

Learning Consensus Opinion: Mining Data from a Labeling Game

Paul N. Bennett
Microsoft Research
One Microsoft Way
Redmond, WA 98052
pauben@microsoft.com

David Maxwell
Chickering
Microsoft Live Labs
One Microsoft Way
Redmond, WA 98052
dmax@microsoft.com

Anton Mityagin
Microsoft Live Labs
One Microsoft Way
Redmond, WA 98052
mityagin@microsoft.com

ABSTRACT

We consider the problem of identifying the consensus ranking for the results of a query, given preferences among those results from a set of individual users. Once consensus rankings are identified for a set of queries, these rankings can serve for both evaluation and training of retrieval and learning systems. We present a novel approach to collecting the individual user preferences over image-search results: we use a collaborative game in which players are rewarded for agreeing on which image result is best for a query. Our approach is distinct from other labeling games because we are able to elicit directly the preferences of interest with respect to image queries extracted from query logs. As a source of relevance judgments, this data provides a useful complement to click data. Furthermore, the data is free of positional biases and is collected by the game without the risk of frustrating users with non-relevant results; this risk is prevalent in standard mechanisms for debiasing clicks. We describe data collected over 34 days from a deployed version of this game that amounts to about 18 million expressed preferences between pairs. Finally, we present several approaches to modeling this data in order to extract the consensus rankings from the preferences and better sort the search results for targeted queries.

Categories and Subject Descriptors

H.3.3 [Information Storage and Retrieval]: Information Search and Retrieval—*selection process*

General Terms

Algorithms, Performance

Keywords

preference judgments, learning preferences

1. INTRODUCTION

In many information retrieval applications, we desire to rank items in relation to an information need or task. The purpose may be to display directly the ranked items to a user of the system or to use the ranking as an intermediate

step in another algorithm. Applications of ranking prediction include image search, ad hoc web search, recommender systems, ordering e-mail and a host of other problems. Often in these tasks, little attention is paid to the challenge of how to reliably identify what ranking should serve as ground truth for training and evaluation purposes. That is, what ranking of the retrieved set would best satisfy all users?

While the standard approach to ranking has been to rank items according to their relevance [14], and in particular *topical* relevance [4], an item may be preferred by a user based on a variety of other characteristics including quality, authoritativeness, and readability; these characteristics may be seen as defining a broader notion of relevance that incorporates the context of a task. For image search—which is the primary ranking task considered in this paper—a user may prefer an image because of its focus, composition, artistry, or a variety of other dimensions.

Ranking systems are typically constructed with data consisting of explicit relevance judgments for a set of training items. One approach to getting this training data is to elicit relevance judgments from a small number of “expert” judges or editors in a controlled setting. A difficulty of this approach is that, due to query ambiguity and personal preferences, it may be difficult for any single person except the query issuer to accurately judge the relevance of results.

An alternative approach to data collection is to sample judgments from a large population resembling the user base. This has motivated much of the research on mining click data. While click data has undeniable value, most notably because it comes from the issuer of the query, it also has potential weaknesses. In particular, items that have not been displayed cannot be clicked, and the lack of a click is often not informative because the search page itself may satisfy the information need. This latter problem is even more of a concern in image search because the search result page typically consists of actual images (potentially scaled down). Furthermore, if the designers of a ranking system experiment with the live system by (*e.g.*) swapping items or placing potentially non-relevant results in the top of the list [12, 13], there is a risk of frustrating the user and prompting him to switch search engines.

In this work, we attempt to solve the data-acquisition problem in the domain of image search with a social labeling game [19]. The game, which we describe in detail in Section 3.1, pairs two participants on the internet who are shown a sequence of queries and corresponding sets of image results; they are both asked to choose the best image for each query,

Copyright is held by the International World Wide Web Conference Committee (IW3C2). Distribution of these papers is limited to classroom use, and personal use by others.

WWW 2009, April 20–24, 2009, Madrid, Spain.

ACM 978-1-60558-487-4/09/04.

and every time they agree they are awarded credits. After collecting enough credits, they may turn the credits in for prizes. Assuming participants are primarily seeking credits or prizes, the incentive mechanism encourages users to give their actual opinion of the best image because it is a good way to achieve agreement with high likelihood.

We study several probabilistic models that can be used to convert the preference data that results from game play into a set of relevance scores for the items; these relevance scores can be used for evaluating an existing system, training a new ranking algorithm, or simply improving the retrieval set for the particular queries in the data.

In the remainder of the paper, we first give an overview of related work, including a more in-depth discussion of why current social labeling games fall short, and a survey of preference-learning models of particular relevance to this work. We then present a detailed description of our labeling game. Next, we present several methods for mining the resulting preference data to learn rankings. Then, we present our results and discuss their implications. Finally, we conclude with a discussion of interesting future directions.

2. RELATED WORK

There have been a number of human-computation games developed for the purpose of collecting data. Arguably the most successful such game is the ESP Game [19], where as a result of game play, participants produce descriptions of images. Others include Peekaboom [20], where the players produce data about where objects are located in an image, and TagATune [10] where players produce descriptions of sounds and music.

The image-description data produced by the ESP game could in principle be used to help infer ranking quality; for example, if players annotate one image with “dog” more often or sooner than a second image, we might reason that the first image is a better result for the query “dog”. One problem with this approach is that the descriptions that are associated with an image are determined by the players and not the game designer. A second problem is that, because players can earn more points in the ESP game if they match quickly, there is incentive to provide very short and obvious descriptions that may not be very informative (*e.g.*, prevalent colors in the images). That is, it is not clear that the task of trying to get a match with a partner on tags (which leads one to give common tags) maps closely to asking for the tags for which this image is one of the most relevant results. Instead, the ESP game is most suited for addressing the indexing aspect of search (identifying a large set of relevant items) while our approach targets the ranking aspect (how to only present a small subset of very relevant items out of a large number of matching items).

While obtaining *absolute* relevance judgments over a particular image/query pair may seem like a better way to obtain judgments, Carterette et al. [2] demonstrate that absolute judgments are less reliable in terms of agreement with other assessors than *pairwise preferences*, *i.e.*, relative comparisons between two items. In Section 4.2, we discuss several methods that are based on pairwise preferences.

In terms of preference learning, Carterette et al. [2] address how to infer a ranking given only a linear number of preferences from a single assessor but do not address how to use preferences across a population of users. Fürnkranz and Hüllermeier [7] present a study on breaking label pref-

erence ranking (*e.g.*, predicting order of preferred classes in a multiclass classification problem) into a series of binary classification tasks [7]. The primary issue that concerns us here, however, is one of instance-preference ranking. Chu and Ghahramani [3] propose a Gaussian-process model applicable to both instance and label preference learning. The model is most similar to the Go model we present but does not have some of the flexibility to model player abilities.

If we consider the challenge proposed here as learning to predict the preferences of all users given a small sample of users, one natural solution is to consider averaging the preferences of the users in the training set and then extracting the optimal ranking for the averaged preferences to predict on the testing set. Cohen, Schapire, and Singer [5] demonstrated that the problem of extracting the optimal ranking is NP-complete, and whereas they provide useful approximation algorithms for this, we are more focused on the second question of how well this generalizes to the full population. Similarly, as alternate models to those presented here, we could consider this a task of rank aggregation [6] and apply rank aggregation methods to extract the consensus ranking. Whereas a number of these approaches would require rankings from users instead of preferences, these rankings could easily be extracted from the preferences first (*e.g.* run quick-sort for each user). We note that in our setting, we do not ask directly for rankings because the increased complexity in the task both increases noise in response and interferes with the fast-paced excitement of the game. Finally, Joachims [9] proposes ranking SVMs for preference learning and applies them to preference learning using click data. Our work is most similar in flavor to this, but we deviate in both the source of our data and the scale.

In contrast to previous work on social labeling games, our work is the first to use queries drawn from search logs to drive image rankings. As a result, the language distribution actually seen on the search engine drives the labeling. Furthermore, ours is the first attempt to study the suitability of large-scale preference data as a signal to augment other sources of relevance information such as click data and assessor judgments. In particular, our approach is unique in its ability to obtain information like click data but with positional biases removed and with no risk of frustrating users when given non-relevant results. Finally, from a learning perspective, we contribute a systematic study of a number of preference models.

3. PREFERENCE COLLECTION AS GAME

Picture This is a collaborative game where pairs of players are rewarded for agreeing on the best result for image-search queries. In this section, we describe the game in detail. Currently the game can be accessed online via http://club.live.com/Pages/Games/GameList.aspx?game=Picture_This.

3.1 Game Play

When a user starts the *Picture This* game he is first paired with a random partner. If no human player is available within a few seconds of the player starting the game, the player is assigned a robot as a partner; we describe robots in detail in Section 3.3. The two players then proceed through a set of rounds, working together for either two minutes or until they reach 100 points, whichever comes first. The players are synchronized such that the next round starts only when both players have completed the current round.




Figure 1: The *Picture This* game interface. The query shown at the top is “Microphone” with eight candidate images. The green boxes in the grid at the top right show how many of the 100 possible points the player has earned; the orange boxes with diamonds show how many points will be earned on agreement (*i.e.*, k). In this case, the partner has made his selection (indicated in lower right corner) and is waiting for this player to choose.

In each round, the two players are shown a query and two different random permutations of the same set of k images. The images are permuted both to eliminate positional bias in the preference data and, as we discuss in Section 3.4, to mitigate against fraud. k is initially set to three and changes throughout the game. Each player selects the image that, in his opinion, best matches the query. If the two players agree on the best image, they are both awarded k points and k is incremented by one if $k < 9$. If the players disagree on the best image, they are awarded zero points and k is decremented by one if $k > 2$. Thus the number of images displayed varies from 2 to 9 depending on player actions: when users agree, the number of images displayed (and the difficulty of the game) increases, and when users disagree, the number of images displayed decreases. Adapting the game difficulty to a player’s performance allows us to both (1) take advantage of discerning players by effectively getting more preference judgments per click and (2) make the game more entertaining.

Players have the option to choose “no good image” to indicate that none of the images are a good result for the query. If either player chooses this option, the number of images k is incremented or decremented as usual depending on agreement, but no points are awarded for agreement. We discuss this decision further in Section 3.4. Players can also flag individual images as being “bad”. If two players flag the same image as “bad”, then they are awarded a time bonus of five seconds if they also agree on the best image. Flag matches are not rewarded if the players agree on “no good image”.

In Figure 1 we show what the game interface looks like when a player is choosing an image. Note that the interface

indicates in the lower right corner that the partner has already chosen. After the player makes his own selection, his partner’s choice will be revealed.

After two minutes has passed or the pair of players has attained 100 points, the game ends. For players signed into the game site, these points are converted into a currency that can be spent on various items including t-shirts, computer hardware, computer software, and music. Players then have three options: they can choose to play again with the same partner, they can choose to play again with a new partner, or they can quit the game. If one player requests to play again with the same partner, that partner is given a choice to either accept or reject the invitation.

3.2 Query and Image Selection

The queries shown for each round in a game of *Picture This* are chosen randomly from a specific list. This list is updated over time, but the data that we analyze in this paper results from game play when the list contained a specific set of 427 queries that were both of interest to the image-search team and were appropriate for a public game (*e.g.*, they were not likely to return offensive images).

The candidate images for the query in each round are the top 50 results extracted from Live Image Search. One could, of course, use different values of n as the number of top images from which to select. We were partially guided by the differences in image search from web search. In a variety of commercial search engines, the default is to display 21-35 images as the first page of results, organized into a matrix-like format rather than a single list. Thus, we chose an n large enough to find relevant images which should be in the first page but are not currently, as well as keeping n small enough to easily get ranking feedback as to the desired full ordering of the images.

With each image, we store the fraction of rounds in which that image was chosen best, and we use this fraction as a simple estimate of image quality; we discuss this model further in Section 4.2.1. On each round, we either choose the k images uniformly at random from the entire set of 50 results, or we choose a random “base” image and select uniformly at random from the $2k$ images that have the smallest difference in quality estimate from the base image. Rounds in which images are chosen to be close to a base image are typically both more difficult for the user and more informative to the system. We choose between the “hard” method and the “easy” method randomly, where the “easy” method is currently applied in roughly a third of the rounds. We did this to balance making fine-grained distinctions between images that are close to the same level of relevance with learning the overall ordering. Note that this is a simple form of active learning and depending on what the goals of the system are, a different active learning mechanism can be employed.

We emphasize that both the query and image pool can be constructed according to a variety of choices. For example, if no base search engine was available, one might choose to display a large number of images by default to quickly find some relevant images, and then compare winners of those rounds to get a refined ranking. Likewise, if some images are not currently in the top set, one can inject them into the result set to find out if they are in fact highly relevant. One might choose to adopt this strategy if the game is being used to explore new features for ranking and the injected images could be images that would be highly ranked if that

feature is given increased weight. In addition to our data-collection needs, the choice of the image-selection algorithm was informed both by architectural constraints in the system and by the enjoyment of game play.

3.3 Partner Robots

For players who could not be matched with a human partner within a small amount of time after starting the game, we implemented partner robots. These robots choose images based on preference estimates in the data, essentially using the method of Section 4.2.1. Unlike the ESP game [19], human behavior is not easily mimicked by simply replaying the choices (and timing) from previous game transcripts because the number of choices that are available for a query will vary from game to game. For example, the best image chosen by a player when two images are shown may be an obviously bad choice when there are additional more relevant images.

3.4 Incentive Structure and Fraud Mitigation

Because game points have real-world value, there is incentive for game players to cheat. In fact, there are entire discussion boards dedicated to the topic of cheating games on the Microsoft Live Search Club site. In this section, we describe a number of design decisions we made in order to mitigate against this threat.

Because points are awarded based on agreement, most cheating methods require coordination between two players. Our first line of defense against these methods is the random assignment of partners, although a fraudster can attack the game with several bots that start at the same time in order to increase the probability of getting a coordinated pair of players.

As described above, the two players are shown the candidate images in different orders and therefore strategies such as always choosing the first image will be no better than random selection. Because we do not award points for agreement on the “no good image” selection, the strategy of always making this choice does not pay off.

We were particularly concerned about defending against a random-selection strategy. Recall from Section 3.1 that the number of points earned for agreement is equal to the number of images shown; this arithmetic progression of the rewards helps to some degree because—relative to normal play—the expected value per round is low for random play. But if the system is attacked by two random-voting (and quickly voting) bots that are paired together, the total payoff could be significant. We have three major lines of defense against random-voting bots. First, all players are occasionally challenged with a puzzle (a CAPTCHA) that is difficult for an automated algorithm to pass. Second, the game site monitors all user-interaction events of the games and applies a bot-detection algorithm to this data. Finally, we have implemented a number of “honeypot” queries in the game for which a human has verified that there is an obvious best image; if a player chooses the wrong image for these queries, he is flagged as a potential cheater.

4. PREFERENCE MODELS

4.1 Data Representation

As described in Section 3.1, players can choose “no good image” instead of a best image, and they can also flag individual images as being bad. We model these choices in our

preference data by including a virtual “neutral” image for every query in the data. In particular, whenever a player selects “no good image”, this is recorded as a preference for the neutral image over all the candidate images. Similarly, whenever a player flags an image as bad, we record this as a preference for the neutral image to that bad image. Finally, every time a player chooses an image instead of choosing “no good image” or flagging that image as bad, this is interpreted as a preference for the chosen image over the neutral image. Thus, the neutral image is implicitly a part of every comparison set. From a modeling perspective, this has the advantage of making the neutral image serve as an anchor and ensures that every image that is chosen at least once or part of a comparison when “no good image” is chosen is connected in a preference graph.

Now we abstract the representation of a collection of data to a more general learning setting. For a given query, we assume there is a set of items \mathcal{I} with $I = |\mathcal{I}|$ elements; in our domain, these are all images displayed during the use of the query in game play plus the virtual “neutral” image described above. We will typically refer to the i th element of \mathcal{I} as item i or image i . We are given a set of labeled data pertaining to the query $\mathcal{D} = \{o_1 = \langle r_1, \langle c_{1,1}, \dots, c_{1,\mathbb{S}(1)} \rangle \rangle, \dots, o_N = \langle r_N, \langle c_{N,1}, \dots, c_{N,\mathbb{S}(N)} \rangle \rangle\}$ consisting of $N = |\mathcal{D}|$ observation records, o_k , where $\mathbb{S}(k) = \text{length}(\langle c_{k,1}, \dots, c_{k,\mathbb{S}(k)} \rangle)$, $r_k, c_{k,l} \in \{1, \dots, I\}$, and $r_k \in \{c_{k,1}, \dots, c_{k,\mathbb{S}(k)}\}$. Each observation corresponds to the action of one player in a round. The $c_{k,l}$ are the items being compared in the k th observation while r_k is the item selected from those compared in the k th observation. For shorthand, we will refer to the tuple of comparison items in an observation as $C_k = \langle c_{k,1}, \dots, c_{k,\mathbb{S}(k)} \rangle$.

4.2 Methods

In this section, we describe three different methods for modeling the preference data. We chose these models because of their progression from the most naïve probabilistic approach one might take with the grossest independence assumptions to increasingly tailored approaches with fewer independence assumptions. In particular, the frequency model simply models the global win probability of the image when it is displayed and does not account for any interaction between items. The pairwise probability model accounts for pairwise interactions, but not interactions among the entire comparison group. Finally, the Go model accounts for interactions among the comparison set by conditioning on the current example when learning and predicting. In all cases, the final set of model parameters are linear in the number of images. We restrict our attention to probabilistic models because the probabilities are useful for future improvement of our active-learning and image-selection components.

4.2.1 Frequency Model

Our simplest model for the relevance score s_i for each image i is to use the fraction of times image i wins out of the number of times it appears. We refer to this as the *frequency* model. In practice, smoothing is necessary and we use a Bayesian m-estimate with a uniform prior over chosen or not chosen and two virtual observations. This simply reduces to:

$$P(i \text{ chosen} \mid i \text{ shown}) = \frac{|\{C_k \mid i \in C_k, i = r_k\}| + 1}{|\{C_k \mid i \in C_k\}| + 2} \quad (1)$$

To use this model to predict the probability that i wins

for some particular comparison set, you can compute:¹

$$P(i = r_k | C_k) = \frac{P(i \text{ chosen} | i \text{ shown})}{\sum_{j \in C_k} P(j \text{ chosen} | j \text{ shown})} \quad (2)$$

This results in a model where the predicted winner depends on how high above the competition the relevant item is predicted to be. Given a large enough number of random comparisons, we expect this model to perform well because, in the long run, any particular image will have been compared against all other images a large number of times. Given non-random samples (*e.g.*, always comparing against very relevant images would yield a disproportionate number of losses) or a small number of samples relative to the n^2 pairs, the model does not have the structure to distribute relevance scores in a transitive manner. Thus, it may not prove the most ideal choice when the goal is to quickly learn the ranking. Without empirical study, it is unclear what number of comparisons is needed for it to perform well and whether more complicated models are justified.

4.2.2 A Pairwise Probability Model

Because we are measuring responses by humans, it is natural to look to psychometric theory for insights into this area. Thurstone [16, 17, 18] formulated a general law of comparative judgment that hypothesizes a relationship of the human ability (which has been repeatedly born out in empirical observations) to make pairwise comparisons between items based on a number of factors such as the actual difference between the quantities of the two items being judged (*e.g.*, actual difference in weight when trying to determine which item is heavier), the variance of the distribution, and the magnitude of each quantity involved. Under certain assumptions, and using the logistic function as the underlying distribution, this can be used to derive the Bradley-Terry-Luce (BTL) pairwise comparison model [1, 11].

In the BTL model, assuming a set of location parameters s_i (in our case, “relevance score” parameters), the probability of selecting one item as being greater (“more relevant”) than another is modeled as:

$$P(i \succ j) = \frac{\exp\{s_i - s_j\}}{1 + \exp\{s_i - s_j\}} \quad (3)$$

To interpret the data as pairwise preferences, we assume an observation $\langle r_k, \langle c_{k,1}, \dots, c_{k,S(k)} \rangle \rangle$ amounts to $S(k) - 1$ preferences expressing the response item is preferred to each remaining item, $\{r_k \succ j \mid j \in C_k, r_k \neq j\}$. For a pairwise probability model, modeling item i is preferred to item j (denoted $i \succ j$), we assume no mass is reserved for a tie and that $P(i \succ j) = 1 - P(j \succ i)$. Given a pairwise probability model, $P(i \succ j)$, we assume an observation’s probability is the product of all expressed pairwise preferences. That is, the $P(o_k) = \prod_{j \in C_k, r_k \neq j} P(r_k \succ j)$. To set the parameters by maximum likelihood, we must maximize the following function (see Appendix A):

$$\sum_{i=1}^I \sum_{j=i+1}^I [B_{i,j}^+(s_i - s_j) - n_{i,j} \log[1 + \exp\{s_i - s_j\}]] \quad (4)$$

¹The other obvious use of these terms such as $P(i = r_k | C_k) = P(i \text{ chosen} | i \text{ shown}) \prod_{j \in C_k - i} [1 - P(j \text{ chosen} | j \text{ shown})]$ would yield a deficient probability model in that it allocates non-zero probability to zero probability events because exactly one item must be chosen; we could consider this model where equality is replaced with proportionality.

where $n_{i,j}$ is the number of times i and j were compared and i or j won and $B_{i,j}^+$ is the number of times i and j were compared and i won.

Equation 3 is simply a linear model with a “1” on the winner and a “-1” on the loser, and it turns out that we can use a standard logistic-regression algorithm to optimize the s_i values given training data. In particular, it can be shown (see Appendix A for a full derivation) that the model is identical to a two-class ($c \in \{0, 1\}$) logistic regression model that does not have a bias term.² To represent the pairwise model as a logistic regression problem, each pair of preferences derived from a comparison set is encoded as a feature vector with the same number of features as images. All feature values are 0 except for the two features corresponding to the preference being expressed. The feature with the lower *index* is set to 1 and the higher to -1. Finally, if the image corresponding to the feature with the lower index was preferred, the class is set to 1 and otherwise to 0.

Because of the easy availability of tested code implementing logistic regression, we chose to implement the optimization in this way. In both the original pairwise formulation and in the reduction to logistic regression, it may be advantageous to introduce a regularizer or prior to deal with sparse data. We use a L2 regularizer (Gaussian prior) which will pull the relevance scores of the images to zero unless the data dictate otherwise. The weight (λ) of the regularizer is 1. Early experiments in which we investigated optimizing the regularizer weight over validation data indicated the results were not particularly sensitive to changes in the regularizer weight.

Note that Equation 3 can be rewritten as:

$$P(i \succ C_k) = \frac{\exp(s_i)}{\sum_{j \in C_k} \exp(s_j)} \quad (5)$$

where $i \in C_k$ for those comparison sets containing two images. We can also use this equation directly to predict for comparison sets of *more* than two images. Assuming the resulting prediction model when training instead of Equation 3 would require the use of a conditional model (conditioning on the current comparison set). Since Section 4.2.3 presents an alternative conditional model, we present this to gain an understanding of whether the conditional model is needed to gain the majority of the predictive power.

4.2.3 Go Model

The final model was first applied to learning to predict moves in the board game Go [15] and is related in some aspects to the TrueSkillTM model used to rank and match players in online gaming [8]. The Go model is a Bayesian system with conditional online updates that can be understood as a general preference learning algorithm which learns to predict when the winning item (originally which Go move is selected) is preferred from a set of items (all Go moves available given the current board).

There are several interesting properties of the Go model from a preference modeling point of view. First, the Go model does not “tie” the parameters of non-winning items. In contrast, many preference models implicitly or explicitly treat each preference vote as evidence that the losing items have the same relevance. The Go model is a *conditional*

²A non-zero bias yields probabilities such that $P(i \succ j) \neq 1 - P(j \succ i)$.

