The mathematics of ranking

The American Institute of Mathematics

The following compilation of participant contributions is only intended as a lead-in to the AIM workshop "The mathematics of ranking." This material is not for public distribution.
Corrections and new material are welcomed and can be sent to `workshops@aimath.org`
Version: Thu Aug 12 02:28:46 2010

Table of Contents

A. Participant Contributions 3

1. Arrow, Kenneth
2. Gleich, David
3. Hochbaum, Dorit
4. Holmes, Susan
5. Kondor, Risi
6. Langville, Amy
7. Levitt, Michael
8. Lim, Lek-Heng
9. Mackey, Lester
10. Orrison, Michael
11. Rudin, Cynthia
12. Saari, Don
13. Saaty, Thomas
14. Sarkar, Purnamrita
15. Shader, Bryan
16. Shiu, Anne
17. Singer, Yoram
18. Small, Kevin
19. Vayatis, Nicolas
20. Zhang, Tong

A.1 Arrow, Kenneth

My interest is in the reasonable conditions on aggregating multiple preferences into a single one. I am not doing any active research in this area.

A.2 Gleich, David

Most of my previous work on ranking involves Google's PageRank system, which associates a real-valued quantity with each node in a graph to reflect the node's importance. I hope to learn about other types of numerical ranking schemes for graphs as well as other types of data and non-numerical ranking techniques. Moreover, I've often wondered when we should and shouldn't rank large quantities of data. Are there times when only subsets of the data are rankable? Finally, fast data streams like Twitter and Facebook raise questions about how we should rank items in rapidly changing databases of information.

A.3 Hochbaum, Dorit

The problem of group ranking, a.k.a. rank aggregation, has been studied in contexts varying from sports, to multi-criteria decision making, to machine learning, to ranking web pages, and to behavioral issues. The dynamics of the group aggregation of individual decisions has been a subject of central importance in decision theory. We present here a new paradigm using an optimization framework that addresses major shortcomings that exist in current models of group ranking. Moreover, the framework provides a specific performance measure for the quality of the aggregate ranking as per its deviations from the individual decision makers' rankings.

The new model for the group ranking problem presented here is based on rankings provided with intensity – that is, the degree of preference is quantified. The model allows for flexibility in decision protocols and can take into consideration imprecise beliefs, less than full confidence in some of the rankings, and differentiating between the expertise of the reviewers. Our approach relaxes frequently made assumptions of: certain beliefs in pairwise rankings; homogeneity implying equal expertise of all decision makers with respect to all evaluations; and full list requirement according to which each decision maker evaluates and ranks all objects. The option of preserving the ranks in certain subsets is also addressed in the model here. Significantly, our model is a natural extension and generalization of existing models, yet it is solvable in polynomial time. The group rankings models are linked to network flow techniques.

A.4 Holmes, Susan

I work the statistics of ranking data, often trying to find underlying hidden variables that explain the rankings.

A.5 Kondor, Risi

I am interested in machine learning problems with a combinatorial or algebraic flavor. I am increasingly interested in ranking, especially since it ties in with a series of papers I wrote recently on harmonic analysis on the symmetric group (the group of permutations). Marconi Barbosa and I just had a paper at COLT on using the representation theory of the symmetric group to efficiently compute kernels between partial rankings.

I would be interested in exploring whether the kernels approach really is a promising way to approach ranking problems. More generally, I would like to know how much the algebra of permutations impact real-world ranking problems.

A.6 Langville, Amy

I originally became interested in ranking around 2001 after reading about the famous PageRank algorithm for ranking webpages. Since then, I written half a dozen papers on webpage ranking and one book, which was published in 2006. I subsequently extended my ranking research into other application areas, most often sports, given that that data is easily accessible and plentiful. For the past 3-5 years, I have been studying ranking in general and adding material to my next book, "Who's #1? The Science of Ranking items from movies to webpages to sports teams," which is due out next year.

Overall, my work involves creating new linear algebra based algorithms for ranking items, creating new measures for comparing several ranked lists, and analyzing the sensitivity of ranking vectors. Most recently, my tools have been optimization techniques and sensitivity analysis. See the preprints and reprints I've included for the AIM library.

I am very excited about this workshop, and in particular, the opportunity to discuss and work in person with colleague and attendee David Gleich. I am interested in hearing about the work of others in this field, learning about various journal outlets (wouldn't a mathematics of ranking journal be nice?), and networking and making new collaborations. It seems like the attendees come from a broad range of fields, so the inter-disciplinary exchange will be great, especially if presenters do a good job with introductory or tutorial sessions. Some other questions: what are the open problems and major challenges in the field? What questions are of interest to the industrial attendees? Why is ranking important to these applied folks (aside from the inherent fun of the mathematical problem)?

A.7 Levitt, Michael

Ranking as a practical activity that seems to carry remarkable weight in society yet be based on very poor procedures.

A.8 Lim, Lek-Heng

Will return to fill out this portion at a later date.

A.9 Mackey, Lester

Coming from a statistical machine learning background, I have a particular interest in the statistical properties of different ranking procedures. In particular, when is a ranking algorithm risk consistent for a loss function of interest? When is it statistically efficient, and what are the rates of convergence? What final sample guarantees can be established for a given ranking loss? There is a rich literature on the risk consistency of surrogate loss minimization in the setting of binary classification; my colleagues and I have been investigating these issues in the ranking setting.

I am also generally interested in learning more about the diversity of ranking techniques and perspectives employed by different communities.

A.10 Orrison, Michael

For the past few years, my students and I have been using algebraic ideas and techniques to develop something that might be called "algebraic voting theory." In particular, we have been using the representation theory of the symmetric group to better understand positional voting, which basically occurs when voters return a (full or partial) ranking of the candidates in an election, and the candidates are then given points based on the positions they occupy in the voters' rankings.

Because of my research interests, I would be particularly interested in conversations at the workshop concerning algebraic approaches to rankings. I am also interested in algorithmic aspects of rankings (e.g., doing "generalized spectral analysis" of ranked data) and the statistical analysis of ranked data (e.g., linear rank tests of uniformity). As such, I would also enjoy learning more about machine learning and Markov chain approaches to rankings.

A.11 Rudin, Cynthia

For the last few years I have been using ranking methods in machine learning to rank manholes on the NYC power grid in order of vulnerability to fires and explosions. I work on theoretical aspects of machine learning ranking problems, such as algorithmic convergence and generalization analysis, and also algorithm design. I am interested in finding out new algorithmic and theoretical approaches, and will enjoy talking with Tong, Risi, Yoram, Nicolas, Shivani and others.

A.12 Saari, Don

In trying to fulfill this request from AIM, let me start with some of my background. I am a mathematician where my earlier (and continuing) interests centered around celestial mechanics – the mathematics of the dynamics of the Newtonian N-body problem. Here I worked on problems such as the evolution of the universe, the structure and likelihood of collisions, (recently, dark matter), etc. After a Post-Doc in the Yale Astronomy Department, I spent the next three decades in the Northwestern University math department. Now, it turns out that the math department is located next door to the building that houses the Northwestern economics department and the NU Kellogg School of Business. This last statement alone should explain how and why I became interested in mathematical economics along with issues of rankings that come from voting, statistics, economics, sports, and on and on.

In fact, I became sufficiently interested in the mathematical questions arising from rankings, and, more generally, concerns from the social and behavioral sciences, that this general area evolved into my "day job," while the dynamics of the physical sciences dropped to become a second interest, but still above being an avocation.

In the late 1990s, Duncan Luce invited me to spend a term at the Institute for Mathematical Behavioral Sciences (IMBS) at the University of California, Irvine. Upon my arrival, Duncan started recruiting me. It took a couple of years for me to fully appreciate the wisdom of his offer, so in July, 2000, my wife and I moved to UCI. Currently I am the director of the IMBS, and I am having a delightful time!

So, what has been my take on the "mathematics of ranking"? Stated simply, it is to develop mathematical approaches that will be able to explain all of those troubling paradoxes that can occur with ranking approaches. So far, I have had success with voting rules;

e.g., by extracting lessons from "chaotic dynamics," it now is possible to find all possible ranking paradoxes that could ever occur (with any number of candidates and over all possible profiles); by using orbits of symmetry groups, it now is possible to explain how and why all of these paradoxes occur and to create any number of examples; by using concepts from dynamical systems, it now is possible to explain a variety of issues such as strategic behavior, etc.. (The "publication" that I will attach to the workshop's webpage is an expository description of a portion of this work.)

With graduate students (the most recent being Anna Bargagliotti), I started to extend my algebraic approach so that it could answer questions about non-parametric statistics. But now that Anna and Mike Orrison have joined forces and are nicely developing this topic, I decided to move on to other issues.

As for the workshop; I am interested in "ranking concerns" that arise anywhere and everywhere.

Another developing interest of mine is to explore how the hard won lessons developed in social choice and other ranking areas extend to explain mysteries that arise in other disciplines. As an example, not only has "Arrow's possibility theorem" been central to voting theory, but it has been extended to explain problems that arise in "consensus theory" and other areas. But, does it offer even more? Motivated by discussions with UCI colleagues working in nanotechnology, I wondered whether some of Arrow's insights could explain basic concerns that occur in engineering and the physical sciences. They can. (A first paper in this direction with an emphasis on engineering should appear in the *Journal of Mechanical Design* around the time of our workshop.)

Other extensions of notions developed in the study of "rankings" are being made explain some issues in psychology.

A.13 Saaty, Thomas

On the measurement of intangibles: Scales,comparisons, eigenvalues and eigenvectors in deriving priorties to rank altenatives.

The AHP has four axioms, (1) reciprocal judgments, (2) homogeneous elements, (3) hierarchic or feedback dependent structure, and (4) rank order expectations.

Assume that one is given n stones, A1, ..., An, with known weights w1, ..., wn, respectively, and suppose that a matrix of pairwise ratios is formed whose rows give the ratios of the weights of each stone with respect to all others. Thus one has the equation: Aw = nw where A has been multiplied on the right by the vector of weights w. The result of this multiplication is nw. Thus, to recover the scale from the matrix of ratios, one must solve the problem Aw = nw or (A - nI)w = 0. This is a system of homogeneous linear equations. It has a nontrivial solution if and only if the determinant of A-nI vanishes, that is, n is an eigenvalue of A. Now A has unit rank since every row is a constant multiple of the first row. Thus all its eigenvalues except one are zero. The sum of the eigenvalues of a matrix is equal to its trace, the sum of its diagonal elements, and in this case the trace of A is equal to n. Thus n is an eigenvalue of A, and one has a nontrivial solution. The solution consists of positive entries and is unique to within a multiplicative constant.

To make w unique, one can normalize its entries by dividing by their sum. Thus, given the comparison matrix, one can recover the scale. In this case, the solution is any column of A normalized. Notice that in A the reciprocal property aji = 1/aij holds; thus, also aii = 1. Another property of A is that it is consistent: its entries satisfy the condition ajk =

aik/aij. Thus the entire matrix can be constructed from a set of n elements which form a chain across the rows and columns.

In the general case, the precise value of wi/wj cannot be given, but instead only an estimate of it as a judgment. For the moment, consider an estimate of these values by an expert who is assumed to make small perturbations of the coefficients. This implies small perturbations of the eigenvalues. The problem now becomes A'w' = maxw' where max is the largest eigenvalue of A'. To simplify the notation, we shall continue to write Aw = maxw, where A is the matrix of pairwise comparisons. The problem now is how good is the estimate of w. Notice that if w is obtained by solving this problem, the matrix whose entries are wi/wj is a consistent matrix. It is a consistent estimate of the matrix A. A itself need not be consistent. In fact, the entries of A need not even be transitive; that is, A1 may be preferred to A2 and A2 to A3 but A3 may be preferred to A1. What we would like is a measure of the error due to inconsistency. It turns out that A is consistent if and only if max = n and that we always have max d n.

Since small changes in aij imply a small change in max, the deviation of the latter from n is a deviation from consistency and can be represented by (max - n)/(n-1), which is called the consistency ratio (C.I.). When the consistency has been calculated, the result is compared with those of the average value of the same index computed for many randomly generated reciprocal matrices from the scale 1 to 9, with reciprocals forced. This index is called the random index (R.I.). The following gives the order of the matrix (first row) and the R.I. (second row) as computed by generating the R.I. for tens of thousands of matrices of the given order and averaging:

The ratio of C.I. to the average R.I. for the same order matrix is called the consistency ratio (C.R.). A consistency ratio of 0.10 or less is positive evidence for informed judgment.

The relations aji = 1/aij and aii = 1 are preserved in these matrices to improve consistency. The reason for this is that if stone 1 is estimated to be k times heavier than stone 2, one should require that stone 2 be estimated to be 1/k times the weight of the first. If the consistency ratio is significantly small, the estimates are accepted; otherwise, an attempt is made to improve consistency by obtaining additional information. The things that contribute to the consistency of a judgment are: (1) the homogeneity of the elements in a group, that is, not comparing a grain of sand with a mountain; (2) the sparseness of elements in the group, because an individual cannot hold in mind simultaneously the relations of many more than a few objects; and (3) the knowledge and care of the decision maker about the problem under study.

A.14 Sarkar, Purnamrita

I have worked on ranking nodes in a graph using random walk based measures. These are interesting, and intuitive heuristics which are widely used. While there are some interesting algorithmic questions, e.g. how to compute these fast in large graphs, there are also questions about why they are useful. The first one often exploits a variety of algorithms and tools and it would be fun to find out more about these techniques.

The second question is: how do we theoretically justify the common trends in the predictive performance of different path based heuristics? Recently there has been some work on axiomatic formulations of random walk based measures, which also addresses this. Given a set of properties I need a similarity measure to have, which is the right measure to pick? Or, is there at all any measure which satisfies all of them simultaneously?

Apart from graph based proximity measures, I am interested in other formulations of ranking problems, which for example, attempt at learning the ranking directly instead of focusing on the underlying measure.

A.15 Shader, Bryan

I'm working with a group of zoologist's who are trying to understand the correlation between the social interaction of young male manakins and the male's placement in the mating queue. My initial interest in ranking is associated with the study of the spectral radius of tournament matrices.

A.16 Shiu, Anne

My work on rank test was motivated by the problem of finding periodic gene expression profiles in time-course microarray data. Periodic genes were inferred to be related to somitogenesis, the biological process during early embryonic development in which somites (the precursors to the backbone segments and related tissues) are created. Our approach to this problem was based on the cyclohedron test, which is a rank test inspired by recent advances in algebraic combinatorics. In particular, we focused on the statistics and combinatorics that underlie the cyclohedron test and its implementation, especially in light of multiple hypothesis testing concerns. During this workshop, I hope to learn about new mathematical problems in this research area.

A.17 Singer, Yoram

Statistical machine learning algorithms for large preference data with emphasis on approaches that promote structural sparsity of ranking models.

A.18 Small, Kevin

My research interests in regards to ranking problems primarily lie in the areas of learning ranking functions and learning to aggregate rankings, particularly in unsupervised and interactive learning scenarios. Application domains which have contributed to much of my relevant research include natural language processing applications (e.g. aggregation of predicted structures) and rapid deployment of targeted search engines (e.g. vertical search engines). Of specific interest to my current work includes deriving distance functions between rankings which are position sensitive (e.g. emphasizing higher rankings for aggregating predicted permutations/structures) and approximation methods for estimating probability distributions parametrized by such distance functions.

A.19 Vayatis, Nicolas

My particular interest in the workshop concerns the statistical aspects of ranking methods. One of the major issues addressed by statistical learning theory after the pioneering work of Vapnik (1995, 1998) was to explain the generalization ability of efficient algorithms used for prediction purposes on the basis of complex (high dimensional) data. Over the last decade, several advances were made in this direction when dealing with classification and regression problems. Concepts like consistency, complexity control, the margin of a classifier, convex risk minimization principles, optimal penalty calibration, fast rates of convergence, sparsity, have become central in the understanding and the theoretical analysis of learning algorithms. Now ranking can be seen as an intermediate problem, more difficult than classification but easier than regression. However, as a global learning problem, ranking presents some special features and when considering for instance the study of consistency or convergence rates for ranking methods, the same questions may appeal to different answers. Hence, further theoretical developments need to be undertaken and I expect that the workshop will have a fruitful impact in this respect.

A.20 Zhang, Tong

Learning to rank. Statistical theory. Scalability.