AGGREGATION OF IMPLICIT FEEDBACKS FROM SEARCH ENGINE LOG FILES

Ashok Veilumuthu and Parthasarathy Ramachandran
Department of Management Studies, Indian Institute of Science, Bangalore 560 012, India

Keywords: Ranking Models, Feedback Aggregation, Implicit Feedbacks, Click Sequence, Partial Ordering.

Abstract: The current approaches to information retrieval from the search engine depends heavily on the web linkage structure which is a form of relevance judgment by the page authors. However, to overcome spamming attempts and language semantics, it is important to also incorporate the user feedback on the documents' relevance for a particular query. Since users can be hardly motivated to give explicit/direct feedback on search quality, it becomes necessary to consider implicit feedback that can be collected from search engine logs. Though there are number of implicit feedback measures proposed to improve the search quality, there is no standard methodology proposed yet to aggregate those implicit feedbacks meaningfully to get a final ranking of the documents. In this article, we propose an extension to the distance based ranking model to aggregate different implicit feedbacks based on their expertise in ranking the documents. The proposed approach has been tested on two implicit feedbacks, namely click sequence and time spent in reading a document from the actual log data of AlltheWeb.com. The results were found to be convincing and indicative of the possibility of expertise based fusion of implicit feedbacks to arrive at a single ranking of documents for the given query.

1 INTRODUCTION

Search engines are information retrieval systems intended to help the users locate their needed information from internet by querying. These queries are usually constructed of keywords to express the users' information needs. Initially, traditional keyword similarity based techniques were used to retrieve the relevant documents from the web. These techniques rely completely on the content of the documents upon which only its authors have sole control. This allows them to manipulate the search results by tampering the documents' content. This lead to the failure of the traditional keyword based techniques. Later in the late 90s, the research community realized the importance of utilizing the linkage structure that exists within the web documents in the form of hyperlinks to improve the search results. Two seminal works by (Brin and Page, 1998) and (Kleinberg, 1999) have used the linkage structure of the web as the human annotation about the quality of the documents. The web linkage structure captures the importance of pages to a large extent, though, it captures only the collective relevance judgment given by the authors of the webpages and not of the readers/users. It is true that the collective judgement of authors is a reliable information, but, the end users are more eligible to judge the credibility of the documents presented to them. Therefore, the users' feedback information would be a valuable source to help improve the search results further.

Multitude of studies have proposed ways of obtaining the user relevance feedback information and methods to incorporate them into the retrieval engine. The feedback information could be either explicit or implicit. In the explicit feedback based methods the users are explicitly asked to register their feedback on the documents presented to them. Such a strategy would impose an increased burden and cognitive load on the users. Further, many users may not be motivated to provide this information (White et al., 2002). In the case of implicit feedbacks, the users interaction with the search system will be recorded in the form of a log file and its entries can be suitably interpreted to infer the users' relevance judgement on the documents presented. There was a debate going on for almost half a decade on substituting the explicit feedback with the implicit ones. Later from experiments it has been concluded that implicit feedbacks can be used as viable alternative for their explicit counterparts by (White et al., 2002).

Veilumuthu A. and Ramachandran P..
AGGREGATION OF IMPLICIT FEEDBACKS FROM SEARCH ENGINE LOG FILES.
DOI: 10.5220/0003096502690274
In *Proceedings of the International Conference on Knowledge Discovery and Information Retrieval* (KDIR-2010), pages 269-274
ISBN: 978-989-8425-28-7

Since then, many attempts have been made to grasp every possible user behavior during the search and use them as a proxy to their relevance feedback later. See, for example (Kim et al., 2000; Kelly and Belkin, 2004; Ramachandran, 2005; Agichtein et al., 2006; Veilumuthu and Ramachandran, 2007). It is understandable that every feedback will have its own advantage and disadvantage and therefore, it would be more useful, if there was a way to aggregate these feedbacks based on their expertise in achieving the ideal ranking. Such an aggregation would help extracting the advantages from each one of the feedbacks and will result in a more accurate and unbiased ranking of documents.

A number of studies have been conducted in the information retrieval literature by borrowing the theories from various fields to aggregate the rankings from various sources. First of that kind was proposed by (Dwork et al., 2001). They proposed a markov chain based formulation to aggregate the individual ranking produced by multiple search engines, and they studied the impact of local kemenization in reducing the spams. In this work, all the rankers are weighed equally and hence it ignores the importance of the better ranker. (Lebanon and Lafferty, 2002) used the distance based ranking models to propose a formalism for the supervised ensemble learning and analyzed the results for partial rankings. An unsupervised rank aggregation approach using the distance based model has been given by (Klementiev et al., 2008). These two works hinted the possibility of incorporating the expertise of the rankers while combining their input rankings. All these studies intended to solve the supervised as well as unsupervised rank aggregation problems in the context of metasearch where the top k-list is fixed, no study has been done in the context of relevance feedback aggregation where the partial orderings can be of any length. This motivated us to extend the distance based ranking models to combine the various feedback rankings based on their expertise in achieving the unknown ideal ranking.

In this paper, we propose a framework to aggregate multiple feedbacks obtained in the form of partial rankings from various user sessions into a single consensus ranking. We extend the distance based ranking models proposed by Mallows to make this unsupervised aggregation more meaningful. The proposed aggregation framework has been examined over the two implicit feedbacks namely, click sequence and time spent reading a document, extracted from the actual log data. The results are found be encouraging and it also ensured the feasibility of achieving such an expertise based aggregation of feedback rankings. In this study, though we discuss only the aggregation of implicit feedbacks, the construction doesn't prevent one from using it for the explicit counterparts. The only requirement is that the easy convertibility of the feedbacks in to partial orders without much information loss.

2 DISTANCE BASED RANKING MODELS

Given a set of elements, any meaningful ranking scheme will assume the existence of an ideal ordering of elements π_0 and will tend to arrange the elements in an order closer to π_0. Therefore, it is highly preferable for a ranking scheme to produce a ranking closer to π_0 than a ranking farther from it. That is, the probability of getting a permutation π should decrease as its distance from π_0 increases. This is the basic intuition behind most of the ranking models proposed in the literature. In this paper, we use the family of distance based ranking models first proposed by (Mallows, 1957) and extended to partial orders by (Fligner and Verducci, 1986). The two main features that motivated this selection are: (1) It gives a distributional view of data, hence, an effective way of representation, and (2) easily interpretable distributional parameters, more importantly the dispersion parameter can be interpreted as the expertise of the ranking scheme (Lebanon and Lafferty, 2002).

Let $\mathcal{X} = \{x_1, \ldots, x_k\}$ be the set of items to be ranked by the judges, identified with the indexes $1, \ldots, k$. We denote the ranking given by the judges with the permutation $\pi = (\pi(1), \ldots, \pi(k))$, where $\pi(i)$ is the rank given to the item x_i and $\pi^{-1}(i)$ is the index of the item assigned to rank i. If x_i is preferred to x_j, then $\pi(i) < \pi(j)$.

We will use π and π^{-1} as vectors similar to (Lebanon and Lafferty, 2002) whose $i-$th component is $\pi(i)$ and $\pi^{-1}(i)$ respectively. Thus π and π^{-1} are the vectors representing the ranking and ordering over the set $\mathcal{X}$ respectively. If the ranking is over the entire set $\mathcal{X}$ then it forms a full ranking π, but if the judge ranks only $p < k$ items in $\mathcal{X}$ then the resultant ordering will be a partial ordering $\pi^{-1^*} = (\pi^{-1}(1), \cdots, \pi^{-1}(p))$. For brevity we represent π^{-1^*} as π^*.

2.1 Generalized Mallows' Model

According to the generalized Mallows' model, for a given dispersion parameter $-\theta$ and location parameter π_0, the judges are assumed to generate their rankings π from

$$P(\pi|\theta,\pi_0) \quad = \frac{exp\{-\theta d(\pi,\pi_0)\}}{\sum\limits_{\pi\in\Omega} exp\{-\theta d(\pi,\pi_0)\}} \quad \pi\in\Omega,\ \theta\in\mathbb{R} \tag{1}$$

where, $d(.,.)$ is a right invariant distance metric, π_0 is a fixed ranking, θ is the dispersion parameter and $\psi(\theta)$ is the normalizing constant. When $\theta > 0$, the fixed ranking π_0 is the modal ranking and when θ approaches infinity mass gets concentrated at the single ranking π_0. When $\theta = 0$ the distribution is uniform and for $\theta < 0$, π_0 is an antimode. The $\theta \geq 0$ can be interpreted as the expertise of the judges. The probability of the ranking π decreases exponentially with increase in the distance from the modal ranking π_0.

(Fligner and Verducci, 1986) extended the above Mallows model for the presence of partial orders by viewing it as a multi-stage ranking process. In the case of partial orderings, a judge reports only his top $p < k$ preferences, denoted by $\pi^* = (\pi^{-1}(1),\cdots,\pi^{-1}(p))$ and let the set of all partial orders be Ω^*. They considered the partial ordering π^* as a censored observation from the Mallows distribution (1) and modeled the probability of observing π^* as the probability of getting a full ranking from the coset $\mathcal{S}_{k-p}\pi$ of all $\pi\in\Omega$ consistent with π^*.

Let, V_j be the number of adjacent transpositions taken in the order of π_0 required to place the item $\pi_0^{-1}(j)$ in j^{th} position. For example, if π_0 is an identity permutation then, V_j is the number of adjacent transpositions required to place the item j in j^{th} position. $V_1,\cdots,V_p$ depend on $\pi\in S_{k-p}\pi$ only through π^*. The remaining vectors $V_{p+1},\cdots,V_{k-1}$ takes all its $(k-p)!$ possible values, thereby independent of π^* and a function of p and θ alone. The induced model in partial ranking can be expressed as (Fligner and Verducci, 1986):

$$P(\pi^*|\theta,\pi_0) = \frac{exp\left\{-\theta\sum\limits_{j=1}^{p} V_j(\pi^*,\pi_0)\right\}}{\prod\limits_{j=1}^{p}\left\{\frac{1-exp\{-(k-j+1)\theta\}}{1-exp(-\theta)}\right\}} \tag{2}$$

3 AGGREGATION OF IMPLICIT FEEDBACKS

As stated earlier, the feedback could be anything that impose an ordering over a subset of documents for a particular query. There are lots of implicit and explicit feedbacks proposed in the literature, and anything that can be converted into ranking with little effort would be considered as feedback for our purpose. The users are under no obligation to register their feedback over the entire document list and therefore, it will be a partial order. Since it is under the users' discretion, they might give their feedback on document sets of varying size. This demands a modification to the model (2) where the length of the partial orderings generated is fixed.

3.1 Model for Partial Orders of Varying Length

For every length p, there exists a probability model (2). Since the length of the partial orders given in user sessions vary from one session to the another, it follows a probability distribution P and let us assume it to be known. Let $p(.)$ be the function which maps the partial ordering π^* to its length. Then the extended model can be written as:

$$P(\pi^*|\theta,\pi_0) = \frac{exp\left\{-\theta\sum\limits_{j=1}^{p(\pi^*)} V_j(\pi^*,\pi_0)\right\}P(p(\pi^*))}{\prod\limits_{j=1}^{p(\pi^*)}\left\{\frac{1-exp\{-(k-j+1)\theta\}}{1-exp(-\theta)}\right\}} \tag{3}$$

$P(p(\pi^*))$ is the probability of getting an ordering of length $p(\pi^*)$. Note that this mass function $P(\cdot)$ is independent of π_0 and θ.

Consider that there are n judges giving their feedback on each of the m feedback measures, then the ordering ${\pi^*}_s^{(r)}$ obtained from the feedback r given by judge s can be considered to be generated from the extended model (3). The loglikelihood function of the model will be as follows:

$$\begin{aligned}\mathcal{L}(\pi_0,\theta^{(r)}) = \sum_{s=1}^{n}\Big\{ & p({\pi^*}_s^{(r)})ln\left[1-exp(-\theta^{(r)})\right] \\ & +ln\left[P(p({\pi^*}_s^{(r)}))\right] - \theta^{(r)}\sum_{j=1}^{p({\pi^*}_s^{(r)})} V_j({\pi^*}_s^{(r)},\pi_0) \\ & -\sum_{j=1}^{p({\pi^*}_s^{(r)})} ln\left[1-exp\{-\theta^{(r)}[k-j+1]\}\right]\Big\}\end{aligned} \tag{4}$$

Estimation of $\theta^{(r)}$: The dispersion parameter $\theta^{(r)}$ for a fixed π_0 can be estimated by solving the following equation:

$$\begin{aligned}\frac{\partial\mathcal{L}}{\partial\theta^{(r)}} = \quad \sum_{s=1}^{n}\Bigg\{ & \frac{p({\pi^*}_s^{(r)})e^{(-\theta^{(r)})}}{1-e^{(-\theta^{(r)})}} - \sum_{j=1}^{p({\pi^*}_s^{(r)})} V_j({\pi^*}_s^{(r)},\pi_0) \\ & - \sum_{j=1}^{p({\pi^*}_s^{(r)})}\frac{(k-j+1)e^{(-\theta^{(r)}(k-j+1))}}{1-e^{(-\theta^{(r)}(k-j+1))}}\Bigg\} = 0\end{aligned} \tag{5}$$

Estimation of π_0: Since the π_0 has been fixed in the estimation of $\theta^{(r)}$, the estimator $\hat{\theta}^{(r)}$ will be a function of π_0, denoted by $\hat{\theta}^{(r)}(\pi_0)$. Therefore, the ideal ranking estimate $\hat{\pi_0}$ can be obtained by iterating the above equation (5) for all $\pi_0 \in \Omega$ and by subsequently substituting the $\hat{\theta}^{(r)}$ and $\hat{\pi_0}$ values in the log likelihood function (4) to get the pair that maximizes it.

$$\hat{\pi_0} = \underset{\pi_0}{\operatorname{argmax}} \mathcal{L}(\pi_0, \hat{\theta}^{(r)}(\pi_0)) \qquad (6)$$

3.2 Multiple Feedback Aggregation

Despite the unavailability of ideal ordering of documents π_0, we would assume its existence for all the queries and we argue that it is possible to effectively estimate it through the partial orders $\{\pi^{*(r)} : r = 1, \cdots, m\}$ induced by the observable feedbacks from user sessions. These partial orders are proxy information of the ideal ordering but with difference in their precision $\{\theta^{(r)} : r = 1, \cdot, m\}$. The precision not only changes with feedbacks but also with queries. This is because of the feedbacks' ability to achieve ideal ranking changes amongst themselves as well as with queries. Let $\pi^* = (\pi^{*(1)}, \cdots, \pi^{*(m)})$ be the vector of all m partial orders given in an user session, and each component $\pi^{*(r)}$ of this vector follows an extended model (3) with dispersion parameter $\theta^{(r)}$ and ideal ordering π_0. Let us denote the vector of all these m dispersion parameters as $\theta = (\theta^{(1)}, \cdots, \theta^{(m)})$.

Formally, for a given ideal ordering π_0, the partial orders $\{\pi^{*(r)} : r = 1, \cdots, m\}$ extracted from m feedbacks given in an user session are assumed to be conditionally independent. Hence, the probability of getting π^* for a given π_0 and θ is given by:

$$P(\pi^*|\theta, \pi_0) \quad = \quad \prod_{r=1}^{m} P(\pi^{*(r)}|\theta^{(r)}, \pi_0) \qquad (7)$$

3.3 Model Benefits

The main advantages of using the distance based modeling framework for feedback aggregation are the following:

1) Since the distance used in the proposed model (3) is the Kendall distance, the maximum likelihood estimate of π_0 will be Kemney optimal and it will enjoy the important rank aggregation properties such as neutrality and consistency in social choice literature, widely known as Condorcet property (Dwork et al., 2001). 2) It has been reported in (Fligner and Verducci, 1988) that the maximum likelihood estimator for π_0 can be obtained by arranging the elements based on the vector of average ranks $\bar{\pi}$. This makes the model computationally viable for larger n. 3) Since sorting based on the average ranking $\bar{\pi}$ being the unbiased estimate of π_0, it may seem similar to Borda's method of rank aggregation, in the case of single implicit feedback. But, if there are more than one feedback ranking that need to be jointly aggregated, then the proposed method will have an edge over the Borda's method, where the weights that need to be given to the individual ranking schemes are not so obvious. Being an unsupervised aggregation framework, it will estimate the expertise of the individual feedbacks based on the given data, rather than getting it externally.

4 MODEL EVALUATION

Since the experimental evaluation of the implicit feedback is severely hampered due to the absence of an independent evaluation of the documents, the experiments are aimed at establishing the benefits of the proposed framework stated in Section 3.3.

In this experiment, a 24 hours log data recorded on 6^{th} February 2001 by AlltheWeb.com has been used. This dataset has been previously used by (Jansen and Spink, 2005) to study the emerging trends in web searching and later by (Veilumuthu and Ramachandran, 2007) to verify the existence of the incremental information in using the click sequence and time based implicit feedback measures.

Each tuple in the dataset corresponds to a click event made by an user. The log contains the userID (masked IP), clickTime (i.e., the time at which the click has been made), the query (masked) posed and the URL (masked) to which the click has been made. In the present study, only the log entries pertaining to two implicit feedbacks namely, (1) Click sequence based ordering $\pi_s^{(o)}$ and (2) Time based based ordering, $\pi_s^{(t)}$, stated in (Veilumuthu and Ramachandran, 2007) are used. For more details on order extraction from log entries, the readers are referred to (Veilumuthu and Ramachandran, 2007). Top 10 non-trivial queries (omitting queries like "google") that had sufficient number of sessions (≥ 30), have been chosen for the study. These selected queries formed the query set Q. It is a known fact that the majority of the sessions will be of length lesser than 3. Therefore, we picked only the top 5 of the document list formed by the documents that appear in the top 3 positions in either of the rankings. This is under the acceptable assumption that the documents that are ranked higher in any of the ranking will represent the data much better than the others.

Table 1: Model parameters of the proposed and Borda's aggregation models.

Query	Order based			Time based			Unified
	Proposed		Borda's	Proposed		Borda's	
ID	π_0	$\hat{\theta}^{(o)}$	π_0	π_0	$\hat{\theta}^{(t)}$	π_0	π_0
Q1	1,2,3,5,4	1.23	1,2,3,5,4	1,2,3,5,4	0.94	1,2,3,5,4	1,2,3,5,4
Q2	2,1,3,4,5	0.25	2,1,4,3,5	2,1,4,3,5	0.21	2,1,4,3,5	2,1,3,4,5
Q3	2,5,3,1,4	0.65	2,5,3,1,4	5,2,3,1,4	0.48	5,2,3,1,4	2,5,3,1,4
Q4	4,5,2,1,3	0.51	4,5,2,1,3	4,2,5,1,3	0.47	4,2,5,1,3	4,5,2,1,3
Q5	5,3,4,1,2	0.59	5,3,4,1,2	5,3,4,1,2	0.48	5,3,4,1,2	5,3,4,1,2
Q6	5,4,3,1,2	0.43	5,4,3,1,2	4,5,3,2,1	0.24	3,5,4,1,2	5,4,3,1,2
Q7	1,4,2,3,5	0.28	1,4,2,3,5	1,4,2,3,5	0.33	1,4,2,3,5	1,4,2,3,5
Q8	4,5,1,2,3	0.36	4,5,1,2,3	4,5,1,3,2	0.26	4,5,3,1,2	4,5,1,2,3
Q9	2,5,3,4,1	0.72	2,5,3,4,1	2,5,3,4,1	0.69	2,5,3,4,1	2,5,3,4,1
Q10	2,3,4,1,5	0.52	2,3,4,5,1	2,3,4,5,1	0.48	2,3,5,4,1	2,3,4,5,1

4.1 Results and Discussion

The sessionwise feedback data for the click sequence based and time based feedbacks have been aggregated through, the proposed method and the Borda's method and the results are tabulated in Table 1. From Table 1, it can be observed that the π_0 obtained through the proposed method for click sequence based and time based feedbacks are same as that of their respective Borda's aggregations (8 out of 10 cases in both the schemes). This is supportive of the fact that the maximum likelihood estimator for π_0 would converge to the ordering obtained by arranging the elements based on the vector of average ranks $\bar{\pi}$, (Fligner and Verducci, 1988).

The partial orders obtained for a feedback from various sessions are all observations from model (3). For a particular feedback, sessionwise observations are all from the same population and therefore, they all get equal weights in the aggregation process, resulting in an ordering same as that of Borda's aggregation. But in contrast, while aggregating the observations from different feedbacks, since they are all from population differing in their dispersion parameter (precision), fusioning with equal weights would be misleading. Borda's method instead gives equal weights to the ranking schemes ignoring the differences in accuracies. Eventhough the weighted Borda's method proposes to use precision based weights, it appears less attractive, because, the precision values are not readily available and need to be fed externally which is impractical. But in the proposed method, the weights (intuitively dispersion parameters) are assigned inherently by the very construction of the model and hence, the feedback rankings are combined in to final ranking of documents based on their expertise.

From Table 1, it can be observed that the $\hat{\theta}^{(t)}$s of time based aggregation are found to be lesser when compared to the respective $\hat{\theta}^{(o)}$s of click sequence based aggregation. This can be supported by the argument that the click sequence based information is ordinal and the time based information is continuous. Therefore, the time based ranking will be more sensitive, i.e., even a second difference in the reading time will force a document to be ranked differently. Though it will have a modal ranking π_0, the mass won't be concentrated heavily on one ranking. This very fact can be seen in the frequency Table 2 for a typical query. Since the time based rankings are widely dispersed, its aggregation parameter $\theta^{(t)}$ is expected to be less when compared to the aggregation parameter $\theta^{(o)}$ of highly skewed (biased) click sequence based ranking. This motivates the need of the proposed unified aggregation model (Section 3.2) which can make trade off between the two aggregations and fully extract the advantages of both the feedbacks.

Table 2: Order-based and time-based rank frequency distribution of the documents for a typical query.

URLs	Order-based					Time-based				
	1	2	3	4	5	1	2	3	4	5
URL31	6	4	1	1	0	7	3	1	1	0
URL32	1	5	1	1	1	5	2	1	1	0
URL33	4	12	4	0	0	9	8	3	0	0
URL34	12	6	0	0	0	8	9	0	0	1
URL35	16	2	1	0	0	10	7	2	0	0

From Table 1, it can be seen that for relatively lower values of $\hat{\theta}^{(t)}$ (when compared to $\hat{\theta}^{(o)}$) the unified aggregation model produces π_0 equal to that of the click sequence based aggregation and for relatively higher values of $\hat{\theta}^{(t)}$ it produces a π_0 which forms the trade off between the two implicit feedback rankings considered. As the loglikelihood of combined aggregation model is the summation of loglikelihoods of its component rankers, it could break the ties that emerge more frequently while estimating π_0s from the maximum loglikelihood. The preliminary results hinted the possibility of using the unified aggregation of feedbacks with comparably better performance. But, this need to be tested with a larger dataset for better understanding and assurance.

5 CONCLUSIONS

In this article, we proposed a generalized framework to incorporate multiple feedbacks on page quality from search engine log files to improve the result quality. Specifically, we considered the click sequence and the time spent by the users in reading a document as measures of a documents' importance for a query. We proposed an extension to the distance based raking method to jointly aggregate the feedbacks based on their expertise. This attempt is a precursor to demonstrate the feasibility of unsupervised fusion of feedbacks in the form of partial orders in to a single ranking of documents taking into account their varying levels of accuracies. The experiments were conducted on the actual search log data from AlltheWeb.com to demonstrate the strength of the proposed model in meaningfully combining the feedbacks.

REFERENCES

Agichtein, E., Brill, E., and Dumais, S. (2006). Improving web search ranking by incorporating user behavior information. In *Procs. SIGIR '06*, pages 19–26, New York, NY, USA. ACM.

Brin, S. and Page, L. (1998). The anatomy of a large-scale hypertextual web search engine. *Computer Networks and ISDN Systems*, 30(1-7):107–117.

Dwork, C., Kumar, R., Naor, M., and Sivakumar, D. (2001). Rank aggregation methods for the web. In *Procs WWW '01*, pages 613–622, New York, NY, USA. ACM.

Fligner, M. A. and Verducci, J. S. (1986). Distance based ranking models. *Journal of the Royal Statistical Society. Series B (Methodological)*, 48(3):359–369.

Fligner, M. A. and Verducci, J. S. (1988). Multistage ranking models. *Journal of the American Statistical Association*, 83(403):892–901.

Jansen, B. J. and Spink, A. (2005). An analysis of web searching by european alltheweb.com users. *Inf. Process. Manage.*, 41(2):361–381.

Kelly, D. and Belkin, N. J. (2004). Display time as implicit feedback: understanding task effects. In *ACM SIGIR*, pages 377–384.

Kim, J., Oard, D., and Romanik, K. (2000). Using implicit feedback for user modeling in internet and intranet searching. Technical report, College of Library and Information Services, University of Maryland at College Park.

Kleinberg, J. M. (1999). Authoritative sources in a hyperlinked environment. *J. ACM*, 46(5):604–632.

Klementiev, A., Roth, D., and Small, K. (2008). Unsupervised rank aggregation with distance-based models. In *Procs. ICML '08*, pages 472–479, New York, NY, USA. ACM.

Lebanon, G. and Lafferty, J. D. (2002). Cranking: Combining rankings using conditional probability models on permutations. In *Procs. ICML '02*, pages 363–370, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.

Mallows, C. L. (1957). Non-null ranking models. i. *Biometrika*, 44(1/2):114–130.

Ramachandran, P. (2005). Discovering user preferences by using time entries in click-through data to improve search engine results. In *Discovery Science*, pages 383–385.

Veilumuthu, A. and Ramachandran, P. (2007). Discovering implicit feedbacks from search engine log files. In *Discovery Science*, pages 231–242.

White, R., Ruthven, I., and Jose, J. M. (2002). The use of implicit evidence for relevance feedback in web retrieval. In *Proceedings of the 24th BCS-IRSG European Colloquium on IR Research*, pages 93–109. Springer-Verlag.