently.

Given a pairwise comparison matrix X, the triple $v(X), m(X), h(X)$ are outputs of three ranking algorithms: Perron Rank [86,87], Tropical Rank [31] and HodgeRank [20,55]. These are the three most commonly used algorithms in pairwise comparison ranking. Perron Rank plays a fundamental role behind the Analytic Hierarchical Process [86], a ranking procedure extensively applied in decision making. Meanwhile, HodgeRank is closely related to many

pairwise ranking algorithms in the broader literature of rank learning in computer science [39,55]. A number of papers have been devoted to their comparisons [30,31,35,55,85,86,101]. The previous two theorems imply that Perron Rank can be viewed as a family of ranking methods $V_k : X \mapsto v(X^{(k)})^{1/k}$ for any $k \in (0, \infty)$, with HodgeRank and Tropical Rank appearing as limiting cases. The choice of k reflects the freedom in choosing an exponent base, or equivalently, the freedom in choosing the unit of additive measurement, when one converts additive comparison to multiplicative comparison data. This conversion is not uncommon, as many datasets come naturally in additive comparison form in computer science applications [39, 49, 55].

In Chapter 4, we compare these three methods and their mathematical connections. We prove Theorem 1.13 and its geometric version, stated in terms of the one-dimensional real positive Kalman variety, a term coined by Ottaviani and Sturmfels [73]. These theorems show that the three methods mathematically belong to a parametrized family of ranking algorithms. On the other hand, we also show that they can produce arbitrarily different rank order on the same input. To be precise, for any two of the three methods, and for any pair of rankings of at least four items, there exists a comparison matrix for the items such that the rankings found by the two methods are the prescribed ones. We then consider random ranking models, and prove Theorem 4.19, which states that Perron Rank and HodgeRank asymptotically almost surely give the same ranking when pairwise scores are independent and identically distributed (i.i.d) from some general distributions. Thus, the method choice may not matter in large scale applications.

Chapter 2

Tropical Eigenvector: Cones of Linearity

This chapter is based on the paper 'Combinatorial types of tropical eigenvectors', joint work with Bernd Sturmfels [97], published in the Bulletin of the London Mathematical Society 45, 27-36, 2013. The main results are generalized in Chapter 3. We work with the max-plus algebra $(\mathbb{R}, \oplus, \odot)$.

2.1 Introduction

Recall Theorem 1.3 which states that the tropical eigenvalue $\lambda(A)$ of a matrix $A \in \mathbb{R}^{n\times n}$ is unique and is equal to the maximum cycle mean in the graph of A. Our exposition shows that the map $A \mapsto \lambda(A)$ is the support function of the normalized cycle polytope $\mathcal{C}_n$ defined in (1.5). The normal fan $\mathcal{N}_n$ of $\mathcal{C}_n$ partitions $\mathbb{R}^{n\times n}$ into polyhedral cones on which the eigenvalue map $A \mapsto \lambda(A)$ is linear. The closed cones of full dimension in $\mathcal{N}_n$ are naturally indexed by cycles in the graph on n vertices. Closed cones of lower dimensions in this fan are thus indexed by graph unions of cycles. Therefore, we can speak of the *combinatorial type* of the tropical eigenvalue, or in short, its *eigenvalue type*.

Definition 2.1. The *eigenvalue type* of a matrix $A \in \mathbb{R}^{n\times n}$ is the closed cone in the normal fan $\mathcal{N}_n$ of the polytope $\mathcal{C}_n$ that contains A.

Example 2.2. Let $n = 3$. There are eight cycles, two of length 3, three of length 2 and three of length 1, and hence eight eigenvalue types corresponding to the full-dimensional cones of the fan $\mathcal{N}_3$. The polytope $\mathcal{C}_3$ is six-dimensional: it is the threefold pyramid over the bipyramid formed by the 3-cycles and 2-cycles. □

Our goal in this chapter is to refine $\mathcal{N}_n$ into cones of linearity for the eigenspace map $A \mapsto Eig(A)$. We shall prove the existence of a partition $\mathcal{P}_n$ of $\mathbb{R}^{n\times n}$ into cones of linearity of this map, whose closed full-dimensional cones are indexed by *connected functions*.

Proposition 2.3. *There exists a partition $\mathcal{P}_n$ of matrix space $\mathbb{R}^{n\times n}$ into finitely many convex polyhedral cones such that each matrix in the interior of a full-dimensional cone has a unique eigenpair (λ, x) in $\mathbb{R} \times \mathbb{TP}^{n-1}$. Moreover, on each full-dimensional cone in that partition, the* eigenpair map *$A \mapsto (\lambda(A), x(A))$ is represented by a unique linear function $\mathbb{R}^{n\times n} \to \mathbb{R} \times \mathbb{TP}^{n-1}$.*

The eigenvalue is always unique, but by Theorem 1.11 the projective tropical eigenspace can be of dimension anywhere between 0 and $n-1$. The proposition implies that the set of matrices with more than one eigenvector lies in the finite union of subspaces of codimension one, and hence a generic $n \times n$ matrix has a unique eigenpair. Furthermore, cones of linearity of the eigenpair map refines the normal fan of the cycle polytope. Our main result is as follows.

Theorem 2.4. *All full-dimensional cones of the partition $\mathcal{P}_n$ are linearly isomorphic to $\mathbb{R}^n \times \mathbb{R}_{>0}^{n(n-1)}$. These cones are indexed by the connected functions $\phi : [n] \to [n]$, so their number is*

$$\sum_{k=1}^{n} \frac{n!}{(n-k)!} \cdot n^{n-k-1}. \tag{2.1}$$

For $n \geq 3$, these cones do not form a fan, that is, two cones may intersect in a non-face.

Here a function ϕ from $[n] = \{1, 2, \ldots, n\}$ to itself is called *connected* if its graph is connected as an undirected graph. The count in (2.1) is the sequence A001865 in [81]. In Section 2.2, we explain this combinatorial representation and we prove both Proposition 2.3 and Theorem 2.4. For $n = 3$, the number (2.1) equals 17, and our cone decomposition is represented by a 5-dimensional simplicial complex with f-vector $(9, 36, 81, 102, 66, 17)$. The locus in $\mathbb{R}^{n\times n}$ where the cone decomposition fails to be a fan consists precisely of the matrices A whose eigenspace is positive-dimensional. We explain the details in Section 2.3. In Section 2.4 we restrict to matrices A that are skew-symmetric, in symbols: $A = -A^T$. Tropical eigenvectors of skew-symmetric matrices arise in *pairwise comparison ranking*, in the approach that was pioneered by Elsner and van den Driessche [31, 32]. We shall discuss this connection in Chapter 4. Chapter 3 provides a complete combinatorial description of all cones in $\mathcal{P}_n$ and that of its natural fan refinement, the fan of linearity of the polytrope map.

Notation

We shall retain this notation for Chapter 3. The *graph of a matrix $A \in \mathbb{R}^{n\times n}$* is the weighted graph with edge weights A_{ij}. For a path $P_{ii'}$ from i to i', let $|P_{ii'}|$ denote the number of edges, $A(P_{ii'})$ denote the sum of all edge weights along the path in the graph of A. We say that two paths are disjoint if they do not share any edges. For disjoint paths $P_{ii'}, P_{i'j}$, we write $P_{ii'} + P_{i'j}$ for their concatenation.

2.2 The partition $\mathcal{P}_n$ and its full-dimensional cones

Recall that a *connected function* on $[n]$ is a function from $[n]$ to itself such that is graph is connected. We shall also use this term to refer to its graph. We shall need some equivalent characterizations of connected functions. These can be checked by direct computation.

Lemma 2.5. *The following are equivalent.*

1. *G is (the graph of) a connected function.*
2. *$G = T \cup e$, where e is an edge, and T is an in-directed spanning tree.*
3. *G has a unique cycle C, and for each node $\ell \in C$, the spanning tree T_ℓ rooted at ℓ is unique, and $G = T_\ell \cup C$.*

By Proposition 1.6 and Theorem 1.11, each eigenvector x is associated with a collection of trees $\mathcal{T}(x)$ of longest paths rooted at a critical node. Suppose A lies in the interior of a full-dimensional cone of $\mathcal{N}_n$, and hence has a unique critical cycle C. For any $T \in \mathcal{T}(x)$, one can use the cycle C to re-root at a different critical node. Since the root nodes lie in the same critical cycle, the re-rooted tree is a tree of longest path of x by Theorem 1.11. Therefore, by Lemma 2.5, each eigenvector can be associated with a collection of *connected functions*, each records a tree of longest path from any node to one of the root nodes.

Proof of Proposition 2.3. Since the eigenvector type necessarily refines the eigenvalue type, it is sufficient to construct the refinement of each eigenvalue type in the full-dimensional cones of $\mathcal{N}_n$. Let A lie in the interior of a cone with critical cycle C. Suppose A has an extreme eigenvector x associated with a connected function G. Fix any critical node ℓ, let T be the unique spanning tree of $[n]$ in G rooted at ℓ. Since the eigenvalue map is linear, for any path $P_{i\ell}$ the quantity $\bar{A}(P_{i\ell})$ is given by a unique linear form in the entries of A. A path $Q_{i\ell}$ is maximal if and only if $\bar{A}(Q_{i\ell}) - \bar{A}(P_{i\ell}) \geq 0$ for all paths $P_{i\ell} \neq Q_{i\ell}$. Thus the coordinate x_i of the eigenvector is given by a unique linear function in the entries of A (up to choices of ℓ) if and only if

$$\bar{A}(Q_{i\ell}) - \bar{A}(P_{i\ell}) > 0 \quad \text{for all paths } \; P_{i\ell} \neq Q_{i\ell}. \tag{2.2}$$

As previously noted, by Theorem 1.11, considering any other tree in G rooted at other nodes will yield the same eigenvector. Thus as linear functions in the entries of A, the linear forms in (2.2) are independent of the choice of ℓ. We have shown that A with critical cycle C has an extreme eigenvector x with tree $T \in \mathcal{T}(x)$, or equivalently, with connected function $G = T \cup C$, if and only if it satisfies a collection of linear inequalities on its entries, which define a full-dimensional polyhedral cone. Thus matrices in the interior of these cones are precisely those with a unique eigenvector x associated with a unique connected function with cycle C. Any other matrix necessarily lies at the intersections of these cones. The union of

the full-dimensional cones defined by (2.2) over all connected functions with cycle C defines a partition of the cone of $\mathcal{N}_n$ indexed by the cycle C. Thus their union over all cycles define the partition $\mathcal{P}_n$ of $\mathbb{R}^{n\times n}$, which refine $\mathcal{N}_n$ by construction. This proves Proposition 2.3. □

Two points should be noted in the proof of Proposition 2.3. Firstly, in the interior of each eigenpair cone (2.2), the unique eigenvector is associated with a unique connected function G on $[n]$. Secondly, the number of facet defining equations for these cones are potentially as large as the number of distinct paths from each $i \in [n]$ to a root node ℓ. In Theorem 2.4 we shall show that there are only $n^2 - n$ facets.

Proof of Theorem 2.4. We saw that each full-dimensional cone of $\mathcal{P}_n$ is associated with a connected function. We next argue that every connected function $\phi : [n] \to [n]$ is the critical graph of some generic matrix $A \in \mathbb{R}^{n\times n}$. If ϕ is surjective then ϕ is a cycle and we take any matrix A with the corresponding eigenvalue type. Otherwise, we may assume that n is not in the image of ϕ. By induction we can find an $(n-1) \times (n-1)$-matrix A' with critical graph $\phi\backslash\{(n, \phi(n))\}$. We enlarge A' to the desired $n\times n$-matrix A by setting $A_{n,\phi(n)} = 0$ and all other entries very negative. Then A has ϕ as its critical graph. We conclude that, for every connected function ϕ on $[n]$, the set of all $n \times n$-matrices that have the critical graph ϕ is a full-dimensional convex polyhedral cone Ω_ϕ in $\mathbb{R}^{n\times n}$, and these are the full-dimensional cones, characterized in (2.2), on which the eigenpair map is linear.

We next show that these cones are linearly isomorphic to $\mathbb{R}^n \times \mathbb{R}^{n(n-1)}_{\geq 0}$. Recall that the n-dimensional subspace W_n in (1.7) is the orthogonal complement to the affine span of the cycle polytope $\mathcal{C}_n$. The normal cone at each vertex of $\mathcal{C}_n$ is the sum of W_n and a pointed cone of dimension $n(n-1)$. We claim that the subcones Ω_ϕ inherit the same property. Let $\bar{\Omega}_\phi$ denote the image of Ω_ϕ in the quotient space $\mathbb{R}^{n\times n}/W_n$. This is an $n(n-1)$-dimensional pointed convex polyhedral cone, so it has at least $n(n-1)$ facets. As shown in the proof of Proposition 2.3, $\bar{\Omega}_\phi$ is defined by linear inequalities of the form (2.2), independent of the choice of the critical vertex ℓ. Fix some critical vertex ℓ. To show that it has precisely $n(n-1)$ facets, we claim that exactly $n(n-1)$ of such inequalities are active constraint, namely

$$\Omega_\phi \quad = \quad \big\{\, A \in \mathbb{R}^{n\times n} \,:\, \bar{A}_{ij} \,\leq\, \bar{A}(\phi_{i\ell}) - \bar{A}(\phi_{j\ell}) \,:\, (i,j) \in [n]^2\backslash\phi \,\big\}. \tag{2.3}$$

In this formula, $\phi_{i\ell}$ denotes the unique cycle-free directed path from i to ℓ in the graph of ϕ, with the convention $\phi_{\ell\ell} = \emptyset$, hence $\bar{A}(\phi_{\ell\ell}) = 0$. Existence of such a path follows from Lemma 2.5.

There are $n(n-1)$ distinct non-edges $(i,j) \in [n]^2\backslash\phi$, hence the inequalities in (2.3) are linearly independent. Furthermore, each of them clearly is of the form (2.2). Therefore, the cone defined by the right-hand side of (2.3) is linearly isomorphic to $\mathbb{R}^n \times \mathbb{R}^{n(n-1)}_{\geq 0}$, and it contains Ω_ϕ. To show that it is equal to Ω_ϕ, let A be in the cone of the right-hand side of (2.3). For each node $i \in [n]$, choose any path $P_{i\ell} \neq \phi_{i\ell}$. Without loss of generality, we can

assume that $i \to j$ is the first edge that differ between $P(i\ell)$ and $\phi_{i\ell}$. Clearly this edge is not in ϕ. Then

$$\bar{A}(P_{i\ell}) = \bar{A}_{ij} + \bar{A}(P_{j\ell}) \leq \bar{A}_{ij} + \bar{A}(\phi_{j\ell}) \leq \bar{A}(\phi_{i\ell}).$$

Thus (2.2) holds, that is, $A \in \Omega_\phi$. This proves (2.3), concluding the proof of the first sentence in Theorem 2.4.

For the second sentence we note that the number of connected functions in (2.1) is the sequence A001865 in [81]. Finally, it remains to be seen that our simplicial cones do not form a fan in $\mathbb{R}^{n^2}/W_n$ for $n \geq 3$. We shall demonstrate this explicitly in Example 2.8. □

2.3 Eigenpair cones and failure of the fan property

Let $(x_\phi, \lambda_\phi) : \mathbb{R}^{n\times n} \to \mathbb{TP}^{n-1} \times \mathbb{R}$ denote the unique linear map which takes any matrix A in the interior of the cone Ω_ϕ to its eigenpair $(x(A), \lambda(A))$. Of course, this linear map is defined on all of $\mathbb{R}^{n\times n}$, not just on Ω_ϕ. The following lemma is a useful characterization of Ω_ϕ in terms of the linear map (x_ϕ, λ_ϕ) which elucidates its product structure as $\mathbb{R}^n \times \mathbb{R}^{n(n-1)}_{\geq 0}$.

Lemma 2.6. *For a matrix $A \in \mathbb{R}^{n\times n}$, we abbreviate $x := x_\phi(A)$, $\lambda := \lambda_\phi(A)$, and we set $C = (c_{ij}) = (A_{ij} - x_i + x_j - \lambda)$. Then A is in the interior of the cone Ω_ϕ if and only if*

$$C_{i\phi(i)} \,=\, 0 \quad \textit{for all } i \in [n] \quad \textit{and} \quad C_{ij} < 0 \textit{ otherwise.} \tag{2.4}$$

Proof. Since the matrix $(x_i - x_j + \lambda)$ is in the linear subspace W_n, the matrices A and C lie in the same eigenpair cone Ω_ϕ. Since $C_{ij} \leq 0$ for all $i, j = 1, \ldots, n$, the conditions (2.4) are thus equivalent to

$$(C \odot [0, \ldots, 0]^T)_i \,=\, \max_{k\in[n]} C_{ik} \,=\, C_{i\phi(i)} \,=\, 0 \quad \text{for all } i \in [n].$$

In words, the matrix C is a normalized version of A which has eigenvalue $\lambda(C) = 0$ and eigenvector $x(C) = [0, \ldots, 0]^T$. The condition (2.4) is equivalent to that in (2.3), with strict inequalities for $\{(i, j) : j \neq \phi(i)\}$, and it holds if and only if C is in the interior of Ω_ϕ. □

The linear map $A \mapsto (C_{ij} : j \neq \phi(i))$ defined in Lemma 2.6 realizes the projection from the eigenpair cone Ω_ϕ onto its pointed version $\bar{\Omega}_\phi$. Thus, the simplicial cone $\bar{\Omega}_\phi$ is spanned by the images in $\mathbb{R}^{n\times n}/W_n$ of the incidence vectors $-\chi_{ij}$ that are indexed by the $n(n-1)$ non-edges:

$$\Omega_\phi \quad = \quad W_n \,+\, \mathbb{R}_{\geq 0}\big\{-\chi_{ij} \,:\, (i, j) \in [n]^2 \backslash \phi\big\} \qquad \simeq \quad \mathbb{R}^n \times \mathbb{R}^{n(n-1)}_{\geq 0}. \tag{2.5}$$

At this point, we find it instructive to work out the eigenpair cone Ω_ϕ explicitly for a small example, and to verify the equivalent representations (2.3) and (2.5) for that example.

Example 2.7 ($n = 3$)**.** Fix the connected function $\phi = \{12, 23, 31\}$. Its eigenvalue functional is $\lambda := \lambda_\phi(A) = \frac{1}{3}(A_{12} + A_{23} + A_{31})$. The eigenpair cone Ω_ϕ is 9-dimensional and is characterized by $3 \cdot 2 = 6$ linear inequalities, one for each of the six non-edges (i, j), as in (2.3). For instance, consider the non-edge $(i, j) = (1, 3)$. Using the identities $B(\phi_{13^*}) = b_{12} + b_{23} = A_{12} + A_{23} - 2\lambda$ and $B(\phi_{33^*}) = b_{31} + b_{12} + b_{23} = A_{31} + A_{12} + A_{23} - 3\lambda$, the inequality $b_{13} \leq B(\phi_{13^*}) - B(\phi_{33^*})$ in (2.3) translates into $A_{13} \leq 2\lambda - A_{31}$ and hence into

$$A_{13} \;\leq\; \frac{1}{3}(2A_{12} + 2A_{23} - A_{31}).$$

Similar computations for all six non-edges of ϕ give the following six linear inequalities:

$$A_{11} \leq \frac{1}{3}(A_{12} + A_{23} + A_{31}), \quad A_{22} \leq \tfrac{1}{3}(A_{12} + A_{23} + A_{31}), \quad A_{33} \leq \frac{1}{3}(A_{12} + A_{23} + A_{31}),$$
$$A_{13} \leq \frac{1}{3}(2A_{12} + 2A_{23} - A_{31}), \quad A_{32} \leq \tfrac{1}{3}(2A_{12} - A_{23} + 2A_{31}), \quad A_{21} \leq \frac{1}{3}(-A_{12} + 2A_{23} + 2A_{31}).$$

The eigenpair cone Ω_ϕ equals the set of solutions in $\mathbb{R}^{3\times 3}$ to this system of inequalities.

According to Lemma 2.6, these same inequalities can also derived from (2.3). We have

$$x := x_\phi(A) \;=\; \big[\, A_{12} + A_{23} - 2\lambda,\; A_{23} - \lambda,\; 0 \,\big]^T.$$

The equations $c_{12} = c_{23} = c_{31} = 0$ in (2.5) are equivalent to

$$A_{12} = x_1 - x_2 + \lambda, \quad A_{23} = x_2 - x_3 + \lambda, \quad A_{31} = x_3 - x_1 + \lambda,$$

and the constraints $c_{11}, c_{13}, c_{21}, c_{22}, c_{32}, c_{33} < 0$ translate into the six inequalities above. □

To describe the combinatorial structure of the eigenpair types, we introduce a simplicial complex Σ_n on the vertex set $[n]^2$. The facets (= maximal simplices) of Σ_n are the complements $[n]^2 \backslash \phi$ where ϕ runs over all connected functions on $[n]$. Thus Σ_n is pure of dimension $n^2 - n - 1$, and the number of its facets equals (2.1). To each simplex σ of Σ_n we associate the simplicial cone $\mathbb{R}_{\geq 0}\{\bar{e}_{ij} : (i, j) \in \sigma\}$ in $\mathbb{R}^{n\times n}/W_n$. We have shown that these cones form a decomposition of $\mathbb{R}^{n\times n}/W_n$ in the sense that every generic matrix lies in exactly one cone. The last assertion in Theorem 2.4 states that these cones do not form a fan. We shall now show this for $n = 3$ by giving a detailed combinatorial analysis of Σ_3.

Example 2.8. [$n = 3$] We here present the *proof of the third and final part of Theorem 2.4.* The simplicial complex Σ_3 is 5-dimensional, and it has 9 vertices, 36 edges, 81 triangles, etc. The f-vector of Σ_3 is $(9, 36, 81, 102, 66, 17)$. The 17 facets of Σ_3 are, by definition, the set complements of the 17 connected functions ϕ on $[3] = \{1, 2, 3\}$. For instance, the connected function $\phi = \{12, 23, 31\}$ in Example 2.6 corresponds to the facet $\{11, 13, 21, 22, 32, 33\}$ of Σ_3. This 5-simplex can be written as $\{11, 22, 33\} * \{21, 32, 13\}$, the join of two triangles, so it appears as the central triangle on the left in Figure 2.1.

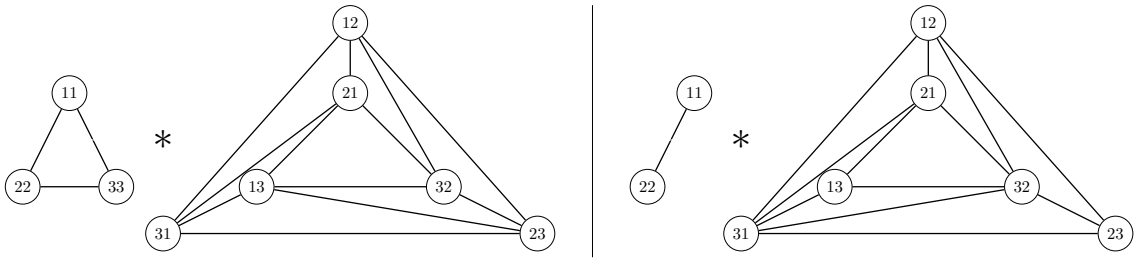

Figure 2.1: The simplicial complex Σ_3 of connected functions $\phi : [3] \to [3]$. Fixed-point free ϕ are on the left and functions with $\phi(3) = 3$ on the right.

Figure 2.1 is a pictorial representation of the simplicial complex Σ_3. Each of the drawn graphs represents its clique complex, and $*$ denotes the join of simplicial complexes. The eight connected functions ϕ whose cycle has length ≥ 2 correspond to the eight facets on the left in Figure 2.1. Here the triangle $\{11, 22, 33\}$ is joined with the depicted cyclic triangulation of the boundary of a triangular prism. The other nine facets of Σ_3 come in three groups of three, corresponding to whether $1, 2$ or 3 is fixed by ϕ. For instance, if $\phi(3) = 3$ then the facet $[3]^2 \backslash \phi$ is the join of the segment $\{11, 22\}$ with one of the three tetrahedra in the triangulation of the solid triangular prism on the right in Figure 2.1.

In the geometric realization given by the cones Ω_ϕ, the square faces of the triangular prism are flat. However, we see that both of their diagonals appear as edges in Σ_3. This proves that the cones covering these diagonals do not fit together to form a fan. □

Naturally, each simplicial complex Σ_n for $n > 3$ contains Σ_3 as a subcomplex, and this is compatible with the embedding of the cones. Hence the eigenpair types fail to form a fan for any $n \geq 3$. For the sake of concreteness, we note that the 11-dimensional simplicial complex Σ_4 has f-vector $(16, 120, 560, 1816, 4320, 7734, 10464, 10533, 7608, 3702, 1080, 142)$.

The failure of the fan property is caused by the existence of matrices that have disjoint critical cycles. Such a matrix lies in a lower-dimensional cone in the normal fan of $\mathcal{C}_n$, and it has two or more eigenvectors in $\mathbb{TP}^{n-1}$ that each arise from the unique eigenvectors on neighboring full-dimensional cones. These eigenvectors have distinct critical graphs ϕ and ϕ' and the cones Ω_ϕ and $\Omega_{\phi'}$ do not intersect along a common face. In other words, the failure of the fan property reflects the discontinuity in the eigenvector map $A \mapsto x(A)$.

For concrete example, consider the edge connecting 13 and 23 on the left in Figure 2.1 and the edge connecting 31 and 32 on the right in Figure 2.1. These edges intersect in their relative interiors, thus violating the fan property. In this intersection we find the matrix

$$A \;=\; \begin{pmatrix} 0 & 0 & -1 \\ 0 & 0 & -1 \\ -1 & -1 & 0 \end{pmatrix}, \tag{2.6}$$

whose eigenspace is a tropical segment in $\mathbb{TP}^2$. Any nearby generic matrix has a unique eigenvector, and that eigenvector lies near one of the two endpoints of the tropical segment. A diagram like Figure 2.1 characterizes the combinatorial structure of such discontinuities.

2.4 Skew-symmetric matrices

In this section we are interested in the combinatorial types of eigenpairs when restricted to the space $\wedge_2\mathbb{R}^n$ of skew-symmetric matrices. These appeared in the application of tropical eigenvectors to the statistical problem of inferring rankings from pairwise comparison matrices pioneered by Elsner and van den Driessche [31, 32]. Working on the additive scale, any pairwise comparison matrix $A = (A_{ij})$ is *skew-symmetric*, i.e. it satisfies $A_{ij} + A_{ji} = 0$ for all $1 \leq i, j \leq n$. The set $\wedge_2\mathbb{R}^n$ of all skew-symmetric matrices is a linear subspace of dimension $\binom{n}{2}$ in $\mathbb{R}^{n\times n}$. The input of the *tropical ranking algorithm* is a generic matrix $A \in \wedge_2\mathbb{R}^n$ and the output is the permutation of $[n] = \{1, \ldots, n\}$ given by sorting the entries of the eigenvector $x(A)$. See Chapter 4 for a comparison with other ranking methods.

For this section, we shall study the decomposition of the space $\wedge_2\mathbb{R}^n$ into the convex polyhedral cones $\Omega_\phi \cap \wedge_2\mathbb{R}^n$ where ϕ runs over connected functions on $[n]$. Note that, $\lambda(A) \geq 0$ for all $A \in \wedge_2\mathbb{R}^n$, and the equality $\lambda(A) = 0$ holds if and only if $A \in W_n$. Hence the intersection $\Omega_\phi \cap \wedge_2\mathbb{R}^n$ is trivial for all connected functions ϕ whose cycle has length ≤ 2. This motivates the following definition.

We define a *kite* to be a connected function ϕ on $[n]$ whose cycle has length ≥ 3. By restricting the sum in (2.1) accordingly, we see that the number of kites on $[n]$ equals

$$\sum_{k=3}^{n} \frac{n!}{(n-k)!} \cdot n^{n-k-1}. \tag{2.7}$$

Thus the number of kites for $n = 3, 4, 5, 6, 7, 8$ equals $2, 30, 444, 7320, 136590, 2873136$. The following result is the analogue to Theorem 2.4 for skew-symmetric matrices.

Theorem 2.9. *The full-dimensional cones in $\wedge_2\mathbb{R}^n$ on which the tropical eigenpair map for skew-symmetric matrices is represented by distinct and unique linear functions are $\Omega_\phi \cap \wedge_2\mathbb{R}^n$ where ϕ runs over all kites on $[n]$. Each cone has $n(n-3)$ facets, so it is not simplicial, but it is linearly isomorphic to $\mathbb{R}^{n-1}$ times the cone over the standard cube of dimension $n(n-3)/2 = \binom{n}{2} - n$. This collection of cones does not form a fan for $n \geq 6$.*

Proof. It follows from our results in Section 2.2 that each cone of linearity of the map

$$\wedge_2\mathbb{R}^n \to \mathbb{R} \times \mathbb{TP}^{n-1}\,,\ A \mapsto (\lambda(A), x(A))$$

has the form $\Omega_\phi \cap \wedge_2\mathbb{R}^n$ for some kite ϕ. Conversely, let ϕ be any kite on $[n]$ with cycle $(1 \to 2 \to \ldots \to k \to 1)$. We must show that $\Omega_\phi \cap \wedge_2\mathbb{R}^n$ has non-empty interior (inside $\wedge_2\mathbb{R}^n$). We

shall prove the statement by induction on $n-k$. Note that this would prove distinctiveness, for the matrices constructed in the induction step lie strictly in the interior of each cones. The base case $n-k=0$ is easy: here the skew-symmetric matrix $A=\sum_{i=1}^{n}(\chi_{i\phi(i)}-\chi_{\phi(i)i})$ lies in the interior of Ω_ϕ.

For the induction step, suppose that A lies in the interior of $\Omega_\phi \cap \wedge_2\mathbb{R}^n$, and fix an extension of ϕ to $[n+1]$ by setting $\phi(n+1)=1$. Our task is to construct a suitable matrix $A \in \wedge_2\mathbb{R}^{n+1}$ that extends the old matrix and realizes the new ϕ. To do this, we need to solve for the n unknown entries $A_{i,n+1}=-A_{n+1,i}$, for $i=1,2,\ldots,n$.

By (2.3), the necessary and sufficient conditions for A to satisfy $\phi(n+1)=1$ are

$$\begin{aligned} A_{(n+1)\,j} &\le \lambda(A)+\bar{A}(\phi_{(n+1)\,j^*})-\bar{A}(\phi_{jj^*}),\\ A_{j\,(n+1)} &\le \lambda(A)+\bar{A}(\phi_{jj^*})-\bar{A}(\phi_{1j^*})-\bar{A}(\phi_{(n+1),1}). \end{aligned}$$

Let $|\phi_{jj^*}|$ denote the number of edges in the path ϕ_{jj^*} Since $A_{ij}=-A_{ji}$, rearranging gives

$$\begin{aligned} A_{1\,(n+1)}+A_{(n+1)\,j} &\le A(\phi_{1j^*})-A(\phi_{jj^*})+(|\phi_{jj^*}|-|\phi_{1j^*}|)\lambda(A),\\ A_{1\,(n+1)}+A_{(n+1)\,j} &\ge A(\phi_{1j^*})-A(\phi_{jj^*})+(|\phi_{jj^*}|-|\phi_{1j^*}|)\lambda(A)-2\lambda(A). \end{aligned}$$

The quantities on the right hand side are constants that do not depend on the new matrix entries we seek to find. They specify a solvable system of upper and lower bounds for the quantities $A_{1\,(n+1)}+A_{(n+1)\,j}$ for $j=2,\ldots,n$. Fixing these $n-1$ sums arbitrarily in their required intervals yields $n-1$ linear equations. Working modulo the 1-dimensional subspace of W_{n+1} spanned by $\sum_{j=1}^{n}(\chi_{n+1,j}-\chi_{j,n+1})$, we add the extra equation $\sum_{j=1}^{n}A_{j\,(n+1)}=0$. From these n linear equations, the missing matrix entries $A_{1(n+1)},A_{2(n+1)},\ldots,A_{n(n+1)}$ can be computed uniquely. The resulting matrix $A\in\wedge_2\mathbb{R}^{n+1}$ strictly satisfies all the necessary inequalities, so it is in the interior of the required cone Ω_ϕ.

The quotient of $\wedge_2\mathbb{R}^n$ modulo its n-dimensional subspace $W_n\cap\wedge_2\mathbb{R}^n$ has dimension $n(n-3)/2$. The cones we are interested in, one for each kite ϕ, are all pointed in this quotient space. From the inductive construction above, we see that each cone $\Omega_\phi\cap\wedge_2\mathbb{R}^n$ is characterized by upper and lower bounds on linearly independent linear forms. This proves that this cone is linearly isomorphic to the cone over a standard cube of dimension $n(n-3)/2$. If $n=4$ then the cubes are squares, as shown in Figure 4.1.

Failure to form a fan stems from the existence of disjoint critical cycles, as discussed at the end of Section 3. For $n\ge 6$, we can fix two disjoint triangles and their adjacent cones in the normal fan of $\mathcal{C}_n$. By an analogous argument to that given in Example 2.7, we conclude that the cones $\Omega_\phi\cap\wedge_2\mathbb{R}^n$, as ϕ runs over kites, do not form a fan for $n\ge 6$. $\square$

2.5 Summary

In this chapter we examined the division of the space of all (skew-symmetric) $n \times n$-matrices into open polyhedral cones that represent distinct combinatorial types of tropical eigenpairs. Since that partition is not a fan, interesting phenomena happen for special matrices A, i.e. those not in any of the full-dimensional cones Ω_ϕ. For such matrices A, the eigenvalue λ is still unique but the tropical polytope $Eig(A)$ may have dimension ≥ 1.

Each tropical vertex of an eigenpolytope $\mathrm{Eig}(A)$ can be represented as the limit of eigenvectors $x(A_\epsilon)$ where (A_ϵ) is a sequence of generic matrices lying in the cone Ω_ϕ for some fixed connected function ϕ. This means that the combinatorial structure of the eigenpolytope $\mathrm{Eig}(A)$ is determined by the connected functions ϕ that are adjacent to A. For example, let us revisit the (inconsistently subdivided) square $\{13, 32, 23, 31\}$ in Figure 2.1. By Lemma 2.6, any 3×3-matrices that correspond to the points on that square can be taken to have the form

$$A \quad = \quad \begin{pmatrix} 0 & 0 & a \\ 0 & 0 & b \\ c & d & 0 \end{pmatrix}, \quad \text{where } a, b, c, d < 0.$$

One particular instance of this was featured in (2.6). The eigenpolytope $\mathrm{Eig}(A)$ of the above matrix is the tropical line segment spanned by the columns of A, and its two vertices are limits of the eigenvectors coming from the two adjacent facets of Σ_3.

In Chapter 3, we complete our study of $\mathcal{P}_n$ by providing a natural fan refinement $\mathcal{F}_n$, and combinatorial descriptions for all of their closed cones. The indexing set for closed cones of $\mathcal{P}_n$ is the collection of *compatible sets of connected relations*. As indicated in the above discussions, these generalize the notion of 'neighboring connected functions'.

Chapter 3

Polytropes and Tropical Eigenspaces: Cones of Linearity

This chapter is based on the submitted paper with the same title, available as an arXiv eprint [102]. We now address all the open questions raised in the previous chapter. Like in the previous chapter, we shall work with the max-plus algebra $(\mathbb{R}, \oplus, \odot)$ and utilize the same notation on graphs and paths. However, as in the archived version, we aim to provide a self-contained treatment.

3.1 Introduction

In this chapter we completely characterize conditions under which the piecewise linear maps $A \mapsto Eig(A)$ and $A \mapsto Pol(A)$ are given by linear functionals in the entries of A. Our results in this section parallel and extend Proposition 2.3 and Theorem 2.4. We first state the theorems. The notation will be fully defined later.

Theorem 3.1. *There exists a fan $\mathcal{F}_n$ such that in the relative interior of its cones, the polytrope map $A \mapsto Pol(A)$ is given by a unique set of linear functionals in the entries of A. Furthermore, $\mathcal{F}_n$ is the normal fan of an $n(n-1)$-dimensional polytope in $\mathbb{R}^{n \times n}$, whose face lattice is isomorphic to the lattice of complete sets of connected relations, denoted $\mathcal{CF}[n]$.*

Our proof provides an explicit formula for $\mathcal{F}_n$, see (3.1). Theorem 3.1 addresses all the unresolved issues in Theorem 2.4, namely that of a combinatorial indexing for the lower dimensional cones of $\mathcal{P}_n$, and the existence of a natural fan refinement which preserves information on the eigenspace.

Corollary 3.2. *The fan $\mathcal{F}_n$ in Theorem 3.1 refines the partition $\mathcal{P}_n$ appearing in Proposition 2.3. The coarsening to $\mathcal{P}_n$ of the lattice isomorphism in Theorem 3.1 produces an order-preserving bijection from the closed cones of $\mathcal{P}_n$ to the poset of compatible sets of connected relations, denoted $\mathcal{CP}[n]$.*

Theorem 3.1 converts the study of cones in $\mathcal{F}_n$ and $\mathcal{P}_n$ to the study of $\mathcal{CF}[n]$ and $\mathcal{CP}[n]$, which essentially are collections of graphs with constraints. From these sets of graphs, one can recover the defining equations and inequalities of the corresponding cones by applying the lattice isomorphism, which is a simple computation given in Algorithm 2. Cone intersections, for example, correspond to taking the join of two elements in $\mathcal{CF}[n]$ as described in Algorithm 1. Coarser information on a cone such as its codimension is given by an explicit formula (Proposition 3.24).

It follows from [29, Proposition 18] that all polytropes arise as $Pol(A)$ for some matrix A. This was also known to Sergeev, Schneider and Butkovic [92], and further studied in [91]. Full-dimensional polytropes, also known as tropical simplices, are of great importance to tropical polyhedral geometry, see [57] and references therein.

The face lattice of polytropes provides a natural combinatorial classification which coincides with Green's D-relation on the semigroup of square tropical matrices [50]. On the other hand, one can declare two polytropes to be equivalent if the corresponding matrices have finite entries and they lie in the relative interior of the same cone in $\mathcal{F}_n$. We term this the *graphical type* of a polytrope, and the former its *combinatorial type*. They are not directly comparable. We do not know if there exists an indexing set equivalent to $\mathcal{CF}[n]$ for the combinatorial type of polytropes. Such a set may shed light on the semigroup structure of tropical matrices, a topic intimately tied in with the 'tropical' (ie: Brazilian) origins of tropical geometry in the work of Imre Simon [95].

Organization

In Section 3.2 we explicitly construct $\mathcal{F}_n$ and prove that it is the desired polytopal fan. In Section 3.3 we define $\mathcal{CF}[n]$ and $\mathcal{CP}[n]$, the complete and compatible sets of connected relations, respectively. Section 3.4 proves the remaining statement of Theorem 3.1 and Corollary 3.2 by constructing the lattice anti-isomorphism Ψ from $\mathcal{CF}[n]$ to the face lattice of $\mathcal{F}_n$. We give an algorithm to compute Ψ in Section 3.5, prove a formula which gives the co-dimension of the cone from its graphical encoding, and discuss symmetries of the fan $\mathcal{F}_n$. We conclude with examples in Section 3.6.

Notation

Throughout this chapter, a *graph* is a directed graph, allowing self-loops but not multiple edges. By a *connected* graph we mean weakly connected. The *subgraph rooted at a node u* is the set of nodes and edges belonging to paths which flow into u. A collection of nodes S in a graph G is a *strong component* if the induced subgraph on S is strongly connected. A strong component S is a *sink component* if there are no directed edges from S to its complement. We use the term *multigraph* to mean a directed graph with multiple distinct arcs, that is, it consists of an ordered 4-tuple (V, E, s, t) where V is the set of nodes, E is the set of edges, $s : E \to V$ is a map assigning to each edge its source node, and $t : E \to V$ is a map assigning

to each edge its target node. The *contraction* $\tilde{G}$ of a connected graph G is a multigraph whose nodes are indexed by the sink components of G, and whose edge set, source and target maps are induced by the edges of G. An *in-directed tree* is a tree whose edges are oriented towards the root. For two fans $\mathcal{F}_1, \mathcal{F}_2$, let $\mathcal{F}_1 \wedge \mathcal{F}_2$ denote their common refinement.

3.2 Construction of $\mathcal{F}_n$

We now give an explicit construction of $\mathcal{F}_n$ and show that it is polytopal. We first need a small result in polyhedral geometry which gives one way of constructing a polytopal fan.

Definition 3.3. Let $\mathcal{F}$ be a fan refinement of a convex, full-dimensional, pointed polyhedral cone C in $\mathbb{R}^n$, $v \in \mathbb{R}^n$ a vector in its interior. A cone of $\mathcal{F}$ is an *outer cone* if it is a subset of the boundary of C. The *flattening of $\mathcal{F}$ along v* is the fan in $\mathbb{R}^{n-1}$ whose cones are the images of the outer cones of $\mathcal{F}$ projected onto the orthogonal complement of v.

Lemma 3.4. *Suppose $\mathcal{F}$ is the normal fan of some convex polyhedron P in $\mathbb{R}^n$ with a pointed, full-dimensional recession cone. Let v be a vector in the interior of $\mathcal{F}$. Then the flattening of $\mathcal{F}$ along v is a polytopal fan in $\mathbb{R}^{n-1}$.*

Proof. Implicitly the lemma claimed that $\mathcal{F}$ is a fan refinement of a full-dimensional, pointed polyhedral cone. To see this, let P be the polyhedron in the lemma. Write $P = Q + C$ where C is its recession cone, Q a polytope. By definition, the polar C° is the normal fan of C. Since C is full-dimensional, the facet normals of Q lie in the interior of C°. Thus the normal fan of P is a fan refinement of the polyhedral cone C°.

An $(n-1)$-dimensional unbounded face of P is a face of Q plus a face of C, hence its normal vector is an outer ray of $\mathcal{F}$. An $(n-1)$-dimensional bounded face of P is necessarily a face of Q, whose facet normal must lie in the interior of C°. Thus the outer rays of $\mathcal{F}$ are precisely the facet normals of the unbounded faces of P.

We now explicitly construct the polytope in question. Let $\mathcal{H}$ denote the orthogonal complement of v. Under the convention that the facet normals point outwards, the linear functional $x \mapsto \langle v, x\rangle$ on $\mathbb{R}^n$ is bounded above on P, and bounded below on Q. Thus there exists M sufficiently small such that $\langle v, x\rangle > M$ for all $x \in Q$. Let $\mathcal{H}_M = \mathcal{H} + M$. View $P_M = P \cap \mathcal{H}_M$ as a polytope in the $(n-1)$-dimensional affine subspace $\mathcal{H}_M$. We claim that P_M is the polytope needed. Indeed, by construction, the i-dimensional faces of P_M are the unbounded $(i+1)$-dimensional faces of P intersected with $\mathcal{H}_M$. Thus it is sufficient to show that the projected outer rays of $\mathcal{F}$ along v are the facet normals of the corresponding faces in P_M. Let x be a vector in an $(n-2)$-dimensional face of P_M, f the corresponding outer ray of $\mathcal{F}$, and Πf the orthogonal projection of f onto $\mathcal{H}$. Since $f - \Pi f$ is orthogonal to $\mathcal{H}$, $\langle x, f - \Pi f\rangle = 0$. Thus $\langle x, \Pi f\rangle = \langle x, f\rangle = 0$. Therefore, the rays Πf are the facet normals of $P \cap \mathcal{H}_M$. □

Recall that the polytrope map $A \mapsto \bar{A}^*$ is the composition of the map $A \mapsto \bar{A} = A \odot (-\lambda(A))$ with the Kleene star map $A \mapsto A^*$. Thus to study the polytrope map, we need to understand cones of linearity of the eigenvalue map $A \mapsto \lambda(A)$ and the Kleene star map. The former, as we now know, is exactly the normal fan $\mathcal{N}_n$ of $\mathcal{C}_n$.

Let us revisit the results in Chapter 1 on the construction of $\mathcal{N}_n$ in view of Lemma 3.4. Recall the notation Γ_n for the set of simple directed cycles on the complete graph, χ_C for the incidence vector of a cycle $C \in \Gamma_n$. Let $\mathcal{J}_n$ be the polyhedral cone of matrices in $\mathbb{R}^{n\times n}$ with no positive cycles

$$\mathcal{J}_n = \{A \in \mathbb{R}^{n\times n} : A \cdot \chi_C \leq 0, C \in \Gamma_n\}.$$

Its lineality space V_n has dimension $n-1$, consisting of matrices with zero-cycle sum.

$$V_n = \{A \in \mathbb{R}^{n\times n} : A \cdot \chi_C = 0, C \in \Gamma_n\},$$

which we have identified to be the space of gradient flows (1.8).

Proposition 3.5. *Let $\partial\mathcal{J}_n$ be the flattening of $\mathcal{J}_n$ along $-\mathbf{1}\mathbf{1}^T$. Define*

$$\mathcal{N}_n = \partial\mathcal{J}_n + \operatorname{span}\{\mathbf{1}\mathbf{1}^T\} + V_n,$$

where $+$ denotes direct sum of subspaces. Then $\mathcal{N}_n$ is a polytopal fan (the normal fan of a polytope) in $\mathbb{R}^{n\times n}$. Its cones are precisely the cones of linearity of the eigenvalue map.

Proof. The rays of the polar $\mathcal{J}_n^\circ$ are the incidence vectors $\{\chi_C : C \in \Gamma_n\}$. Since all faces of $\mathcal{J}_n$ are unbounded except the 0 vertex, we can choose any constant $M < 0$ for the flattening, say, $M = -1$. Now $\mathcal{J}_n^\circ \cap \mathcal{H}_{-1}$ is the normalized cycle polytope

$$\mathcal{C}_n = \operatorname{conv}\left\{\frac{1}{|C|}\chi_C : C \in \Gamma_n\right\}.$$

By Lemma 3.4, $\partial\mathcal{J}_n$ is the normal fan of $\mathcal{C}_n$ viewed as an (n^2-n)-dimensional polytope in the affine subspace $\mathcal{H}_{-1}$. One can check that $\mathbb{R}^{n\times n} = \mathcal{H}_{-1} + \operatorname{span}\{\mathbf{1}\mathbf{1}^T\} + V_n$, thus $\mathcal{N}_n$ is the normal fan of $\mathcal{C}_n$ in $\mathbb{R}^{n\times n}$. By [21], the eigenvalue function $A \mapsto \lambda(A)$ is the support function of $\mathcal{C}_n$ in $\mathbb{R}^{n\times n}$. Thus $\mathcal{N}_n$ is precisely the fan of linearity of this map. □

We now consider the Kleene star map $A \mapsto A^*$. Fix a column i^*. Recall the linear program of single-target longest path (1.9) to which the map $A \mapsto A^*_{\cdot i^*}$ is the solution. This linear program is bounded if and only if $A \in \mathcal{J}_n$. Its constraint set is a polyhedron in $\mathbb{R}^{n\times n}$ with lineality space V_n, whose vertices are in bijection with in-directed spanning trees on n nodes with root i^* [1]. Thus the cones of linearity of $A \mapsto A^*_{\cdot i^*}$ form a fan partition of $\mathcal{J}_n$.

Denote this fan $\mathcal{K}_{n,i^*}$. Take the common refinement over all columns i^*, we obtain the fan

$$\mathcal{K}_n = \bigwedge_{i^*=1}^{n} \mathcal{K}_{n,i^*}.$$

Corollary 3.6. *$\mathcal{K}_n$ is the fan of linearity of the Kleene star map on $\mathcal{J}_n$. Furthermore, it is the normal fan of a polyhedron in $\mathbb{R}^{n\times n}$.*

Proof. The first statement follows from Proposition 1.8. For each i^*, $\mathcal{K}_{n,i^*}$ is the normal fan of a polyhedron. Thus $\mathcal{K}_n$ is the normal fan of the Minkowski sum of the corresponding polyhedra, see [66, §2.3]. □

Proof of Theorem 3.1, part (a). Let $\partial\mathcal{K}_n$ be the flattening of $\mathcal{K}_n$ along $-\mathbf{1}\mathbf{1}^T$. Define

$$\mathcal{F}_n = \partial\mathcal{K}_n + \operatorname{span}\{\mathbf{1}\mathbf{1}^T\} + V_n. \tag{3.1}$$

By Proposition 3.5 and Corollary 3.6, $\mathcal{F}_n$ is the fan of linearity of the polytrope map $A \mapsto Pol(A)$. By Lemma 3.4, it is a polytopal fan in $\mathbb{R}^{n\times n}$. □

The fan $\mathcal{F}_n$ refines $\mathcal{N}_n$, and inherits its lineality space $\operatorname{span}\{\mathbf{1}\mathbf{1}^T\} + V_n$. We often identify $\mathcal{F}_n$ with $\partial\mathcal{K}_n$, $\mathcal{N}_n$ with $\partial\mathcal{J}_n$, as in Example 3.7 and Example 3.26.

Example 3.7 ($n = 2$)**.** Here $\mathcal{K}_2$ is the trivial fan, and $\mathcal{N}_2 = \mathcal{P}_2 = \mathcal{F}_2$. The fan $\mathcal{N}_2$ has three cones of codimension 0, three cones of codimension 1, and one cone of codimension 2. Express the matrix A as the vector (x, y, z, w), where $x = A_{11}, y = A_{22}, z = A_{12} + A_{21}, w = A_{12} - A_{21}$. Then $V_2 = \{(x, y, z, w) : x = y = z = 0\}$, $\mathcal{J}_2 = \{x, y, z \le 0, w = 0\}$ is the negative orthant, and $\partial\mathcal{J}_2$ is the image of the three quadrants $\{x = w = 0, y, z \le 0\}$, $\{y = w = 0, x, z \le 0\}$ and $\{z = w = 0, x, y \le 0\}$ projected onto the subspace defined by $\{w = 0, x + y + z = 0\}$. This projection results in a fan partition of $\mathbb{R}^2$ seen in Figure 3.1.

We shall revisit this example in Section 3.6 when discussing complete sets of connected relations. For the rest of this paper we study how the sets of graphs indexing cones in Figure 3.1 can be generalized to indexing sets for cones of $\mathcal{P}_n$ and $\mathcal{F}_n$ for $n \ge 3$.

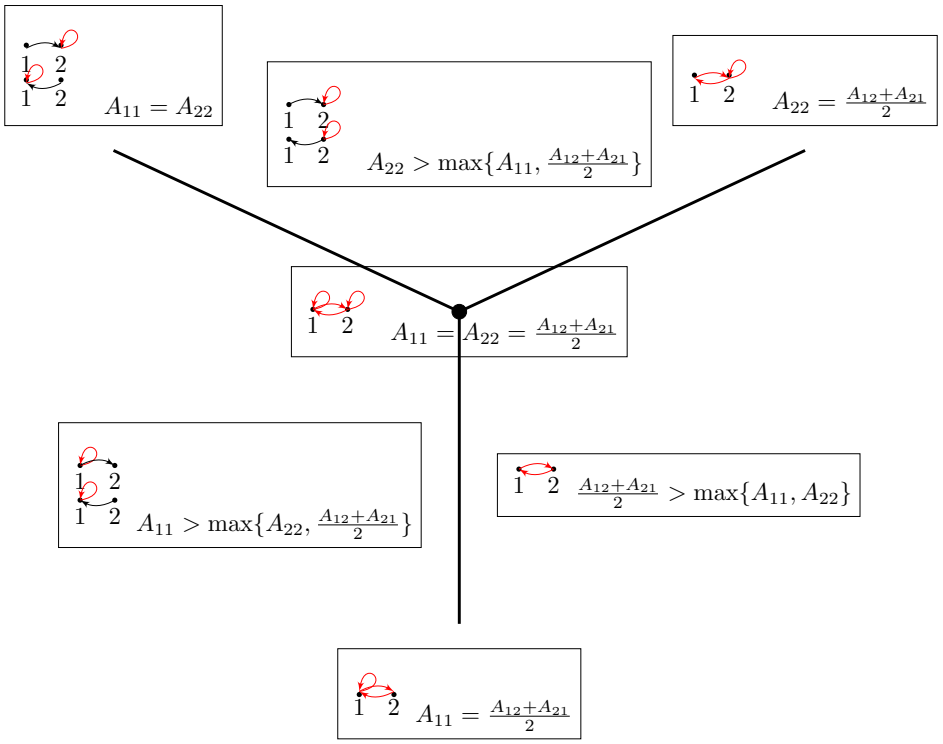

Figure 3.1: Cones in the partition $\mathcal{N}_2 = \mathcal{P}_2 = \mathcal{F}_2$ of $\mathbb{R}^{2\times 2}$, their defining equations and inequalities, and their associated maximal paths. Figure is drawn modulo lineality space, and thus represents the fan $\partial\mathcal{J}_2 = \partial\mathcal{K}_2$ of $\mathbb{R}^2$. Edges in the critical cycles (cycles with maximal mean) are shown in red.

3.3 Complete and compatible sets of connected relations

In this section we give intuition behind necessary properties for the indexing set of $\mathcal{P}_n$ and $\mathcal{F}_n$. We then define these sets, which are the two terms in the title, and conclude with a few examples.

For fixed i^*, the full-dimensional cones of $\mathcal{K}_{n,i^*}$ are in bijection with in-directed spanning trees on n nodes with root i^*, whose edges form the optimal paths from j to i^* for all nodes $j \in \{1, \ldots, n\}, j \neq i^*$ [1]. One can show that the full-dimensional cones of $\mathcal{K}_n$ are in bijection with collections of n in-directed spanning trees on n nodes, one rooted at each node, which satisfy the following *path compatibility* condition.

Lemma 3.8 ([1]). *If the path $i_1 \to i_2 \to \ldots \to i_k$ is a longest path from node i_1 to node i_k, then for every $q = 1, \ldots, k-1$, the subpath $i_q \to i_{q+1} \to \ldots \to i_k$ is a longest path from node i_q to node i_k.*

Cones of $\mathcal{J}_n$ with codimension one are in bijection with cycles on the complete graph. Their intersections with the full-dimensional cones of $\mathcal{K}_n$ form the full-dimensional cones of $\mathcal{F}_n$. Thus, the later are in bijection with *complete connected functions* (see Definition 3.10). Informally, a complete connected function consists of n trees unioned with a cycle, one rooted at each node, which satisfy the path compatibility condition.

Many properties of the full-dimensional cones of $\mathcal{N}_n$ also hold for those of $\mathcal{F}_n$. For example, it can be shown that neighboring full-dimensional cones of $\mathcal{F}_n$ differ by one edge. However, due to the path compatibility condition, it is less obvious what pairs of edges can be swapped to move to a neighboring cone. Similarly, one can compute cone intersection by taking the graph union between pairs of trees with the same root node. However, in doing so, new paths (such as a cycle) may appear, which by the path compatibility condition should also be present in other trees.

Therefore, one needs an efficient indexing set for all closed cones of $\mathcal{F}_n$, and a join operation that corresponds to cone intersections. With the hindsight of Theorem 3.1, we chose to start off by defining the desired indexing set, the complete sets of connected relations. We first introduce connected relations, the building block for our indexing set. In relation to [97], they are generalizations of connected functions.

Definition 3.9. A *circled tree* on $[n]$ is the graph union of a cycle and an in-directed spanning tree on n nodes, such that it has exactly one cycle. A *connected relation* on $[n]$ is a non-empty graph which can be written as the graph union of circled trees on $[n]$.

Let $\mathcal{R}[n]$ denote the set of all connected relations on $[n]$ adjoined with the empty set. It is a lattice ordered by subgraph inclusion, generated by its join-irreducibles, which are circled trees. Connected functions, for example, are strongly connected circled trees, or equivalently, circled trees in which the root node belongs to the cycle. Note that every connected relation has a unique strongly connected sink component. For $G \in \mathcal{R}[n]$, denote its unique sink component by $s(G)$ and the induced subgraph $S(G)$.

Definition 3.10. Let $\mathcal{G} = \{G_1, \ldots, G_k\}$ be a set of connected relations, $G_i \in \mathcal{R}[n]$. For node u, let $\mathcal{G}(u)$ denote the part $G_i \in \mathcal{G}$ containing u in its sink. The list $\mathcal{G}$ is a *complete set of connected relations* if:

(a) The sinks $s(G_j), j = 1, \ldots, k$ form a partition of $[n]$.

(b) For each pair $i, j \in \{1, \ldots, k\}$, let S_{ij} be the subgraph of the contraction $\widetilde{G}_j$ of G_j rooted at the node indexed by $s(G_i)$. Then S_{ij} equals the induced subgraph in $\widetilde{G}_i$.

(c) For every triple of distinct nodes $u, v, w \in \{1, \ldots, n\}$, if there exists paths P_{wv}, Q_{wu} in $\mathcal{G}(v)$ such that Q_{wu} is not a subpath of P_{wv}, then P_{wv} cannot appear in $\mathcal{G}(u)$.

Let $\mathcal{CF}[n]$ denote the collection of completed sets of connected relations.

Condition (b) is the analogue of the path compatibility condition in Lemma 3.8. Note that it implies that the contractions $\widetilde{G}_i$ have the same node set, and they are indexed by the set of sinks $s(G_j), j = 1, \ldots k$. The need for condition (c) is discussed after the proof of Lemma 3.19.

Definition 3.11. Write $\mathcal{G} = (\mathcal{D}, \mathcal{E})$ when we decompose $\mathcal{G}$ into the set $\mathcal{D}$ of connected relations whose sinks each contain at least one edge, and the set $\mathcal{E}$ of connected relations whose sinks each contain exactly one node and no edges. We say that $\mathcal{D}$ is a *compatible set of connected relations* and let $\mathcal{CP}[n]$ denote all such $\mathcal{D}$. If $\mathcal{D}$ consists of exactly one connected function, we say that $\mathcal{G}$ is a *complete connected function*.

Note that $\mathcal{E}$ can be an empty set but not $\mathcal{D}$, since at least one strongly connected component of each G_i necessarily contains a cycle. We shall use capital letters D and E instead of the calligraphic ones when the corresponding sets $\mathcal{D}$ and $\mathcal{E}$ are singletons.

Example 3.12. Figure 3.2 shows an incomplete set of connected relations $\mathcal{G} = (D, E)$ on 3 nodes. Here $s(D) = \{2\}$, $s(E) = \{3\}$. Since $\mathcal{G}$ is missing a part with the strong component $\{1\}$ as its sink, it is not complete.

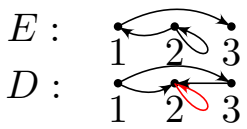

Figure 3.2: An incomplete set of connected functions $\mathcal{G} = (D, E)$.

Example 3.13. We can complete $\mathcal{G}$ in Example 3.12 by adding in a connected relation with sink $\{1\}$ that does not violate condition (b) of Definition 3.10. Figure 3.3 shows all such compatible connected relations, and Figure 3.4 shows some incompatible connected relations. In the first case, the subgraph rooted at 3 consists of the edge $2 \to 3$, which is not a subgraph of E. The second violates condition (a) since the sinks do not form a partition of $[n]$. The third violates condition (b) as D and E do not contain the self-loop at 3.



Figure 3.3: All three possible connected relations that can complete the set $\mathcal{G}$ of Figure 3.2.



Figure 3.4: Some connected relations incompatible to $\mathcal{G}$ of Figure 3.2.

3.4 The lattice isomorphism

In this section we prove the rest of Theorem 3.1, namely that the face lattice of $\mathcal{F}_n$ is anti-isomorphic to the lattice $\mathcal{CF}[n]$. The proof consists of three steps. First, we show that $\mathcal{CF}[n]$ is a join semilattice generated by its join irreducibles, which are the complete connected functions. Second, we produce a semilattice homomorphism Ψ from $\mathcal{CF}[n]$ to a collection of cones in $\mathbb{R}^{n\times n}$ ordered by subset inclusion. Third, we show that the full-dimensional cones of the fan $\mathcal{F}_n$ defined in (3.1) are precisely the images of the complete connected functions under Ψ. Since $\mathcal{F}_n$ is a fan, its face lattice ordered by subset inclusion is generated by meet-irreducibles, which are the full-dimensional cones of $\mathcal{F}_n$. Together with the first step, this implies that Ψ is a lattice anti-isomorphism between $\mathcal{CF}[n]$ and cones of $\mathcal{F}_n$.

The semilattice of $\mathcal{CF}[n]$.

We now show that $\mathcal{CF}[n]$ is a join semilattice generated by its join irreducibles, which are the complete connected functions. For $\mathcal{G} \in \mathcal{CF}[n]$, let $\lambda(\mathcal{G})$ denote the partition of $[n]$ induced by the sinks of its parts. For a pair $\mathcal{G}, \mathcal{H} \in \mathcal{CF}[n]$, write $\mathcal{H} \geq \mathcal{G}$ if for any pair of nodes u, v, the set of paths from u to v in $\mathcal{G}(v)$ is a subgraph of $\mathcal{H}(v)$. It is clear that this defines a partial ordering on $\mathcal{CF}[n]$.

We now define the join of two elements $\mathcal{G}, \mathcal{H} \in \mathcal{CF}[n]$ to be the output of Algorithm 1. This is an elaborated version of taking graph union and updating the strongly connected sink components at the same time. Broadly speaking, the algorithm iterates over three steps. First, we take the minimal amount of graph union required to satisfy condition (a) of Definition 3.10. In doing so, new cycles may form which may connect two sink components, requiring them to be merged. Thus in step two, we check for all such cycles. In step three, we find all triples violating condition (c) of Definition 3.10, and resolve them by taking the graph union of the corresponding connected relations. If 'actions' happened in steps two and three, we go back to step one. Otherwise, the algorithm terminates and outputs its current set of connected relations.

Note that after each iteration of line 21 the partition λ is unique and becomes coarser. Thus the algorithm is well-defined and terminates. By construction, the output satisfies Definition 3.10, and hence it is a complete set of connected relations.

Lemma 3.14. *$\mathcal{G} \vee \mathcal{H}$ is the join of $\mathcal{G}$ and $\mathcal{H}$ in $\mathcal{CF}[n]$.*

Algorithm 1 The join operation on $\mathcal{CF}[n]$.

1: Input: a pair $\mathcal{G}, \mathcal{H} \in \mathcal{CF}[n]$.
2: Output: a set of connected relations called $\mathcal{G} \vee \mathcal{H}$.
3: **procedure** JOIN($\mathcal{G}, \mathcal{H}$)
4: Set $B = \emptyset$.
5: Compute λ, the finest partition of $[n]$ that $\lambda(\mathcal{G}), \lambda(\mathcal{H})$ both refine.
6: **Step 1:**
7: **for** each component λ_i of λ **do**
8: Compute

$$I_i = \bigcup_{u \in \lambda_i} \mathcal{G}(u) \cup \mathcal{H}(u) \cup B \tag{3.2}$$

9: **end for**
10: **Step 2:** Set B' to be the union over all i of all cycles $C \subset I_i$ where $C \not\subset S(I_j)$ for any j.
11: **Step 3:**
12: **for** each ordered triple of distinct nodes $u, v, w \in [n]$ **do**
13: **for** each pairs of paths P_{wv}, Q_{wu} in $\mathcal{G}(v)$ that violates Definition 3.10(c) **do**
14: Add to B' the cycles formed by the union of all paths P_{uv} in $\mathcal{G}(v)$ and all paths Q_{vu} in $\mathcal{G}(u)$.
15: **end for**
16: **end for**
17: **if** $B' = \emptyset$ **then**
18: **return** $\mathcal{I} = (I_1, \ldots, I_{|\lambda|})$ ▷ The output is the list $\mathcal{I}$
19: **else**
20: Set $B = B \cup B'$
21: Compute λ, the finest partition of $[n]$ such that each component of $\lambda(I_j)$, $j = 1, \ldots, |\lambda|$ and each cycle C in B' is contained in exactly one component of λ.
22: Go to line 7.
23: **end if**
24: **end procedure**

Proof. From construction it is clear that $\mathcal{G} \vee \mathcal{H} \in \mathcal{CF}[n]$ and $\mathcal{I} \geq \mathcal{G}, \mathcal{H}$. To see that it is the smallest such element in $\mathcal{CF}[n]$, let $\mathcal{I}^{(1)}, \mathcal{I}^{(2)}, \ldots \mathcal{I}^{(k)} = \mathcal{I}$ be the sequence of intermediate candidate sets produced after the end of line 9 for each time this line is reached.

Let $\mathcal{I}' \in \mathcal{CF}[n]$ be such that $\mathcal{I}' \geq \mathcal{G}, \mathcal{H}$. Then each component of $\lambda(\mathcal{G}), \lambda(\mathcal{H})$ is necessarily contained in exactly one component of $\lambda(\mathcal{I}')$. Thus any path in $\mathcal{I}^{(1)}(v)$ is a subgraph of $\mathcal{I}'(v)$, and in particular, any cycle C in B' must be a subgraph of some part $\mathcal{I}'_i$. Since C is a cycle, by Definition 3.10, it is contained in exactly one component of $\lambda(\mathcal{J})$. Thus any path in $\mathcal{I}^{(2)}(v)$ must be a subgraph of $\mathcal{I}'(v)$. Iterating this argument up the sequence $\mathcal{I}^{(j)}$, $j = 1, \ldots, k$, we see that any path in $\mathcal{I}(v)$ must be a subgraph of $\mathcal{I}'(v)$, so $\mathcal{I}' \geq \mathcal{I}$. □

By the same argument, one can check that the operation $\vee$ is idempotent, commutative and associative. We make use of associativity in the proof of the following proposition.

Proposition 3.15. *$\mathcal{CF}[n]$ is a join semilattice generated by complete connected functions, which are precisely its sets of join-irreducibles.*

Proof. The join operation of Lemma 3.14 turns $\mathcal{CF}[n]$ into a join semilattice. It is clear that every complete connected function is join-irreducible. Thus it is sufficient to show that they generate $\mathcal{CF}[n]$.

For $\mathcal{G} \in \mathcal{CF}[n]$, let L be the collection of complete connected function $\mathcal{H}$ such that $\mathcal{H} \leq \mathcal{G}$, and define $\mathcal{I} := \bigvee_{\mathcal{H} \in L} \mathcal{H}$. By definition $\mathcal{I} \leq \mathcal{G}$. To show that $\mathcal{G} \leq \mathcal{I}$, for any pair of nodes $u, v \in [n]$, fix a path P_{uv} in $\mathcal{G}(v)$. We shall construct a complete connected function $\mathcal{H}$ such that $\mathcal{H}(r) \subseteq \mathcal{G}(r)$ for any node $r \in [n]$, and $P_{uv} \subseteq \mathcal{H}(v)$. This implies $\mathcal{H} \in L$, $P_{uv} \subseteq \mathcal{I}(v)$, and hence $\mathcal{G} \leq \mathcal{I}$ by definition of the partial order on $\mathcal{CF}[n]$. We consider two cases.

Case 1: $|\lambda(\mathcal{G})| = 1$. That is, $\mathcal{G}$ consists of only one part G whose graph is strongly connected. It is sufficient to construct the contraction $\mathcal{H}(v)$, thus we can assume without loss of generality that G contains a self-loop at v. Order the nodes so that $v = 1$. Let $\mathcal{H}$ be the output of the following algorithm.

1: Let $G_1, G_2, \ldots, G_n$ be copies of G, to be modified in line 5 below.
2: **for** $i = 1, \ldots, n$ **do**
3: Let H_i be an in-directed spanning tree rooted at i of G_i. If $i = v = 1$, require that H_1 contains P_{uv}.
4: **for** $j = i + 1, \ldots, n$ **do**
5: Replace the induced subgraph on G_j by the subtree of H_i rooted at j. ▷ Since ST is a spanning tree, the new G_j is still connected.
6: **end for**
7: **end for**
8: Set $\mathcal{H} = (H_1, \ldots, H_n)$.

By construction the parts of $\mathcal{H}$ are in-directed spanning trees on distinct root nodes, each containing exactly one cycle which is the self-loop at node $v = 1$. Thus Conditions (a) and

(c) of Definition 3.10 are satisfied. Condition (b) follows from line 5. Thus $\mathcal{H}$ is a complete connected function, and $\mathcal{H} \leq \mathcal{G}$, $P_{uv} \subseteq \mathcal{H}(v)$ by construction.

Case 2: $|\lambda(\mathcal{G})| > 1$. Again, we can assume without loss of generality that there exists a self-loop at v. We first apply the above algorithm to the contractions $\tilde{G}_1, \ldots, \tilde{G}_k$, producing k compatible in-directed spanning trees $T_1, \ldots, T_k$ on the strong components of the G_i's. Each edge in these trees fix a source and target node between two strong components. Now we apply the above algorithm to each strong component to split them into compatible spanning trees on singleton sinks. Finally, for each T_i, we replace each of its node (which currently is a strong component) by the corresponding spanning tree of that strong component, whose sink is specified by the edges of T_i. By construction, the collection of resulting n trees on the n singleton sinks is the desired complete connected function. □

From connected relations to equations and inequalities

We now define the map which takes an element of $\mathcal{R}[n]$ to a cone in $\mathbb{R}^{n \times n}$. We then extend its domain to complete sets of connected relations, resulting in a candidate for the lattice anti-isomorphism of Theorem 3.1.

Definition 3.16. Let $\psi : G \mapsto \psi(G)$ be the map from $\mathcal{R}[n]$ to cones in $\mathbb{R}^{n \times n}$, where $\psi(G)$ is the cone defined by the linear equations and inequalities output by Algorithm 2.

Algorithm 2 Definition of $\psi(G)$.

1: Input: $G \in \mathcal{R}[n]$
2: Output: a collection of non-redundant linear equations and inequalities defining a cone $\psi(G)$ in $\mathbb{R}^{n \times n}$
3: **procedure** ψ(G)
4: Pick an in-directed spanning tree T of G, let r be its root. Let T_u denote the path from u to the root r in the tree.
5: Pick some cycle C in G. Let $\lambda(A) = \frac{A(C)}{|C|}$, $\bar{A} = A \odot (-\lambda(A))$.
6: **for** each edge $u \to v$ not in T **do**
7: **if** edge $u \to v$ is in G **then**
8: **print** $\bar{A}_{uv} + \bar{A}(T_v) = \bar{A}(T_u)$. ($\bar{A}(T_r) = 0$ by convention)
9: **else print** $\bar{A}_{uv} + \bar{A}(T_v) \leq \bar{A}(T_u)$.
10: **end if**
11: **end for**
12: **end procedure**

Lemma 3.17. *Algorithm 2 is well-defined: its output is independent of the choice of C and the spanning tree T.*

Proof. For any cycle $C' \in G$, the algorithm outputs $\bar{A}(C') = 0$, and hence $A(C)/|C| = A(C')/|C'|$ for all pairs of cycles $C, C' \in G$. Thus the algorithm output is independent of the choice of the cycle C.

To show that the algorithm output is independent of T, it is sufficient to prove that for any path P_{ur} from u to r, we have $\bar{A}(P_{ur}) = \bar{A}(T_u)$. Indeed, suppose that this is true. Let T' be another in-directed spanning tree of G with root r'. Then r, r' are strongly connected, hence $\bar{A}(T_u) = \bar{A}(T'_u) + \bar{A}(T_{r'})$. Thus the linear inequalities and equalities outputted are the same, since the term $\bar{A}(T_{r'})$ cancels on both sides.

We now prove the above claim. Fix a path P_{ur}, and without loss of generality, suppose that $P_{ur} \neq T_u$, and that the first edge of the two paths differ, that is, the edge $u \to v$ on the path P_{ur} is not in T. Thus

$$\bar{A}(T_u) = \bar{A}_{uv} + \bar{A}(T_v) = \bar{A}(P_{ur}(u \to v)) + \bar{A}(T_v).$$

Therefore it is sufficient to prove that $\bar{A}(T_v) = \bar{A}(P_{ur}(v \to r))$ for the induced path from v to r in P_{ur}. Repeating this argument, noting that at each step the length of the remaining path is decreasing, we see that $\bar{A}(P_{ur}) = \bar{A}(T_u)$. □

Lemma 3.18. *Algorithm 2 outputs the minimal set of equations and inequalities defining the cone $\psi(G)$. This is a closed, non-trivial cone, and its codimension is between 0 and $n(n-1)$. The cone $\psi(G)$ has codimension 0 if and only if G is a connected function, in which case it is the cone Ω_G of $\mathcal{P}_n$ as defined in (2.3).*

Proof. The output of Algorithm 2 is minimal since the $n(n-1)$ equations and inequalities in the output are linearly independent. Indeed, each edge $u \to v$ in G gives rise to the unique equation or inequality in the output that contains the entry A_{uv}.

The n-dimensional lineality space of $\mathcal{N}_n$ is contained in $\psi(G)$ for any $G \in \mathcal{R}[n]$. Thus $\psi(G)$ is a closed, non-trivial cone, with codimension between 0 and $n(n-1)$. Algorithm 2 outputs no equality if and only if G is a circled tree with tree T and cycle C. In this case, the output of Algorithm 2 is precisely the $n(n-1)$ facet-defining inequalities of the cone Ω_G found in [97]. □

Unfortunately ψ does not completely carry over the lattice structure of $\mathcal{R}[n]$ to its image set of cones. The following lemma explains the necessity for condition (c) in Definition 3.10 of complete set of connected relations.

Lemma 3.19. *Suppose $G, H \in \mathcal{R}[n]$. Then $\psi(G \cup H) \subseteq \psi(G) \cap \psi(H)$. Furthermore, $\psi(G \cup H) = \psi(G) \cap \psi(H)$ if and only if either*

(a) The sink components are not *disjoint, or*

(b) There exists a triple of distinct nodes u, v, w, $u \in s(G), v \in s(H)$ and paths P_{wv}, Q_{wu} in both G and H such that $Q_{wu} \neq P_{wv}(w \to u)$.

Proof. The first statement follows from the definition of ψ. We prove the second statement by considering three disjoint cases: when condition (a) holds, when (a) does not hold but (b) holds, and when neither (a) nor (b) holds.

Suppose that (a) holds, that is, $s(G) \cap s(H) \neq \emptyset$. Any inequality in $\psi(G \cup H)$ is also an inequality in both $\psi(G)$ and $\psi(H)$. Thus for $A \in \psi(G) \cap \psi(H)$, we need to show that A satisfies the extra equations in $\psi(G \cup H)$. Let $r \in s(G) \cap s(H)$. For $u \in [n]$, consider the two paths from u to r: G_{ur} and H_{ur}, coming from G and H respectively. Since $A \in \psi(G)$, $\bar{A}(G_{ur}) \geq \bar{A}(H_{ur})$. Since $A \in \psi(H)$, $\bar{A}(G_{ur}) \leq \bar{A}(H_{ur})$. Thus $\bar{A}(G_{ur}) = \bar{A}(H_{ur})$ for all $u \in [n]$, so $A \in \psi(G \cup H)$.

Suppose $s(G) \cap s(H) = \emptyset$ and condition (b) holds. By induction on the number of nodes, we reduce to the case $n = 3$ with $s(G) = \{1\}, s(H) = \{2\}$, and and the triple $1, 2, 3$ witnesses condition (b). That is, G contains the edges $2 \to 1, 3 \to 1, 3 \to 2$, and H contains the edges $1 \to 2, 3 \to 1, 3 \to 2$. Let $A \in \psi(G) \cap \psi(H)$. By definition of ψ, $A_{11} = A_{22} = \lambda(A)$. Since $A \in \psi(G)$, $\bar{A}_{31} = \bar{A}_{32} + \bar{A}_{21}$. But $A \in \psi(H)$, hence $\bar{A}_{32} = \bar{A}_{31} + \bar{A}_{21}$. Thus $\bar{A}_{21} + \bar{A}_{12} = 0$, that is, the cycle $C = 1 \to 2 \to 1$ is also a critical cycle of A. Thus

$$\psi(G) \cap \psi(H) = \psi(G \cup C) \cap \psi(H \cup C) = \psi((G \cup C) \cup (H \cup C)) = \psi(G \cup H)$$

where the last equality follows from the first case.

Suppose neither conditions (a) nor (b) hold. By induction on the number of nodes, we can reduce to the case $n = 2$, where $s(G) = \{1\}, s(H) = \{2\}$. By enumerating all possibilities, we can check that regardless of the critical cycle, the matrix

$$A = \begin{bmatrix} 0 & -1 \\ -1 & 0 \end{bmatrix}$$

lies in $\psi(G) \cap \psi(H)$, but not $\psi(G \cup H)$, for example. □

We now show that one can extend ψ to a map Ψ from $\mathcal{CF}[n]$ to cones in $\mathbb{R}^{n \times n}$ satisfying $\Psi(\mathcal{G} \vee \mathcal{H}) = \Psi(\mathcal{G}) \cap \Psi(\mathcal{H})$. This will be our lattice anti-isomorphism between $\mathcal{CF}[n]$ and the fan $\mathcal{F}_n$.

Definition 3.20. Let $\mathcal{G} = (G_1, \ldots, G_k)$ where each $G_i \in \mathcal{R}[n]$. Define a cone in $\mathbb{R}^{n \times n}$ by

$$\Psi(\mathcal{G}) = \bigcap_{i=1}^{k} \psi(G_i).$$

Proposition 3.21. *For $\mathcal{G}, \mathcal{H} \in \mathcal{CF}[n]$,*

$$\Psi(\mathcal{G} \vee \mathcal{H}) = \Psi(\mathcal{G}) \cap \Psi(\mathcal{H}).$$

Proof. The '$\subseteq$' direction follows from Lemma 3.19. For the other direction, consider the sequence $(B'_1, B'_2, \ldots, B'_k = \emptyset)$ at line 17 of Algorithm 1. Let $(\lambda_1, \ldots, \lambda_k)$ be the correspond-

ing sequence of partitions used at the start of line 7. We shall prove the '$\supseteq$' direction by induction on k.

If $k = 1$, $B_1' = \emptyset$. Since the parts of $\mathcal{G}$ and $\mathcal{H}$ are connected, this cannot happen when the union of line 7 was taken over components with disjoint sinks. Thus $\lambda(\mathcal{G}) = \lambda(\mathcal{H}) = \lambda(\mathcal{G} \vee \mathcal{H}) = \lambda$. By Lemma 3.19,

$$\Psi(\mathcal{G} \vee \mathcal{H}) = \bigcap_{i=1}^{|\lambda|} \psi \left(\bigcup_{u \in \lambda_i} \mathcal{G}(u) \cup \mathcal{H}(u) \right) = \bigcap_{i=1}^{|\lambda|} \bigcap_{u \in \lambda_i} \psi(\mathcal{G}(u)) \cap \psi(\mathcal{H}(u)) = \Psi(\mathcal{G}) \cap \Psi(\mathcal{H}).$$

If $k = 2$, each cycle $C \in B_1'$ was generated by a union of parts in $\mathcal{G}$ and $\mathcal{H}$ that either have non-disjoint sinks, or have the triple-node situation. These are precisely the two cases covered by Lemma 3.19. Thus for any $A \in \Psi(\mathcal{G}) \cap \Psi(\mathcal{H})$, $\bar{A}(C) = 0$, and the cone defined by $\Psi(\mathcal{G}) \cap \Psi(\mathcal{H})$ is unchanged when we adjoin each cycle $C \in B_1'$ to each of the parts of $\mathcal{G}$ and $\mathcal{H}$. Call this new set $\mathcal{G}'$ and $\mathcal{H}'$, that is, define $\mathcal{G}'(u) = \mathcal{G}(u) \bigcup_{C \in B_1'} C, \mathcal{H}'(u) = \mathcal{H}(u) \bigcup_{C \in B_1'} C$. Now suppose that there are nodes $u, v \in C$ which lie in different components of λ_1. But $u, v \in s(\mathcal{G}'(u)) \cap s(\mathcal{H}'(u))$, thus

$$\Psi(\mathcal{G}(u)) \cap \Psi(\mathcal{H}(u)) = \Psi(\mathcal{G}'(u)) \cap \Psi(\mathcal{H}'(u)) = \Psi(\mathcal{G}'(u) \cup \mathcal{H}'(u)).$$

Therefore $\Psi(\mathcal{G}') \cap \Psi(\mathcal{H}') = \Psi(\mathcal{G}) \cap \Psi(\mathcal{H})$. Since $B_2' = \emptyset$, it follows that

$$\Psi(\mathcal{G} \vee \mathcal{H}) = \Psi(\mathcal{G}') \cap \Psi(\mathcal{H}').$$

Combining the last two equations proves the case $k = 2$. Since the algorithm terminates in at most n^2 steps, $k \leq n^2$, and induction completes the proof. □

Proposition 3.22. *Ψ defines a bijection from the join-irreducibles of $\mathcal{CF}[n]$ to the closed, full-dimensional cones of $\mathcal{F}_n$.*

Proof. By construction of $\mathcal{F}_n$, its closed, full-dimensional cone are in bijection with a collection of n trees of longest paths on the n roots, each union with a cycle C. Note that trees with roots in C coincide. Thus the set of such trees are precisely the complete connected function, which are the join-irreducibles of $\mathcal{CF}[n]$. Finally, Ψ maps a complete connected function to the corresponding cone in $\mathcal{F}_n$ by definition of ψ. □

Proof of Theorem 3.1 (lattice anti-isomorphism statement). Propositions 3.15 and 3.22 show that Ψ is a bijection from the join-irreducibles of the finitely generated semilattice $\mathcal{CF}[n]$ to the meet-irreducibles of the face lattice of $\mathcal{F}$. It follows from Proposition 3.21 that Ψ is a lattice anti-isomorphism. This concludes the proof of Theorem 3.1. □

Proof of Corollary 3.2. Let $\mathcal{G} = (G_1, \ldots, G_k)$ be a complete set of connected relations. Let $A \in \Psi(\mathcal{G})$. By Proposition 1.8, the sinks $s(G_i), i = 1, \ldots, k$ index distinct tropical vertices of $Pol(A)$. Write $\mathcal{G} = (\mathcal{D}, \mathcal{E})$, where $\mathcal{D}$ is the collection of complete set of connected

relations whose sink contain a critical cycle. Thus the corresponding columns of $Pol(A)$ are precisely the extreme tropical eigenvectors of $Eig(A)$ [5, 16]. Therefore, compatible set of connected relations index cones of linearity of the tropical eigenvector map, which are closed cones of $\mathcal{P}_n$. □

3.5 Codimension, facet-defining equations and symmetries

It follows from the proof of Theorem 3.1 that one can obtain the defining equations and inequalities of cones in $\mathcal{F}_n$ from their complete connected relations $\mathcal{G}$ by computing $\Psi(\mathcal{G})$. One method is to apply Algorithm 2 to obtain $\psi(G_i)$ for each part G_i of $\mathcal{G}$, then compute their intersections. Here we present a more self-contained algorithm and show how one can quickly compute the codimension of $\Psi(\mathcal{G})$. The proof of minimal representation similar to that of Algorithm 2 and hence omitted.

Algorithm 3 Minimal representation of $\Psi(\mathcal{G})$.

Input: $\mathcal{G} = (G_1, \ldots, G_k) \in \mathcal{CF}[n]$
Output: $\Psi(\mathcal{G})$
procedure MIN-$\Psi(G_1, \ldots, G_k)$
    for each $i = 1, \ldots, k$ **do**
        print $\psi(S(G_i))$ ▷ Compute $\psi(S(G_i))$ using Algorithm 2
        Compute a spanning tree T of G_i. Let r denote its root.
        for each edge $u \to v$ not in T, $u \in s(G_j), v \in s(G_i)$, $j \neq i$ **do**
            if $u \to v$ in $\widetilde{G}_j$ **then**
                print $\bar{A}_{uv} + \bar{A}(T(v \to r)) = \bar{A}(T(u \to r))$.
            else
                print $\bar{A}_{uv} + \bar{A}(T(v \to r)) \leq \bar{A}(T(u \to r))$.
            end if
        end for
    end for
end procedure

Lemma 3.23. *Algorithm 3 outputs a minimal set of equations and inequalities defining the cone $\Psi(G)$, and is independent of the choice of the spanning tree T.*

Proposition 3.24. *Let $\mathcal{G} = (G_1, \ldots, G_k) \in \mathcal{CF}[n]$. The codimension of the cone $\Psi(\mathcal{G})$ is*

$$codim(\Psi(\mathcal{G})) = \sum_{i=1}^{k} (e_i - v_i + E_i - V_i)$$

where v_i, e_i are the number of nodes and edges of the graph $S(G_i)$, V_i is the number of nodes in the contraction $\tilde{G}_i$ with an outgoing edge to the node indexed by $s(G_i)$, and E_i is the total out-degree of such nodes. In particular, the maximal codimension of $\mathcal{G}$ is $n(n-1)$, and this happens when G is the complete graph on n nodes with n self-loops.

Proof. The codimension of $\Psi(\mathcal{G})$ is the number of equalities returned by Algorithm 3. As in the proof of Proposition 3.15, we consider two cases.

Case 1: $k = 1$**.** Here $E = V = 1$ and $\Psi(\mathcal{G}) = \psi(G)$. The spanning tree T of G consists of $v - 1$ edges, each other edge in G contributes one equality in the output of Algorithm 2, except for the special edge $r \to v^*$. Thus $codim(\Psi(\mathcal{G})) = e - (v-1) - 1 = e - v$.

Case 2: $k > 1$**.** By case 1, the number of equalities produced from computing $\psi(S(G_i))$ is $e_i - v_i$. Thus it is sufficient to show that lines 6 to 13 of Algorithm 3 yield $\sum_{i=1}^{k}(E_i - V_i)$ equalities. Suppose $s(G_j)$ has an outgoing edge to $s(G_i)$ in $\tilde{G}_i$. If this is its unique outgoing edge, then it must be in the spanning tree T, yielding no equality. Otherwise, since $\tilde{G}_i$ is weakly connected, each outgoing edge from $s(G_j)$ yields a unique alternative path from $s(G_j)$ to $s(G_i)$. Hence the number of equalities each such $s(G_j)$ contributes is exactly its number of outgoing edges minus 1. Taking the sum over all such j, we obtain $\sum_{i=1}^{k}(E_i - V_i)$. □

Symmetries of the fan $\mathcal{F}_n$

The automorphism group of the complete directed graph on n vertices consists of vertex permutations and reversing edge orientations. Vertex permutations induce symmetries on the fan $\mathcal{F}_n$, mapping one full-dimensional cone to another. For edge reversals, note that $Pol(A) = Pol(A^T)$ up to vertex permutations [29, 57], thus edge reversals coincide with some vertex permutations. In terms of $\mathcal{CF}[n]$, edge reversals correspond to tracking all-pairs longest paths by the sources rather than the sinks. If we take $\mathcal{G} \in \mathcal{CF}[n]$, reverse all arrows, and then group the paths by their sink, we obtain another element $\mathcal{H} \in \mathcal{CF}[n]$, and $\mathcal{G} = \mathcal{H}$ up to vertex permutations. This is a non-obvious symmetry of the fan $\mathcal{F}_n$. See Example 3.26 below.

3.6 Examples

Example 3.25 ($n = 2$, continued)**.** In this case $\mathcal{N}_2 = \mathcal{P}_2 = \mathcal{F}_2$, and the cones indexed by complete sets of connected functions are shown in Figure 3.1. We shall refer to the cones by their critical cycles. Along the cone $C_{11,22}$ defined by $A_{11} = A_{22} > (A_{12} + A_{21})/2$, $dim(Eig(A)) = 1$, and the two extreme tropical eigenvectors has types shown in the upper-left box of Figure 3.1. Even in this small example we can see Lemma 3.19 in action: matrices in the relative interior of the cone $C_{11,22}$ are precisely the matrices in $\psi(C_{11}) \cap \psi(C_{22})$ but not $\psi(C_{11} \cup C_{22})$, and this is precisely due to the lack of uniqueness of the tropical eigenvector for matrices in this region.

Example 3.26 ($n = 3$)**.** The f-vector of $\mathcal{F}_3$ is $(68, 207, 267, 186, 72, 14, 1)$. As a sanity check, note that $\mathcal{F}_3$ has a lineality space of dimension 3. Identify it with a complete pointed fan in $\mathbb{R}^6$, we see that the Euler characteristic of the f-vector of $\mathcal{F}_3$ excluding the point of codimension 6 should be that of a 5-sphere, and it is indeed 0.

Figure 3.5 shows $N(k, \lambda, p, c)$, the number of cones in $\mathcal{F}_3$ of a given dimension k whose indexing complete set of connected relations $\mathcal{G}$ satisfy $\lambda(\mathcal{G}) = \lambda$, with p number of equalities coming from paths, and c number of equalities coming from cycles. For example, $N(2, (1, 1, 1), 1, 1) = 90$ is the number of cones in $\mathcal{F}_3$ which has critical cycles of length at most 1, codimension 2 and thus 2 defining equations, out of which 1 comes from a pair of distinct paths to a critical vertex, and 1 comes from the existence of two critical cycles. For $\lambda(\mathcal{G}) = (3)$, that is, $\mathcal{G}$ has a 3-cycle, we omit p and c since any equality can be regarded as a cycle equality.

k = 0			1		2		3		4		5		6
λ=(1,1,1)	(p,c) = 0,0	54	1,0 0,1	90 54	2,0 1,1 0,2	45 90 18	3,0 2,1 1,2	6 45 30	3,1 2,2	6 15	3,2	2	
(2,1)	0,0	12	1,0 0,1	12 36	2,0 1,1 0,2	3 36 36	2,1 1,2 0,3	9 36 12	2,2 1,3	9 12	2,3	3	
(3)		2		15		39		48		30		9	1
Total		68	207		267		186		72		14		1

Figure 3.5: Table of $N(k, \lambda, p, c)$, the number of cones in $\mathcal{F}_3$ of a given dimension k whose indexing complete set of connected relations G satisfy $\lambda(G) = \lambda$, with p the number of equalities coming from paths, and c the number of equalities coming from cycles.

The 54 cones with partition $\lambda = (1, 1, 1), p = 0, c = 0$ refine the three cones of $\mathcal{N}_3$ corresponding to the self-loops at $1, 2$ and 3. Ignoring the self-loops, we have 18 cones, with 5 equivalence classes up to permutations (and edge reversals). These equivalence classes and their sizes are shown in Figure 3.6 below.

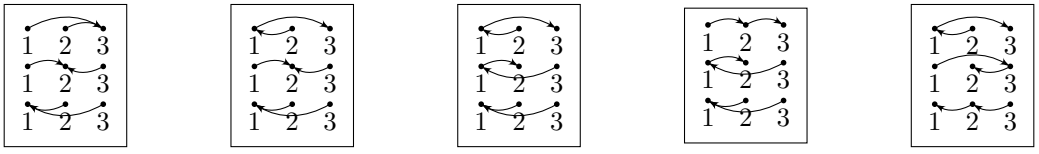

$\mathcal{G}_1, 1$ $\mathcal{G}_2, 6$ $\mathcal{G}_3, 3$ $\mathcal{G}_4, 6$ $\mathcal{G}_5, 2$

Figure 3.6: The five equivalence classes and their sizes, $\lambda = (1, 1, 1), p = c = 0$

Suppose we reverse all edges of $\mathcal{G}_3$. Then we obtain three out-directed trees. Re-organize the paths in these trees by their sinks, we obtain back $\mathcal{G}_3$. Thus edge reversal acts trivially on the orbit of $\mathcal{G}_3$. For $\mathcal{G}_4$, edge reversal coincides with the permutation $(1\,2)$.

Example 3.27 (A non-simplicial full-dimensional cone of $\mathcal{F}_3$)**.** Consider the full-dimensional cone in $\mathcal{F}_3$ indexed by the complete connected relation G shown in the middle box of Figure 3.7. This cone has 8 facets, thus it cannot be a simplicial cone. This example can be generalized to higher dimensions. Thus not all full-dimensional cones in $\mathcal{F}_n$ are simplicial for $n \geq 3$.

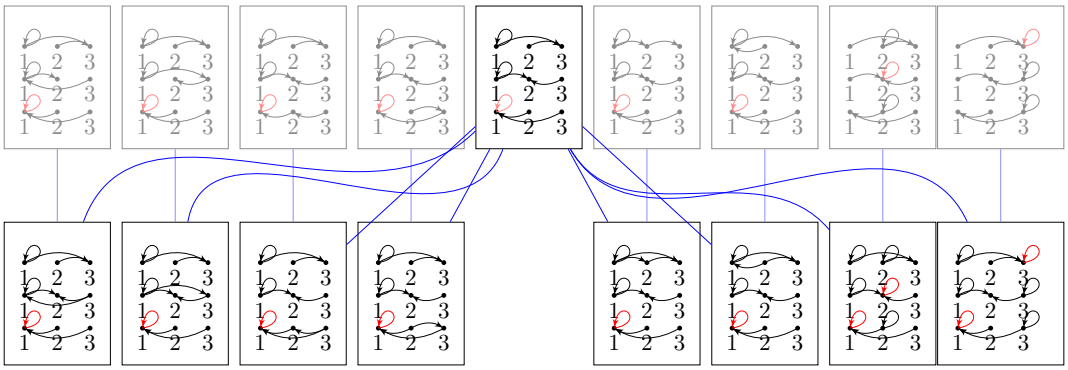

Figure 3.7: Facets of a non-simplicial full-dimensional cone in $\mathcal{F}_3$ labeled by $\mathcal{G}$, the complete connected relation shown in the middle box.

Example 3.28 (Face lattice of a full-dimensional cone)**.** We enumerate the face lattice of a full-dimensional cone in $\mathcal{F}_3$ indexed by the complete connected function $\mathcal{G}$ shown in the top solid box of Figure 3.8. The f-vector of this cone is $1, 6, 15, 20, 15, 6, 1$. It is a cone over a 5-dimensional simplex with a three-dimensional lineality space. Algorithm 3 gives the following set of defining equations and inequalities for $\Psi(\mathcal{G})$. For clarity we express them as formulas in $\bar{A}$ instead of A

$$\begin{aligned} &\bar{A}_{12} + \bar{A}_{21} \geq \bar{A}_{33}, && \bar{A}_{23} \geq \bar{A}_{21} + \bar{A}_{13}, && \\ &\bar{A}_{12} + \bar{A}_{21} \geq \bar{A}_{11}, && \bar{A}_{22} \leq \bar{A}_{12} + \bar{A}_{21}, && \bar{A}_{31} \geq \bar{A}_{32} + \bar{A}_{21}. \end{aligned}$$

The face lattice of $\Psi(\mathcal{G})$ indexed by complete connected relations is displayed as solid graphs in Figure 3.8. Blue arrows indicate subset inclusion. Red edges are those which belong to a cycle in the sink of the corresponding part. The six full-dimensional cones of $\mathcal{F}_3$ adjacent to $\Psi(\mathcal{G})$ are shown in lighter print next to $\mathcal{G}$.

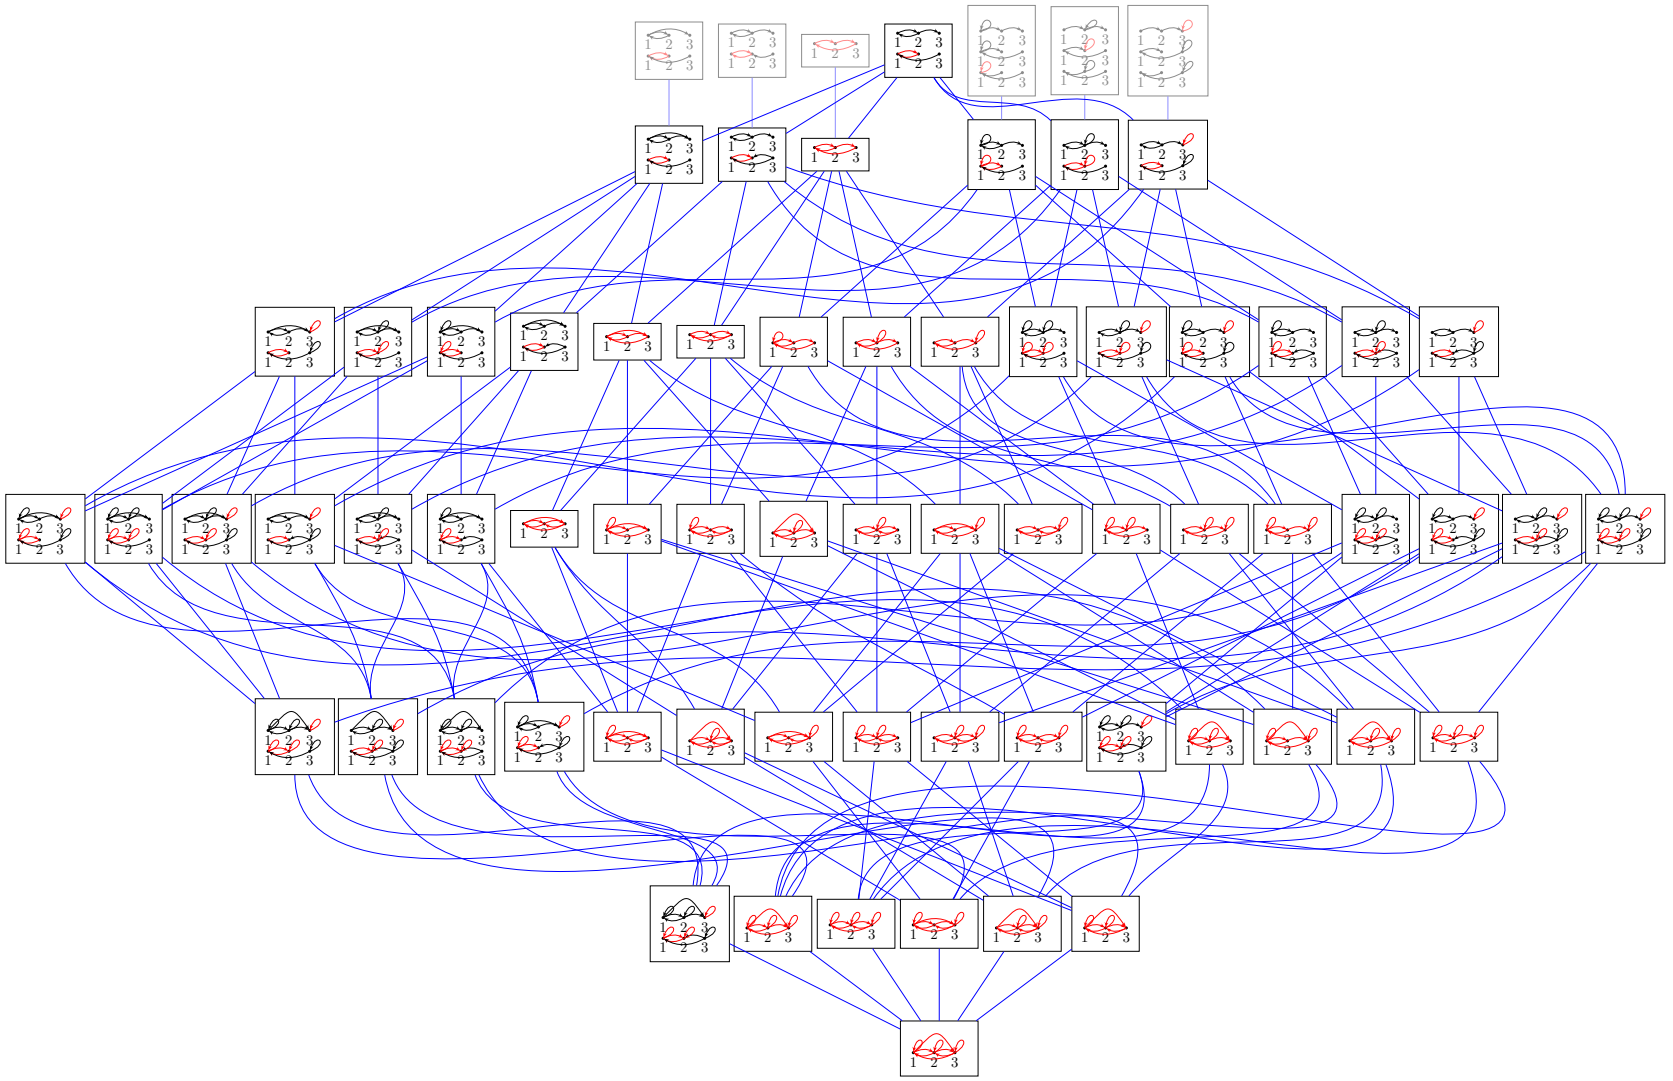

Figure 3.8: The face lattice of $\Psi(G)$.

Chapter 4

TropicalRank

The first five sections of this chapter made up the paper titled 'Pairwise ranking: choice of method can produce arbitrarily different orders', published in Linear Algebra and its Applications, 438(3), 1012-1024. The last two are based on a submitted paper titled 'HodgeRank is the limit of PerronRank', available as arXiv eprint [100]. Section 4.7 is a new result which will be incorporated in future versions of this paper. Compared to the published version, we have changed some notation and order of presentation to be consistent with the rest of the dissertation.

4.1 Introduction

In this chapter we examine three methods for ranking by pairwise comparison: PerronRank (Principal Eigenvector) [86], HodgeRank [55] and TropicalRank [31, 32]. The first two are most widely used in business applications and decision making. Both of these methods are closely related to TropicalRank, proposed by Elsner and Van Driessche. A number of papers have been devoted to their comparison [30, 35, 36, 85, 86]. We have two main contributions. On the one hand, we show that they mathematically belong to a parametrized family of ranking methods. On the other hand, we show that they can produce arbitrarily different rank order on the same input. To be precise, for any two of the three methods, and for any pair of rankings of at least four items, there exists a comparison matrix for the items such that the rankings found by the two methods are the prescribed ones. We discuss the implications of this result in practice, study the geometry of the methods, and state some open problems. We also examine the connections between these methods as the number of rank items tends to infinity.

Organization

In Section 4.3 we describe the three methods, their known properties, and their connections. Section 4.4 studies the geometry of TropicalRank. We compute an explicit formula for the

tropical eigenvector of generic skew-symmetric square matrices when $n = 4$. Section 4.5 contains the proof of Theorem 4.1. In Section 4.6, we prove Theorem 1.13 stated in Chapter 1, which states that the three methods in consideration belong to the same parametrized family of ranking models. We then use the proof techniques to study the i.i.d random ranking model in Section 4.7, and conclude with open problems in Section 4.8.

We recall some notation used in Section 1.6. Let $X = [X_{ij}]$ be an $n \times n$ matrix, $\mathcal{K} := \mathbb{R}_+^{n \times n}$ be the open cone of elementwise positive matrices in $\mathbb{R}^{n \times n}$ (we shall suppress the index n on $\mathcal{K}$). For $X, Y \in \mathcal{K}$, let $X^{(k)}$ denote the kth Hadamard power $X_{ij}^{(k)} = X_{ij}^k$, and $X \circ Y$ denote their Hadamard product $(X \circ Y)_{ij} = X_{ij} Y_{ij}$.

4.2 Pairwise ranking

Let n be the number of items to be ranked. Suppose one observes an $n \times n$ *multiplicative comparison* or *symmetrically reciprocal* matrix [22, 31, 35, 86], which is a matrix in the set

$$\mathcal{X}_n = \{X \in \mathcal{K} : X_{ij} = 1/X_{ji}\}.$$

Here X_{ij} measures the multiplicative preference of i over j. If $X_{ji} = 2$, for example, then we prefer j twice as much as i, which means we prefer i half as much as j, hence $X_{ij} = 1/X_{ji}$. The matrix X can come from direct measurement, or by taking elementwise exponential of an *additive comparison* matrix A in which A_{ij} measures the additive preference of i over j. An additive comparison matrix is skew-symmetric, that is, it lies in

$$ACM = \wedge_2 \mathbb{R}^n = \{A \in \mathbb{R}^{n \times n} : A_{ij} = -A_{ji}\}.$$

One has an analogous ranking interpretation: if i dominates j by $+2$ points, then j dominates i by -2 points (that is, j loses to i).

A ranking algorithm takes X (resp. A) as input and returns a score vector $s \in \mathbb{PR}^{n-1}$ (resp. $a \in \mathbb{TP}^{n-1}$) by which the items are ranked. Following Saari [85], we say that an additive comparison matrix A is *strongly transitive* if $A_{ik} = A_{ij} + A_{jk}$ for all $i, j, k \in \{1, \ldots, n\}$, A is *transitive* if $A_{ij}, A_{jk} > 0 \Rightarrow A_{ik} > 0$. Similar definitions hold for MCMs. That is, a multiplicative comparison matrix X is strongly transitive if $X_{ik} = X_{ij}/X_{jk}$ for all $i, j, k \in \{1, \ldots, n\}$, and X is transitive if $X_{ij}, X_{jk} > 1 \Rightarrow X_{ik} > 1$. It is trivial to retrieve an ordinal ranking from strongly transitive and transitive comparison matrices. Furthermore, it is immediate to check that strong transitivity implies the existence of a unique cardinal ranking up to constants [86]:

$$A \text{ is strongly transitive ACM} \Leftrightarrow A_{ij} = a_i - a_j$$
$$X \text{ is strongly transitive MCM} \Leftrightarrow X_{ij} = s_i/s_j.$$

In particular, the space of strongly transitive ACM is precisely the space of gradient flows V_n defined in (1.8). In applications one often does not observe strongly transitive nor transitive matrices. The comparison matrix A or X may come from averaging the opinion of multiple individuals, or may be corrupted by noise. Thus finding a score vector is equivalent to approximating the observed matrix by a strongly transitive matrix.

The triple $v(X), m(X), h(X)$ to be defined below are the outputs of the three ranking algorithms: PerronRank [86, 87], TropicalRank [31] and HodgeRank [20, 55]. PerronRank plays a fundamental role behind the Analytic Hierarchical Process [86], a ranking procedure extensively applied in decision making. Meanwhile, HodgeRank is closely related to many pairwise ranking algorithms in the broader literature of rank learning in computer science [39, 55]. The main result of this section states that these three methods are, in some sense, fundamentally different.

Theorem 4.1. *For $n \geq 4$, for any two of the three methods HodgeRank, PerronRank and TropicalRank, given any pairs of rank order (σ_1, σ_2), there exists a pairwise comparison matrix X such that the first method applied to X gives the ranking σ_1, and the second method applied to X gives the ranking σ_2 on the n items.*

This result is somewhat surprising and may remind the reader of Arrow's Impossibility Theorem in social choice theory, see [85] for further discussions on consistency between methods in pairwise ranking. In particular, it implies that these methods are fundamentally different, and the choice of methods in practice deserves further study.

4.3 Definitions and basic properties

Definition 4.2 (PerronRank)**.** For multiplicative input $X \in \mathcal{X}_n$, PerronRank outputs the (classical) principal eigenvector $v(X)$ of X, that is, the unique (up to scaling) positive vector $\mathbb{R}^n_+$ satisfying

$$X \cdot v(X) = \mu \cdot v(X) \quad \text{for some } \mu \in \mathbb{R}_+.$$

PerronRank is not defined for additive inputs.

As the existence of $v(X)$ only relies on the non-negativity of the matrix entries, this method has also found applications in other ranking contexts, such as in Google's PageRank algorithm. The specific application of principal eigenvector to MCMs first appeared in the work of Saaty [86]. This method has a Bayesian interpretation: given a prior score vector $s \in \mathbb{R}^n_+$, we can think of $X \cdot s$ as the posterior score, since

$$(X \cdot s)_i = \sum_j X_{ij} s_j = \text{sum over all } j \text{ of how much } i \text{ wins } j, \text{ weighted by our prior } s_j.$$

Then the principal eigenvector $v(X)$ is the Bayesian solution to the ranking problem, since up to scaling it is invariant under further updates by X.

Definition 4.3 (HodgeRank)**.** For multiplicative input $X \in \mathcal{X}_n$, HodgeRank outputs the row geometric mean $h(X)$ of X, defined as

$$h(X)_i = \left(\prod_{j=1}^{n} X_{ij} \right)^{1/n} \qquad \text{for } 1 \leq i \leq n.$$

For additive input $A \in \wedge_2 \mathbb{R}^n$, HodgeRank outputs the average row sum

$$h(A)_i = \frac{1}{n} \sum_{j=1}^{n} A_{ij} \text{ for } 1 \leq i \leq n,$$

The output of HodgeRank for multiplicative input is also known as the log least square estimator, or row geometric mean, first defined by Crawford [20]. Borda count [85] is a closely related method, which also outputs the row sum but the input comparison matrix have the form $A_{ij} = 1 - A_{ji}$ as opposed to the skew-symmetry condition. The name 'HodgeRank' was proposed by Jiang-Lim-Yao-Ye [55], who approached the problem of pairwise comparison from combinatorial Hodge theory. In the additive input setting, the entries A_{ij} are viewed as edge flows from node i to node j, and the unknown score vector measures the nodes' potentials. Entries of $h(A)$ can be interpreted as the average outflow from each node, which is an intuitive measure of their potentials. This method has a number of appealing properties both computationally and mathematically [55]. For example, it is the maximum likelihood estimator for the true score vector under certain statistical models of ranking [20].

Definition 4.4 (TropicalRank)**.** For multiplicative input $X \in \mathcal{X}_n$, TropicalRank outputs the tropical max-times eigenvector $m(X)$ of X if it is unique. For additive input $A \in \wedge_2 \mathbb{R}^n$, TropicalRank outputs the tropical max-plus eigenvector $m(A)$ of A, if it is unique. If $m(X)$ (resp. $m(A)$) is not unique, this algorithm does not produce a ranking.

TropicalRank was pioneered by Elsner and van den Driessche [31,32]. To make TropicalRank well-defined, we shall restrict our study to matrices with unique tropical eigenvectors for the rest of our paper. This is not a severe assumption, since any 'generic' matrix in $\mathcal{X}_n$ (resp. $\wedge_2 \mathbb{R}^n$) will have a unique tropical max-times eigenvector. That is, the set of matrices with non-unique max-times (resp. max-plus) tropical eigenvector has measure 0 under any non-atomic measure on $\mathcal{X}_n$ (resp. $\wedge_2 \mathbb{R}^n$). This follows from Proposition 2.4.

Basic properties and consistency of definitions

Lemma 4.5 (Properties of the three methods)**.** *Let $X \in \mathcal{X}_n$ be a multiplicative matrix with unique tropical max-times eigenvector $m(X)$.*

1. *If $P \in \{0,1\}^{n \times n}$ is a permutation matrix, then*

$$v(PXP^\top) = v \cdot h(X), \quad h(PXP^\top) = P \cdot h(X), \quad m(PXP^\top) = P \cdot m(X).$$

That is, the output of the three methods are invariant under item relabeling.

2. *For any $k > 0$, for all $1 \le i \le n$, $h(X^{(k)})_i = h(X)_i^k$, $m(X^{(k)})_i = m(X)_i^k$. In general this does not hold for $v(X)$, that is, $v(X^{(k)})_i \ne v(X)_i^k$.*

3. *If $X, X' \in \mathcal{X}_n$, then for $1 \le i \le n$,*

$$h(X \circ X')_i = h(X)_i h(X')_i.$$

If in addition X is strongly transitive, that is, $X_{ij} = s_i/s_j$ for some $s \in \mathbb{R}^n_+$, then for $1 \le i \le n$,

$$m(X \circ X')_i = m(X)_i m(X')_i, \quad v(X \circ X')_i = v(X)_i v(X')_i.$$

It follows from Lemma 4.5 that the definitions of HodgeRank and TropicalRank for additive and multiplicative inputs respect the natural log/exponential isomorphism between ACM and MCM. In particular, consider the family of isomorphisms $\wedge_2\mathbb{R}^n \to \mathcal{X}_n$, where

$$A \mapsto X := (X_{ij}) = (\exp_b(A_{ij})) = \exp(\log(b) \cdot A_{ij})$$

for some $b > 0$ referred to as the *exponent base*. For all possible choices of exponent bases b, a matrix $A \in \wedge_2\mathbb{R}^n$ is mapped to a family of matrices $\mathcal{E}(A) = \{X^{(k)} : X_{ij}^{(k)} = \exp(kA_{ij}), k > 0\}$. Since there is no *a priori* choices for b, one may prefer ranking algorithms which are invariant under choices of b. By Lemma 4.5 part (2), HodgeRank and TropicalRank satisfy this property. That is, for an additive comparison matrix A and any $k > 0$, if $X = (X_{ij}) = (\exp(k \cdot A_{ij}))$, then $h(A) = \log h(X)$ and $m(A) = \log m(X)$, where the last two equalities are up to additive constants.

However, by Lemma 4.5, $v(X^{(k)})_i \ne v(X)_i^k$. Thus one cannot extend PerronRank to handle additive comparison matrices in a consistent way. This does not imply that the method depends on the measurement unit. In fact, the entries of X are unit-free by definition, since they measure the *ratio* between two scores. This lack of invariance of PerronRank simply means that the method is not designed for additive pairwise comparison data. In general, a pairwise ranking method on the set of additive comparison matrices is invariant under all monotone transformations of the input entries A_{ij} if and only if it only depends on their signs, that is, one only need to consider $A_{ij} \in \{\pm 1, 0\}$. In this case, it only makes sense to ask for an ordinal and not cardinal ranking. Thus any cardinal ranking method lacks the invariance property under *some* monotone transformations of the data. One should only apply PerronRank to multiplicative inputs, while HodgeRank and TropicalRank can be applied to both multiplicative and additive inputs.

An optimization perspective

Recall the interpretation of ranking algorithms as finding the best approximation to the observed additive comparison matrix over V_n, the space of strongly transitive matrices. The following theorem states that HodgeRank and TropicalRank solve these problems where 'best' means minimizing ℓ_2 and ℓ_∞ distance, respectively.

Theorem 4.6 ([20, 55], [31])**.** *Let $A \in \wedge_2\mathbb{R}^n$ be an additive comparison matrix with unique tropical max-plus eigenvector. Let $H(A)$ denote the unique ℓ_2-minimizer from A to the subspace of strongly transitive matrices V_n, $H_\infty(A)$ denote the set of ℓ_∞-minimizers from A to V_n. Then the HodgeRank vector $h(A)$ is the score vector of the matrix $H(A)$, and $Pol(A)$, the polytrope of A, is precisely the set of score vectors of matrices in $H_\infty(A)$. In particular, this implies that the TropicalRank vector $m(A)$ is the score vector of some matrix in $H_\infty(A)$.*

It follows that translating A by a strongly transitive matrix does not change the relative difference between $m(A)$ and $h(A)$. From either a geometric proof using Theorem 4.6, or a direct computation, we obtain a restatement of Lemma 4.5 part (3), for additive comparison matrices.

Corollary 4.7. *Let $A \in \wedge_2\mathbb{R}^n$ be an additive comparison matrix with unique tropical max-plus eigenvector. Let $S \in V_n$ be a strongly transitive matrix with score vector s. Then $m(A - S) = m(A) - s, h(A - S) = h(A) - s$.*

4.4 Geometry of TropicalRank

It follows from Theorem 4.6 that for a skew-symmetric matrix $A \in \wedge_2\mathbb{R}^n$, $\lambda(A)$ has the geometric meaning of being the ℓ_∞ distance from A to V_n. We now study the geometric meaning of its tropical eigenvector $m(A)$ and HodgeRank vector $h(A)$, and connect it with our discussion of cones of linearity of the tropical eigenvector map in Section 2.4. In particular, we explicitly compute the partition $\mathcal{P}_4 \cap \wedge_2\mathbb{R}^4$ consisting of cones of linearity of the tropical eigenvector map defined in Section 2.4. Our computation verifies Theorem 2.9, and does not rely on results in Chapter 2. In other words, one can derive the tropical spectral theory for skew-symmetric matrices purely using Proposition 4.8 below.

Proposition 4.8. *Fix $A \in \wedge_2\mathbb{R}^n$, viewing it as a vector in $\mathbb{R}^{\binom{n}{2}}$. The projection of the standard cube centered at A onto V_n is the $(n-1)$-permutahedron, scaled by $1/n$ and centered at $H(A)$, the matrix with score vector $h(A)$. Its dual zonotope $Z_\perp$, which is the projection of the standard cube centered at A onto $V_n^\perp$, is a convex symmetric polytope centered at $A - H(A)$, whose facets are indexed by cycles on the complete graph of length at least 3. The distance in $V_n^\perp$ from the center of $Z_\perp$ to 0 is the tropical max-plus eigenvalue of A, $\lambda(A)$. Furthermore, the facet(s) of the scaled polytope $\lambda(A) \cdot Z_\perp$ which contain(s) 0 are precisely indexed by the critical cycle(s) of A.*

Proof. The first statement follows from a standard construction of the $(n-1)$-permutahedron as an affine projection of the cube in dimension $\binom{n}{2}$, see [110, Example 7.15]. The zonotope $Z_\perp$ is the cographic zonotope associated with the complete graph on n nodes (see [110]), thus its faces can be indexed by cycles of length at least 3. By definition of ℓ_∞-minimization, the distance in $V_n^\perp$ from the center of $Z_\perp$ to 0 is precisely the ℓ_∞ distance from A to V_n, which is its tropical max-plus eigenvalue $\lambda(A)$ by Theorem 4.6. The last statement follows from the aforementioned properties of $Z_\perp$ and the maximum cycle mean interpretation of tropical max-plus eigenvalue. □

It follows from the results of Chapter 3 that $Z_\perp$ is precisely the normal fan $\mathcal{N}_n$ of the cycle polytope $\mathcal{C}_n$ studied in Section 2.1, intersected with the subspace $\wedge_2\mathbb{R}^n$ and the condition $\lambda(A) = 1$. In particular, $Z_\perp$ has a subdivision induced by the subdivision of $\mathcal{P}_n$ of $\mathcal{N}_n$. We compute this subdivision explicitly for $n = 4$ in Example 4.11. In this case, the critical cycles has to have length 3 or 4, thus the tropical eigenvector is always unique, and $\mathcal{P}_4 \cap \wedge_2\mathbb{R}^4$ is a *fan partition* of $\mathbb{R}^{4\times 4}$. As found in Theorem 2.9, full-dimensional cones of this fan are linearly isomorphic to V_n times the standard cube of dimension $\binom{4}{2} - 2 = 2$. See Figure 4.1.

From a ranking perspective developed in [55, 85], the strongly transitive subspace V_n can be viewed as the 'signal' space, and its orthogonal complement $V_n^\perp$ is the 'noise level'. Thus it is beneficial to work with a basis of $\wedge_2\mathbb{R}^n$ which respects the orthogonal decomposition $\wedge_2\mathbb{R}^n = V_n \oplus V_n^\perp$. Saari [85] proposed a natural basis for V_n and identified a spanning set for $V_n^\perp$. For implications of this theorem to ranking applications, see [85]. Here we extend his work by showing that there is also a natural choice of basis for $V_n^\perp$. First we state his result in our notation. Identify each matrix $A \in \wedge_2\mathbb{R}^n$ with the vector of its upper diagonal entries $(A_{ij}, i > j) \in \wedge_2\mathbb{R}^n$. For a subgraph G of the complete directed graph on n nodes, write $\bar{\chi}_G \in \{\pm 1, 0\}^{\binom{n}{2}}$ for its *signed* incidence vector, where

$$\bar{\chi}_G(ij) = \begin{cases} 0 & \text{if} \quad \text{G does not contain the edge } i \to j \\ 1 & \text{if} \quad \text{G contains the edge } i \to j \text{ and } i < j \\ -1 & \text{if} \quad \text{G contains the edge } i \to j \text{ and } i > j \end{cases} .$$

Proposition 4.9 ([85]). *For $i \in [n]$, let G_i be the star graph at i on n vertices, with all edges directed towards i. That is, G_i contains all edges of the form $j \to i$ for all $j \in [n], j \neq i$. Then any subset of $n-1$ vectors from the set $\{\chi_{G_1}, \ldots, \chi_{G_n}\}$ defines a basis for V_n. Furthermore, the set $S = \{\bar{\chi}_C : C \in \Gamma_n, |C| \geq 3\}$ spans $V_n^\perp$.*

The cardinality of S is the number of k-cycles for $3 \leq k \leq n$. This is much greater than $\binom{n-1}{2}$, the dimension of $V_n^\perp$. Clearly any linearly independent subset of size $\binom{n-1}{2}$ would be a basis. We note that there is a natural choice: the set of 3-cycles involving a fixed vertex v.

Proposition 4.10. *Fix a vertex v. Define S_v to be the set of signed incidence vectors of all three cycles containing the node v, that is,*

$$S_v := \{\bar{\chi}_C : C \in \Gamma_n, |C| = 3, v \in C\}$$

Then S_v has $2 \times \binom{n-1}{2}$ vectors which occur in pairs of the form $(\bar{\chi}_C, -\bar{\chi}_C)$. Pick one representative for each pair. Then the resulting set is a basis of $V_n^\perp$.

Proof. The cardinality of S_v is $2 \times \binom{n-1}{2}$ since there are two choices for the other two nodes to complete an undirected 3-cycle. Note that any k-cycle C containing v can be graphically represented by a planar k-gon, which we shall call the *graph of C*. Since A as a matrix is skew-symmetric,

$$\bar{\chi}_C = \sum_{j=1}^{k-2} \bar{\chi}_{\tau_j}$$

where $\{\tau_j \in S_v : j = 1, \ldots, k-2\}$ is a triangulation of the graph of π obtained by putting in chords going through the node v. Thus S_v spans the set of all k-cycles containg v. On the other hand, the graph of C can also be triangulated by fixing another node v'. Since the choice of cycle is arbitrary, any other 3-cycle involving v' is also in the span of S_v. An inductive argument shows that S_v spans S, completing the proof. □

The output of TropicalRank purely depends on $Z_\perp$. This zonotope exhibits many symmetries, since relabeling of the n items correspond to rotations of $Z_\perp$. We now compute $Z_\perp$ for the case $n = 4$ using the basis of 4-cycles which highlights these symmetries in this special case. In particular, by studying the geometry of $Z_\perp$, we are able to give explicit formulas for the tropical max-plus eigenvector and tropical eigenvalue, without relying on results in Chapter 3.

Example 4.11 (n = 4)**.** Here $\dim(V_n) = \dim(V_n^\perp) = 3$. We shall equip $V_n^\perp$ with a basis consisting of 4-cycles. There are six 4-cycles in $\mathbb{S}_4$ and they occur in pairs of the form (σ, σ^{-1}). Choose a representative for each pair. The resulting set S_π consists of three independent vectors in $\mathbb{R}^6$, and they form a basis for $V_n^\perp$. For concreteness, we shall choose S_π to be the set $\{f_1, f_2, f_3\}$, with

$$\begin{aligned} f_1 &= \bar{\chi}_{(1\,2\,3\,4)} = (1, 0, -1, 1, 0, 1) \\ f_2 &= \bar{\chi}_{(1\,3\,4\,2)} = (-1, 1, 0, 0, -1, 1) \\ f_3 &= \bar{\chi}_{(1\,4\,2\,3)} = (0, -1, 1, 1, -1, 0). \end{aligned}$$

If we regard A as the $\mathbb{R}^6$ vector $(A_{12}, A_{13}, A_{14}, A_{23}, A_{24}, A_{34})$, then

$$\langle f_1, A\rangle = A_{12} - A_{14} + A_{23} + A_{34} = A_{12} + A_{23} + A_{34} + A_{41}$$

which is the value of the cycle (1 2 3 4). The computation for f_2 and f_3 can be checked similarly. Figure 4.1 shows the zonotope $Z_\perp$ expressed in this basis.

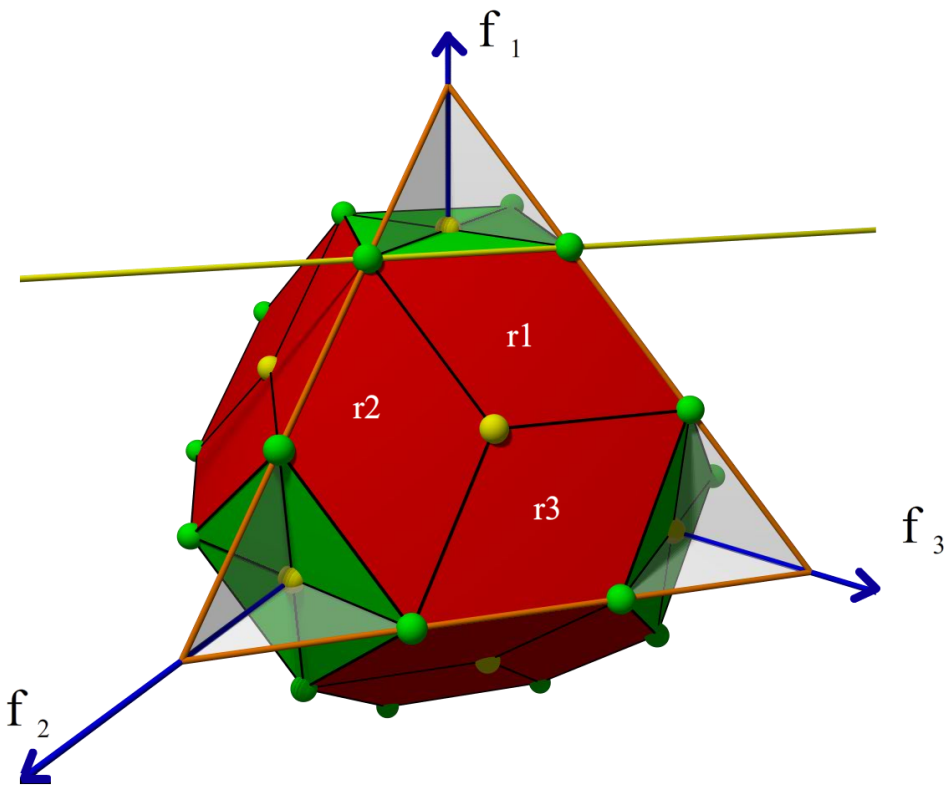

Figure 4.1: The zonotope $Z_\perp$ and its subdivision to cones of linearity of the tropical eigenvector map.

Note that $Z_\perp$ is also a permutahedron; this is due to the complete graph on 4 elements being self-dual, therefore the graphic and cographic zonotope of Proposition 4.8 are isomorphic [110]. $Z_\perp$ has eight hexagonal facets colored red, corresponding to the eight 3-cycles, and six square facets colored green, corresponding to the six 4-cycles. The value and critical cycle of the tropical max-plus eigenvalue λ is determined by entirely by the facet of $\lambda \cdot Z_\perp$ that contains 0. The hexagonal facets are further subdivided into regions corresponding to different formulas for the tropical eigenvector. Three such regions are labeled $r1, r2, r3$ in Figure 4.1. This subdivision is due to the fact that the tropical eigenvector is sensitive to permutation of the labeling of the items that fixes the critical cycle.

The subdivision on the square is an artifact of the graphics, since we subdivided $\mathbb{R}^3$ into the eight quadrants with respect to the f_1, f_2, f_3 coordinates. The square facets are in fact

not subdivided. They correspond to cycles of length four. The hexagonal facets correspond to critical cycles of length three. They are subdivided into three regions, corresponding to the three types of connected functions on four nodes with a cycle of length three. These connected functions are equivalent up to permutations. Overall, up to permutations, there are only two formulas for the tropical eigenvalue and tropical eigenvector, one for each facet type.

To obtain their explicit formulas from Figure 4.1, we only need to know the equations defining boundaries of the facets and their subdivisions. The equations

$$\langle f_1, A\rangle = 0, \langle f_2, A\rangle = 0, \langle f_3, A\rangle = 0$$

are three gray planes with orange boundary, defining the three edges shared by adjacent hexagons. They are the equations that artificially induced a subdivision of the square facets. The equations

$$2(\langle f_2, A\rangle + \langle f_3, A\rangle) = \langle f_1, A\rangle, 2(\langle f_1, A\rangle + \langle f_3, A\rangle) = \langle f_2, A\rangle, 2(\langle f_1, A\rangle + \langle f_2, A\rangle) = \langle f_3, A\rangle$$

define the other three edges of the hexagon. In Figure 4.1 we only show, in yellow, the edge cut out by the first equation. The black lines subdividing the frontal hexagon into three regions are

$$2\langle f_1, A\rangle = \langle f_2, A\rangle + \langle f_3, A\rangle, 2\langle f_2, A\rangle = \langle f_1, A\rangle + \langle f_3, A\rangle, 2\langle f_3, A\rangle = \langle f_1, A\rangle + \langle f_2, A\rangle.$$

One can compute and see that the frontal hexagon corresponds to the 3-cycle $(2\,3\,4)$, and the three green squares perpendicular to the axes f_1, f_2, f_3 correspond to the three 4-cycles defining these vectors. One can also read off the formula for $\lambda(A)$ in terms of the cycle values and the inequalities defining the regions, as demonstrated below.

Example 4.12 (Formulas for the tropical eigenvalue of skew-symmetric matrices for $n = 4$.)**.** Suppose $\langle f_1, A\rangle, \langle f_2, A\rangle, \langle f_3, A\rangle > 0$. This means 0 lies in the triangular region bounded by the orange lines. If

$$\begin{aligned} &\langle f_1, A\rangle \geq 2 \cdot (\langle f_2, A\rangle + \langle f_3, A\rangle) \quad \text{or} \\ &\langle f_2, A\rangle \geq 2 \cdot (\langle f_1, A\rangle + \langle f_3, A\rangle) \quad \text{or} \\ &\langle f_3, A\rangle \geq 2 \cdot (\langle f_1, A\rangle + \langle f_2, A\rangle), \end{aligned}$$

which implies that 0 lies in one of the three small triangles outside of the hexagon, then the 4-cycle corresponding to f_1 (or f_2, or f_3, respectively) is a critical cycle of A, and

$$\lambda_{max}(A) = \frac{1}{4}(\langle f_1, A\rangle) \text{ or } \frac{1}{4}(\langle f_2, A\rangle) \text{ or } \frac{1}{4}(\langle f_3, A\rangle), \text{ respectively.}$$

If

$$\begin{aligned}
\langle f_1, A\rangle &\leq 2\cdot(\langle f_2, A\rangle + \langle f_3, A\rangle) \quad \text{and}\\
\langle f_2, A\rangle &\leq 2\cdot(\langle f_1, A\rangle + \langle f_3, A\rangle) \quad \text{and}\\
\langle f_3, A\rangle &\leq 2\cdot(\langle f_1, A\rangle + \langle f_2, A\rangle),
\end{aligned}$$

which implies that 0 lies in the hexagon, then the 3-cycle $(2\,3\,4)$ is a critical cycle of A, and

$$\lambda_{max}(A) = \frac{1}{6}(\langle f_1, A\rangle + \langle f_2, A\rangle + \langle f_3, A\rangle).$$

Example 4.13 (Formulas for the tropical eigenvector of skew-symmetric matrices for $n = 4$)**.** The tropical eigenvector can also be read off from $Z_\perp$. Here we give explicit formulas for the red region marked $r1$ and the green region above it.

- The red region $r1$: this region is defined by $\langle f_1, A\rangle > \langle f_2, A\rangle, \langle f_3, A\rangle > 0$, $2(\langle f_2, A\rangle + \langle f_3, A\rangle) > \langle f_1, A\rangle$, $2f_2 < \langle f_1, A\rangle + \langle f_3, A\rangle$, and $2f_1 > \langle f_2, A\rangle + \langle f_3, A\rangle$. In this case,

$$m(A) = h(A) + \frac{1}{12}\begin{bmatrix} 0 & 0 & 0\\ -1 & 5 & 2\\ -2 & 7 & 1\\ -3 & 6 & 3 \end{bmatrix}\begin{bmatrix} \langle f_1, A\rangle\\ \langle f_2, A\rangle\\ \langle f_3, A\rangle \end{bmatrix}$$

- The green region above $r1$: this region is defined by $\langle f_1, A\rangle > \langle f_2, A\rangle, \langle f_3, A\rangle > 0$, and $2(\langle f_2, A\rangle + \langle f_3, A\rangle) < \langle f_1, A\rangle$. In this case,

$$m(A) = h(A) + \frac{1}{12}\begin{bmatrix} 0 & 0 & 0\\ 0 & 3 & 0\\ 0 & 3 & -3\\ 0 & 0 & -3 \end{bmatrix}\begin{bmatrix} \langle f_1, A\rangle\\ \langle f_2, A\rangle\\ \langle f_3, A\rangle \end{bmatrix}$$

Formulas for the tropical eigenvalue and eigenvector for other regions can be obtained by permuting the labels of the vertices, corresponding to rotating the zonotope. For example, if one acts on the labellings of the items by $\sigma = (4\,2\,3)$, then

$$\begin{aligned}
\sigma(\langle f_1, A\rangle) &= \sigma(A_{12} + A_{23} + A_{34} + A_{41}) = A_{13} + A_{34} + A_{42} + A_{21} = \langle f_2, A\rangle\\
\sigma(\langle f_2, A\rangle) &= \sigma(A_{13} + A_{34} + A_{42} + A_{21}) = A_{14} + A_{42} + A_{23} + A_{31} = \langle f_3, A\rangle\\
\sigma(\langle f_3, A\rangle) &= \sigma(A_{14} + A_{42} + A_{23} + A_{31}) = A_{12} + A_{23} + A_{34} + A_{41} = \langle f_1, A\rangle.
\end{aligned}$$

Therefore, the action of σ on the label of items corresponds to a counter-clockwise 120° rotation of $Z_\perp$ around the center of the hexagonal facet in the picture. The red region $r1$ is

mapped to region $r2$, and the formula for the max-plus eigenvector in this region is

$$m(A) = h(A) + \frac{1}{12}\begin{bmatrix} 0 & 0 & 0 \\ 3 & -3 & 6 \\ 2 & -1 & 5 \\ 1 & -2 & 7 \end{bmatrix}\begin{bmatrix} \langle f_1, A\rangle \\ \langle f_2, A\rangle \\ \langle f_3, A\rangle \end{bmatrix}$$

4.5 Proof of the pairwise inconsistency theorem

We now prove Theorem 4.1. For a generic vector $w \in \mathbb{R}^n$, let $r(w)$ denote its induced ranking. Say that A *induces a ranking triple* $(\sigma_1, \sigma_2, \sigma_3)$ for $(v, h, m) \in \mathbb{R}^3$ if $r(v) = \sigma_1, r(h) = \sigma_2, r(m) = \sigma_3$. Say that a ranking triple $\sigma = (\sigma_1, \sigma_2, \sigma_3)$ is *realizable for* (v, h, m) if there exists a matrix A which induces that ranking triple. Definitions for ranking pairs follow similarly.

Lemma 4.14. *To prove Theorem 4.1 for the pair HodgeRank and TropicalRank, it is sufficient to find one matrix $A \in \wedge_2\mathbb{R}^n$ such that $h(A) = (0, \ldots, 0)$, and $m(A)$ induces a ranking without ties. To prove Theorem 4.1 for the pair TropicalRank and PerronRank, it is sufficient to find one matrix $X \in \mathcal{X}_n$ such that $m(X) = (1, \ldots, 1)$ and $v(X)$ is a ranking without ties.*

Proof. For the first case, suppose A is such a matrix with unique tropical max-plus eigenvector $m(A)$. Fix a ranking σ. It is sufficient to find a matrix B such that $r(m(B)) = r(m(A))$ and $r(h(B)) = \sigma$. Let $W \in V_n$ be a strongly transitive matrix with $r(h(W)) = \sigma$. Note that $h(W) = m(W)$. Define $B(\epsilon) = A + \epsilon \cdot W$ for some $\epsilon > 0$. By Lemma 4.5, $m(B(\epsilon)) = m(A) + \epsilon \cdot m(W)$, and this function is continuous in ϵ. Thus for sufficiently small ϵ, $r(m(B(\epsilon))) = r(m(A))$, while $r(h(B(\epsilon))) = r(h(W)) = \sigma$.

For the second case, suppose X is such a matrix with unique tropical max-plus eigenvector $m(X)$. Let $M \in \mathcal{X}_n$ be a strongly transitive matrix with $r(v(M)) = \sigma$. Note that $v(M) = m(M)$. For $k > 0$, define $Y(k) = X \circ M^{(k)}$. By Lemma 4.5, $m(Y(k)) = m(X) \circ (m(M)^k)$, so $r(m(Y(k))) = r(m(M)) = \sigma$. It is sufficient to find k small enough such that $r(v(Y(k))) = r(v(X))$. Since λ has multiplicity 1, any column of the adjugate matrix $\mathrm{adj}(X - \lambda\mathbf{I})$ is a principal eigenvector, where $\mathbf{I}$ is the identity matrix. Since each entry of this matrix is a cofactor of an entry in $X - \lambda\mathbf{I}$, each v_i is a polynomial in λ and entries of X, and thus is continuous in the entries of X (see, for example, [107]). Let $\epsilon = \min_{i,j} |v_i(X) - v_j(X)|$. Since there are only finitely many v_i's, one can find a $\delta > 0$ such that the ordering of the v_i's are preserved when each upper-diagonal entry of X is perturbed by ϵ^δ. Hence $k = \delta / \log(\|M\|_\infty)$ would do. □

Proof for HodgeRank and TropicalRank. We shall construct the matrix A in Lemma 4.14 as follows: start with a matrix $A' \in \wedge_2\mathbb{R}^n$ such that

- (1) $h(A') = 0$.
- (2) $A'_{12} > A'_{23} > \ldots > A'_{(n-1)\,n} > A'_{n1} > 0$.
- (3) For $\lambda := \dfrac{A'_{12} + A'_{23} + \ldots + A'_{(n-1)\,n} + A'_{n\,1}}{n}$, we have $A'_{i\,(i+1)} \neq \mu$ for all i.

Such a matrix A' exists since $\dim(\wedge_2 \mathbb{R}^n) > n$ for all $n > 3$. Let $A := A' + k \cdot B$, where $B \in \wedge_2 \mathbb{R}^n$ is the matrix with upper diagonal entries

$$B_{i,i+1} = 1 \text{ for all } i \leq n-1, \;\; B_{1,n} = -1, \;\; B_{ij} = 0 \text{ else.}$$

Note that A satisfies properties (1) to (3), and for sufficiently large k, the following also hold.

- (4) For each row i, $\max_j A_{ij} = A_{i\,(i+1)}$.
- (5) The n-cycle $1 \to 2 \to \ldots \to n \to 1$ is a critical cycle of A with value μ.

By Theorem 1.12, property (5) implies that A has a unique tropical max-plus eigenvector, denoted m. Property (4) implies $m_i - m_{i+1} = A_{i\,(i+1)} - \lambda$ for all i. By (2) and (3), there exists $j \in \{1, \ldots, n\}$ such that $A_{i\,(i+1)} > \lambda$ for all $i < j$, and $A_{i\,(i+1)} < \lambda$ for all $i \geq j$. This implies

$$m_j < m_{j-1} < \ldots < m_1 < m_n < \ldots < m_{j+1}.$$

Hence A is the matrix needed. □

Proof for HodgeRank and PerronRank. By Theorem 1.12, this result follows immediately from the above by a limit argument. Specifically, let $X \in \mathcal{X}_n$ be a matrix that induces the ranking (σ_1, σ_2) on the pair HodgeRank and TropicalRank. Let $\epsilon := \min_{i,j} |m(X)_i - m(X)_j|$. By Corollary 1.12, one can choose sufficiently large k such that the matrix $Y := X^{(k)}$ satisfies $|v_i(Y) - m_i(X)| < \epsilon/3$ for all i. So $r(v(Y)) = r(m(X)) = \sigma_2$, and by Lemma 4.5, $r(h(Y)) = \sigma_1$. Hence Y is a matrix that realizes the pair (σ_1, σ_2) for HodgeRank and PerronRank. □

Proof for Tropical and PerronRank. To find the desirable X in Lemma 4.14, we consider a family of matrices called *perturbed comparison matrices*, introduced by Farkas [35]. These are strongly transitive MCMs with the first row and column perturbed by a noise vector. Explicit formulas for the principal eigenvalue and eigenvector are given in [4, 35].

For our case, fix a number $L > 1, L \in \mathbb{Q}$. Define $s \in \mathbb{R}^n$ by

$$s_i = 1 \text{ for all } i < n, \;\; s_n = \frac{1}{L}.$$

Let the upper-diagonal entries of $X \in \mathcal{X}_n$ be

$$X_{ij} = \frac{s_i}{s_j} \text{ if } i \neq 1, \quad X_{1j} = \frac{1}{s_j} \cdot \delta_j$$

where $\delta = (\delta_2, \ldots, \delta_n) \in \mathbb{Q}^{n-1}$ is the noise vector, with

$$\delta_2 < \delta_3 < \ldots < \delta_{n-1} = L, \text{ and } \delta_n = \frac{1}{L^2} \tag{4.1}$$

Then

$$X = \begin{bmatrix} 1 & \delta_2 & \delta_3 & \cdots & L & \frac{1}{L} \\ \frac{1}{\delta_2} & 1 & 1 & \cdots & 1 & L \\ \frac{1}{\delta_3} & 1 & 1 & \cdots & 1 & L \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ \frac{1}{L} & 1 & 1 & \cdots & 1 & L \\ L & \frac{1}{L} & \frac{1}{L} & \cdots & \frac{1}{L} & 1 \end{bmatrix}.$$

By [4], the principal eigenvector v of X is

$$v = \alpha_1 e_1 + \alpha_2 s + \alpha_3 w$$

where

$$e_1 = \begin{bmatrix} 1 \\ 0 \\ 0 \\ \vdots \\ 0 \\ 0 \end{bmatrix}, s = \begin{bmatrix} 1 \\ 1 \\ 1 \\ \vdots \\ 1 \\ \frac{1}{L} \end{bmatrix}, w = \begin{bmatrix} 0 \\ \frac{1}{\delta_2} - 1 \\ \frac{1}{\delta_3} - 1 \\ \vdots \\ \frac{1}{L} - 1 \\ \frac{1}{L}\left(\frac{1}{L^2} - 1\right) \end{bmatrix}$$

and $\alpha := (\alpha_1, \alpha_2, \alpha_3)$ is a column vector of the adjugate matrix $\text{adj}(r\mathbf{I} - Z)$, where $\mathbf{I}$ is the identity matrix, and

$$Z := \begin{bmatrix} 0 & a & b \\ 1 & n & c \\ 1 & 1 & 0 \end{bmatrix},$$

with

$$a = \sum_{i=2}^{n} (\delta_i - 1), b = \sum_{i=2}^{n} (\delta_i - 1)\left(\frac{1}{\delta_i} - 1\right), c = \sum_{i=2}^{n}\left(\frac{1}{\delta_i} - 1\right),$$

and r is the principal eigenvalue of Z, that is, it is the root of maximal modulus of the characteristic polynomial

$$p(t) = t^2(t - n) + b(n - 1) - ac. \tag{4.2}$$

Since $L > 1$ is the maximal element in each row, $m(X) = (1, \ldots, 1)$. By Lemma 4.14, it is sufficient to prove that there exists choices of δ such that $v(X)$ induces a complete ranking. In fact, we claim that any rational δ satisfying Equation (4.1) would do. Indeed, computing an explicit formula for α, we obtain

$$\alpha_1 = (r-n)r - c, \;\; \alpha_2 = r + c, \;\; \alpha_3 = r - n + 1.$$

Then,

$$\begin{aligned}
v_1 &= (r-n)r - c + r + c = (r-n+1)r \\
v_i &= r + c + (r-n+1) \cdot \left(\frac{1}{\delta_i} - 1\right) \text{ for } 2 \leq i < n-1 \\
v_{n-1} &= r + c + (r-n+1) \cdot \left(\frac{1}{L} - 1\right) \\
v_n &= (r+c) \cdot \frac{1}{L} + (r-n+1) \cdot \frac{1}{L}\left(\frac{1}{L^2} - 1\right)
\end{aligned}$$

Note that v_1 is a quadratic in r, while v_i are linear for $i > 1$. Since the δ_i's are distinct, $v_i \neq v_j$ for all $2 \leq i, j \leq n-1$. Suppose $v_1 = v_i$ for some $i \neq 1$, or $v_n = v_i$ for some $i \neq n$. In the first case, r is a root of a quadratic with coefficients in $\mathbb{Q}$. In the second case, r is a root of a linear equation with coefficients in $\mathbb{Q}$. Either cases imply the cubic p in Equation (4.2) must be factorisable into a quadratic times a linear term over $\mathbb{Q}$. That is

$$p(t) = t^3 - nt^2 + b(n-1) - ac = (t^2 + et + 1)(t - g) \text{ for some } e, g \in \mathbb{Q}$$

Equating coefficients, we have $eg = 1$, $e - g = -n$, and $b(n-1) - ac = -g$. The first two equations imply $1/g - g + n = 0$, hence $g = (n - \sqrt{n^2+4})/2$ or $(n + \sqrt{n^2+4})/2$. But $g \in \mathbb{Q}$, therefore $n^2 + 2^2$ must be a square natural number, since $n \in \mathbb{N}$. However, there is no positive Pythagorean triple involving 2, hence no such g exists. Therefore, the v_i's are all distinct. □

4.6 HodgeRank is the limit of PerronRank

Recall that $\mathcal{K} = \mathbb{R}_+^{n \times n}$ is the open cone of elementwise positive matrices in $\mathbb{R}^{n \times n}$. Equip $\mathbb{R}^n$ with the Euclidean norm. Let $\|\cdot\|$ denote the operator norm if the argument is a matrix, and the ℓ_2 norm if the argument is a vector. Let $\mathbf{I}$ be the $n \times n$ identity matrix. In this section we shall prove the Theorem 1.13 on the inverse tropical limit of the principal eigenvector. For the reader's convenience, we restate this theorem.

Theorem 4.15. *For $X \in \mathcal{K}$, define $h(X)_i = (\prod_{j=1}^{n} X_{ij})^{1/n}$. Then*

$$\lim_{k \to 0} v(X^{(k)})^{1/k} = h(X)$$

When X is restricted to the subvariety of multiplicative comparison matrices, Theorem 1.13 and Theorem 1.12 precisely state that HodgeRank and TropicalRank are limits of PerronRank. Ranking methods are often compared and selected via an optimization approach, where one finds methods which maximize an objective function, such as number of items ranked correctly [39, 49]. Therefore, these two theorems convert the comparison of PerronRank and HodgeRank into a parametrized optimization problem. It also contributes towards solving whether there exists an objective function in which PerronRank for $k \neq 0$ or ∞ is the optimal solution over all possible ranking algorithms. This question arose in the literature since HodgeRank and TropicalRank are known to be optimizers under ℓ_2 and ℓ_∞ error, respectively (cf. Theorem 4.6).

Since many optimization problems can be rephrased as geometric questions on the projected level sets of the noise distribution, we provide the following geometric version of Theorem 1.13.

Theorem 4.16. *Define $\tilde{V}_k : \wedge_2 \mathbb{R}^n \to \mathbb{TP}^{n-1}$, $\tilde{V}_k(A) = 1/k \log v([\exp(ka_{ij})])$ to be the log version of the Perron rank V_k. Let $\mathbf{1} = (1, \ldots, 1)^\top$ be the all one vector and $\mathbf{0} = (0, \ldots, 0)^\top$ the all-zero vector. For all $s \in \mathbb{TP}^{n-1}$, as a set, the fibers of the map $\tilde{V}_k$ satisfy $\tilde{V}_k^{-1}(s) = [s_i - s_j] + \tilde{V}_k^{-1}(\mathbf{0})$. The zero-fiber $\tilde{V}_k^{-1}(\mathbf{0})$ can be decomposed as*

$$\tilde{V}_k^{-1}(\mathbf{0}) = \bigoplus_{i=1}^{n} S_i(k) + \mathbb{R} \cdot \mathbf{1}\mathbf{1}^\top,$$

where the $S_i(k)$ for $i = 1, \ldots, n$ are orthogonal, with

$$\begin{aligned}
S_i(k) &:= \{a_{i\cdot} \in \mathbb{R}^n : a_{ij} < 0 \text{ for all } j, \sum_{j=1}^{n} \exp(ka_{ij}) = 1\} \text{ for } k \in (0, \infty), \\
S_i(\infty) &:= \{a_{i\cdot} \in \mathbb{R}^n : a_{ij} \leq 0, a_{ij} = 0 \text{ for at least one } j = 1, \ldots, n\}, \\
S_i(0) &:= \{a_{i\cdot} \in \mathbb{R}^n : \sum_{j=1}^{n} a_{ij} = 0\}
\end{aligned}$$

For large k, the zero-fiber is the real amoeba of the tropical max-plus eigenvector map [66]. For $k = 0$, the zero-fiber is the set of matrices decomposing into the all-one matrix plus a matrix whose column sum is zero, which is the zero-fiber of the HodgeRank map. Thus we obtain a crude version of Theorem 1.13.

4.6.1 Proof of the inverse tropical limit theorem

Our proof of Theorem 1.13 relies on the following proposition, which gives a linear approximation to the 'error term' of the principal eigenvector when X is not far from being a rank one matrix.

Proposition 4.17. *Fix a vector $s \in \mathbb{R}^n_+$. Let $\xi := [\xi_{ij}] := [s_j/s_i \cdot X_{ij}]$. Suppose there exists a constant $\kappa \geq 1$ such that the $n \times n$ matrix $\Xi = \xi - \kappa \mathbf{1}\mathbf{1}^\top - (1-\kappa)\mathbf{I}$ satisfies $\rho := (\frac{2\|\Xi\|}{n\kappa - 2\|\Xi\|})^2 < \frac{1}{2}$. Then*

$$v(X) = s \cdot (\mathbf{1} + \frac{1}{\kappa n}(r - \bar{r}\mathbf{1}) + \epsilon) \tag{4.3}$$

where $r = (\Xi \cdot \mathbf{1})$ is the row sum of Ξ, $\bar{r} = \frac{1}{n}\sum_{j=1}^n r_j$ is its mean, and $\|\epsilon\| < \frac{\rho}{1-\rho} \cdot \frac{\|\Xi\|}{\kappa\sqrt{n}}$.

This proposition is interesting in itself. One could think of s as the true score vector, $[s_i/s_j]$ as the true multiplicative comparison matrix, and ξ as the multiplicative perturbation with centered version Ξ, so defined since $v(\kappa\mathbf{1}\mathbf{1}^\top + (1-\kappa)\mathbf{I}) = \mathbf{1}$ for all $\kappa \geq 1$. If $\Xi \equiv 0$ then the perturbation does not affect the principal eigenvector, and $v(X)$ is exactly s. The proposition states that if the centered perturbation Ξ is small, as measured by ρ and $\frac{\|\Xi\|}{\kappa\sqrt{n}}$, then $v(X)$ differs from the true multiplicative score s only by a linear factor plus a lower order term. Unfortunately this proposition cannot be applied easily in practice since it requires knowledge of κ, which depends on the choice of s and in general neither are easy to find. For the proof of Theorem 1.13 it is sufficient to choose $\kappa = 1$.

Proof of Proposition 4.17. By Lemma 4.5, $v(X) = s \cdot v(\xi)$. Hence one can assume without loss of generality that $s \equiv \mathbf{1}$. Write $\xi = \kappa\mathbf{1}\mathbf{1}^\top + (\Xi + (1-\kappa)\mathbf{I})$. Then $\kappa\mathbf{1}\mathbf{1}^\top$ is a rank one matrix with one non-zero eigenvalue $n\kappa$, corresponding to the normalized eigenvector $\frac{1}{\sqrt{n}}\mathbf{1}$. Let $Y \in \mathbb{R}^{n\times(n-1)}$ be an orthonormal basis of the zero eigenspace of $\kappa\mathbf{1}\mathbf{1}^\top$. From standard results in perturbed linear operators (see, for example, [96, Theorem 2.7]), $\rho < \frac{1}{2}$ implies $v(\xi) = \frac{1}{\sqrt{n}}\mathbf{1} + YP + \frac{\epsilon}{\sqrt{n}}$, where $P := \frac{1}{n\sqrt{n}\kappa}Y^\top(\Xi + (1-\kappa)\mathbf{I})\mathbf{1}$, and the error term is bounded by

$$\|\epsilon\| \leq \frac{\rho}{1-\rho} \cdot \frac{\|\Xi\|}{n\kappa}.$$

Since $YY^\top = \mathbf{I} - \frac{1}{n}\mathbf{1}\mathbf{1}^\top$, and since $v(\xi)$ is only defined up a multiplicative constant, we have

$$\begin{aligned} v(\xi) &= \frac{1}{\sqrt{n}}\left(\mathbf{1} + \frac{1}{n\kappa}(\mathbf{I} - \frac{1}{n}\mathbf{1}\mathbf{1}^\top)(\Xi + (1-\kappa)\mathbf{I})\mathbf{1} + \epsilon\right) \\ &= \mathbf{1} + \frac{1}{n\kappa}(\Xi\mathbf{1} - \frac{1}{n}\mathbf{1}\mathbf{1}^\top\Xi\mathbf{1}) + \epsilon = \mathbf{1} + \frac{1}{n\kappa}(r - \bar{r}) + \epsilon \end{aligned}$$

□

Proof of Theorem 1.13. Under the notation of Proposition 4.17, let $\kappa = 1$, $s \equiv \mathbf{1}$, so $X = \xi$. Define $A = [A_{ij}] = [\log X_{ij}]$. Then

$$\Xi_{ij}^{(k)} = \exp(kA_{ij}) - 1 = k \cdot \left(A_{ij} + \sum_{t \geq 2} \frac{k^{t-1} A_{ij}^t}{t!} \right).$$

Hence for k close to 0, $\|\Xi^k\| = k\|A\| + o(k)$ and $\frac{\rho}{1-\rho} = O(k^2)$, and Proposition 4.17 applies. Note that $r = \Xi^{(k)}\mathbf{1} = k \cdot A\mathbf{1} + O(k^2)$, $\bar{r}\mathbf{1} = k \cdot \mathbf{1}\mathbf{1}^\top A\mathbf{1} + O(k^2)$, and $\epsilon = O(k^3)$. After applying the logarithm to Equation 4.3, Taylor expansion gives

$$\lim_{k \to 0} \frac{1}{k} \log v(X^k) = \lim_{k \to 0} \left(\frac{1}{kn} r - \frac{1}{kn} \bar{r} + O(k^2) \right) = \frac{1}{n} A\mathbf{1} - \frac{1}{n} \mathbf{1}\mathbf{1}^\top A\mathbf{1} = \frac{1}{n} \log h(X)$$

since the HodgeRank vector is defined only up to additive constants. □

We now prove the geometric version of the inverse tropical theorem, Theorem 4.16. It is sufficient to prove equivalent statements on the cone of elementwise positive matrices $\mathcal{K}$ for the map V_k. The fibers of this map are the set of elementwise positive matrices whose principal eigenvector is a fixed vector $w \in \mathbb{R}_+^n$, that is

$$\mathcal{K}(w) := \bigsqcup_{\mu \in (0,\infty)} \mathcal{K}(w, \mu) := \bigsqcup_{\mu \in (0,\infty)} \{X \in \mathcal{K} : v(X) = w, \rho(X) = \mu\}$$

where $\sqcup$ denote disjoint set union, where $\rho(X)$ is the Perron eigenvalue of X. We call this the *one-dimensional real positive Kalman variety*, motivated by the definition of the Kalman variety in Ottaviani and Sturmfels [73]. In general real Kalman varieties are difficult to characterize, however the one-dimensional positive variety has a simple description.

Corollary 4.18. *For each fixed pair (w, μ), let $\Psi_{w,\mu} : \mathcal{K} \to \mathcal{K}$ be the map $[X_{ij}] \mapsto [Y_{ij}] := [\frac{X_{ij} w_j}{\mu w_i}]$. Then*

$$\Psi_{w,\lambda}(\mathcal{K}(w, \mu)) = \bigoplus_{i=1}^{n} (\Delta_{n-1})_i$$

where each $(\Delta_{n-1})_i$ is the interior of the $(n-1)$-dimensional simplex on the $n \cdot (i-1) + 1$ to $n \cdot i$ coordinates of $\mathbb{R}^{n \times n}$.

Proof of Theorem 4.16. The first statement follows by a direct computation. By Corollary 4.18, $V_k^{-1}(\mathbf{1}) = \mathbb{R} \cdot \bigoplus_{i=1}^n (\Delta_{n-1})_i$. By taking log, one obtains the stated result for $\tilde{V}_k^{-1}(\mathbf{0})$. The case for $k = \infty$ follows from the convergence of the real amoebas to the zero set of the tropical eigenvectors map. As k approaches 0, each component $S_i(k)$ is flattened, and a little calculation shows that the each component converges, up to a translation by a large

constant times the $(1, \ldots, 1)$ vector, to its tangent plane at $a_{i1} = \ldots = a_{in}$, which has equation $\sum_{j=1}^{n} a_{ij} = 0$. This proves the theorem. □

4.7 Random ranking models

We mention an application of Proposition 4.17 to random ranking models, that is, when the input matrix has independent and identically distributed (i.i.d) entries, and the number of items is large. The main result of this section is the following.

Theorem 4.19. *Let $s \in \mathbb{R}_+^{\mathbb{N}}$. Consider the infinite triangular array $\{s_i/s_j\xi_{ij}, i, j = 1, 2, \ldots\}$, $\xi_{ij} = 1/\xi_{ji}$, and the entries $\{\xi_{ij}, i < j\}$ are positive, independent and identically distributed as ξ, $\xi \stackrel{d}{=} 1/\xi$, with finite $4k^{th}$ moment $m_{4k} := \mathbb{E}\xi^{4k} < \infty$ for some $k > 0$. Then for sufficiently large n, Equation (4.3) in Proposition 4.17 holds for $v(X^{(k)})$ with $\kappa = m_k$*

We first give the proof idea. Proposition 4.17 is very useful in comparing HodgeRank and PerronRank. However, as stated, it is not clear how to find the perturbation constant κ. In the proof of Theorem 1.13, where k is close to 0, it is sufficient to take $\kappa = 1$. In practice, one may need to utilize known or assumed structures of ξ. The key step is to obtain an upper bound on the operator norm of Ξ. The case of i.i.d noise is closely related to random matrices: since one can assume $s \equiv 1$, if $\{\xi_{ij} : i > j\}$ are i.i.d, the operator norm of ξ behaves almost like that of a matrix with iid entries. Then, it is natural to center at the mean, that is, picking $\kappa = \mathbb{E}(\xi_{ij})$. The main step in our proof is a modification of the combinatorial proof of the Bai-Yin theorem [109], allowing us to bound the operator norm of the centered matrix Ξ.

Proof of Theorem 4.19. Without loss of generality we can assume that $s_i = 1$ for all i and $k = 1$. Under the notation of Proposition 4.17, let $\kappa = m = \mathbb{E}\xi$. Note that $\Xi_{ij} = \xi_{ij} - m$, $\Xi_{ji} = 1/\xi_{ij} - m$ for $i < j$, and $\Xi_{ii} = 0$ for all i. Write

$$\|\Xi_n\| \leq \frac{1}{2}\|\Xi_n + \Xi_n^\top\| + \frac{1}{2}\|\Xi_n - \Xi_n^\top\|.$$

Now, $\Xi_n + \Xi_n^\top$ is a symmetric matrix with mean zero, i.i.d entries with distribution $\xi + 1/\xi - 2m$, finite second and fourth moment. By the Bai-Yin theorem for symmetric matrices [6],

$$\limsup_{n\to\infty} \|\Xi_n + \Xi_n^\top\|_{op}/\sqrt{n} \leq C$$

for some constant C depending on the second moment of ξ. For the second term, note that $(\Xi_n - \Xi_n^\top)_{ij} = \xi_{ij} - 1/\xi_{ij}$. As $\xi \stackrel{d}{=} 1/\xi$, we have $\xi_{ij} - 1/\xi_{ij} \stackrel{d}{=} 1/\xi_{ij} - \xi_{ij} = -(\xi_{ij} - 1/\xi_{ij})$. That is, $B_n := (\Xi_n - \Xi_n^\top)$ is a skew-symmetric matrix with mean zero, i.i.d upper-diagonal entries

whose distribution is symmetric, has finite second and fourth moment. We claim that for any even m,

$$\mathbb{E}\|B_n\|_{op}^m \leq \mathbb{E}tr((B_nB_n^\top)^{m/2}) \leq \mathbb{E}tr(\tilde{B}_n^m)$$

where tr denote the trace of a matrix, and $\tilde{B}_m$ is the i.i.d symmetric matrix with ij-th entry distributed as $\xi - 1/\xi$. If holds, by repeating the combinatorics proof of the Bai-Yin theorem (see [109]), we see that

$$\limsup_{n\to\infty} \|\Xi_n - \Xi_n^\top\|_{op}/\sqrt{n} \leq C'$$

for some constant C'. Indeed, since B_n is skew-symmetric with symmetrically distributed entries, in taking the expectation of $tr((B_nB_n^\top)^{m/2})$, the only terms contributing are those with legs transversed an even number of times as in the case of $tr(\tilde{B}_n^m)$. The two sum differ only in the signs associated to each term: $tr((B_nB_n^\top)^{m/2})$ has some terms with negative signs, whereas all terms in $tr(\tilde{B}_n^m)$ have positive signs. Since the legs appear with even multiplicity, the expectation is always positive, thus all terms are positive. Hence

$$tr((B_nB_n^\top)^{m/2}) \leq tr(\tilde{B}_n^m).$$

Thus this proves the claim. Finally, we have

$$\limsup_{n\to\infty} \|\Xi_n\|/\sqrt{n} \leq C$$

for some constant C. That is, for sufficiently large n, with high probability, $\rho_n = O(n^{-1}) \ll \frac{1}{2}$, and Proposition 4.17 applies, with $\|\epsilon\| = O(n^{-1})$. □

4.8 Summary and open problems

In this chapter we considered the problem of obtaining a cardinal ranking from an $n \times n$ pairwise comparison matrix on additive and multiplicative scales. We studied the mathematical properties and connections of three proposed methods: PerronRank, HodgeRank and TropicalRank. We showed that for $n \geq 4$, all regions of ranking differences can be realized between any pair of these three methods. This indicates that the choice of method in practice deserves further study. We also proved Theorem 1.13 and its geometric version, Theorem 4.16, which show that HodgeRank is the limit of PerronRank. This is an important mathematical connection, and the proof technique shows that HodgeRank is the first-order approximation to PerronRank under the i.i.d random ranking model. Below we list some future research directions.

Realizability of ranking triples

It is natural to ask whether all ranking triples can be realized. For $n = 3$, the answer is no.

Lemma 4.20. *For $n = 3$, if $X \in \mathcal{X}_n$, then $h(X) = m(X) = v(X)$ (up to scaling). That is, all methods produce the same score vector, and hence the same ranking.*

The result is a direct computation: if $X \in \mathcal{X}_n$ is not strongly transitive, there can be only one maximal cycle of length 3, hence the tropical eigenvector is always unique. In this case, one can plug in the formula for the HodgeRank vector (as row geometric mean) to verify that it is indeed also the tropical and principal eigenvector of X. We can also see this geometrically in $\wedge_2\mathbb{R}^n$ for the pair HodgeRank and TropicalRank: in this case $V_n^\perp$, viewed as vectors of the upper-diagonal entry, is spanned by the vector $(1, -1, 1)$. Hence the standard 3-cube centered at A will always hits V_n at either the corner $A + \lambda(A) \cdot (1, -1, 1)$, or the opposite corner $A + \lambda(A) \cdot (-1, 1, -1)$, therefore $h(A) = m(A)$ for all $A \in \wedge_2\mathbb{R}^n$.

The case $n > 3$ is substantially more difficult. Due to their mathematical connections in Theorem 1.12 and Theorem 1.13, one may suspect that not all ranking regions may be realizable, especially when the ranking induced by HodgeRank and TropicalRank agree. This is illustrated in the example below.

Example 4.21. Here our matrix $X \in \mathcal{X}_n$ is

$$X = \begin{bmatrix} 1 & 1.57 & 0.72 & 0.70 \\ 0.63 & 1 & 1.52 & 0.65 \\ 1.38 & 0.65 & 1 & 1.57 \\ 1.45 & 1.52 & 0.63 & 1 \end{bmatrix}$$

The corresponding vectors, normalized to have the first component be 1 and rounded in 3 decimal places, are

$$\begin{aligned} v(X) &= \begin{bmatrix} 1 & 0.991 & 1.191 & 1.151 \end{bmatrix} \Rightarrow \text{ranking: } 3 > 4 > 1 > 2 \\ h(X) &= \begin{bmatrix} 1 & 0.942 & 1.155 & 1.151 \end{bmatrix} \Rightarrow \text{ranking: } 3 > 4 > 1 > 2 \\ m(X) &= \begin{bmatrix} 1 & 0.979 & 0.989 & 0.968 \end{bmatrix} \Rightarrow \text{ranking: } 1 > 3 > 2 > 4 \end{aligned}$$

Both the matrix X and the corresponding scores do not have 'crazy' entries that can intuitively indicate intransitivity. For example, if one uses the consistency index $(\mu(X) - n)/(n-1)$ suggested by Saaty [86], where $\mu(X)$ is the principal eigenvalue of X, then the consistency index for this case is 0.07073, well within the proposed 0.1 recommended cut-off. Thus consistency index may not be a good measure of ranking agreement between methods.

Disagreement for random matrices

Our current construction of matrices realizing a particular ranking pair is somewhat artificial. However, simulations suggested that the true region of disagreement is much larger than those considered in our construction. Thus one may ask the following: given a 'random' pairwise comparison matrix, what is the probability that a given pair of methods disagree?

Disagreement for large number of items

We already saw that for large number of items, under i.i.d noise on each entry, HodgeRank serves as a first order approximation of PerronRank. In this case, one would be less worried about small differences in ranking (say, a difference by one transposition), but more on how the methods compare in terms of statistical optimality and computational efficiency. Can we compare these properties of HodgeRank and PerronRank under the i.i.d model?

Pairwise ranking as an optimization problem

As previously mentioned, Theorem 1.13 and 1.12 converts the problem of comparing the three methods HodgeRank, Principal Eigenvector and Tropical Eigenvector into a parametric optimization problem in a single variable over the positive reals. This leads to two interesting, open questions in the area of ranking by pairwise comparison: how to solve this optimization problem, and how exactly does the principal eigenvector behave as a function of the parameter k and the input matrix X?

Chapter 5

Finite Size Biased Permutation

This chapter appeared as the joint paper with Jim Pitman titled 'Size-biased permutation of a finite sequence with independent and identically distributed terms', submitted and is available as an arXiv eprint [79].

5.1 Introduction

Let $x = (x(1), x(2), \ldots)$ be a positive sequence with finite sum $t = \sum_{i=1}^{\infty} x(i)$. Its size-biased permutation (s.b.p) is the same sequence presented in a random order $(x(\sigma_1), x(\sigma_2), \ldots)$, where $\mathbb{P}(\sigma_1 = i) = \dfrac{x(i)}{t}$, and for k distinct indices $i_1, \ldots, i_k$,

$$\mathbb{P}(\sigma_k = i_k | \sigma_1 = i_1, \ldots, \sigma_{k-1} = i_{k-1}) = \frac{x(i_k)}{t - (x(i_1) + \ldots + x(i_{k-1}))}. \tag{5.1}$$

An index i with bigger 'size' $x(i)$ tends to appear earlier in the permutation, hence the name size-biased. Size-biased permutation of a random sequence is defined by conditioning on the sequence values.

One of the earliest occurrences of size-biased permutation is in social choice theory. For fixed sequence length n, the goal is to infer the $x(i)$ given multiple observations from the random permutation defined by (5.1). Here the $x(i)$ are the relative scores or desirabilities of the candidates, and (5.1) models the distribution of their rankings in an election. Now known as the Plackett-Luce model, it has wide applications [18, 80, 98].

Around the same time, biologists in population genetics were interested in inferring the distribution of alleles in a population through sampling. In these applications, $x(i)$ is the abundance and $x(i)/t$ is the relative abundance of the ith species [33]. Size-biased permutation models the outcome of successive sampling, where one samples without replacement from the population and records the abundance of newly discovered species in the order that they appear. To account for the occurrence of new types of alleles through mutation

and migration, they considered random abundance sequences and did not assume an upper limit to the number of possible types. Species sampling from a random infinite sequence is sometimes known as *size-biased random permutation*, a term coined by Patil and Taillie [75]. The earliest work along this vein is perhaps that of McCloskey [67], who obtained results on the size-biased permutation of ranked jumps in a certain Poisson point process (p.p.p). The distribution of this ranked sequence is now known as the Poisson-Dirichlet distribution $PD(0,\theta)$. The distribution of its size-biased permutation is the $GEM(0,\theta)$ distribution. See Section 5.4.3 for their definitions. This work was later generalized by Perman, Pitman and Yor [76], who studied size-biased permutation of ranked jumps of a subordinator; see Section 5.4.

One can also approach size-biased permutation of finite sequences through the study of partition structure, starting with Ewens' sampling formula [34]. We do not take this route, but would like to offer a quick overview. Ewens' sampling formula is the finite combinatorial version of size-biased sampling from the $PD(0,\theta)$ distribution. Kingman [59] wrote: 'One of the most striking results of recent theoretical research in population genetics is the sampling formula enunciated by Ewens and shown by (others) to hold for a number of different population models'. In studying this formula, Kingman initiated the theory of partition structure [59, 60]. Kingman showed that Ewens' sampling formula defines a particular partition structure by deletion of type. This property is closely connected to invariance under size-biased permutation, see [40, 77]. Subsequent authors have studied partition structures and their representations in terms of exchangeable random partitions, random discrete distributions, random trees and associated random processes of fragmentation and coalescence, Bayesian statistics, and machine learning. See [78] and references therein.

This chapter focuses on *finite i.i.d size-biased permutation*, that is, the size-biased permutation of n independent and identically distributed (i.i.d) random variables from some distribution F on $(0,\infty)$. Our setting is a finite dimensional analogue of the size-biased permutation of ranked jumps of a subordinator studied in [76], as well as a special form of *induced order statistics* [12, 23]. This intersection grants us different tools for deriving distributional properties. Their comparison lead to new results, as well as simpler proofs of existing ones. By considering size-biased permutation of i.i.d triangular arrays, we derive convergence in distribution of the remaining u fraction in a successive sampling scheme. This provides alternative proofs to similar statements in the successive sampling literature. Our main contribution, Theorem 5.28 in Section 5.6, describes the asymptotic distribution of the last few terms in a finite i.i.d size-biased permutation via a Poisson coupling with its few smallest order statistics.

Organization

We derive joint and marginal distribution of finite i.i.d size-biased permutation in Section 5.2 through a Markov chain, and re-derive them in Section 5.3 using induced order statistics. Section 5.4 connects our setting and its infinite version of [76]. As the sequence length tends

to infinity, we derive asymptotics of the last u fraction of finite i.i.d size-biased permutation in Section 5.5, and that of the first few terms in Section 5.6.

Notation

We shall write $gamma(a, \lambda)$ for a Gamma distribution whose density at x is $\lambda^a x^{a-1} e^{-\lambda x}/\Gamma(a)$ for $x > 0$, and $beta(a, b)$ for the Beta distribution whose density at x is $\dfrac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} x^{a-1}(1-x)^{b-1}$ for $x \in (0, 1)$. For an ordered sequence $(Y_n(k), k = 1, \ldots, n)$, let $Y_n^{rev}(k) = Y_n(n-k+1)$ be the same sequence presented in reverse. For order statistics, we write $Y^{\uparrow}$ for the increasing sequence, and $Y^{\downarrow}$ for the decreasing sequence. Throughout this chapter, unless otherwise indicated, we use $X_n = (X_n(1), \ldots, X_n(n))$ for the underlying i.i.d sequence with n terms, and $(X_n[1], \ldots, X_n[n])$ for its size-biased permutation. To avoid having to list out the terms, it is also sometimes convenient to write $X_n^* = (X_n^*(1), \ldots, X_n^*(n))$ for the size-biased permutation of X_n.

5.2 Markov property and stick-breaking

Assume that F has density ν_1. Let $T_{n-k} = X_n[k+1] + \ldots + X_n[n]$ denote the sum of the last $n-k$ terms in an i.i.d size-biased permutation of length n. We first derive joint distribution of the first k terms $X_n[1], \ldots, X_n[k]$.

Proposition 5.1 (Barouch-Kaufman [7])**.** *For $1 \le k \le n$, let ν_k be the density of S_k, the sum of k i.i.d random variables with distribution F. Then*

$$\begin{aligned}
&\mathbb{P}(X_n[1] \in dx_1, \ldots, X_n[k] \in dx_k) \\
&= \frac{n!}{(n-k)!} \left(\prod_{j=1}^{k} x_j \, \nu_1(x_j) \, dx_j \right) \int_0^\infty \nu_{n-k}(s) \prod_{j=1}^{k} (x_j + \ldots + x_k + s)^{-1} \, ds \qquad (5.2) \\
&= \frac{n!}{(n-k)!} \left(\prod_{j=1}^{k} x_j \, \nu_1(x_j) \, dx_j \right) \mathbb{E} \left(\prod_{j=1}^{k} \frac{1}{x_j + \ldots + x_k + S_{n-k}} \right). \qquad (5.3)
\end{aligned}$$

Proof. Let σ denote the random permutation on n letters defined by size-biased permutation as in (5.1). Then there are $\frac{n!}{(n-k)!}$ distinct possible values for $(\sigma_1, \ldots, \sigma_k)$. By exchangeability of the underlying i.i.d random variables $X_n(1), \ldots, X_n(n)$, it is sufficient to consider $\sigma_1 = 1$, …, $\sigma_k = k$. Note that

$$\mathbb{P}\left((X_n(1), \ldots, X_n(k)) \in dx_1 \ldots dx_k, \sum_{j=k+1}^{n} X_n(j) \in ds \right) = \nu_{n-k}(s) \, ds \prod_{j=1}^{k} \nu_1(x_j) \, dx_j.$$

Thus, restricted to $\sigma_1 = 1, \ldots, \sigma_k = k$, the probability of observing $(X_n[1], \ldots, X_n[k]) \in dx_1 \ldots dx_k$ and $T_{n-k} \in ds$ is precisely

$$\frac{x_1\, dx_1}{x_1 + \ldots + x_k + s} \frac{x_2\, dx_2}{x_2 + \ldots + x_k + s} \cdots \frac{x_k\, dx_k}{x_k + s}\, \nu_{n-k}(s) \left(\prod_{j=1}^{k} \nu_1(x_j)\, dx_j \right) ds.$$

By summing over $\frac{n!}{(n-k)!}$ possible values for $(\sigma_1, \ldots, \sigma_k)$, and integrating out the sum T_{n-k}, we arrive at (5.2). Equation (5.3) follows by rewriting. □

Note that $X_n[k] = T_{n-k+1} - T_{n-k}$ for $k = 1, \ldots, n-1$. Thus we can rewrite (5.2) in terms of the joint law of $(T_n, T_{n-1}, \ldots, T_{n-k})$:

$$\mathbb{P}(T_n \in dt_0, \ldots, T_{n-k} \in dt_k) = \frac{n!}{(n-k)!} \left(\prod_{i=0}^{k-1} \frac{t_i - t_{i+1}}{t_i} \nu_1(t_i - t_{i+1}) \right) \nu_{n-k}(t_k)\, dt_0 \ldots dt_k. \tag{5.4}$$

Rearranging (5.4) yields the following result, which appeared as an exercise in [17, §2.3].

Corollary 5.2 (Chaumont-Yor [17]). *The sequence $(T_n, T_{n-1}, \ldots, T_1)$ is an inhomogeneous Markov chain with transition probability*

$$\mathbb{P}(T_{n-k} \in ds | T_{n-k+1} = t) = (n-k+1)\, \frac{t-s}{t} \cdot \nu_1(t-s)\, \frac{\nu_{n-k}(s)}{\nu_{n-k+1}(t)}\, ds, \tag{5.5}$$

for $k = 1, \ldots, n-1$. Together with $T_n \stackrel{d}{=} S_n$, equation (5.5) specifies the joint law in (5.4), and vice versa.

The stick-breaking representation

An equivalent way to state (5.5) is that for $k \geq 1$, conditioned on $T_{n-k+1} = t$, $X_n[k]$ is distributed as the first size-biased pick out of $n-k+1$ i.i.d random variables conditioned to have sum $S_{n-k+1} = t$. This provides a recursive way to generate a finite i.i.d size-biased permutation: first generate T_n (which is distributed as S_n). Conditioned on the value of T_n, generate T_{n-1} via (5.5), let $X_n[1]$ be the difference. Now conditioned on the value of T_{n-1}, generate T_{n-2} via (5.5), let $X_n[2]$ be the difference, and so on. Let us explore this recursion from a different angle by considering the ratio $W_{n,k} := \dfrac{X_n[k]}{T_{n-k+1}}$ and its complement, $\overline{W}_{n,k} = 1 - W_{n,k} = \dfrac{T_{n-k}}{T_{n-k+1}}$. For $k \geq 2$, note that

$$\frac{X_n[k]}{T_n} = \frac{X_n[k]}{T_{n-k+1}} \frac{T_{n-k+1}}{T_{n-k+2}} \cdots \frac{T_{n-1}}{T_n} = W_{n,k} \prod_{i=1}^{k-1} \overline{W}_{n,i}. \tag{5.6}$$

The variables $\overline{W}_{n,i}$ can be interpreted as residual fractions in a *stick-breaking* scheme: start with a stick of length 1. Choose a point on the stick according to distribution $W_{n,1}$, 'break' the stick into two pieces, discard the piece of length $W_{n,1}$ and rescale the remaining half to have length 1. Repeating this procedure k times, and (5.6) is the fraction broken off at step k relative to the original stick length.

Together with $T_n \stackrel{d}{=} S_n$, one could use (5.6) to compute the marginal distribution for $X_n[k]$ in terms of the ratios $\overline{W}_{n,i}$. In general the $W_{n,i}$ are not necessarily independent, and their joint distributions need to be worked out from (5.5). However, when F has gamma distribution, $T_n, W_{n,1}, \ldots, W_{n,k}$ are independent, and (5.6) leads to the following result of Patil and Taillie [75].

Proposition 5.3 (Patil-Taillie [75])**.** *If F has distribution gamma(a, λ) for some $a, \lambda > 0$, then T_n and the $W_{n,1}, \ldots, W_{n,n-1}$ in (5.6) are mutually independent. In this case,*

$$\begin{aligned} X_n[1] &= \gamma_0 \beta_1 \\ X_n[2] &= \gamma_0 \bar{\beta}_1 \beta_2 \\ &\cdots \\ X_n[n-1] &= \gamma_0 \bar{\beta}_1 \bar{\beta}_2 \ldots \bar{\beta}_{n-2} \beta_{n-1} \\ X_n[n] &= \gamma_0 \bar{\beta}_1 \bar{\beta}_2 \ldots \bar{\beta}_{n-1} \end{aligned}$$

where γ_0 has distribution gamma(an, λ), β_k has distribution beta$(a+1, (n-k)a)$, $\bar{\beta}_k = 1 - \beta_k$ for $1 \leq k \leq n-1$, and the random variables $\gamma_0, \beta_1, \ldots, \beta_n$ are independent.

Proof. This statement appeared as a casual in-line statement without proof in [75], perhaps since there is an elementary proof (which we will outline later). For the sake of demonstrating previous computations, we shall start with (5.5). By assumption S_k has distribution *gamma*(ak, λ). One substitutes the density of *gamma*(ak, λ) for ν_k to obtain

$$\mathbb{P}(T_{n-k} \in ds | T_{n-k+1} = t) = C \left(\frac{(t-s)^a}{t} \right) \left(\frac{s^{a(n-k)-1}}{t^{a(n-k)+a-1}} \right) = C \left(1 - \frac{s}{t} \right)^a \left(\frac{s}{t} \right)^{a(n-k)-1}$$

for some normalizing constant C. By rearranging, we see that $\dfrac{T_{n-k}}{T_{n-k+1}}$ has distribution *beta*$(a+1, a(n-k))$, and is independent of T_{n-k+1}. Therefore $W_{n,1}$ is independent of T_n. By the stick-breaking construction, $W_{n,2}$ is independent of T_{n-1} and T_n, and hence of $W_{n,1}$. The final formula follow from rearranging (5.6). $\square$

Here is another direct proof. By the stick-breaking construction, it is sufficient to show that T_n is independent of $W_{n,1} = \dfrac{X_n[1]}{T_n}$. Note that

$$\mathbb{P}(X_n[1]/T_n \in du, T_n \in dt) = n\,u\,\mathbb{P}\left(\frac{X_n(1)}{X_n(1) + (X_n(2) + \ldots + X_n(n))} \in du, T_n \in dt\right). \quad (5.7)$$

Since $X_n(1) \stackrel{d}{=} gamma(a,1)$, $S_{n-1} = X_n(2) + \ldots + X_n(n) \stackrel{d}{=} gamma(a(n-1),1)$, independent of $X_n(1)$, the ratio $\frac{X_n(1)}{X_n(1)+S_{n-1}}$ has distribution $beta(a, a(n-1))$ and is independent of T_n. Thus

$$\begin{aligned}\mathbb{P}(X_n[1]/T_n \in du) &= n\,u\,\frac{\Gamma(a + a(n-1))}{\Gamma(a)\Gamma(a(n-1))}u^{a-1}(1-u)^{a(n-1)-1} \\ &= \frac{\Gamma(a+1+a(n-1))}{\Gamma(a+1)\Gamma(a(n-1))}u^{a}(1-u)^{a(n-1)-1}.\end{aligned}$$

In other words, $X_n[1]/T_n \stackrel{d}{=} beta(a, a(n-1))$. This proves the claim.

Lukacs [65] proved that if X, Y are non-degenerate, positive independent random variables, then $X + Y$ is independent of $\frac{X}{X+Y}$ if and only if both X and Y have gamma distributions with the same scale parameter. Thus one obtains another characterization of the gamma distribution.

Corollary 5.4 (Patil-Taillie converse)**.** *If T_n is independent of $X_n[1]/T_n$, then F is $gamma(a,\lambda)$ for some $a, \lambda > 0$.*

Proof. One applies Lukacs' theorem to $X_n(1)$ and $(X_n(2) + \ldots + X_n(n))$ in (5.7). □

5.3 Size-biased permutation as induced order statistics

When n i.i.d pairs $(X_n(i), Y_n(i))$ are ordered by their Y-values, the corresponding $X_n(i)$ are called the *induced order statistics* of the vector Y_n, or its *concomitants*. Gordon [43] first proved the following result for finite n which shows that finite i.i.d size-biased permutation is a form of induced order statistics. Here we state the infinite sequence version, which is a special case of [76, Lemma 4.4].

Proposition 5.5 (Perman, Pitman and Yor [76])**.** *Let x be a fixed positive sequence with finite sum $t = \sum_{i=1}^{\infty} x(i)$, ϵ a sequence of i.i.d standard exponential random variables, independent of x. Let Y be the sequence with $Y(i) = \epsilon(i)/x(i)$, $i = 1, 2, \ldots$, $Y^{\uparrow}$ its sequence of increasing order statistics. Define $X^*(k)$ to be the value of the $x(i)$ such that $Y(i)$ is $Y^{\uparrow}(k)$. Then $(X^*(k), k = 1, 2, \ldots)$ is a size-biased permutation of the sequence x. In particular, the size-biased permutation of a positive i.i.d sequence $(X_n(1), \ldots, X_n(n))$ is distributed as the*

induced order statistics of the sequence $(Y_n(i) = \epsilon_n(i)/X_n(i), 1 \le i \le n)$ *for an independent sequence of i.i.d standard exponentials* $(\epsilon_n(1), \ldots, \epsilon_n(n))$*, independent of the* $X_n(i)$*.*

Proof. Note that the $Y(i)$ are independent exponentials with rates $x(i)$. Let σ be the random permutation such $Y(\sigma(i)) = Y^{\uparrow}(i)$. Note that $X^*(k) = x(\sigma(k))$. Then

$$\mathbb{P}(\sigma(1) = i) = \mathbb{P}(Y(i) = \min\{Y(j), j = 1, 2, \ldots\}) = \frac{x(i)}{t},$$

thus $X^*(1) \stackrel{d}{=} x[1]$. In general, for distinct indices $i_1, \ldots, i_k$, by the memoryless property of the exponential distribution,

$$\begin{aligned} &\mathbb{P}(\sigma(k) = i_k | \sigma(1) = i_1, \ldots, \sigma(k) = i_{k-1}) \\ =&\mathbb{P}\left(Y(i_k) = \min\{Y(\sigma(j)), j \ge k\} | \sigma(1) = i_1, \ldots, \sigma(k) = i_{k-1}\right) = \frac{x(i_k)}{t - \sum_{j=1}^{k-1} x(i_j)}. \end{aligned}$$

Induction on k completes the proof. □

Proposition 5.5 readily supplies simple proofs for joint, marginal and asymptotic distributions of i.i.d size-biased permutation. For instance, the proof of the following nesting property, which can be cumbersome, amounts to i.i.d thinning.

Corollary 5.6. *Consider a finite i.i.d size-biased permutation* $(X_n[1], \ldots, X_n[n])$ *from a distribution* F*. For* $1 \le m \le n$*, select* m *integers* $a_1 < \ldots < a_m$ *by uniform sampling from* $\{1, \ldots, n\}$ *without replacement. Then the subsequence* $\{X_n[a_j], 1 \le j \le m\}$ *is jointly distributed as a finite i.i.d size-biased permutation of length* m *from* F*.*

In general, the induced order statistics representation of size-biased permutation is often useful in studying limiting distribution as $n \to \infty$, since one can consider the i.i.d pair $(X_n(i), Y_n(i))$ and appeal to tools from empirical process theory. We shall demonstrate this in Section 5.5 and 5.6.

5.3.1 Joint and marginal distribution revisited

We now revisit the results of 5.2 using induced order statistics. This leads to a different formula for the joint distribution, and an alternative proof of the Barouch-Kaufman formula (5.3).

Proposition 5.7. $(X_n[k], k = 1, \ldots, n)$ *is distributed as the first coordinate of the sequence of pairs* $((X_n^*(k), U_n^{\downarrow}(k)), k = 1, \ldots, n)$*, where* $U_n^{\downarrow}(1) \ge \ldots \ge U_n^{\downarrow}(n)$ *is a sequence of uniform order statistics, and conditional on* $(U_n^{\downarrow}(k) = u_k, 1 \le k \le n)$*, the* $X_n^*(k)$ *are independent with*

distribution $(G_{u_k}(\cdot), k = 1, \ldots, n)$*, where*

$$G_u(dx) = \frac{xe^{-\phi^{-1}(u)x}\, F(dx)}{-\phi'(\phi^{-1}(u))}. \tag{5.8}$$

Here ϕ *is the Laplace transform of* X*, that is,* $\phi(y) = \int_0^\infty e^{-yx}\, F(dx)$*,* ϕ' *its derivative and* ϕ^{-1} *its inverse function.*

Proof. Let X_n be the sequence of n i.i.d draws from F, ϵ_n an independent sequence of i.i.d standard exponentials, $Y_n(i) = \epsilon_n(i)/X_n(i)$ for $i = 1, \ldots, n$. Note that the pairs $\{(X_n(i), Y_n(i)), 1 \le i \le n\}$ is an i.i.d sample from the joint distribution $F(dx)[xe^{-yx}\, dy]$. Thus $Y_n(i)$ has marginal density

$$P(Y_n(i) \in dy) = -\phi'(y)\, dy, \quad 0 < y < \infty, \tag{5.9}$$

and its distribution function is $F_Y = 1 - \phi$. Given $\{Y_n(i) = y_i, 1 \le i \le n\}$, the $X_n^*(i)$ defined in Proposition 5.5 are independent with conditional distribution $\widetilde{G}(y_i, \cdot)$ where

$$\widetilde{G}(y, dx) = \frac{xe^{-yx}\, F(dx)}{-\phi'(y)}. \tag{5.10}$$

Equation (5.8) follows from writing the order statistics as the inverse transforms of ordered uniform variables

$$(Y_n^{\uparrow}(1), \ldots, Y_n^{\uparrow}(n)) \stackrel{d}{=} (F_Y^{-1}(U_n^{\downarrow}(n)), \ldots, F_Y^{-1}(U_n^{\downarrow}(1))) \stackrel{d}{=} (\phi^{-1}(U_n^{\downarrow}(1)), \ldots, \phi^{-1}(U_n^{\downarrow}(n))) \tag{5.11}$$

where $(U_n^{\downarrow}(k), k = 1, \ldots, n)$ is an independent decreasing sequence of uniform order statistics. Note that the minus sign in (5.9) results in the reversal of the sequence U_n in the second equality of (5.11). □

Corollary 5.8. *For* $1 \le k \le n$ *and* $0 < u < 1$*, let*

$$f_{n,k}(u) = \frac{\mathbb{P}(U_n(k) \in du)}{du} = n\binom{n-1}{k-1} u^{n-k}(1-u)^{k-1} \tag{5.12}$$

be the density of the k*th largest of the* n *uniform order statistics* $(U_n(i), i = 1, \ldots, n)$*. Then*

$$\frac{\mathbb{P}(X_n[k] \in dx)}{xF(dx)} = \int_0^\infty e^{-xy} f_{n,k}(\phi(y))\, dy. \tag{5.13}$$

Proof. Equation (5.12) follows from known results on order statistics, see ([27]). For $u \in [0, 1]$, let $y = \phi^{-1}(u)$. Then $\dfrac{dy}{du} = \dfrac{1}{\phi'(\phi^{-1}(u))}$ by the inverse function theorem. Apply this change of variable to (5.8), rearrange and integrate with respect to y, we obtain (5.13). □

In particular, for the first and last values,

$$\frac{\mathbb{P}(X_n[1] \in dx)}{xF(dx)} = n \int_0^\infty e^{-xy} \phi(y)^{n-1}\, dy$$
$$\frac{\mathbb{P}(X_n[n] \in dx)}{xF(dx)} = n \int_0^\infty e^{-xy} (1-\phi(y))^{n-1}\, dy.$$

Alternative derivation of the Barouch-Kaufman formula

Write $\phi(y) = \mathbb{E}(e^{-yX})$ for X with distribution F. Then $\phi(y)^{n-1} = \mathbb{E}(e^{-yS_{n-1}})$ where S_{n-1} is the sum of $(n-1)$ i.i.d random variables with distribution F. Since all integrals involved are finite, by Fubini's theorem

$$\begin{aligned}\frac{\mathbb{P}(X_n[1] \in dx)}{xF(dx)} &= n \int_0^\infty e^{-xy} \phi(y)^{n-1}\, dy = n\,\mathbb{E}\left(\int_0^\infty e^{-xy} e^{-yS_{n-1}}\, dy \right) \\ &= n\,\mathbb{E}\left(\frac{1}{x + S_{n-1}} \right),\end{aligned}$$

which is a rearranged version of the Barouch-Kaufman formula (5.3) for $k = 1$. Indeed, one can derive the entire formula from Proposition 5.7. For simplicity we demonstrate the case $k = 2$.

Proof of (5.3) for $k = 2$. The joint distribution of the two largest uniform order statistics $U_n^{\downarrow}(1), U_n^{\downarrow}(2)$ has density

$$f(u_1, u_2) = n(n-1) u_2^{n-1}, \qquad \text{for } 0 \le u_2 \le u_1 \le 1.$$

Conditioned on $U_n^{\downarrow}(1) = u_1, U_n^{\downarrow}(2) = u_2$, $X_n[1]$ and $X_n[2]$ are independent with distribution (5.8). Let $y_1 = \phi^{-1}(u_1), y_2 = \phi^{-1}(u_2)$, so $dy_1 = \dfrac{du_1}{\phi'(\phi^{-1}(u_1))}, dy_2 = \dfrac{du_2}{\phi'(\phi^{-1}(u_2))}$. Let S_{n-2} denote the sum of $(n-2)$ i.i.d random variables with distribution F. Apply this change of

variable and integrate out y_1, y_2, we have

$$\begin{aligned}
\frac{\mathbb{P}(X_n[1] \in dx_1, X_n[2] \in dx_2)}{x_1 x_2 F(dx_1) F(dx_2)} &= n(n-1) \int_0^\infty \int_{y_1}^\infty e^{-y_1 x_1} e^{-y_2 x_2} (\phi(y_2))^{n-2} \, dy_2 dy_1 \\
&= n(n-1) \int_0^\infty \int_{y_1}^\infty e^{-y_1 x_1} e^{-y_2 x_2} \mathbb{E}(e^{-y_2 S_{n-2}}) \, dy_2 dy_1 \\
&= n(n-1) \mathbb{E} \left(\int_0^\infty \int_{y_1}^\infty e^{-y_1 x_1} e^{-y_2 (x_2 + S_{n-2})} \, dy_2 dy_1 \right) \\
&= n(n-1) \mathbb{E} \left(\int_0^\infty e^{-y_1 x_1} \frac{e^{-y_1 (x_2 + S_{n-2})}}{x_2 + S_{n-2}} \, dy_1 \right) \\
&= n(n-1) \mathbb{E} \left(\frac{1}{(x_2 + S_{n-2})(x_1 + x_2 + S_{n-2})} \right),
\end{aligned}$$

where the swapping of integrals is justified by Fubini's theorem, since all integrals involved are finite. □

Example 5.9. Suppose F is $gamma(a, 1)$. Then $\phi(y) = (\frac{1}{1+y})^a$, and $\phi^{-1}(u) = u^{-1/a} - 1$. Hence G_u in (5.8) is

$$G_u(dx) = \frac{x}{a u^{(a+1)/a}} \, e^{-(u^{-1/a} - 1)x} \, F(dx) = \frac{x^a}{\Gamma(a+1)} u^{-(a+1)/a} e^{-x u^{-1/a}}.$$

That is, G_u is $gamma(a + 1, u^{-1/a})$.

Patil-Taillie revisited

When F is $gamma(a, \lambda)$, Lemma 5.7 gives the following result, which is an interesting complement to the Patil-Taillie representation in Proposition 5.3.

Proposition 5.10. *Suppose F is $gamma(a, \lambda)$. Then G_u is $gamma(a + 1, \lambda u^{-1/a})$, and*

$$(X_n[k], k = 1, \ldots, n) \stackrel{d}{=} ([U_n^{\downarrow}(k)]^{1/a} \gamma_k, k = 1, \ldots, n) \tag{5.14}$$

where $\gamma_1, \ldots, \gamma_n$ are i.i.d $gamma(a + 1, \lambda)$ random variables, independent of the sequence of decreasing uniform order statistics $(U_n^{\downarrow}(1), \ldots, U_n^{\downarrow}(n))$. Alternatively, jointly for $k = 1, \ldots, n$

$$\begin{aligned}
X_n^{rev}[1] &= \gamma_1 \beta_{an,1} \\
X_n^{rev}[2] &= \gamma_2 \beta_{an,1} \beta_{an-a,1} \\
&\cdots \\
X_n^{rev}[n-1] &= \gamma_{n-1} \beta_{an,1} \beta_{an-a,1} \cdots \beta_{2a,1} \\
X_n^{rev}[n] &= \gamma_n \beta_{an,1} \beta_{an-a,1} \cdots \beta_{a,1}.
\end{aligned}$$

where the $\beta_{an-ia,1}$ for $i = 0, \ldots, n-1$ are distributed as beta$(an - ia, 1)$, and they are independent of each other and the γ_k.

Proof. The distribution G_u is computed in the same way as in Example 5.9 and (5.14) follows readily from Proposition 5.7. □

A direct comparison of the two different representations in Proposition 5.3 and 5.10 creates n distributional identities. For example, the equality $X_n[1] = X_n^{rev}[n]$ shows that the following two means of creating a product of independent random variables produce the same result in law:

$$\beta_{a+1,(n-1)a}\, \gamma_{an,\lambda} \stackrel{d}{=} \beta_{an,1}\, \gamma_{a+1,\lambda} \tag{5.15}$$

where $\gamma_{r,\lambda}$ and $\beta_{a,b}$ denote random variables with distributions $gamma(r, \lambda)$ and $beta(a, b)$, respectively. Indeed, this identity comes from the usual 'beta-gamma' algebra, which allows us to write

$$\beta_{a+1,(n-1)a} = \frac{\gamma_{a+1,\lambda}}{\gamma_{a+1,\lambda} + \gamma_{a(n-1),\lambda}}, \quad \beta_{an,1} = \frac{\gamma_{an,\lambda}}{\gamma_{an,\lambda} + \gamma_{1,\lambda}}$$

for $\gamma_{a(n-1),\lambda}, \gamma_{1,\lambda}$ independent of all others. Thus (5.15) reduces to

$$\gamma_{a+1,\lambda} + \gamma_{a(n-1),\lambda} \stackrel{d}{=} \gamma_{an,\lambda} + \gamma_{1,\lambda},$$

which is true since both sides have distribution $gamma(an + 1, \lambda)$.

5.4 Limit in distributions of finite size-biased permutations

As hinted in the introduction, our setup is a finite version of the size-biased permutation of ranked jumps of a subordinator, studied in [76]. In this section we make this statement rigorous.

Let $\Delta = \{x = (x(1), x(2), \ldots) : x(i) \geq 0, \sum_i x(i) \leq 1\}$ and $\Delta^{\downarrow} = \{x^{\downarrow} : x \in \Delta\}$ be closed infinite simplices, the later contains sequences with non-increasing terms. Denote their boundaries by $\partial\Delta = \{x \in \Delta : \sum_i x(i) = 1\}$ and $\partial\Delta^{\downarrow} = \{x \in \Delta^{\downarrow}, \sum_i x(i) = 1\}$ respectively. Any finite sequence can be viewed as an element of $\partial\Delta$ after being normalized by its sum and extended with zeros. Thus one can speak of convergence in distribution of sequences in Δ.

We have to consider Δ and not just $\partial\Delta$ because a sequence in $\partial\Delta$ can converge to one in Δ. For example, the sequence $(X_n, n \geq 1) \in \partial\Delta$ with $X_n(i) = 1/n$ for all $i = 1, \ldots, n$ converges to the elementwise zero sequence in Δ. Thus, we need to define convergence in distribution of size-biased permutations in Δ. We shall do this using Kingman's paintbox. In particular, with this definition, convergence of size-biased permutation is equivalent to

convergence of order statistics. Our treatment in Section 5.4.1 follows that of Gnedin [41] with simplified assumptions.

It then follows that size-biased permutation of finite i.i.d sequences with almost sure finite sum converges to the size-biased permutation of the sequence of ranked jumps of a *subordinator*, roughly speaking, a non-decreasing process with independent and homogeneous increments. We give a review of Lévy processes and subordinators in Section 5.4.2. Many properties such as stick-breaking and the Markov property of the remaining sum have analogues in the limit. We explore these in Section 5.4.3.

5.4.1 Kingman's paintbox and some convergence theorems

Kingman's paintbox [59] is a useful way to describe and extend size-biased permutations. For $x \in \Delta$, let s_k be the sum of the first k terms. Note that x defines a partition $\varphi(x)$ of the unit interval $[0,1]$, consisting of *components* which are intervals of the form $[s_k, s_{k+1})$ for $k = 1, 2, \ldots$, and the interval $[s_\infty, 1]$, which we call the *zero component*. Sample points $\xi_1, \xi_2, \ldots$ one by one from the uniform distribution on $[0,1]$. Each time a sample point discovers a new component that is not in $[s_\infty, 1]$, write down its size. If the sample point discovers a new point of $[s_\infty, 1]$, write 0. Let $X^* = (X^*(1), X^*(2), \ldots)$ be the random sequence of sizes. Since the probability of discovery of a particular (non-zero) component is proportional to its length, the non-zero terms in X^* form the size-biased permutation of the non-zero terms in x as defined by (5.1). In Kingman's paintbox terminology, the components correspond to different colors used to paint the balls with labels $1, 2, \ldots$. Two balls i, j have the same paint color if and only if ξ_i and ξ_j fall in the same component. The size-biased permutation X^* records the size of the newly discovered components, or paint colors. The zero component represents a continuum of distinct paint colors, each of which can be represented at most once.

It follows that the size-biased permutation of a sequence x does not depend on the ordering of its terms. In particular, convergence in $\Delta^\downarrow$ implies convergence in distribution of the corresponding sequences of size-biased permutations.

Theorem 5.11 (Convergence of order statistics implies convergence of s.b.p. [41])**.** *Suppose $X^\downarrow$, $X_1^\downarrow, X_2^\downarrow, \ldots$ are random elements of $\Delta^\downarrow$ and $X_n^\downarrow \xrightarrow{f.d.d} X^\downarrow$. Then $(X_n^\downarrow)^* \xrightarrow{f.d.d} (X^\downarrow)^*$.*

Proof. The following proof is due to Gnedin [41, Theorem 1]. By the Skorokhod representation theorem, we can reduce the statement to the case $X_n^\downarrow$ converges almost surely. Let $S_n^\downarrow(k), S^\downarrow(k)$ be the sum of the k terms in $X_n^\downarrow$ and $X^\downarrow$, respectively. Then $S_n^\downarrow(k) \xrightarrow{a.s.} S^\downarrow(k)$ for all k. Generate the size-biased permutations from the same sequence of sample points $\xi_1, \xi_2, \ldots$ in Kingman's paintbox. Thus if ξ_i discovers a new interval I of the partition $\varphi(X^\downarrow)$, for large n, ξ_i also discovers a new interval I_n of $\varphi(X_n^\downarrow)$, whose size tends to I as $n \to \infty$ almost surely. If I is the zero component of $\varphi(X^\downarrow)$, then the sequence of intervals I_n either are the the zero components of $\varphi(X_n^\downarrow)$ for large n, or their length tends to 0. In either cases, the corresponding terms in $X^*(n)$ tend to 0. Therefore, $(X_n^\downarrow)^* \xrightarrow{f.d.d} (X^\downarrow)^*$. □

We can speak of convergence of size-biased permutation on Δ without having to pass to order statistics. However, convergence of random elements in Δ implies neither convergence of order statistics nor of their size-biased permutations. To achieve convergence we need to keep track of the sum of components. This prevents large order statistics from 'drifting' to infinity.

Lemma 5.12. *Suppose $x, x_1, x_2, \ldots \in \Delta$, $x_n \to x$, $\sum_i x_n(i) \to \sum_i x(i)$. Then $x_n^{\downarrow} \to x^{\downarrow}$.*

Proof. We reproduce the proof in [41]. Fix k and assume $x^{\downarrow}(k) > 0$. Let s, s_n be the sequence of partial sums of x, x_n, $s(m) = \sum_{i=1}^{m} x(i)$, with the convention $s(\infty) = \sum_i x(i)$. Choose m sufficiently large such that $x_1^{\downarrow}, \ldots, x_k^{\downarrow}$ are among the first m terms $x(1), \ldots, x(m)$, and $s(\infty) - s(m) = \sum_{i=m+1}^{\infty} x(i) < x^{\downarrow}(k)/2$. As $s_n(m) \to s(m)$ and $s_n(\infty) \to s(\infty)$, the corresponding tail sum $s_n(\infty) - s_n(m)$ also converge to $s(\infty) - s(m)$. Thus for large n, the k largest values of $x(n)$ are among its first m terms. Since $(x_n(1), \ldots, x_n(m))$ converges to $(x(1), \ldots, x(m))$, the corresponding order statistics converge, in particular, the largest k. So the first k terms of $x_n^{\downarrow}$ converges to the first k terms of $x^{\downarrow}$. □

Using Lemma 5.12, we can upgrade Theorem 5.11 to the following.

Theorem 5.13 ([41]). *Suppose $X, X_1, X_2, \ldots$ are random elements of Δ and*

$$\left(X_n, \sum_i X_n(i)\right) \stackrel{f.d.d}{\to} \left(X, \sum_i X(i)\right).$$

Then $X_n^* \stackrel{f.d.d}{\to} X^*$ and $X_n^{\downarrow} \stackrel{f.d.d}{\to} X^{\downarrow}$.

Theorem 5.14 (Convergence of s.b.p implies convergence of order statistics [41]). *Suppose $X_1, X_2, \ldots$ are random elements of Δ and $X_n^* \stackrel{f.d.d}{\to} Y$ for some $Y \in \Delta$. Then $Y^* \stackrel{d}{=} Y$, and $X_n^{\downarrow} \stackrel{f.d.d}{\to} Y^{\downarrow}$.*

Proof. If the components of X_n, Y are element-wise positive, $\sum_i X_n(i) = \sum_i Y(i) = 1$ for all n, then one has a short proof using the induced order statistics. One generates $(X_n^*)^*, Y^*$ from the same sequence of i.i.d standard exponentials ϵ in Proposition 5.5. Since $X_n^* \stackrel{f.d.d}{\to} Y$, $(X_n^*)^* \stackrel{f.d.d}{\to} Y^*$. But $(X_n^*)^* \stackrel{d}{=} X_n^*$, so $Y^* \stackrel{d}{=} Y$. The second statement comes from Theorem 5.13. See [41] for an analytic proof of the general case. □

5.4.2 Finite i.i.d size-biased permutation and ranked jumps of a subordinator

Let $(X_n), n \geq 1)$ be an i.i.d positive triangular array, that is, $X_n = (X_n(1), \ldots, X_n(n))$, where $X_n(i)$ are i.i.d and a.s. positive. Write T_n for $\sum_{i=1}^{n} X_n(i)$. We ask for conditions

under which the size-biased permutation of the sequence $(X_n, n \geq 1)$ converges to the size-biased permutation of some infinite sequence X. Since we can only make sense of size-biased permutation of sequences with finite sum, let us restrict to the case $T_n \xrightarrow{d} T$ for some $T < \infty$ a.s. A classical result in probability states that $T_n \xrightarrow{d} T$ if and only if $T = \tilde{T}(1)$ for some *Lévy process* $\tilde{T}$, which in this case is a *subordinator*. For self-containment, we gather some necessary facts about Lévy processes and subordinators below. See [58, §15] for their proofs, [11] for a thorough treatment of subordinators.

Definition 5.15. A *Lévy process* $\tilde{T}$ in $\mathbb{R}$ is a stochastic process with right-continuous left-limits paths, stationary independent increments, and $\tilde{T}(0) = 0$. A *subordinator* $\tilde{T}$ is a Lévy process, with real, finite, non-negative increments.

Following [58], we do not allow the increments to have infinite value. We suffer no loss of generality, since subordinators with jumps of possibility infinite size do not contribute to our discussion of size-biased permutation. Let $\tilde{T}$ be a subordinator, $T = \tilde{T}(1)$. For $t, \lambda \geq 0$, using the fact that increments are stationary and independent, one can show that

$$\mathbb{E}(\exp(-\lambda \tilde{T}(t))) = \exp(-t\Phi(\lambda)),$$

where the function $\Phi : [0, \infty) \to [0, \infty)$ is called the *Laplace exponent* of $\tilde{T}$. It satisfies the *Lévy-Khinchine formula*

$$\Phi(\lambda) = \mathsf{d}\,\lambda + \int_0^\infty (1 - e^{-\lambda x})\Lambda(dx), \quad \lambda \geq 0,$$

where $\mathsf{d} > 0$ is the *drift coefficient*, and Λ a unique measure on $(0, \infty)$ with $\Lambda([1, \infty)) < \infty$, called the *Lévy measure of* $\tilde{T}$. Assume $\int_0^1 x\,\Lambda(dx) < \infty$, which implies a.s. $\tilde{T}(1) = T < \infty$. Then over $[0, 1]$, $\tilde{T}$ can be decomposed as the sum of a deterministic drift plus a Poisson point process with i.i.d jumps

$$(\tilde{T})(t) = \mathsf{d}\,t + \sum_i X(i)\mathbf{1}_{\{\sigma(i) \leq t\}}$$

for $0 \leq t \leq 1$, where $\{(\sigma(i), X(i)), i \geq 1\}$ is the countable random set of points of a Poisson point process (p.p.p) on $(0, \infty)^2$ with intensity measure $dt\,\Lambda(dx)$. The $X(i)$ are the *jumps* of the subordinator $\tilde{T}$. We say that the sequence $X = (X(i) : \sigma(i) \leq 1)$ are the jumps of T. Finally, we need a classical result on convergence of i.i.d positive triangular arrays to subordinators (see [58, §15]).

Theorem 5.16. *Let $(X(n), n \geq 1)$ be an i.i.d positive triangular array, $T_n = \sum_{i=1}^n X_n(i)$. Then $T_n \xrightarrow{d} T$ for some random variable T, $T < \infty$ a.s. if and only if $T = \tilde{T}(1)$ for some subordinator $\tilde{T}$ whose Lévy measure Λ satisfies $\int_0^1 x\,\Lambda(dx) < \infty$. Furthermore, let μ_n be the*

measure of $X_n(i)$. Then on $\mathbb{R}_+$, the sequence of measures $(n\mu_n)$ converges vaguely to Λ, written

$$n\mu_n \xrightarrow{v} \Lambda.$$

That is, for all $f : \mathbb{R}_+ \to \mathbb{R}_+$ continuous with compact support, $n\mu_n(f) = n\int_0^\infty f(x)\,\mu_n(dx)$ converges to $\Lambda(f) = \int_0^\infty f(x)\,\Lambda(dx)$. In particular, if μ_n, Λ have densities ρ_n, ρ, respectively, then we have pointwise convergence for all $x > 0$

$$n\rho_n(x) \to \rho(x).$$

Proposition 5.17. *Let $(X_n, n \geq 1)$ be an i.i.d positive triangular array, $T_n = \sum_{i=1}^n X_n(i)$. Suppose $T_n \xrightarrow{d} T$ for some T a.s. finite. Let X be the sequence of ranked jumps of T arranged in any order, $(X/T)^*$ be the size-biased permutation of the sequence (X/T) as defined using Kingman's paintbox, $(X^*)' = T \cdot (X/T)^*$. Then*

$$(X_n)^* \xrightarrow{f.d.d} (X^*)'$$

Proof. The sequence of decreasing order statistics $X_n^\downarrow$ converges in distribution to $X^\downarrow$ [58]. Since $T_n, T > 0$ a.s. and $T_n \xrightarrow{d} T$, $X_n^\downarrow/T_n \xrightarrow{f.d.d} X^\downarrow/T$. Theorem 5.11 combined with multiplying through by T prove the claim. □

For subordinators without drift, $\mathsf{d} = 0$, $\sum_i X(i) = T$, hence $(X^*)' = X^*$. When $\mathsf{d} > 0$, the sum of the jumps $\sum_i X_i$ is strictly less than T, so $(X^*)' \neq X^*$. In this case, there is a non-trivial zero component at infinity, which comes from an accumulation of mass at 0 of $n\mu_n$ in the limit as $n \to \infty$. At each finite, large n, we have a significant number of jumps with 'microscopic' size.

The case without drift was studied by Perman, Pitman and Yor in [76] with the assumption $\Lambda(0, \infty) = \infty$ to ensure that the sequence of jumps has infinite length. We shall re-derive some of their results as limits of results for finite i.i.d size-biased permutation using Theorem 5.16 in the next section. One can obtain another finite version of the Perman-Pitman-Yor setup by setting $\Lambda(0, \infty) < \infty$, but this can be reduced to finite i.i.d size-biased permutation by conditioning. Specifically, $\tilde{T}$ is now a compound Poisson process, where the subordinator waits for an exponential time with rate $\Lambda(0, \infty)$ before making a jump X, whose length is independent of the waiting time and distributed as $\mathbb{P}(X \leq t) = \Lambda(0, t]/\Lambda(0, \infty)$ [11]. If $(X(1), X(2), \ldots)$ is the sequence of successive jumps of $(\tilde{T}_s, s \geq 0)$, then $(X(1), X(2), \ldots, X(N))$ is the sequence of successive jumps of $(\tilde{T}_s, 0 \leq s \leq 1)$, where N is a Poisson random variable with mean $\Lambda(0, \infty)$, independent of the jump sequence $(X(1), X(2), \ldots)$. For $N > 0$, properties of the size-biased permutation of $(X(1), \ldots, X(N))$ can be deduced from those of a finite i.i.d size-biased permutation by conditioning on N.

5.4.3 Markov property in the limit

We now indicate how some results in [76] can be obtained as limits of those in Section 5.2, including the Markov property and the stick-breaking representation. Consider a subordinator with Lévy measure Λ, drift $\mathsf{d} = 0$. Let $\tilde{T}_0$ be the subordinator at time 1. Assume $\Lambda(1,\infty) < \infty$, $\Lambda(0,\infty) = \infty$, $\int_0^1 x\Lambda(dx) < \infty$, and $\Lambda(dx) = \rho(x)\,dx$ for some density ρ. Note that $\tilde{T}_0 < \infty$ a.s, and it has a density determined by ρ via its Laplace transform, which we denote ν. Let $\tilde{T}_k$ denote the remaining sum after removing the first k terms of the size-biased permutation of the sequence $X^{\downarrow}$ of ranked jumps.

Proposition 5.18 ([76])**.** *The sequence $(\tilde{T}_0, \tilde{T}_1, \ldots)$ is a Markov chain with stationary transition probabilities*

$$\mathbb{P}(\tilde{T}_1 \in dt_1 | \tilde{T}_0 = t) = \frac{t - t_1}{t} \cdot \rho(t - t_1)\,\frac{\nu(t_1)}{\nu(t)}\,dt_1.$$

Note the similarity to (5.5). Starting with (5.4) and send $n \to \infty$, for any finite k, we have $\nu_{n-k} \to \nu$ pointwise, and by Theorem 5.16, $(n-k)\nu_1 \to \rho$ pointwise over $\mathbb{R}$, since there is no drift term. Thus the analogue of (5.4) in the limit is

$$\mathbb{P}(\tilde{T}_0 \in dt_0, \ldots, \tilde{T}_k \in dt_k) = \left(\prod_{i=0}^{k-1} \frac{t_i - t_{i+1}}{t_i}\rho(t_i - t_{i+1})\right)\nu(t_k)\,dt_0 \ldots dt_k.$$

Rearranging gives the transition probability in Proposition 5.18.

Conditionally given $\tilde{T}_0 = t_0, \tilde{T}_1 = t_1, \ldots, \tilde{T}_n = t_n$, the sequence of remaining terms in the size-biased permutation $(X[n+1], X[n+2], \ldots)$ is distributed as $(X^{\downarrow}(1), X^{\downarrow}(2), \ldots,)$ conditioned on $\sum_{i\geq 1} X^{\downarrow}(i) = t_n$, independent of the first n size-biased picks [76, Theorem 4.2]. The stick-breaking representation in (5.6) now takes the form

$$\frac{X[k]}{\tilde{T}_0} = W_k \prod_{i=1}^{k-1} \overline{W}_i, \tag{5.16}$$

where $X[k]$ is the kth size-biased pick, and $W_i = \frac{X[i]}{\tilde{T}_{i-1}}$, $\overline{W}_i = 1 - W_i = \frac{\tilde{T}_i}{\tilde{T}_{i-1}}$. Proposition 5.3 and Corollary 5.4 parallel the following result.

Proposition 5.19 (McCloskey [67] and Perman-Pitman-Yor [76])**.** *The random variables $\tilde{T}_0$ and $W_1, W_2, \ldots$ in (5.16) are mutually independent if and only if $\tilde{T}_0$ has distribution gamma(a, λ) for some $a, \lambda > 0$. In this case, the W_i are i.i.d with distribution beta$(1, a)$ for $i = 1, 2, \ldots$.*

5.4.4 Invariance under size-biased permutation

We take a small detour to explain some results related to Propositions 5.3 and 5.19 on characterization of size-biased permutations. For a random discrete distribution prescribed by its probability mass function $P \in \partial\Delta$, let P^* be its size-biased permutation. (Recall that Δ is the closed infinite simplex, $\partial\Delta$ is its boundary. These are defined at the beginning of Section 5.4). Given $Q, P \in \partial\Delta$, one may ask when is $Q = P^*$. Clearly $(P^*)^* = P^*$ for any $P \in \partial\Delta$, thus this question is equivalent to characterizing random discrete distributions on $\mathbb{N}$ which are invariant under size-biased permutation (ISBP). Pitman [77, Theorem 4] gave a complete answer in terms of symmetry of a certain function of the finite dimensional distributions. Furthermore, Pitman proved a complete characterization of ISBP when P_k can be written as the right hand side of (5.16) with $W_1, W_2, \ldots$ independent.

Theorem 5.20 ([77]). *Let $P \in \partial\Delta$, $P(1) < 1$. and $P(n) = W_1 \cdots W_{n-1}\overline{W}_n$ for independent W_i. Then $P = P^*$ if and only if one of the four following conditions holds.*

1. *$P(n) \geq 0$ a.s. for all n, in which case the distribution of W_n is*

$$beta(1-\alpha, \theta + n\alpha)$$

   for every $n = 1, 2, \ldots$, for some $0 \leq \alpha < 1$, $\theta > -\alpha$.

2. *For some integer constant m, $P(n) \geq 0$ a.s. for all $1 \leq n \leq m$, and $P(n) = 0$ a.s. otherwise. Then either*

   (a) *For some $\alpha > 0$, W_n has distribution $beta(1+\alpha, m\alpha - n\alpha)$ for $n = 1, \ldots, m$; or*

   (b) *$W_n = 1/(m-n+1)$ a.s., that is, $P(n) = 1/m$ a.s. for $n = 1, \ldots, m$; or*

   (c) *$m = 2$, and the distribution F on $(0,1)$ defined by $F(dw) = \bar{w}\mathbb{P}(W_1 \in dw)/\mathbb{E}(\bar{W}_1)$ is symmetric about $1/2$.*

The McCloskey case of Proposition 5.19 is case 1 with $\alpha = 0, \theta > 0$, and Patil-Taillie case of Proposition 5.3 is case 2(a). These two cases often re-written in the form W_i has distribution $beta(1-\alpha, \theta + i\alpha)$, $i = 1, 2, \ldots$ for pairs of real numbers (α, θ) satisfying either $(0 \leq \alpha < 1, \theta > -\alpha)$ (case 1), or $(\alpha < 0, \theta = m\alpha)$ for some $m = 1, 2, \ldots$ (case 2(a)). In both settings, such a distribution P is known as the $GEM(\alpha, \theta)$ distribution. The abbreviation GEM was introduced by Ewens, which stands for Griffiths-Engen-McCloskey. If P is $GEM(\alpha, \theta)$ then $P^{\downarrow}$ is called a Poisson-Dirichlet distribution with parameters (α, θ) [76]. They have applications in Bayesian statistics and machine learning, see [78].

5.5 Asymptotics of the last u fraction of the size-biased permutation

In this section we derive Glivenko-Cantelli and Donsker-type theorems for the distribution of the last u fraction of terms in a finite i.i.d size-biased permutation. It is especially convenient to work with the induced order statistics representation since we can appeal to tools from empirical process theory. In particular, our results are special cases of more general statements which hold for arbitrary induced order statistics in d dimensions (see Section 5.5.1). Features pertaining to i.i.d size-biased permutation are presented in Lemma 5.21. The proof is a direct computation. We first discuss the interesting successive sampling interpretation of Lemma 5.21, quoting some results needed to make the discussion rigorous. We then derive the aforementioned theorems and conclude with a brief historical account of induced order statistics.

Lemma 5.21. *Suppose F has support on $[0,\infty)$ and finite mean. For $u \in (0,1)$, define*

$$F_u(dx) = \frac{e^{-x\phi^{-1}(u)}}{u} F(dx) \tag{5.17}$$

and extend the definition to $\{0,1\}$ by continuity, where ϕ is the Laplace transform of F as in Proposition 5.7. Then F_u is a probability distribution on $[0,\infty)$ for all $u \in [0,1]$, and G_u in (5.8) satisfies

$$G_u(dx) = xF_u(dx)/\mu_u \tag{5.18}$$

where $\mu_u = \int xF_u(dx) = \frac{-\phi'(\phi^{-1}(u))}{u}$. Furthermore,

$$\int_0^u G_s(dx)\, ds = F_u(dx) \tag{5.19}$$

for all $s \in [0,1]$. In other words, the density

$$f(u,x) = F_u(dx)/F(dx) = u^{-1}e^{-x\phi^{-1}(u)}$$

of F_u with respect to F solves the differential equation

$$\frac{d}{du}[uf(u,x)] = \frac{-xf(u,x)}{\mu_u} \tag{5.20}$$

with boundary condition $f(1,x) \equiv 1$.

For any distribution F with finite mean μ and positive support, $xF(dx)/\mu$ defines its *size-biased distribution*. If F is the empirical distribution of n positive values $x_n(1), \ldots, x_n(n)$, for example, one can check that $xF(dx)/\mu$ is precisely the distribution of the first size-biased

pick $X_n[1]$. For continuous F, the name size-biased distribution is justified by the following lemma.

Lemma 5.22. *Consider an i.i.d size-biased permutation $(X_n[1], \ldots, X_n[n])$ from a distribution F with support on $[0, \infty)$ and finite mean μ. Then*

$$\lim_{n\to\infty} \mathbb{P}(X_n[1] \in dx) = \frac{xF(dx)}{\mu}.$$

Since the $\lfloor nu \rfloor$-th smallest out of n uniform order statistics converge to u as $n \to \infty$, G_u is the limiting distribution of $X_n^{rev}[\lfloor nu \rfloor]$, the size-biased pick performed when a u-fraction of the sequence is left. By (5.18), G_u is the size-biased distribution of F_u. Thus, F_u can be interpreted as the limiting distribution of the remaining u-fraction of terms in a successive sampling scheme. This intuition is made rigorous by Corollary 5.23 below. In words, it states that F_u is the limit of the empirical distribution function (e.d.f) of the last $u > 0$ fraction in a finite i.i.d size-biased permutation.

Corollary 5.23. *For $u \in (0, 1]$, let $F_{n,u}(\cdot)$ denote the empirical distribution of the last $\lfloor nu \rfloor$ values of an i.i.d size-biased permutation with length n. For each $\delta \in (0, 1)$, as $n \to \infty$,*

$$\sup_{u\in[\delta,1]} \sup_{\mathbf{I}} |F_{n,u}(\mathbf{I}) - F_u(\mathbf{I})| \xrightarrow{a.s.} 0, \tag{5.21}$$

where $\mathbf{I}$ ranges over all subintervals of $(0, \infty)$.

Therefore in the limit, after removing the first $1 - u$ fraction of terms in the size-biased permutation, we are left with an (infinitely) large sequence of numbers distributed like i.i.d draws from F_u, from which we do a size-biased pick, which has distribution $G_u(dx) = xF_u(dx)/\mu_u$ as specified by (5.18).

Since $X_n^{rev}[\lfloor nu \rfloor]$ converges in distribution to G_u for $u \in [0, 1]$, Corollary 5.23 lends a sampling interpretation to Lemma 5.21. Equation (5.20) has the heuristic interpretation as characterizing the evolution of the mass at x over time u in a successive sampling scheme. To be specific, consider a successive sampling scheme on a large population of N individuals, with species size distribution H. Scale time such that at time u, for $0 \le u \le 1$, there are Nu individuals (from various species) remaining to be sampled. Let H_u denote the distribution of species sizes at time u, and fix the bin $(x, x + dx)$ of width dx on $(0, \infty)$. Then $NuH_u(dx)$ is the number of individuals whose species size lie in the range $(x, x + dx)$ at time u. Thus $\frac{d}{du}NuH_u(dx)$ is the rate of individuals to be sampled from this range of species size at time u. The probability of an individual whose species size is in $(x, x + dx)$ being sampled at time u is $\frac{xH_u(dx)}{\int_0^\infty xH_u(dx)}$. As we scaled time such that $u \in [0, 1]$, in time du we sample $N\,du$ individuals. Thus

$$\frac{d}{du}NuH_u(dx) = -N\,\frac{xH_u(dx)}{\int_0^\infty xH_u(dx)}.$$

Let $f(u,x) = H_u(dx)/H(dx)$, then as a function in u, the above equation reduces to (5.20).

Example 5.24. Suppose F puts probability p at a and $1-p$ at b, with $a < b$. Let $p(u) = F_u(a)$ be the limiting fraction of a left when proportion u of the sample is left. Then the evolution equation (5.20) becomes

$$\begin{aligned} p'(u) &= u^{-1}\left(\frac{a}{ap(u)+b(1-p(u))} - 1\right)p(u) \\ &= u^{-1}\left(\frac{a-ap(u)-b(1-p(u))}{b-(b-a)p(u)}\right)p(u) \\ &= u^{-1}\frac{(1-p(u))p(u)(a-b)}{b-(b-a)p(u)} \quad ,0 \le u \le 1, \end{aligned}$$

with boundary condition $p(0) = p$. To solve for $p(u)$, let y solve $u = pe^{-ay} + (1-p)e^{-by}$. Then $p(u) = pe^{-ay}/u$.

A Glivenko-Cantelli theorem

We now state a Glivenko-Cantelli-type theorem which applies to size-biased permutations of finite deterministic sequences. Versions of this result are known in the literature [13,43,51,84], see discussions in Section 5.5.1. We offer an alternative proof using induced order statistics.

Theorem 5.25. *Let $(x_n, n = 1,2,\ldots)$ be a deterministic triangular array of positive numbers with corresponding c.d.f sequence $(E_n, 1 \le n)$. Suppose*

$$\sup_x |E_n(x) - F(x)| \to 0 \text{ as } n \to \infty \tag{5.22}$$

for some distribution F on $(0,\infty)$. Let $u \in (0,1]$. Let $E_{n,u}(\cdot)$ be the empirical distribution of the first $\lfloor nu \rfloor$ terms in a size-biased permutation of the sequence x_n. Then for each $\delta \in (0,1)$,

$$\sup_{u\in[\delta,1]} \sup_{\mathbf{I}} |E_{n,u}(\mathbf{I}) - F_u(\mathbf{I})| \xrightarrow{a.s.} 0 \text{ as } n \to \infty, \tag{5.23}$$

where $\mathbf{I}$ ranges over all subintervals of $(0,\infty)$.

Proof. Define $Y_n(i) = \epsilon_n(i)/x_n(i)$ for $i = 1,\ldots,n$ where $\epsilon_n(i)$ are i.i.d standard exponentials as in Proposition 5.5. Let H_n be the empirical distribution function (e.d.f) of the $Y_n(i)$,

$$H_n(y) := \frac{1}{n}\sum_{i=1}^{n} \mathbf{1}_{\{Y_n(i)<y\}}.$$

Let J_n denote the e.d.f of $(x_n(i), Y_n(i))$. By Proposition 5.7,

$$E_{n,u}(\mathbf{I}) = \frac{n}{\lfloor nu \rfloor} \frac{1}{n} \sum_{i=1}^{n} \mathbf{1}_{\{x_n(i) \in \mathbf{I}\}} \mathbf{1}_{\{Y_n(i) < H_n^{-1}(1-u)\}} = \frac{n}{\lfloor nu \rfloor} J_n(\mathbf{I} \times [0, H_n^{-1}(1-u)]). \tag{5.24}$$

Fix $\delta \in (0,1)$, and let $u \in [\delta, 1]$. Let ϕ be the Laplace transform of F and J the joint law of $(X, \epsilon/X)$, where X is a random variable with distribution F, and ϵ is an independent standard exponential. Note that $\frac{1}{u} J(\mathbf{I} \times [0, \phi^{-1}(u)]) = F_u(\mathbf{I})$. Thus

$$\begin{aligned} \mathbb{E}_{n,u}(\mathbf{I}) - F_u(\mathbf{I}) = &\left(\frac{n}{\lfloor nu \rfloor} J_n(\mathbf{I} \times [0, H_n^{-1}(1-u)]) - \frac{n}{\lfloor nu \rfloor} J_n(\mathbf{I} \times [0, \phi^{-1}(u)]) \right) \\ &+ \left(\frac{n}{\lfloor nu \rfloor} J_n(\mathbf{I} \times [0, \phi^{-1}(u)]) - \frac{1}{u} J(\mathbf{I} \times [0, \phi^{-1}(u)]) \right). \end{aligned} \tag{5.25}$$

Let us consider the second term. Note that

$$J_n(\mathbf{I} \times [0, \phi^{-1}(u)]) = \int_0^\infty e^{-t\phi^{-1}(u)} \mathbf{1}_{\{t \in \mathbf{I}\}} \, E_n(dt).$$

Since E_n converges to F uniformly and $e^{-t\phi^{-1}(u)}$ is bounded for all $t \in (0, \infty)$ and $u \in [\delta, 1]$,

$$\sup_{u \in [\delta,1]} \sup_{\mathbf{I}} \left| \frac{n}{\lfloor nu \rfloor} J_n(\mathbf{I} \times [0, \phi^{-1}(u)]) - \frac{1}{u} J(\mathbf{I} \times [0, \phi^{-1}(u)]) \right| \xrightarrow{a.s.} 0 \text{ as } n \to \infty.$$

Let us consider the first term. Since J_n is continuous in the second variable, it is sufficient to show that

$$\sup_{u \in [\delta,1]} |H_n^{-1}(1-u) - \phi^{-1}(u)| \xrightarrow{a.s.} 0 \text{ as } n \to \infty. \tag{5.26}$$

To achieve this, let A_n denote the 'average' measure

$$A_n(y) := \frac{1}{n} \sum_{i=1}^{n} \mathbb{P}(Y_n(i) < y) = 1 - \int_0^\infty e^{-xy} \, dE_n(x).$$

A theorem of Wellner [106, Theorem 1] states that if the sequence of measures $(A_n, n \geq 1)$ is tight, then the Prohorov distance between H_n and A_n converges a.s. to 0 and $n \to \infty$. In this case, since E_n converges to F uniformly, A_n converges uniformly to $1 - \phi$. Thus H_n converges uniformly to $1 - \phi$, and (5.26) follows. □

Proof of Corollary 5.23. When E_n is the e.d.f of n i.i.d picks from F, then (5.22) is satisfied a.s. by the Glivenko-Cantelli theorem. Thus Theorem 5.25 implies Corollary 5.23. □

Proof of Lemma 5.22. Let ϕ be the Laplace transform of F. For $y > 0$,

$$\frac{d\phi(y)^n}{dy} = n\phi(y)^{n-1}\phi'(y).$$

By Corollary 5.8, we have

$$\frac{\mathbb{P}(X_n[1] \in dx)}{xF(dx)} = n\int_0^\infty e^{-xy}\phi(y)^{n-1}dy = \int_0^\infty \frac{e^{-xy}}{\phi'(y)}\frac{d}{dy}(\phi(y)^n)\,dy.$$

Apply integration by parts, the constant term is

$$\frac{e^{-xy}}{\phi'(y)}\phi(y)^n\Big|_0^\infty = -\frac{1}{\phi'(0)} = \frac{1}{\mu}.$$

The integral term is

$$\int_0^\infty \frac{d}{dy}(e^{-xy}\phi'(y))(\phi(y))^n\,dy.$$

The integrand is integrable for all n, thus

$$\lim_{n\to\infty}\int_0^\infty \frac{d}{dy}(e^{-xy}\phi'(y))(\phi(y))^n\,dy = \int_0^\infty \lim_{n\to\infty}\frac{d}{dy}(e^{-xy}\phi'(y))(\phi(y))^n\,dy = 0$$

Since $n \to \infty$, $\phi(y)^n \to 0$ for all $y > 0$. Therefore, $\lim_{n\to\infty}\dfrac{\mathbb{P}(X_n[1] \in dx)}{xF(dx)} = \dfrac{1}{\mu}$. □

The functional central limit theorem

The functional central limit theorem for i.i.d size-biased permutation is a special case of the general result of Davydov and Egorov [25] for induced order statistics.

Theorem 5.26 (Davydov and Egorov, [25])**.** *Suppose the first two moments of F are finite. For a distribution H, let $\mu(H)$, $\sigma(H)$ denote its mean and standard deviation. Let $X_n = (X_n(1), \ldots, X_n(n))$ sequence of i.i.d draws from F, ϵ_n be an independent sequence of n i.i.d standard exponentials, $Y_n = (\epsilon_n(i)/X_n(i), i = 1, \ldots, n)$. For $u \in (0, 1]$, define*

$$\xi_n(u) = \frac{1}{n}\sum_{j=1}^{\lfloor nu\rfloor} X_n^{rev}[j]$$

$$\eta_n(u) = \frac{1}{n}\sum_{j=1}^{n} X_n(j)\mathbf{1}_{\{Y_n(j)<\phi^{-1}(u)\}}$$

$$m(u) = \int_0^u \mu(G_s)\,ds = \mu(F_u).$$

Then $\xi_n \to m$, $\eta_n \to m$ *uniformly on* $[0,1]$, *and*

$$\sqrt{n}(\xi_n - m) \Rightarrow \alpha, \qquad \sqrt{n}\,(\eta_n - m) \Rightarrow \eta$$

in the Skorokhod topology, where

$$\alpha(u) = \int_0^u \sigma(G_s) dW(s), \quad \eta(u) = \alpha(u) + \int_0^u m(s)\, dV(s).$$

W *is a standard Brownian motion, and* V *is a Brownian bridge, independent of* W.

Note that ξ_n is a rescaled version of the sum of the last u-fraction of terms in a size-biased permutation. The difference in ξ_n and η_n is the fluctuation of empirical uniform order statistics around its limit. The proof of Theorem 5.26 can be found in [25], together with similar statements on the functional law of the iterated logarithm for the processes η_n, ξ_n.

5.5.1 Historical notes on induced order statistics and successive sampling

Induced order statistics were first introduced by David [23] and independently by Bhattacharya [12]. Typical applications stem from modeling an indirect ranking procedure, where subjects are ranked based on their Y-attributes although the real interest lies in ranking their X-attributes, which are difficult to obtain at the moment where the ranking is required.[1] For example in cattle selection, Y may represent the genetic makeup, for which the cattle are selected for breeding, and X represents the milk yields of their female offspring. Thus a portion of this literature focuses on comparing distribution of induced order statistics to that of usual order statistics [24, 42, 71, 108]. The most general statement on asymptotic distributions is obtained by Davydov and Egorov [25], who proved the functional central limit theorem and the functional law of the iterated logarithm for the process $S_{n,u}$ under tight assumptions. Their theorem translates directly into Theorem 5.26 for finite i.i.d size-biased permutation. Various versions of results in Section 5.5, including Theorem 5.26 are also known in the successive sampling community [13, 43, 51, 84, 90]. For example, Bickel, Nair and Wang [13] proved Theorem 5.25 with convergence in probability when E_n and F have the same discrete support on finitely many values.

[1] One often uses X for the variable to be ordered, and Y for the induced variable, with the idea that Y is to be predicted. Here we use X for the induced order statistics since $X_n[k]$ has been used for the size-biased permutation. The role of X and Y in our case is interchangeable, as evident when one writes $X_n(i)Y_n(i) = \epsilon_n(i)$.

5.6 Poisson coupling of size-biased permutation and order statistics

Comparisons between the distribution of induced order statistics and order statistics of the same sequence have been studied in the literature [24, 42, 71, 108]. However, finite i.i.d size-biased permutation has the special feature that there exists an explicit coupling between these two sequences as described in Proposition 5.5. Using this fact, we now derive Theorem 5.28, which gives a Poisson coupling between the *last* k size-biased terms $X_n^{rev}[1], \ldots, X_n^{rev}[k]$ and the k smallest order statistics $X_n^{\uparrow}(1), \ldots X_n^{\uparrow}(k)$ as $n \to \infty$. The existence of a Poisson coupling is not surprising, since the increasing sequence of order statistics $(X_n^{\uparrow}(1), X_n^{\uparrow}(2), \ldots)$ converges to points in a Poisson point process (p.p.p) whose intensity measure depends on the behavior of F near the infimum of its support, which is 0 in our case. This standard result in order statistics and extreme value theory dates back to Rényi [82], and can be found in [27].

5.6.1 Random permutations from Poisson scatter

Let $N(\cdot)$ be a Poisson scatter on $(0, \infty)^2$. Suppose $N(\cdot)$ has intensity measure m such that for all $s, t \in (0, \infty)$

$$m((0, s) \times (0, \infty)) < \infty, m((0, \infty) \times (0, t)) < \infty.$$

Then one obtains a random permutation of $\mathbb{N}$ from ranking points $(x(i), y(i))$ in N according to either the x or y coordinate. Let x^* and y^* denote the induced order statistics of the sequence x and y obtained by ranking points by their y and x values in increasing order, respectively. For $j, k = 1, 2, \ldots$, define sequences of integers $(K_j), (J_k)$ such that $x^{\uparrow}(J_k) = x^*(k)$, $y^{\uparrow}(K_j) = y^*(j)$; see Figure 5.1.

For $j \geq 1$, conditioned on $x(j) = s, y^*(j) = t$,

$$K_j - 1 \stackrel{d}{=} Poisson\left(m((s, \infty) \times (0, t))\right) + Binomial\left(j - 1, \frac{m((0, s) \times (0, t))}{m((0, s) \times (0, \infty))}\right), \quad (5.27)$$

where the two random variables involved are independent. Similarly, for $k \geq 1$, conditioned on $x_k^* = s$, $y(k) = t$,

$$J_k - 1 \stackrel{d}{=} Poisson\left(m((0, s) \times (t, \infty))\right) + Binomial\left(k - 1, \frac{m((0, s) \times (0, t))}{m((0, \infty) \times (t, \infty))}\right), \quad (5.28)$$

where the two random variables involved are independent. When m is a product measure, it is possible to compute the marginal distribution of K_j and J_k explicitly for given $j, k \geq 1$.

Random permutations from Poisson scatters appeared in [76, §4]. When $X^{\downarrow}$ is the sequence of ranked jumps of a subordinator, these authors noted that one can couple the

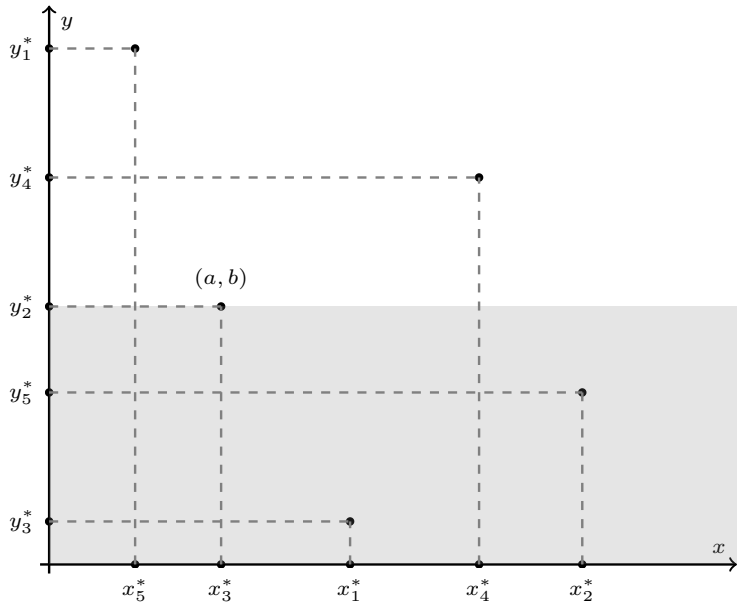

Figure 5.1: A point scatter on the plane. Here $J_1 = 5, J_2 = 3, J_3 = 1, J_4 = 4, J_5 = 2$, and $K_1 = 3, K_2 = 5, K_3 = 2, K_4 = 4, K_5 = 1$. The permutations J and K are inverses. Conditioned on $x^{\uparrow}(2) = a$ and $y^*(2) = b$, the number of points lying in the shaded region determines $J_2 - 1$.

size-biased permutation with the order statistics via the following p.p.p

$$N(\cdot) := \sum_{k \geq 1} \mathbf{1}[(X[k], Y^{\uparrow}(k)) \in \cdot] = \sum_{k \geq 1} \mathbf{1}[(X^{\downarrow}(k), Y(k)) \in \cdot], \tag{5.29}$$

where Y is an independent sequence of standard exponentials. Thus $N(\cdot)$ has measure $m(dxdy) = xe^{-xy}\,\Lambda(dx)dy$. The first expression in (5.29) defines a scatter of (x, y) values in the plane listed in increasing y values, and the second represents the same scatter listed in decreasing x values. Since $\sum_{i \geq 1} X^{\downarrow}(i) < \infty$ a.s., the x-marginal of the points in (5.29) has the distribution of the size-biased permutation X^*, since it prescribes the joint distribution of the first k terms $X[1], \ldots, X[k]$ of X^* for any finite k. Perman, Pitman and Yor used this p.p.p representation to generalize size-biased permutation to h-biased permutation, where the 'size' of a point x is replaced by an arbitrary strictly positive function $h(x)$; see [76, §4].

5.6.2 Poisson coupling in the limit

Our theorem states that in the limit, finite i.i.d size-biased permutation is a form of random permutation obtained from a Poisson scatter with a certain measure, which, under a change of coordinate, is given by (5.35). Before stating the theorem we need some technical results. The distribution of the last few size-biased picks depends on the behavior of F near 0, the

infimum of its support. We shall consider the case where F has 'power law' near 0, like that of a Gamma distribution.

Lemma 5.27. *Suppose F is supported on $(0,\infty)$ with Laplace transform ϕ. Let $u = \phi(y)$, X_u a random variable distributed as $G_u(dx)$ defined in (5.8). For $\lambda, a > 0$,*

$$F(x) \sim \frac{\lambda^a x^a}{\Gamma(a+1)} \text{ as } x \to 0, \tag{5.30}$$

if and only if,

$$\phi(y) \sim \lambda^a / y^a \quad \text{ as } y \to \infty. \tag{5.31}$$

Furthermore, (5.30) implies

$$u^{-1/a} X_u \xrightarrow{d} gamma(a+1, \lambda) \quad \text{ as } u \to 0. \tag{5.32}$$

Proof. The equivalence of (5.30) and (5.31) follows from a version of Karamata Tauberian Theorem [14, §1.7]. Assume (5.30) and (5.31). We shall prove (5.32) by looking at the Laplace transform of the non-size-biased version X'_u, which has distribution F_u. For $\theta \geq 0$,

$$E(\exp(-\theta X'_u)) = \int_0^\infty u^{-1} \exp(-yx - \theta x)\, F(dx) = u^{-1}\phi(y+\theta) = \frac{\phi(y+\theta)}{\phi(y)}. \tag{5.33}$$

Now as $y \to \infty$ and $u = \phi(y) \to 0$, for each fixed $\eta > 0$, (5.31) implies

$$E(\exp(-\eta u^{-1/a} X'_u)) = \frac{\phi(y + \eta\phi(y)^{-1/a})}{\phi(y)} \sim \frac{\lambda^a (y + \eta\lambda^{-1} y)^{-a}}{\lambda^a y^{-a}} = \left(\frac{\lambda}{\lambda+\eta}\right)^a.$$

That is to say

$$u^{-1/a} X'_u \xrightarrow{d} gamma(a, \lambda). \tag{5.34}$$

Since ϕ is differentiable, (5.33) implies $E(X'_u) = \phi'(y)/\phi(y)$. Now ϕ has an increasing derivative ϕ', thus (5.31) implies $\phi'(y) \sim a\lambda^a / y^{a+1}$ as $y \to \infty$. Therefore,

$$u^{-1/a} E(X'_u) = \frac{\phi'(y)}{\phi(y)^{1+1/a}} \to \frac{a}{\lambda},$$

which is the mean of a $gamma(a, \lambda)$ random variable. Thus the random variables $u^{-1/a} X'_u$ are uniformly integrable, so for any bounded continuous function h, we can compute

$$E(h(u^{-1/a} X_u)) = \frac{E[(u^{-1/a} X'_u) h(u^{-1/a} X'_u)]}{u^{-1/a} E(X'_u)} \to \frac{E[\gamma_{a,\lambda} h(\gamma_{a,\lambda})]}{E(\gamma_{a,\lambda})} = E(h(\gamma_{a+1,\lambda}))$$

where $\gamma_{b,\lambda}$ is a $gamma(b,\lambda)$ random variable. This proves (5.32). □

We now present the analogue of (5.42) for the last few size-biased picks $X_n^{rev}[1],\ldots,X_n^{rev}[k]$ and the promised Poisson coupling.

Theorem 5.28. *Suppose that (5.30) holds for some $\lambda, a > 0$. Let $N(\cdot)$ be a Poisson scatter on $(0,\infty)^2$ with intensity measure*

$$\frac{\mu(ds\,dt)}{ds\,dt} = \frac{1}{a}\Gamma(a+1)^{1/a}(s/t)^{1/a}\exp\{-(\Gamma(a+1)s/t)^{1/a}\}. \tag{5.35}$$

By ranking points in either increasing T or S coordinator, one can write

$$N(\cdot) = \sum_k \mathbf{1}[(S(k),T^{\uparrow}(k)) \in \cdot] = \sum_j \mathbf{1}[(S^{\uparrow}(j),T(j)) \in \cdot]. \tag{5.36}$$

Define $\Psi_{a,\lambda}(s) = s^{1/a}\Gamma(a+1)^{1/a}/\lambda$. Define a sequence of random variables ξ via

$$\xi(k) = \Psi_{a,\lambda}(S^{\uparrow}(k)), \tag{5.37}$$

and let ξ^* be its re-ordering defined $\xi^*(k) = \Psi_{a,\lambda}(S(k))$. Then jointly as $n \to \infty$,

$$n^{1/a}X_n^{\uparrow} \xrightarrow{f.d.d.} \xi, \tag{5.38}$$

$$n^{1/a}(X_n^*)^{rev} \xrightarrow{f.d.d} \xi^*. \tag{5.39}$$

In particular, for each n, let $J_n = (J_{nk}, 1 \le k \le n)$ be the permutation of $\{1,\ldots,n\}$ defined by $X_n^{rev}[k] = X_n(J_{nk})$. As $n \to \infty$,

$$(J_{nk}, 1 \le k \le n) \xrightarrow{f.d.d.} (J_k : 1 \le k < \infty), \tag{5.40}$$

where J_k is the random permutation of $\{1,2,\ldots,\}$, defined by

$$\xi^*(k) = \xi(J_k) \tag{5.41}$$

for $k = 1,2,\ldots,$ and the f.d.d. convergence in (5.38), (5.39), (5.40) all hold jointly.

In other words, the Poisson point process $N(\cdot)$ defined in (5.36) with measure (5.35) defines a random permutation (J_k) of $\mathbb{N}$ and its inverse (K_j). Theorem 5.28 states that (J_k) is precisely the limit of the random permutation induced by the size-biased permutation of a sequence of n i.i.d terms from F. Furthermore, to obtain the actual sequence of size-biased permutation, one only needs to apply the *deterministic* transformation $\Psi_{a,\lambda}$ to the sequence of s-marginals of points in $N(\cdot)$, ranked according to their t-values. The sequence of increasing order statistics can be obtained by applying the transformation $\Psi_{a,\lambda}$ to the s-marginals ranked in increasing order.

Proof of Theorem 5.28. By Lemma 5.27, it is sufficient to prove the theorem for the case F is $gamma(a, \lambda)$. First we check that the sequence on the right hand side of (5.38) and (5.39) have the right distribution. Indeed, by standard results in order statistics [27, Theorem 2.1.1], as $n \to \infty$, the sequence $n^{1/a} X_n^{\uparrow}$ converges (f.d.d) to the sequence $\tilde{\xi}$, where

$$\tilde{\xi}(k) = (S^{\uparrow}(k))^{1/a} \Gamma(a+1)^{1/a} / \lambda = \Psi_{a,\lambda}(S^{\uparrow}(k)) \tag{5.42}$$

where $S^{\uparrow}(k) = \epsilon_1 + \ldots + \epsilon_k$ for ϵ_i i.i.d standard exponentials. Similarly, by Proposition 5.10 and law of large numbers, the sequence $n^{1/a}(X_n^*)^{rev}$ converges (f.d.d) to the sequence $\widetilde{\xi^*}$, where

$$\widetilde{\xi^*}(k) = (T^{\uparrow}(k))^{1/a} \gamma_k / \lambda,$$

where $T^{\uparrow}(k) = \epsilon'_1 + \ldots + \epsilon'_k$ for i.i.d standard exponentials ϵ'_i, and γ_k, $k = 1, \ldots, n$ are i.i.d $gamma(a+1, 1)$, independent of the $T(k)$. By direct computation, we see that $\tilde{\xi} \stackrel{d}{=} \xi$ and $\widetilde{\xi^*} \stackrel{d}{=} \xi^*$. The dependence between the two sequences S and T comes from Proposition 5.5, which tells us that $S(k)$ is the term $S^{\uparrow}(j)$ that is paired with $T^{\uparrow}(k)$ in our Poisson coupling. Observe $\Psi_{a,\lambda}$ has inverse function $\Psi_{a,\lambda}^{-1}(x) = \lambda^a x^a / \Gamma(a+1)$. Thus applying (5.42), we have

$$S(k) = \Psi_{a,\lambda}^{-1}(\widetilde{\xi^*}(k)) = \lambda^a [\widetilde{\xi^*}(k)]^a / \Gamma(a+1) = T(k) \gamma_k^a / \Gamma(a+1). \tag{5.43}$$

Comparing (5.42) and (5.43) gives a pairing between $S(k)$ and $S^{\uparrow}(k)$, and hence $T^{\uparrow}(k)$ and $S^{\uparrow}(k)$, via $\tilde{\xi}(k)$ and $\widetilde{\xi^*}(k)$. Hence we obtain another definition of J_k equivalent to (5.41):

$$S(k) = S^{\uparrow}(J_k).$$

Let $T(j)$ be the T value corresponding to the order statistic $S^{\uparrow}(j)$ of the sequence S. That is,

$$T(j) = T^{\uparrow}(K_j)$$

where (K_j) is a random permutation of the positive integers. By (5.43), (J_k) is the inverse of (K_j). Together with (5.38) and (5.39), this implies (5.40), proving the last statement. The intensity measure μ comes from direct computation. □

Marginal distributions of the random permutation (J_k) and its inverse (K_j) are given in (5.27) and (5.28). Note that for $k = 1, 2, \ldots$,

$$S(k) = T^{\uparrow}(k) \gamma_k^a / \Gamma(a+1)$$

for i.i.d γ_k distributed as $gamma(a+1, 1)$, independent of the sequence $T^{\uparrow}$, and

$$T(k) = \Gamma(a+1) S^{\uparrow}(k) \tilde{\epsilon}_k^{-a}$$

for i.i.d standard exponentials ϵ_k, independent of the sequence $(S(k))$ but not of the γ_k. Since the projection of a Poisson process is Poisson, the s and t-marginal of μ is just Lebesgue measure, as seen in the proof. Thus by conditioning on either $S^{\uparrow}(k)$ or $T^{\uparrow}(k)$, one can evaluate (5.27) and (5.28) explicitly. In particular, by a change of variable $r = \Gamma(a+1)^{1/a}(s/t)^{1/a}$, one can write μ in product form. This leads to the following.

Proposition 5.29. *For $j \geq 1$, conditioned on $S^{\uparrow}(j) = s, T(j) = \Gamma(a+1)\, s\, r^{-a}$ for some $r > 0$, $K_j - 1$ is distributed as*

$$Poisson(m(s,r)) + Binomial(j-1, p(s,r)) \tag{5.44}$$

with

$$m(s,r) = asr^{-a} \int_r^{\infty} x^{a-1} e^{-x}\, dx, \tag{5.45}$$

and

$$p(s,r) = as^{2/a-2} r^{a-2} \int_0^{r} x^{a-1} e^{-x}\, dx, \tag{5.46}$$

where the Poisson and Binomial random variables are independent. Similarly, for $k \geq 1$, conditioned on $T^{\uparrow}(k) = t, S(k) = t\, r^a/\Gamma(a+1)$ for some $r > 0$, $J_k - 1$ is distributed as

$$Poisson(m'(t,r)) + Binomial(k-1, p'(t,r)) \tag{5.47}$$

with

$$m'(t,r) = t\, \left(\frac{r^a + a \int_r^{\infty} x^{a-1} e^{-x}\, dx}{\Gamma(a+1)} - 1 \right), \tag{5.48}$$

and

$$p'(t,r) = \Gamma(a+1)^{1-2/a}\, at^{2/a-2} r^{a-1/a} \int_0^{r} x^{a-1} e^{-x}\, dx, \tag{5.49}$$

where the Poisson and Binomial random variables are independent.

Proposition 5.30 (Marginal distributions of K_1 and J_1)**.** *Suppose that (5.30) holds for some $\lambda > 0$ and $a = 1$. Then the distribution of K_1, the k such that $\xi(1) = \xi^*(k)$, is a mixture of geometric distributions, and so is that for J_1, the j such that $\xi^*(1) = \xi(j)$. In particular,*

$$P(K_1 = k) = \int_0^{\infty} p_r q_r^{k-1} e^{-r}\, dr \tag{5.50}$$

where $p_r = r/(r + e^{-r})$, $q_r = 1 - p_r$, and

$$P(J_1 = j) = \int_0^{\infty} \tilde{p}_r \tilde{q}_r^{j-1} r e^{-r}\, dr \tag{5.51}$$

where $\tilde{p}_r = 1/(r + e^{-r})$, $\tilde{q}_r = 1 - \tilde{p}_r$.

Proof. When $a = 1$, $\int_r^\infty t^{a-1}e^{-t}\,dt = e^{-r}$. Substitute to (5.45) and (5.48) give

$$m(s,r) = sr^{-1}e^{-r}, \quad m'(t,r) = t(r - 1 + e^{-r}).$$

By a change of variable, (5.35) becomes

$$\frac{\mu(ds\,dr)}{ds\,dr} = se^{-r}, \quad \frac{\mu(dt\,dr)}{dt\,dr} = tre^{-r}.$$

Thus, *conditioned on* s *and* r, $K_1 - 1$ is distributed as the number of points in a p.p.p with rate $r^{-1}e^{-r}$ before the first point in a p.p.p with rate 1. This is the geometric distributions on $(0, 1, \ldots)$ with parameter $p_r = 1/(1 + r^{-1}e^{-r})$. Since the marginal density of r is e^{-r}, integrating out r gives (5.50). The computation for the distribution of J_1 is similar. □

One can check that each of (5.50) and (5.51) sum to 1. We conclude with a 'fun' computation. Suppose that (5.30) holds for some $\lambda > 0$ and $a = 1$. That is, F behaves like an exponential c.d.f near 0. By Proposition 5.30, $E(J_1) = 9/4$ and $E(K_1) = \infty$. That is, the last size-biased pick is expected to be almost the second smallest order statistic, while the smallest order statistic is expected to be picked infinitely earlier on in a successive sampling scheme(!). The probability that the last species to be picked in a successive sampling scheme is also the one of smallest species size is

$$\begin{aligned}\lim_{n\to\infty} P(X_n^{rev}[1] = X_n(1)) = P(\xi^*(1) = \xi(1)) = P(J_1 = 1) = P(K_1 = 1) \\ = \int_0^\infty \frac{re^{-r}}{r + e^{-r}}\,dr = 1 - \int_0^1 \frac{u}{u - \log u}\,du \approx 0.555229\end{aligned}$$

5.7 Summary

In this chapter we focused on the size-biased permutation of n independent and identically distributed (i.i.d) positive random variables. Our setting is a finite dimensional analogue of the size-biased permutation of ranked jumps of a subordinator studied in Perman-Pitman-Yor [76], as well as a special form of *induced order statistics* [12, 23]. This intersection grants us different tools for deriving distributional properties. Their comparisons lead to new results, as well as simpler proofs of existing ones. Our main contribution, Theorem 5.28 in Section 5.6, describes the asymptotic distribution of the last few terms in a finite i.i.d size-biased permutation via a Poisson coupling with its few smallest order statistics.

Chapter 6

Exponential Binary Storage in Little-Hopfield Networks

An extended abstract version of this chapter with weaker main theorem appeared in the joint work with Chris Hillar and Kilian Koepsell, available as an arXiv eprint [48]. Here we strengthen the robustness theorem and supply a proof.

Overview

The Little-Hopfield network is an auto-associative computational model of neural memory storage and retrieval. This model is known to robustly store collections of binary patterns as stable states, or fixed points, of the network dynamics. Researchers are particularly interested in the *network encoding problem*: given a collection of target binary patterns, find a Little-Hopfield network which has them as robust stable states. That is, one trains a Little-Hopfield network to 'remember' the target patterns, so that when exposed to a corrupted version of a target pattern, the network can correctly retrieve the original by running the network dynamics. The more corruption the network can tolerate while guaranteeing exact recovery, the more robust the storage. The idea generalizes to a coding-decoding scheme, where one maps the desired message to a stable state of a secret Little-Hopfield network (which is a binary vector), codes it by corrupting some bits, and decodes using the same network. Thus there is a need for constructing Little-Hopfield networks with a large number of known robust stable states.

While theoretical and experimental evidence suggest that a Little-Hopfield network usually has exponentially many stable states, identifying them and measuring their robustness have been difficult. Current algorithms can only construct robust Little-Hopfield networks in which the number of known stable states scales linearly in the number of neurons. It has been a long-standing open problem whether robust exponential storage of binary patterns was possible in such a network memory model. In this chapter, we design families of Little-Hopfield networks that solve this problem affirmatively.

6.1 Introduction

Inspired by early work of McCulloch-Pitts [68] and Hebb [46], the Little-Hopfield model [52,64] is a distributed neural network architecture for binary memory storage and denoising. In [52], Hopfield showed experimentally, using the *outer-product learning rule* (OPR), that $.15n$ binary patterns (generated uniformly at random) can be robustly stored in such an n-node network if some fixed percentage of errors in a recovered pattern were tolerated. Later, it was verified that this number was a good approximation to the actual theoretical answer [3]. However, pattern storage without errors in recovery using OPR is provably limited to $n/(4 \log n)$ patterns [70, 105]. Since then, improved methods to fit Little-Hopfield networks more optimally have been developed [15, 56, 104], with the most recent being [47]. Independent of the method, however, arguments of Cover [19] can be used to show that the number of randomly generated patterns storable in a Little-Hopfield network with n neurons is at most $2n$, although the exact value is not known (it is $\approx 1.6n$ from experiments in [47]).

Nonetheless, theoretical and experimental evidence suggest that Little-Hopfield networks usually have exponentially many stable states. For instance, choosing weights for the model randomly (from a normal distribution) produces an n-node network with $\approx 1.22^n$ fixed-points asymptotically [45, 69, 99]. However, stable states of randomly networks are difficult to determine from the network parameters. Furthermore, a stored pattern corrupted by only a few bit errors does not typically converge under the network dynamics to the original. Fulan [38] constructed a deterministic Little-Hopfield network with identical weights, and showed that the stable states contain all binary vectors with exactly half of their bits equal. Thus, for a network with $2n$ neurons, the number of such stable states is $\binom{2n}{n} \approx \frac{4^n}{\sqrt{\pi n}}$, which scales exponentially in n. However, such a network is not able to denoise a single bit of corruption. Very recently, more sophisticated (non-binary) discrete networks have been developed [61, 88] that give exponential memory storage. However, the storage in these networks is not known to be robust (in the sense of Definition 6.6 below). Moreover, determining or prescribing the network parameters for storing these exponentially many memories is sophisticated, involving expander graphs and solving linear equations over the integers.

In this chapter, we design a simple two-parameter family of Little-Hopfield networks that robustly stores an exponential number of binary patterns. Our construction uses cliques in graphs and actions of the symmetric groups, resulting in networks with explicit entries. Section 6.2 defines Little-Hopfield networks, their stable states and robust storage. We state and prove our main result (Theorem 6.7) in Section 6.3. In Section 6.4 we discuss and prove properties of our two-parameter Little-Hopfield networks. We conclude with applications and open problems in Section 6.5.

6.2 Little-Hopfield networks and robust storage

Let $H : \mathbb{R} \to \{0, 1\}$ denote the Heaviside function

$$H(y) = \begin{cases} 1 & \text{if } y > 0 \\ 0 & \text{else.} \end{cases}$$

A Little-Hopfield network on n nodes $\mathcal{H} = (J, \theta)$ consists of a real symmetric $n \times n$ *weight matrix* $J \in \mathbf{S}_2\mathbb{R}^n$ with zero diagonal, and a threshold vector $\theta \in \mathbb{R}^n$ which defines the *threshold function* $f_\theta : \mathbb{R}^n \to \{0, 1\}^n$, where

$$H_\theta(\sigma)_i = H(\sigma_i - \theta_i).$$

The possible *states* of the network are all length n binary strings $\{0, 1\}^n$, which we represent as binary column vectors $\sigma = (\sigma_1, \ldots, \sigma_n)^\top$, each $\sigma_i \in \{0, 1\}$ indicating the state σ_i of node i. A Little-Hopfield network defines a deterministic dynamical system on the states. Starting at state σ, one *synchronous update of the network* returns

$$\sigma' = H_\theta(J\sigma),$$

that is, simultaneously for all $i = 1, \ldots, n$,

$$\sigma'_i = H_\theta(J\sigma)_i = H(\sum_{j \neq i} J_{ij}\sigma_j - \theta_i). \tag{6.1}$$

In the context of neuroscience, the i-th neuron (node) takes a weighted sum of its neighbors signals, and it 'fires' (returns 1) if and only if its potential is above the threshold θ_i. Variations of the update rules include *asynchronous update*, where the nodes are updated one-by-one in either fixed or random order according to (6.1).

If σ^* satisfies $\sigma^* = H_\theta(J\sigma^*)$, then it is a fixed-point of (6.1), also called a *stable state*, or a *stored memory*. We say that the network *converges to* σ^* starting from σ, or σ^* is an *attractor* of σ, if the sequence of outputs of the network dynamic $(\sigma^{(0)}, \sigma^{(1)}, \ldots)$ converges pointwise to σ^*. Since there are only finitely many states, starting at a given state σ, the sequence of outputs either converges to a fixed point, or enters a finite cycle. Note that in asynchronous network, a point can have multiple attractors, and convergence depends on the update order. An attractor σ^* of σ is *universal* if, starting at σ, the network converges to σ^* regardless of the node orderings.

Example 6.1. For $n = 2$, let $J = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$, $\theta = (0, 0)^\top$. We start the network at $\sigma^{(0)} = (0, 1)^\top$. After one synchronous update, one obtains $\sigma^{(1)} = (1, 0)^\top$. Another update returns $\sigma^{(2)} = \sigma^{(0)}$. Thus the synchronous network cycles through $\{\sigma^{(0)}, \sigma^{(1)}\}$ infinitely

often. The asynchronous network, however, converges after one step either to $(0,0)^\top$ or $(1,1)^\top$, depending on the update order of the nodes.

Theorem 6.2 ([52]). *In an asynchronous Little-Hopfield network, starting at any state $\sigma \in \{0,1\}^n$, the network dynamic converges in finitely many step to a stable state.*

Proof. Define the *energy* of a pattern σ as

$$E_{J,\theta}(\sigma) := -\frac{1}{2}\sigma^\top J\sigma + \theta^\top\sigma = -\sum_{i<j}\sigma_i\sigma_j J_{ij} + \sum_{i=1}^{n}\theta_i\sigma_i. \tag{6.2}$$

Suppose the network updates node i by (6.1). Let σ' denote the new state. Then the change in energy is

$$E_{J,\theta}(\sigma') - E_{J,\theta}(x) = -(\sigma_i' - \sigma_i)(\sum_{j\neq i} J_{ij}\sigma_j - \theta_i).$$

Thus each network update does not increase the energy (6.2). In particular, if the network updates node i to a different value, that is, $\sigma_i' \neq \sigma_i$, then the energy is strictly decreasing, unless if $\sum_{j\neq i} J_{ij}\sigma_j - \theta_i = 0$, in which case the new state has to be 0. We then have two cases: either the network converges to a strict local minimum of the energy function $E_{J,\theta}$, or it cycles through finitely many states, each differ from the previous by switching a node from 1 to 0. Since there are only finitely many states, the later situation cannot possibly form a cycle. Hence the network converges to a strict local minimum on $\{0,1\}^n$ of the energy function $E_{J,\theta}$. □

For this reason, from now on we shall only consider Little-Hopfield networks with asynchronous updates. We now formalize the notion of robust storage, which comes in two flavors: deterministic (r-stable) and probabilistic (α-robust).

Definition 6.3 (r-stable)**.** Equip the state space $\{0,1\}^n$ with the Hamming distance, which is the minimal number of bit flips between two binary strings. For an integer $r \geq 1$, σ^* is *r-stable* if it is a universal attractor for all states within Hamming distance r of σ^*.

In words, if a state σ is r-stably stored, the network is guaranteed to converge to σ when exposed to any corrupted version not more than r bit flips away from σ. This is a fairly restrictive definition of stability. It could be that we only need the state to have stability 'on average', or for 'most randomly corrupted versions'. We make the last concept rigorous by introducing p-corruption.

Definition 6.4 (p-corruption)**.** Let $\sigma \in \{0,1\}^n$ be a state, $p \in [0,1]$. The *p-corruption of σ* is the random pattern Σ_p, obtained by replacing x_i by $1-x_i$ with probability p, independent for each node $i = 1,\ldots,n$.

We now define α-robustness for a set of states S, which says that for each $\sigma \in S$, with high probability, the network guarantees to recover σ when exposed to a p-corruption of σ for all $p < \alpha$.

Definition 6.5 (α-robust)**.** Let $\mathcal{H}$ be a Little-Hopfield network, $S \subset \{0,1\}^n$. Fix $\epsilon \in [0,1]$. For $p \in [0,1], \sigma \in S$ let Σ_p be the p-corruption of σ. The robustness $\alpha(S, \epsilon)$ of S in $\mathcal{H}$ is

$$\operatorname*{argmax}_{p\in[0,1]}\{\min_{\sigma\in S} \mathbb{P}(\sigma \text{ is a universal attractor of } \Sigma_p) \geq 1-\epsilon\}.$$

Finally, the robustness index of a sequence of states in a corresponding sequence of Little-Hopfield networks is their asymptotic α-robustness. In this sense, the networks mentioned in the introduction do not have robust storage (that is, their robustness index is 0), because the number of bits of corruption tolerated in memory recovery does not increase with the number of nodes.

Definition 6.6 (Robustness index)**.** Let $(\mathcal{H}_n)$ be a sequence of Little-Hopfield networks, (S_n) be a sequence of sets of states. We say that the sequence S_n is *robustly stored by* $\mathcal{H}_n$ *with robustness index* α if for some $\alpha > 0$,

$$\liminf_{\epsilon\to 0, n\to\infty} \alpha(S_n, \epsilon) \to \alpha.$$

Note that the p-corruption of a state differs from the original by np bit flips on average. Clearly the higher p is, the more difficult to recover the original. In particular, if $p = 1/2$, then the random pattern Σ_p is independent of the input, hence consistent recovery is impossible for non-trivial networks. Thus the robustness index α is bounded above by $1/2$.

6.3 Main theorem

Theorem 6.7. *For each positive integer $k > 3$, there exists an explicit two-parameter family of Little-Hopfield networks on $n = \binom{2k}{2}$ nodes that stores $\binom{2k}{k}$ known stable states. In particular, there exists an explicit sequence of Little-Hopfield networks $(\mathcal{H}_n)$ such that asymptotically, $\mathcal{H}_n$ robustly stores approximately $\frac{2^{\sqrt{2n}+\frac{1}{4}}}{n^{1/4}\sqrt{\pi}}$ known stable states with robustness index $\alpha = \frac{1}{2}$.*

In particular, for each $n = \binom{2k}{2}$, the set of known stable states of $\mathcal{H}_n$ in Theorem 6.7 is the set of all k-cliques in the complete graph on $2k$ vertices, and the weight matrix J of $\mathcal{H}_n$ assigns weights to pairs of edges in the graph. Another interpretation of Theorem 6.7 is that these n-node networks have a large number of states that converge under the network dynamics to a targeted memory, which is a k-clique. In other words, the networks have "large basins of attraction" around these stored cliques, as cartoonized in Figure 6.1 below.

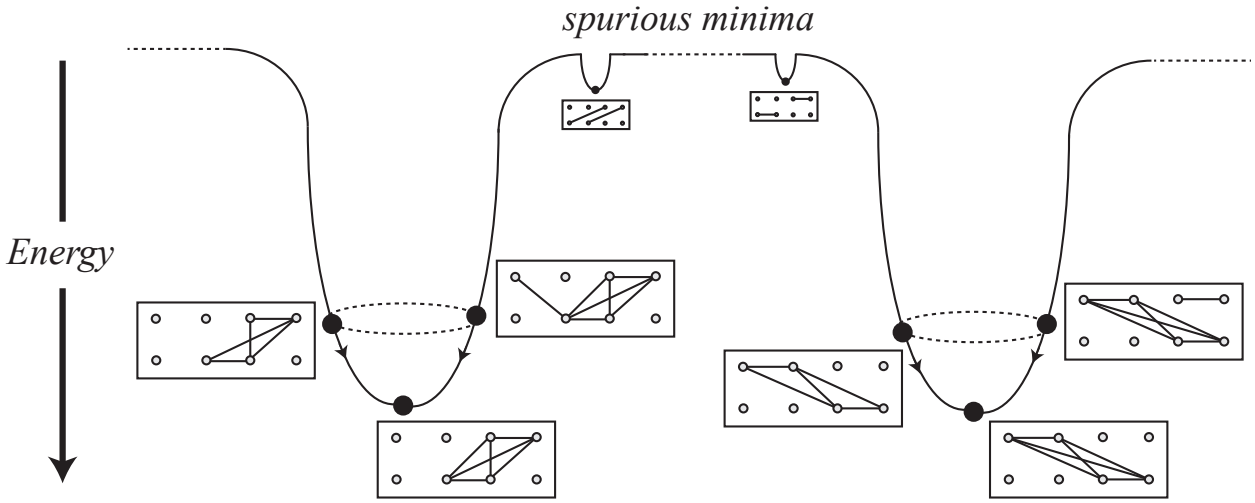

Figure 6.1: Illustration of the energy landscape of a Little-Hopfield network depicting the robust storage of all 4-cliques in graphs on 8 vertices. The network dynamics sends a graph that is almost a clique to a graph with smaller energy, until finally converging to the underlying 4-clique as a stable point.

As previously discussed, α in Definition 6.6 is bounded above by $1/2$. Thus our robustness index is the best possible for any family of Little-Hopfield networks. The number of robust stable states, however, only takes up a vanishingly small fraction of 2^n, the size of the state space. Furthermore, one can show that the fraction of states guaranteed to converge to one of the k-cliques is also vanishingly small. Thus in our Little-Hopfield network, there are numerous other stable states, some of which possibly have very low error tolerance. However, as a model of memory storage and retrieval, this may be more of a feature than a bug: one could argue that our brain stores some memory more vividly than others. From a mathematical perspective, the number of known robust stable states in our Little-Hopfield network is by no means optimal. It would be interesting to identify other stable states of our network. See Conjecture 6.14 for our 'educated guess', supported by simulations.

While the robustness index is defined to be an asymptotic quantity, we can construct finite-sized Little-Hopfield networks which store k-cliques robustly. In Figure 6.2, we demonstrate that the exponential number of stored cliques in our networks have large basins of attraction. For each vertex size $v = 50, 75, 100, 125$ and 150,(with $k = 25, 37, 50, 63$ and 75), we constructed a Little-Hopfield network storing all k-cliques as fixed-points of the dynamics using the parameters identified in Corollary 6.9. Each such k-clique is represented as a binary vector of length $(2k-1)k$. In each experiment, we chose a k-clique uniformly at random, corrupted each of the $\binom{v}{2}$ edges independently with probability p, and ran the network dynamics until convergence. We did 100 of such experiments, and plotted the fraction of the 100 cliques that were correctly recovered (exactly) as a function of p. For example, a network

with $v = 100$ vertices robustly stores $\binom{100}{50} \approx 10^{29}$ memories (i.e., all 50-cliques in a 100-node graph) using binary vectors of length 4950, each having $\binom{50}{2} = 1225$ nonzero coordinates. The figure shows that a 50-clique memory represented with 4950 bits is recovered by the dynamics $\approx 90\%$ of the time after flipping on average $\frac{1}{10} \cdot 4950 = 495$ at random.

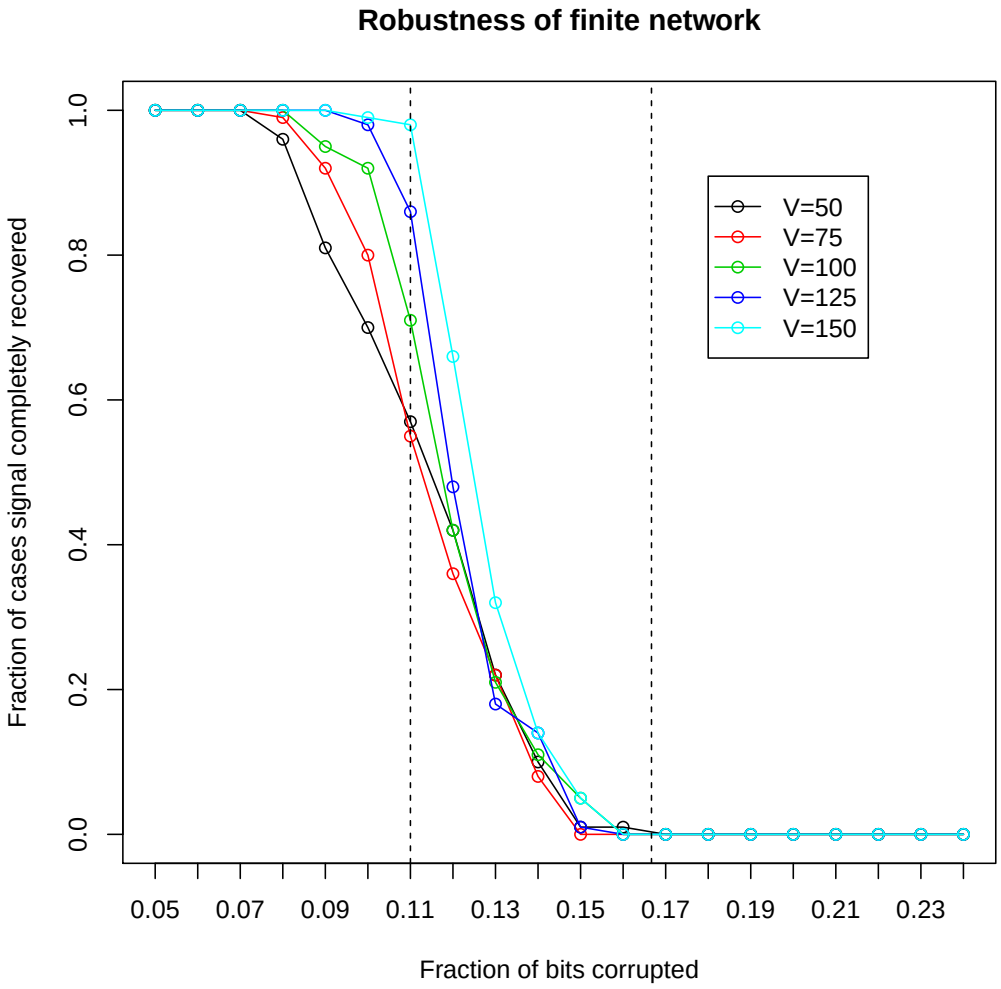

Figure 6.2: Robustness of finite networks with $v = 2k = 50, 75, 100, 125$ and 150 vertices with parameters given in Corollary 6.9. The cut-off for p, the probability of an edge being corrupted, occurs around $p = 0.11$. The asymptotic cut-off for networks with these parameters as $v \to \infty$ is $p = 1/6 \approx 0.167$. These two cut-offs are indicated by vertical dotted lines.

6.3.1 The existence proof

Fix positive integers $v \geq k > 3$, let $n = \binom{v}{2}$. Let K_v denote the complete graph with v vertices. Fix an ordering on the edges of K_v. The set $\{0,1\}^n$ of states of a Little-Hopfield network on n nodes are in bijection with subgraphs of K_v, where for $\sigma \in \{0,1\}^n$, the graph it encodes has edge e if and only if $\sigma_e = 1$. A *k-clique* of K_v is a complete subgraph on k vertices. Let $\theta(z) = z\,\mathbf{1}^\top$, where $\mathbf{1}$ is the all-one vector in $\mathbb{R}^n$. Consider the following two-parameter family of symmetric weight matrices

$$J(x,y)_{ef} = \begin{cases} x & \text{if} \quad |e \cap f| = 1 \\ y & \text{if} \quad |e \cap f| = 0 \end{cases} \tag{6.3}$$

for some $x, y \in \mathbb{R}$, where $|e \cap f|$ is the number of vertices that edges e and f share.

Lemma 6.8. *Fix $k > 3, 1 \leq r < k$. The Little-Hopfield network $(J(x,y), \theta(z))$ stores all k-cliques as r-stable states if and only if the parameters $x, y, z \in \mathbb{R}$ satisfy the following linear inequalities*

$$\begin{bmatrix} -4(k-2)+2r & -(k-2)(k-3) & 2 \\ -4(k-2) & -(k-2)(k-3)-2r & 2 \\ 2(k-1)+2r & (k-1)(k-2) & -2 \\ 2(k-1) & (k-1)(k-2)-2r & -2 \end{bmatrix} \cdot \begin{bmatrix} x \\ y \\ z \end{bmatrix} < \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}.$$

Furthermore, a pattern within Hamming distance r of a k-clique takes one iteration of the network dynamics to converge.

Proof. For fixed r and a state σ of k-cliques, there are 2^r possible patterns within Hamming distance r of σ. Each of these pattern defines a pair of linear inequalities on the parameters x, y, z. However, only the inequalities from the following two extreme cases are active constraints.

- Case 1: r edges are added to a node i not in the clique.
- Case 2: r edges in the clique with a common node i are removed.

In case 1, consider the edges of the form (i,j) for all nodes j in the clique. Such an edge has $r+k-1$ neighboring edges, and $\binom{k-1}{2}$ non-neighboring edges. Thus, such an edge will be labeled as 0 after one network update if and only if x, y satisfy

$$2(r+k-1)x + (k-1)(k-2)y \leq 1. \tag{6.4}$$

Edges of the form (j,j') for all nodes j, j' in the clique have $2(k-2)$ neighbors and $\binom{k-2}{2} + r$ non-neighbors. Thus they impose the following linear constraint

$$4(k-2)x + ((k-2)(k-3)+2r)y > 1. \tag{6.5}$$

In case 2, the edges at risk of being mislabeled are edges of the form (i, j) for all nodes j in the clique. Such an edge has $2(k-2) - r$ neighbors, and $\binom{k-2}{2}$ non-neighbors. This results in another linear constraint for x, y, namely

$$2(2k - r - 4)x + (k-2)(k-3)y > 1. \tag{6.6}$$

Consider edges of the form $(\bar{i}, j)$ for all nodes $\bar{i} \neq i$ in the clique, and j not in the clique. Assume $r < k-1$, such an edge has at most $k-1$ neighbors and $\binom{k-1}{2} - r$ non-neighbors. Thus it imposes the following constraint

$$2(k-1)x + ((k-1)(k-2) - 2r)y \leq 1. \tag{6.7}$$

□

The equations of Lemma 6.8 cut out a cone in $\mathbb{R}^3$. Thus without loss of generality, we can assume either $z = 1/2$ or $z = 0$. Suppose $z = 0$. Without loss of generality, we can choose $y = -1$. This forces $x > 0$, and the second and fourth constraints are dominated by the first and third. Thus we need x which solve

$$(4k - 2r - 8)x > (k-2)(k-3), \quad \text{and} \quad (2k + 2r - 2)x < (k-1)(k-2).$$

To have a feeling of how large r can be relative to k, let us look at the asymptotics as $k \to \infty$, and approximate each coefficient of the matrix by its first-order term. Suppose $r = ck$ for some positive constant c. For large k, we have

$$(4 - 2c)kx > k^2 + O(k) \quad \text{and} \quad (2 + 2c)kx < k^2 + O(k).$$

Thus we need $\frac{k}{2(2-c)} + O(1) < x < \frac{k}{2(1+c)} + O(1)$, which is feasible for all $c < 1/2$, and infeasible for all $c > 1/2$. Thus we can expect to set $r = \lfloor k/2 \rfloor$. By direct computation, we obtain the following.

Corollary 6.9. *Fix $z = 0, y = -1$. For $k = 2r + 1$, the set of feasible solutions is*

$$\frac{2r^2 - 3r + 1}{3r - 2} < x < \frac{2r - 1}{3},$$

which is always non-empty for $r \geq 1$. In particular, $x = \frac{2r-1.5}{3} = \frac{k-2.5}{3}$ is a solution for all $r > 1, k > 3$.

That is, by setting $x = \frac{k-2.5}{3}, y = -1, z = 0$, we obtain a family of Little-Hopfield networks in which all k-cliques are stored as $\lfloor k/2 \rfloor$-stable states. Thus this proves the first part of Theorem 6.7.

6.3.2 The robustness index proof

In the previous section, we have tuned the parameters (x, y, z) to guarantee that the k-cliques stable states are $\lfloor k/2 \rfloor$-stable states. We now chose (x, y, z) to guarantee robust storage.

Let $v = 2k$. Again, fix $y = -1, z = 0$. For $p \in (0, 1)$, let Σ_p be a p-corruption of a k-clique chosen uniformly at random. For each node i, let i_{in}, i_{out} denote the number of edges from i to other clique and non-clique nodes, respectively. We need the following lemma.

Lemma 6.10. *Let Y be an $n \times n$ symmetric matrix with zero diagonal, where Y_{ij} are i.i.d Bernoulli(p) random variables for $i < j$. For each $i = 1, \ldots, n$, let $Y_i = \sum_j Y_{ij}$ be the i-th row sum. Let $M_n = \max_i Y_i, m_n = \min_i Y_i$. Then for any constant $c > 0$,*

$$\mathbb{P}(|m_n - np| > c\sqrt{n}\log(n)) \to 0, \;\; \mathbb{P}(|M_n - np| > c\sqrt{n}\log(n)) \to 0 \text{ as } n \to \infty.$$

That is,

$$|m_n - np|, |M_n - np| = o(\sqrt{n}\log(n)).$$

Proof. Fix a constant $c > 0$. By Bernstein's inequality [10], for each i and for any $\epsilon > 0$,

$$\mathbb{P}(Y_i - np > n\epsilon) \leq \exp\left(-\frac{n\epsilon^2}{2 + 2\epsilon/3}\right).$$

Applying the union bound and choose $\epsilon = \frac{c\log(n)}{\sqrt{n}}$, we have

$$\mathbb{P}(\max_i Y_i - np > n\epsilon) \leq \exp\left(-\frac{n\epsilon^2}{2 + 2\epsilon/3} + \log(n)\right) \leq \exp(-\frac{\log(n)^2}{3} + \log(n)),$$

and the last bound converges to 0 as $n \to \infty$, proving the claim for M_n. Since the distribution of Y_i is symmetric about np, a similar inequality holds for m_n. □

Corollary 6.11. *Let $M_{in} = \max_{i \in C_k} i_{in}$, $m_{in} = \min_{i \in C_k} i_{in}$, $M_{out} = \max_{i \notin C_k} i_{out}$, $m_{out} = \min_{i \notin C_k} i_{out}$, $M_{between} = \max_{i \notin C_k} i_{in}$. Then $M_{in} - k(1-p), m_{in} - k(1-p), M_{out} - kp, m_{out} - kp$ and $M_{between} - kp$ are all of order $o(\sqrt{k}\log(k))$ as $k \to \infty$ almost surely.*

Let T be the total number of edges in Σ_p. Then T is the sum of two independent binomials

$$T \stackrel{d}{=} Binom\left(\binom{k}{2}, 1 - p\right) + Binom\left(\binom{2k}{2} - \binom{k}{2}, p\right).$$

For an edge e, let $N(e), \bar{N}(e)$ be the number of neighbors and non-neighbors of e. Note that $\bar{N}(e) = T - N(e)$. For every e in the clique,

$$N(e) \geq 2m_{in} + 2m_{out} \sim 2k + o(\sqrt{k}\log(k)) \quad \text{w.h.p.}$$

Since the standard deviation is of order k, by a union bound, we have

$$\bar{N}(e) \leq T - (2m_{in} + 2m_{out}) \sim k^2(1/2 + p) + O(k) \quad \text{w.h.p.}$$

To guarantee that $\sigma_e = 1$ for all edge e in the clique after one update iteration, we need $N(e)\,\sigma - \bar{N}(e) > 0$. Therefore, we need

$$x > (\frac{1+2p}{4})\,k + o(\sqrt{k}\log(k)). \tag{6.8}$$

Now let f be an edge with only one vertex in the clique. Then

$$N(f) \leq M_{in} + M_{out} + 2M_{between} \sim k(1+2p) + o(\sqrt{k}\log(k)) \quad \text{w.h.p.}$$

By the same argument as above, we have

$$\bar{N}(f) \geq T - (M_{in} + M_{out} + 2M_{between}) \sim k^2(1/2 + p) + O(k) \quad \text{w.h.p.}$$

To guarantee that $\sigma_f = 0$ for all such edge f after one iteration, we need $N(f)\sigma - \bar{N}(f) < 0$. Thus we obtain the following constraint for x

$$x < \frac{1}{2}\,k + o(\sqrt{k}\log(k)). \tag{6.9}$$

Thus, for some small $\delta \in (0, 1/2)$, if $p = p(k) \sim \frac{1}{2} - k^{\delta - 1/2}$, choosing $x = x(k) = \frac{1}{2}k(1 - k^{\delta/2 - 1/4})$ would guarantee that for large k, both equations (6.8) and (6.9) are simultaneous satisfied. In this case, $\lim_{k\to\infty} p(k) = 1/2$, and thus the family of two-parameters Little-Hopfield networks with $x(k) = \frac{1}{2}k(1 - k^{\delta/2-1/4})$, $y = -1$, $z = 0$ have robustness index $1/2$. This completes the proof of Theorem 6.7.

6.4 Properties of the two-parameter family

This project stem from a question of Chris Hillar as to whether it is possible to use Little-Hopfield networks to solve the hidden clique problem. Consider a graph with v vertices, where a random subset of k vertices has been made into a clique, and the remaining edges are chosen independently with probability $1/2$. The hidden clique problem is to design an algorithm that finds the k-clique in polynomial time with high probability. See [28] for most recent results.

We have explicitly constructed a solution for cliques of size exactly $v/2$. However, our two-parameter family of Little-Hopfield networks cannot store the k-cliques robustly for $k = o(v)$. In practice, our Little-Hopfield network algorithm does no better than the naïve algorithm of picking the k vertices with largest degree. In fact, one synchronous iteration of the two-parameter Little-Hopfield network is equivalent to labeling the edges using an affine

linear hyperplane in $\mathbb{R}^2$. Edges on one side of the hyperplane are labeled $+1$, and those on the other side are labeled 0. In other words, edges are separated into two groups using a linear function in their attributes, in this case the number of neighbors and non-neighbors in the input graph. In statistics, this is known as a *linear classifier*.

Proposition 6.12. *One step of synchronous update using the two-parameter Little-Hopfield network $(J(x,y),\theta(z))$ is equivalent to doing a linear classifier on the edges in $\mathbb{R}^2$ using an affine hyperplane.*

Proof. Associated each edge e with the vector $(n_e, \bar{n}_e)$ which counts the number of neighbors and non-neighbors of e. For $(x, y, z \in \mathbb{R})$, consider the affine hyperplane in $\mathbb{R}^2$

$$xn_e + y\bar{n}_e + z = 0,$$

which separates the points into the positive and negative half-plane. After one synchronous update using $(J(x,y),\theta(z))$, an edge e is labeled $+1$ if and only if the point $(n_e, \bar{n}_e)$ lies on the positive half-plane. Thus we have the claim. □

It follows from Lemma 6.8 that Little-Hopfield networks which store k-cliques always have $x > 0, y < 0$, assuming $z \geq 0$. In other words, every neighbor of e has positive energy contribution x, and every non-neighbor of e has negative energy contribution y. The higher the energy of an edge, the more neighbors, and hence the more 'likely' that it is a clique edge. Starting from some configuration σ, if an edge e has been removed (that is, switched to being labeled 0) in the first iteration of the network dynamics, it means at least one of its two vertices have low degree. The removal of e further reduces its degree, therefore making it 'less likely' for the edge e to ever be labeled 1. This leads to the following conjecture, which is supported by simulations.

Conjecture 6.13. *Let σ be a k-clique, Σ_p be its p-corruption. Staring at Σ_p, with high probability, a two-parameter Little-Hopfield network either converges to σ in one step, or does not converge to σ.*

Simulations also indicate that the other stable states that the network converges to are often either cliques of difference sizes, or a union of star graphs, that is, a graph in which a few nodes have edges to every other node. Both of these cases are intuitive from our heuristic reasoning above.

Conjecture 6.14. *Let $(J(x,y),\theta(z))$ be a two-parameter Little-Hopfield network on $2k$ nodes which stores all k-cliques. Then the other stable states of this network are either k'-cliques for some $k' \neq k$, or a union of star graphs.*

Coming back to the hidden clique problem, it is still interesting to ask whether one can store all k-cliques for a range of k at the expense of robustness. The following theorem gives the largest range of clique size that can be stored by *any* Little-Hopfield networks.

Theorem 6.15. *For each integer $v > 3$, there is a Little-Hopfield network on $n = \binom{v}{2}$ nodes that stores all $\approx 2^v(1 - e^{-Cv})$ k-cliques in the range $k \in R = [\frac{1}{D+2}v, \frac{3D+2}{2(D+2)}v]$ as strict local minima for constants $C \approx .002$ and $D \approx 13.928$. Moreover, this is the best possible range for any Little-Hopfield network. That is, if $\mathcal{H}_n$ has all k-cliques as stable states for all $k \in I$ for some $I \subset \{3, \dots, v\}$ with the same cardinality as R, then $I = R$.*

The above is a corollary of the following more precise version of our range storage result.

Theorem 6.16. *Fix m such that $3 \leq m < v$. For $M \geq m$, there exists a Little-Hopfield network which stores all k-cliques in the range $[m, M]$ if and only if M solves the implicit equation $x_M - x_m < 0$, where*

$$x_m = \frac{-(4m - \sqrt{12m^2 - 52m + 57} - 7)}{2(m^2 - m - 2)}$$
$$x_M = \frac{-(4M + \sqrt{12M^2 - 52M + 57} - 7)}{2(M^2 - M - 2)}.$$

Proof of Theorem 6.15. For large m, M and v, we have the approximations $x_m \approx \frac{\sqrt{12}-4}{2m}$, $x_M \approx \frac{-\sqrt{12}-4}{2M}$ for x_m, x_M defined in Theorem 6.16. Hence $x_M - x_m < 0$ when $M \lesssim \frac{2+\sqrt{3}}{2-\sqrt{3}}m \approx 13.9282m$.

Note that $\binom{v}{k}2^{-v}$ is the fraction of k-cliques in K_v, which is also the probability of a $Binom(v, 1/2)$ equals to k. To store the most number of clique, choose $m = \frac{1}{1+\frac{2+\sqrt{3}}{2-\sqrt{3}}} \approx \frac{1}{15}v$. For large v, approximate the binomial distribution by a Gaussian, then use Mill's ratio [44, p. 98] to approximate the tail of the Gaussian c.d.f, we see that the proportion of cliques storable tends to

$$\Phi(\frac{14}{15}\sqrt{v}) - \Phi(\frac{1}{15}\sqrt{v}) = 1 - 2\Phi(\frac{1}{15}\sqrt{v}) \approx 1 - \exp(-Cv)$$

for some constant $C \approx \frac{1}{2 \cdot 15^2} \approx 0.002$. The range appeared in Theorem 6.15 comes from rounding the fractions involved. □

The 'if and only if' statement in Theorem 6.16 relies on the following simple observation. Let $\mathcal{C}(k, v) \subset \{0, 1\}^n$ denote the set of states whose graphs are k-cliques in the complete graph K_v. Let $\mathbb{S}_v$ be the symmetric group on v letters. Then $\mathbb{S}_v$ induces a natural action on $\mathbf{S}_2\mathbb{R}^n$ by acting on the labeling of the vertices of K_v. In symbols, we have

$$(\tau \cdot J)_{ef} = J_{\tau \cdot e, \tau \cdot f}$$

where an edge $e = (u, v)$ becomes $\tau \cdot e = (\tau(u), \tau(v))$. Let $(\mathbf{S}_2\mathbb{R}^n)^{\mathbb{S}}$ denote the subspace of invariants

$$(\mathbf{S}_2\mathbb{R}^n)^{\mathbb{S}} := \{J \in \mathbf{S}_2\mathbb{R}^n : \tau \cdot J = J \text{ for all } \tau \in \mathbb{S}_v\}.$$

Then $(\mathbf{S}_2\mathbb{R}^n)^{\mathbb{S}}$ is precisely the space of weight matrices in the two-parameter family defined in (6.3). Define $\mathcal{HC}(k) \subset \mathbb{R}^{n\times n}$ to be the cone in $\mathbb{R}^{n\times n}$ of Little-Hopfield networks, identified with their weight matrix and weight vector, which store all k-cliques as stable states. It is a convex polyhedral cone which is closed under the action of $\mathbb{S}_v$. Let $\mathcal{HC}(k)^{\mathbb{S}}$ be the *center* of $\mathcal{HC}(k)$, which is its image under the symmetrization projection

$$J \mapsto \frac{1}{d!}\sum_{\tau\in\mathbb{S}_v} \tau\cdot J, \qquad \theta \mapsto \frac{1}{d!}\sum_{\tau\in\mathbb{S}_v} \tau\cdot \theta.$$

It is clear that $\mathcal{HC}(k) \neq \emptyset$ if and only if $\mathcal{HC}(k)^{\mathbb{S}} \neq \emptyset$. This allows us to reduce all existence statements for $\mathcal{HC}(k)$, their intersections, and similar cones in $\mathbb{R}^{n\times n}$ (such as those with r-stability constraints) to the corresponding statements on their central cone.

Proof of Theorem 6.16. In Lemma 6.8, fix $z = 0.5$ and $r = 0$.[1] Then the cone defined by the inequalities in Lemma 6.8 is in bijection with the polyhedron $\mathbf{I}_k \subseteq \mathbb{R}^2$ cut out by the two halfspaces defined by

$$4(k-2)x + (k-2)(k-3)y - 1 > 0,$$
$$2(k-1)x + (k-1)(k-2)y - 1 < 0.$$

Let R_k be the line $4(k-2)x+(k-2)(k-3)y-1=0$, B_k be the line $2(k-1)x+(k-1)(k-2)y-1=0$. As previously discussed, there exists a Little-Hopfield network which stores all k-cliques in the range $[m, M]$ if and only if $\bigcap_{k=m}^{M} \mathbf{I}_k \neq \emptyset$. For a point $P \in \mathbb{R}^2$, write $x(P)$ for its x-coordinate. Note that for $k \geq 3$, the points $B_k \cap B_{k+1}$ lie on the curve Q implicitly parametrized by k:

$$Q = \{x = \frac{1}{k-1}, y = \frac{-1}{(k-1)(k-2)}\}.$$

Now, when the polytope $\bigcap_{k=m}^{M} \mathbf{I}_k$ is non-empty, its vertices are the following points: $R_M \cap R_m, R_M \cap B_m$, $B_k \cap B_{k+1}$ for $m \leq k \leq M-1$, and $B_M \cap R_m$. This defines a non-empty convex polytope if and only if

$$x_M := x(Q \cap R_M) < x_m := x(Q \cap R_m).$$

Direct computation gives the formulae for x_m, x_M in Theorem 6.16. □

In Figure 6.3 we plot $\mathcal{I}_k$ for $5 \leq k \leq 15$ and shade their intersections. Figure 6.4 shows a similar plot for $1000 \leq k \leq 5500$. Note that appearance of the smooth curve Q enveloping the family B_k in the figure.

[1] Previously we have been considering $z = 0$, which is a different component of the cone, but one can show that the range of cliques recoverable by this component is a subset of the case $z = 0.5$.

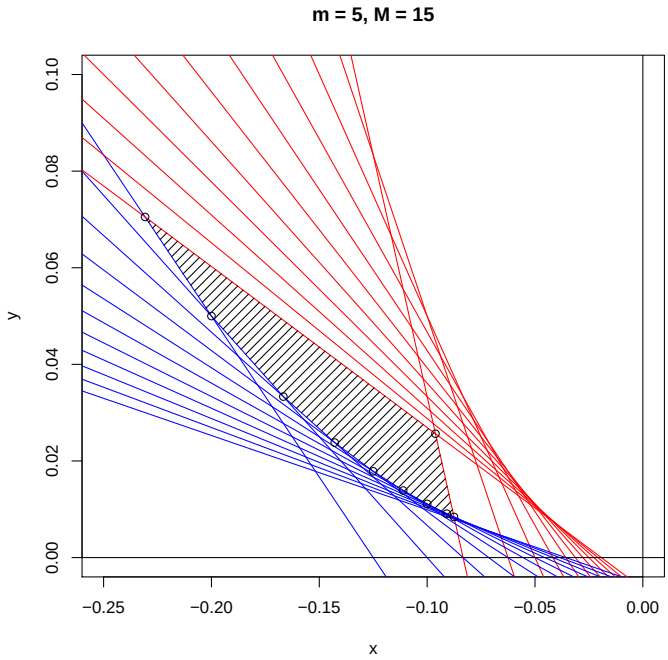

Figure 6.3: Feasible region for network parameters giving exponential storage. The shaded region is the feasible polytope $\bigcap_{k=5}^{15} \mathbf{I}_k$ for network parameters giving clique storage for the range $5 \leq k \leq 15$. Black points are its vertices, the red R_k and blue B_k lines are linear constraints.

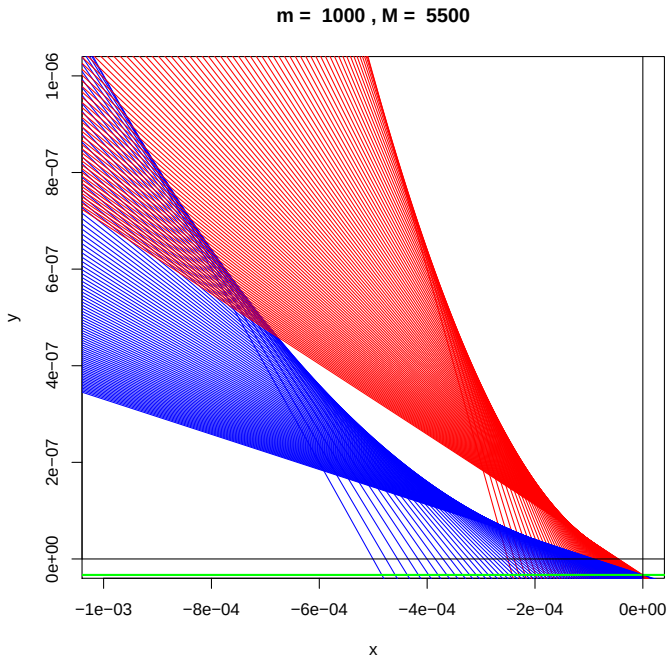

Figure 6.4: Feasible region for larger range of k. Lines R_k (red) and B_k (blue) for $1000 \leq k \leq 5500$.

6.5 Summary

6.5.1 Applications to neuroscience

The Little-Hopfield network is a model of emergent neural computation [52, 53, 64]. One interpretation of the local[2] dynamics in such a model is that by minimizing an energy, the network tries to determine the most probable memory conditioned on a noisy or corrupted version. This concept is in line with arguments of several researchers in theoretical neuroscience [9, 26, 37, 54, 62, 63, 72, 83], and can be traced back to Helmholtz [103]. In addition, recent analyses of spike distributions in neural populations have shown that their joint statistics can sometimes be well-described by the Ising model [8,89,93,94]. The now demonstrated ability of these networks to store large numbers of patterns robustly suggests that the Little-Hopfield architecture should be studied more fully as a possible explanation of neural circuit computation.

6.5.2 Open problems

Some of these problems have been mentioned in the text. We collect them here for the reader's convenience. To our knowledge, the last three problems have not been solved for any non-trivial Little-Hopfield networks. We formulate them in terms of our two-parameter family, with the hope that they provide a good example for tackling these problems.

- (Beating $2^{\sqrt{n}}$). Can one design a family of Little-Hopfield networks with maximal robustness index $\alpha = 1/2$ and have a constant fraction of the state space known to converge to known, robustly stored stable states?

- (Prove Conjecture 6.13). Prove that any two-parameter network $(J(x, y), \theta(z))$ with stable k-clique storage either converges in one step, or does not converge to a k-clique, with high probability. In other words, prove that $(J(x, y), \theta(z))$ cannot do better than a linear classifier in $\mathbb{R}^2$ based on number of neighbors and non-neighbors of edges.

- (Convergence time). Give an upper bound on the number of iterations taken for a two-parameter Little-Hopfield network to converge to any of its stable states.

- (Prove Conjecture 6.14). Characterize all the other stable states of a two-parameter Little-Hopfield network.

[2]The term "local" here refers to the fact that an update (6.1) to a neuron only requires the feed-forward inputs from its neighbors.

Bibliography

[1] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, 1993.

[2] M. Akian, R. Bapat, and S. Gaubert. Asymptotics of the Perron eigenvalue and eigenvector using max-algebra. *Comptes Rendus de l'Acadéimie des Sciences - Series I - Mathematics*, 327(11):927 – 932, 1998.

[3] D.J. Amit, H. Gutfreund, and H. Sompolinsky. Statistical mechanics of neural networks near saturation. *Annals of Physics*, 173(1):30–67, 1987.

[4] P. Astuti and A. D. Garnadi. On eigenvalues and eigenvectors of perturbed comparison matrices. *ITB Journal of Science*, 41 A:69 – 77, 2009.

[5] F. Baccelli, G. Cohen, G.J. Olsder, and J.-P. Quadrat. *Synchronization and Linearity: An Algebra for Discrete Event Systems*. Wiley Interscience, 1992.

[6] Z. Bai and J. W. Silverstein. *Spectral analysis of large dimensional random matrices*. Springer series in statistics, 2010.

[7] E. Barouch and G.M Kaufman. Probabilistic modelling of oil and gas discovery. In *Energy: mathematics and models*, pages 133–150. Philadelphia: SIAM, 1976.

[8] A.K. Barreiro, J. Gjorgjieva, F. Rieke, and E. Shea-Brown. When are feedforward microcircuits well-modeled by maximum entropy methods? *arXiv:1011.2797*, 2010.

[9] P. Berkes, G. Orbán, M. Lengyel, and J. Fiser. Spontaneous Cortical Activity Reveals Hallmarks of an Optimal Internal Model of the Environment. *Science*, 331(6013):83–87, 2011.

[10] S. N. Bernstein. Modification of Chebyshev's inequality and error of Laplace's formula. *Uch. Zap. Nauchno-Issled. Kafedr Ukr., Otd. Mat*, (1):38–48, 1924.

[11] J. Bertoin. Subordinators: examples and applications. In *Lectures on probability theory and statistics (Saint-Flour, 1997)*, volume 1717 of *Lecture Notes in Math.*, pages 1–91. Springer, Berlin, 1999.

[12] P. K. Bhattacharya. Convergence of sample paths of normalized sums of induced order statistics. *Ann. Statist.*, 2:1034–1039, 1974.

[13] P. J. Bickel, V. N. Nair, and P. C. C. Wang. Nonparametric inference under biased sampling from a finite population. *Ann. Statist.*, 20(2):853–878, 1992.

[14] N. H. Bingham, C. M. Goldie, and J. L. Teugels. *Regular Variation*. Encyclopedia of Mathematics and Its Applications. Cambridge University Press, 1987.

[15] AD Bruce, A. Canning, B. Forrest, E. Gardner, and DJ Wallace. Learning and memory properties in fully connected networks. In *AIP Conference Proceedings*, volume 151, page 65, 1986.

[16] P Butkovič. *Max-linear Systems: Theory and Algorithms*. Springer, 2010.

[17] L. Chaumont and M. Yor. *Exercises in Probability: A Guided Tour from Measure Theory to Random Processes, via Conditioning*. Cambridge University Press, 2003.

[18] W. Cheng, K. Dembczyński, and E. Hüllermeier. Label ranking methods based on the plackett-luce model. In Johannes Fürnkranz and Thorsten Joachims, editors, *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pages 215–222, Haifa, Israel, June 2010. Omnipress.

[19] T.M. Cover. Geometrical and statistical properties of systems of linear inequalities with applications in pattern recognition. *IEEE Transactions on Electronic Computers*, 3:326–334, 1965.

[20] G. B. Crawford. The geometric mean procedure for estimating the scale of a judgement matrix. *Mathematical Modelling*, 9:327–334, 1987.

[21] R. A. Cuninghame-Green. Describing industrial processes with interference and approximating their steady-state behaviour. *OR*, 13(1):95–100, 1962.

[22] G. Dahl. A method for approximating symmetrically reciprocal matrices by transitive matrices. *Linear Algebra Appl.*, 403:207–215, 2005.

[23] H. A. David. Concomitants of order statistics. *Bulletin of the International Statistical Institute*, 45:295–300, 1973.

[24] H. A. David, M.J. O'Connell, and S.S. Yang. Distribution and expected value of rank of a concomitant of an order statistic. *Annals of Statistics*, 5(1):216–223, 1977.

[25] Yu. Davydov and V. Egorov. Functional CLT and LIL for induced order statistics. In *Asymptotic methods in probability and statistics with applications (St. Petersburg, 1998)*, Stat. Ind. Technol., pages 333–349. Birkhäuser Boston, Boston, MA, 2001.

[26] P. Dayan, G.E. Hinton, R.M. Neal, and R.S. Zemel. The Helmholtz machine. *Neural computation*, 7(5):889–904, 1995.

[27] L. de Haan and A. Ferreira. *Extreme Value Theory: An Introduction (Springer Series in Operations Research and Financial Engineering)*. Springer, 1st edition. edition, 2010.

[28] Y. Dekel, O. Gurel-Gurevich, and Y. Peres. Finding hidden cliques in linear time with high probability. In *ANALCO11—Workshop on Analytic Algorithmics and Combinatorics*, pages 67–75. SIAM, Philadelphia, PA, 2011.

[29] M. Develin and B. Sturmfels. Tropical convexity. *Doc. Math.*, 9:1–27 (electronic), 2004.

[30] Y. Dong, G. Zhang, W. Hong, and Y. Xu. Consensus models for AHP group decision making under row geometric mean prioritization method. *Decision Support Systems*, 49:281 – 289, 2010.

[31] L. Elsner and P. van den Driessche. Max-algebra and pairwise comparison matrices. *Linear Algebra and its Applications*, 385:47 – 62, 2004.

[32] L. Elsner and P. van den Driessche. Max-algebra and pairwise comparison matrices, II. *Linear Algebra and its Applications*, 432(4):927 – 935, 2010.

[33] S. Engen. *Stochastic Abundance Models with Emphasis on Biological Communities and Species Diversity*. Chapman and Hall Ltd., 1978.

[34] W.J. Ewens. The sampling theory of selectively neutral alleles. *Theor. Popul. Biol.*, 3:87 – 112, 1972.

[35] A. Farkas. The analysis of the principal eigenvector of pairwise comparison matrices. *Acta Polytechnica Hungarica*, 4:99–115, 2007.

[36] A. Farkas, P. Rózsa, and E. Stubnya. Transitive matrices and their applications. *Linear Algebra and its Applications*, 302 - 303(0):423 – 433, 1999.

[37] J. Fiser, P. Berkes, G. Orbán, and M. Lengyel. Statistically optimal perception and learning: from behavior to neural representations. *Trends in Cognitive Sciences*, 14(3):119–130, 2010.

[38] M.N. Fulun. A Hopfield network with exponential storage capability. *Master's Thesis, Ohio State University*, 1988.

[39] J. Fürnkranz and E. Hüllermeier. *Preference Learning*. Springer, 2010.

[40] A. Gnedin, C. Haulk, and J. Pitman. Characterizations of exchangeable partitions and random discrete distributions by deletion properties. In *Probability and mathematical genetics*, volume 378 of *London Math. Soc. Lecture Note Ser.*, pages 264–298. Cambridge Univ. Press, Cambridge, 2010.

[41] A. V. Gnedin. On convergence and extensions of size-biased permutations. *J. Appl. Probab.*, 35(3):642–650, 1998.

[42] P. K. Goel and P. Hall. On the average difference between concomitants and order statistics. *The Annals of Probability*, 22(1):126–144, 1994.

[43] L. Gordon. Successive sampling in large finite populations. *The Annals of Statistics*, 11(2):702–706, 1983.

[44] G. Grimmett and D. Stirzaker. *Probability and Random Processes*. Oxford University Press, 1992.

[45] D.J. Gross and M. Mézard. The simplest spin glass. *Nuclear Physics B*, 240(4):431–452, 1984.

[46] D. O. Hebb. *The Organization of Behavior.* New York Wiley, 1949.

[47] C. Hillar, J. Sohl-Dickstein, and K. Koepsell. Efficient and optimal binary Hopfield associative memory storage using minimum probability flow. arXiv:1204.2916, 2012.

[48] C. Hillar, N. M. Tran, and K. Koepsell. Robust exponential binary pattern storage in Little-Hopfield networks. arXiv:1206.2081v1, 2012.

[49] D. S. Hochbaum. Ranking sports teams and the inverse equal paths problem, 2006.

[50] C. Hollings and M. Kambites. Tropical matrix duality and Green's D relation. *Journal of the London Mathematical Society*, 2010. To Appear.

[51] L. Holst. Some limit theorems with applications in sampling theory. *Ann. Statist.*, 1:644–658, 1973.

[52] J.J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences of the United States of America*, 79(8):2554, 1982.

[53] J.J. Hopfield and D.W. Tank. Neural computation of decisions in optimization problems. *Biological cybernetics*, 52(3):141–152, 1985.

[54] P.O. Hoyer and A. Hyvärinen. Interpreting neural response variability as Monte Carlo sampling of the posterior. In *Advances in Neural Information Processing Systems 15: Proceedings of the 2002 Conference*, page 293. The MIT Press, 2003.

[55] X. Jiang, L-H. Lim, Y. Yao, and Y. Ye. Statistical ranking and combinatorial hodge theory. *Mathematical Programming*, 127:203–244, 2011. 10.1007/s10107-010-0419-x.

[56] M. Jinwen. The asymmetric Hopfield model for associative memory. In *Neural Networks, 1993. IJCNN'93-Nagoya.*, volume 3, pages 2611–2614. IEEE, 1993.

[57] M. Joswig and K. Kulas. Tropical and ordinary convexity combined. *Adv. Geom.*, 10(2):333–352, 2010.

[58] O. Kallenberg. *Foundations of modern probability*. Probability and its Applications (New York). Springer-Verlag, New York, second edition, 2002.

[59] J. F. C. Kingman. Random partitions in population genetics. *Proc. R. Soc. Lond. A.*, 361:1–20, 1978.

[60] J. F. C. Kingman. The representation of partition structures. *J. London Math. Soc.*, 18:374–380, 1978.

[61] K.R. Kumar, A.H. Salavati, and A. Shokrollahi. Exponential pattern retrieval capacity with non-binary associative memory. In *Information Theory Workshop*, pages 80–84. IEEE, 2011.

[62] T.S. Lee and D. Mumford. Hierarchical Bayesian inference in the visual cortex. *JOSA A*, 20(7):1434–1448, 2003.

[63] M. Lewicki and T.J. Sejnowski. Bayesian unsupervised learning of higher order structure. In *Advances in Neural Information Processing Systems 9*. Citeseer, 1996.

[64] W.A. Little. The existence of persistent states in the brain. *Mathematical Biosciences*, 19(1):101–120, 1974.

[65] E. Lukacs. A characterization of the Gamma distribution. *Ann. Math. Statist.*, 26:319–324, 1955.

[66] D. Maclagan and B. Sturmfels. Introduction to Tropical Geometry. American Mathematical Society, 2014, to appear.

[67] J. W. McCloskey. A model for the distribution of individuals by species in an environment. Ph. D. thesis, Michigan State University, 1965.

[68] W.S. McCulloch and W. Pitts. A logical calculus of the ideas immanent in nervous activity. *Bulletin of mathematical biology*, 5(4):115–133, 1943.

[69] R. J. McEliece and E. C. Posner. The number of stable points of an infinite-range spin glass. *JPL Telecomm. and Data Acquisition Progress Report*, 42-83:209–215, 1985.

[70] R. J. McEliece, E. C. Posner, E. R. Rodemich, and S. S. Venkatesh. The capacity of the Hopfield associative memory. *IEEE Trans. Inform. Theory*, 33(4):461–482, 1987.

[71] H. N. Nagaraja and H. A. David. Distribution of the maximum of concomitants of selected order statistics. *The Annals of Statistics*, 22(1):478–494, 1994.

[72] B.A. Olshausen and D.J. Field. Sparse coding with an overcomplete basis set: A strategy employed by V1? *Vision research*, 37(23):3311–3325, 1997.

[73] G. Ottaviani and B. Sturmfels. Matrices with eigenvectors in a given subspace. *Proc. Amer. Math. Soc.*, 141(4):1219–1232, 2013.

[74] L. Pachter and B. Sturmfels. Tropical geometry of statistical models. *Proceedings of the National Academy of Sciences of the United States of America*, 101(46):16132–16137, 2004.

[75] G. P. Patil and C. Taillie. Diversity as a concept and its implications for random communities. *Bull. Int. Stat. Inst.*, XLVII:497 – 515, 1977.

[76] M. Perman, J. Pitman, and M. Yor. Size-biased sampling of poisson point processes and excursions. *Probability Theory and Related Fields*, 92:21–39, 1992. 10.1007/BF01205234.

[77] J. Pitman. Random discrete distributions invariant under size-biased permutation. *Adv. in Appl. Probab.*, 28(2):525–539, 1996.

[78] J. Pitman. Combinatorial stochastic processes - Saint-Flour Summer School of Probabilities XXXII - 2002. In *Combinatorial Stochastic Processes*, volume 1875 of *Lecture Notes in Mathematics*, pages 1+. Springer-Verlag Berlin, 2006.

[79] J. Pitman and N. M. Tran. Size-biased permutation of a finite sequence with independent and identically distributed terms. arXiv:1206.2081v1, 2012.

[80] R. L. Plackett. Random permutations. *Journal of the Royal Statistical Society. Series B (Methodological)*, 30(3):517–534, 1968.

[81] S. Plouffe and N. Sloane. The online encyclopedia of integer sequences. http://oeis.org, 2010.

[82] Alfréd Rényi. On the theory of order statistics. *Acta Math. Acad. Sci. Hungar.*, 4:191–231, 1953.

[83] D.L. Ringach. Spontaneous and driven cortical activity: implications for computation. *Current opinion in neurobiology*, 19(4):439–444, 2009.

[84] B. Rosén. Asymptotic theory for successive sampling with varying probabilities without replacement. I, II. *Ann. Math. Statist.*, 43:373–397; ibid. 43 (1972), 748–776, 1972.

[85] D. A. Saari. A new way to analyze paired comparison rules. Manuscript, 2010.

[86] T. L. Saaty. *The Analytic Hierarchy Process*. McGrawHill, New York, 1980.

[87] T. L. Saaty. Rank according to Perron: A new insight. *Mathematics Magazine*, 60(4):211–213, 1987.

[88] A.H. Salavati, K. Raj Kumar, A. Shokrollahi, and W. Gerstnery. Neural pre-coding increases the pattern retrieval capacity of Hopfield and bidirectional associative memories. In *Information Theory Proceedings (ISIT)*, pages 850–854. IEEE, 2011.

[89] E. Schneidman, M.J. Berry, R. Segev, and W. Bialek. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440(7087):1007–1012, 2006.

[90] P. K. Sen. Invariance principles for the coupon collector's problem: a martingale approach. *Ann. Statist.*, 7(2):372–380, 1979.

[91] S. Sergeev. *Multiorder, Kleene stars and cyclic projectors in the geometry of max cones*, volume 495 of *Contemporary Mathematics*, pages 317–342. Providence: American Mathematical Society, 2009.

[92] S. Sergeev, H. Schneider, and P. Butkovic. On visualization scaling, subeigenvectors and Kleene stars in max algebra. *Linear Algebra and its Applications*, 431(12):2395 – 2406, 2009.

[93] J. Shlens, G.D. Field, J.L. Gauthier, M. Greschner, A. Sher, A.M. Litke, and EJ Chichilnisky. The structure of large-scale synchronized firing in primate retina. *Journal of Neuroscience*, 29(15):5022, 2009.

[94] J. Shlens, G.D. Field, J.L. Gauthier, M.I. Grivich, D. Petrusca, A. Sher, A.M. Litke, and EJ Chichilnisky. The structure of multi-neuron firing patterns in primate retina. *Journal of Neuroscience*, 26(32):8254, 2006.

[95] I. Simon. On semigroups of matrices over the tropical semiring. *RAIRO Inform. Théor. Appl.*, 28(3-4):277–294, 1994.

[96] G. W. Stewart and J. Sun. *Matrix perturbation theory*. Academic Press, Inc, 1990.

[97] B. Sturmfels and N. M. Tran. Combinatorial types of tropical eigenvectors. *Bulletin of the London Mathematical Society*, 45:27–36, 2013.

[98] B. Sturmfels and V. Welker. Commutative algebra of statistical ranking. *Journal of Algebra*, 361(0):264 – 286, 2012.

[99] F. Tanaka and S. F. Edwards. Analytic theory of the ground state properties of a spin glass. I. Ising spin glass. *Journal of Physics F: Metal Physics*, 10(12):2769, 1980.

[100] N. M. Tran. Hodgerank is the limit of Perron rank. arXiv:1201.4632v1, 2012.

[101] N. M. Tran. Pairwise ranking: Choice of method can produce arbitrarily different rank order. *Linear Algebra and its Applications*, 438(3):1012 – 1024, 2013.

[102] N. M. Tran. Polytropes and tropical eigenspaces: cones of linearity. arXiv:1205.3186v2, 2013.

[103] H. von Helmholtz. *Handbuch der Physiologischen Optik*. Voss, Leipzig, 1867.

[104] D. J. Wallace. Memory and learning in a class of neural network models. *Lattice Gauge Theory–A Challenge to Large Scale Computing*, 1986.

[105] G. Weisbuch and F. Fogelman-Soulié. Scaling laws for the attractors of Hopfield networks. *Journal de Physique Lettres*, 46(14):623–630, 1985.

[106] J. Wellner. A Glivenko-Cantelli theorem for empirical measures of independent but non-identically distributed random variables. *Stochastic Processes and their Applications*, 11(3):309 – 312, 1981.

[107] J. H. Wilkinson. *Algebraic Eigenvalue Problem*. Clarendon Press, Oxford, 1965.

[108] S. S. Yang. General distribution theory of the concomitants of order statistics. *Ann. Statist.*, 5(5):996–1002, 1977.

[109] Y. Q. Yin, Z. D. Bai, and P. R. Krishnaiah. On the limit of the largest eigenvalue of the large dimensional sample covariance matrix. *Probability theory and related fields*, 78:509 – 521, 1988.

[110] G. M. Ziegler. *Lectures on Polytopes*. Springer-Verlag, 1995.