

A Probability Ranking Principle for Interactive Information Retrieval

Norbert Fuhr
University of Duisburg-Essen, Germany*

Abstract

The classical Probability Ranking Principle (PRP) forms the theoretical basis for probabilistic Information Retrieval (IR) models, which are dominating IR theory since about 20 years. However, the assumptions underlying the PRP often do not hold, and its view is too narrow for interactive information retrieval (IIR). In this paper, a new theoretical framework for interactive retrieval is proposed: The basic idea is that during IIR, a user moves between situations. In each situation, the system presents to the user a list of choices, about which s/he has to decide, and the first positive decision moves the user to a new situation. Each choice is associated with a number of cost and probability parameters. Based on these parameters, an optimum ordering of the choices can be derived - the PRP for IIR. The relationship of this rule to the classical PRP is described, and issues of further research are pointed out.

1 Introduction

Interactive retrieval systems have become a commodity today. Although there is a significant amount of research on this type of systems, the theoretical foundation for this type of system is still in its infancy. Most work has focused on cognitive issues or usability aspects. Empirical studies of complete systems mostly focus on variations of single components.

*norbert.fuhr@uni-due.de

Given this state of research, the construction of a good interactive IR system is still a task for which there are only some guidelines concerning certain aspects of the system. However, for the core problem, namely performing effective retrieval in such a setting, no solid knowledge is available. The classical probability ranking principle (PRP) [Robertson 77] forms the theoretical basis for optimizing the results of ad-hoc retrieval. On the other hand, experiments in interactive retrieval ([Voorhees & Harman 00, Turpin & Scholer 06]) have shown that systems performing quite differently in the standard retrieval setting (i.e. ad-hoc retrieval for a fixed query) are indistinguishable when being used in an interactive setting. Further studies [Turpin & Hersh 01] pointed out that this is due to the fact that scanning through document lists for identifying the relevant entries is not the most crucial activity in interactive retrieval.

In this paper, we develop a framework for extending probabilistic IR approaches to interactive information retrieval (IIR). For that we develop an abstract view of the functional level of an IIR system, and then derive certain desirable properties that a system should fulfill at this level. Ultimately, this leads us to the formulation of a PRP for IIR.

The remainder of this paper is structured as follows: First, we briefly revisit the classical PRP and point out its shortcomings. In section 3, we describe the basic concepts of our approach, followed by the development of a cost model in section 4. Based on these notions, we are able to derive the PRP for IIR in Section 5. In Section 6, we describe first steps towards applying this theoretical framework. Section 7 gives a survey on related work, before the final section concludes and gives an outlook on further research.

2 Motivation

The classical PRP focuses on the task of retrieving relevant documents for a given, fixed information need. The major assumption in this model is that the relevance of a document to a query is independent of the relevance of other documents the user has seen before. The task addressed by the PRP is the user's scanning through the list of ranked documents.

Both the independence assumptions and the restriction to the scanning task are questionable:

1. In real settings, relevance always depends on documents the user has seen before. Besides the trivial case of duplicates (which happens frequently during Web retrieval), often a user wants to find relevant documents that provide different answers to a given problem (aspectual

recall). Thus, the relevance of any additional relevant document clearly depends on the relevant documents seen before.

2. Interactive retrieval consists of user actions of various types, and scanning through document lists for identifying the relevant entries is not the most crucial activity in interactive retrieval [Turpin & Hersh 01]. In contrast, other activities (like e.g. query reformulation) seem to be more 'expensive' from the user's point of view.

Somewhat related to the first point, there is the empirical finding (see e.g. [O'Day & Jeffries 93]) that user information needs are not static throughout a search, they change in reaction to the information a user has seen already. In light of this result, relevance feedback methods can hardly work, since they try to optimize the query formulation for an information need that is assumed to be static; instead, we are dealing with a moving target.

So we see that the assumptions underlying the classical PRP are not appropriate for interactive retrieval, and its focus on the result list misses the major part of the interaction, thus yielding at best a local optimization.

3 Approach

3.1 Requirements

In order to develop a PRP for interactive IR, we aim at fulfilling the following requirements:

- *Consider the complete interaction process:* Instead of focusing on document ranking, the new approach should cover all kinds of interactions of a user with an IR system, like e.g. browsing through lists of related terms, categories or cluster labels, looking at summaries of varying granularity (e.g. [White et al. 05]), or following links between documents.
- *Allow for different costs and benefits of different activities:* The types of activities in IIR require different effort (e.g. selection of a proposed expansion term may be cheaper than finding a synonym for a search term). Vice versa, the benefit resulting from an action may also vary – modifying a query will often have a bigger effect than declaring a single document to be relevant.
- *Allow for changes of the information need:* Finally, the model should be more dynamic than the classical PRP. In principle, any positive

information a user finds during a search may change his information need — like e.g. in the berrypicking model [Bates 89].

3.2 Assumptions

Like in all probabilistic IR models, our approach refers to the system’s representation of documents and information needs (see e. g. [Fuhr 92]). Since we are dealing with interactive retrieval here, our model refers to the system’s knowledge about the state of the search. Thus, in contrast to cognitive models, which may refer to certain users’ states of mind, our model can only take into account information that is available for the system — either through direct input by the user, or by appropriate sensors (like e. g. an eyetracker — future systems might even observe the user’s face in order to detect satisfaction or disappointment).

Based on the requirements listed above, we formulate the following assumptions underlying our approach:

- *Focus on the functional level of interaction:* Although human-machine interaction involves a variety of usability and visualization issues, we want to restrict here to a purely functional level. That is, the same activity (e.g. selecting expansion terms from a list) may require different effort, depending on the actual design of the interaction. These aspects may affect the values of certain parameters in our model, but we do not consider this issue here.
- *Decisions are the major interaction activity:* As the most important cognitive activity of the user, we focus on decision making. Thus, we assume that the system offers binary choices to the user, who in turn has to decide about these choices. In case the user accepts a choice, we call it a positive decision, otherwise negative. In the positive case, if the user does not want to modify the decision as soon as he learns about its consequences, we call the decision ‘correct’. (Without explicit feedback from the user, however, the system will not be able to distinguish this case from the one where the user found the resulting information relevant and then went back to the original list in order to check the next item.) The evaluation of choices may require cognitive acts of various size (e.g. looking at a single proposed term vs. reading through a full document), which is accounted for by the effort attributed to the actual choice. Creative actions like entering a new term that was not proposed by the system are also regarded as choices here — obviously with a much higher cognitive effort than in the case of a selection from an explicit list of alternatives.

- *Users evaluate choices in linear order:* This means that there is a (explicit or implicit) linear order in which a set of choices is evaluated. Examples could be explicit linear lists, but also the set of links occurring in a text. When there is no such order, we can split up the set of choices so that we have linear orders within each of its subsets, and assume that the user makes the explicit decision to move to another subset. Also, there may be cases where no strict linear order is given (e.g. the 'tag clouds' used by many popular Web 2.0 sites), but the total order considered in this paper can be used for deriving such a partial order. There may be user interface designs where several lists are presented simultaneously (e.g. [White et al. 05]), and the system does not know in what order the the user evaluates these lists — unless we use an eyetracker. As an approximation, one can assume that the user regarded only the list where he made an explicit, positive decision, since the system will use this information for recomputing all the lists currently shown. Further research will be needed for validating this assumption, or deriving better approximations in cases of incomplete knowledge.
- *Only positive, correct decisions are of benefit for a user:* This is the strongest assumption we have to make. There are many non-IR examples of decision-making where both accepting or declining a choice have a certain benefit (because of the usually limited number of choices, rejection implicitly means a restriction to the small set of alternative choices). However, the spectrum of choices in IR typically is rather large, so that the system can conclude hardly anything useful from the rejection of a choice (e.g. even when the user has given negative relevance feedback to all the documents he has seen before, the system has no information on how to improve the query).

3.3 Situations

As an important new concept, we introduce the notion of a *situation*. A situation reflects the system state of the interactive search a user is performing. In terms of our model, a situation consists of a list of choices the user has to evaluate in this situation. The first positive decision by the user will move him to another situation (depending on the choice he selected positively). In order to avoid the user getting stuck in a situation, we assume that there is always a last choice that will move him to another situation with an alternative list of choices (e.g. when the user has found no relevant document, the system might propose terms for modifying the query or browsing of docu-

ment clusters). This 'last choice' is not covered by our model, since we are focusing on the order of choices, which does not affect the 'last choice'.

From a system's point of view, its knowledge about the user's information need does not change during a situation, knowledge is added only when switching to another situation due to a positive decision. Vice versa. we can assume the information need to be static while the user is within the same situation, but a transition to another situation may change the information need. By taking this approach, we implicitly also drop the PRP assumption of the independence of relevance judgments: A positive relevance judgment may change the information need, and thus a previously relevant document now may become irrelevant for the user.

4 A Cost Model for IIR

4.1 Situations, choices and expected benefits

For modeling the interaction, we assume that the user moves from situation to situation. In each situation, the user is presented a list of (binary) choices, about which he decides in sequential order. The first positive decision moves the user to a new situation. A decision requires some effort, and with a certain probability, will be positive. There is some benefit from a positive decision, provided that the decision was correct.

In each situation s_i , we have a set of choices $C_i = \{c_{i1}, c_{i2}, \dots, c_{i,n_i}\}$. Then we define p_{ij} as the probability that a user in situation s_i will accept choice c_{ij} . (The precise specification of the underlying event space is given in the appendix.)

The only independence assumption we now have to make is the following: the probability of a user accepting a choice c_{ij} is independent of the choices he rejected before. In most cases, this supposition will be fairly valid (e.g. ranked list of documents, or list of expansion terms). Please note that this assumption is much weaker than that of the classical PRP, where independence of both positive and negative relevance judgments is assumed. With this presupposition, we exclude any sequence effects, i.e. changing the order of the choices being presented does not affect their probability of being accepted.

Furthermore, let q_{ij} denote the probability that acceptance of this choice is not revised later. In addition, we assume that $p_{ij} > 0$ for $j = 1, \dots, n_i$ (it does not make sense to offer choices a user certainly will reject, and some of the derivations given below are valid for $p_{ij} > 0$ only).

In addition to these probabilistic parameters, we introduce three cost

factors. Since we are interested in maximizing the benefit of a user, we will use the term 'benefit' for referring to negative costs, and specify all parameters as benefits. The decision about the choice c_{ij} requires the effort $e_{ij} < 0$. In case of acceptance, and if the decision was right, the resulting benefit will be b_{ij} ; if the decision was wrong, the additional effort for correction is $g_{ij} \leq 0$.

With these parameters, we can estimate the expected benefit of choice c_{ij} as

$$E(c_{ij}) = e_{ij} + p_{ij} (q_{ij} b_{ij} + (1 - q_{ij}) g_{ij}) \quad (1)$$

Since we are describing a general framework in this paper, we do not address the issue of estimating the parameters p_{ij} , q_{ij} , b_{ij} , e_{ij} , and g_{ij} here; these parameters are specific to the underlying model and the actual design of the user interface. In Section 6, however, we discuss some approaches for parameter estimation.

term	n_i	p_{ij}	b_{ij}	$p_{ij} b_{ij}$	$n_{q'}$	q_{ij}
program	195 mill.	0.67	0.4	0.268	3	-0.5
blend	5 mill.	0.02	4.0	0.08	116	64
island	2 mill.	0.01	4.9	0.049	290	188

Table 1: Example: query refinement and expected benefit

As an illustrating example, assume that a user enters the term t_0 ='Java' in a Web search engine, which yields n_0 =290 mill. hits. Now the system proposes three terms t_i for query refinement along with their number of hits n_i , as shown in table 1¹. As probability of acceptance, we have assumed that $p_{ij} = n_i/n_0$ (i.e. query terms follow the same frequency distributions as document terms); furthermore, as a rough expression of the cost of a choice, we chose $b_{ij} = \log \frac{n_0}{n_i}$ (as an information theoretic measure for the gain when narrowing down from t_0 to t_i). Obviously, the benefit for the less common terms 'blend' and 'island' is much higher than that for 'program'. On the other hand, the expected benefit — approximated here by $p_{ij} b_{ij}$ — is lower for the latter. This outcome seems to be reasonable: most users will be interested in Java programs, thus this choice should be presented first. For a minority of users, however, the other two choices would be very helpful.

4.2 Maximizing expected benefit

In a good IIR system, the expected benefit of the choices presented to the user should be as high as possible. As a first conclusion from eqn (1), we

¹The columns headed $n_{q'}$ and q_{ij} are explained in Section 6.

can say that the expected benefit of any choice presented to the user should be positive — otherwise the user would not gain anything from a choice. This condition already limits the set of choices to be presented to a user. As an implicit consequence of this statement, choices with $p_{ij} = 0$ should not occur in the selection list, since their expected benefit will be negative (due to $e_{ij} < 0$).

Regarding a single choice c_{ij} , our major goal is of course the maximization of its expected benefit. Given that the benefit b_{ij} and the backtracking effort g_{ij} of a decision are fixed, there are three strategies for maximizing $E(c_{ij})$:

1. Minimizing the effort $|e_{ij}|$. However, this may lead to more erroneous decisions, thus reducing the other addends of the expected benefit. So the system should provide enough information for avoiding too many erroneous decisions.
2. Maximizing the 'selection probability' p_{ij} , i.e. the user should choose c_{ij} whenever it is appropriate. At the same time, however, the 'success probability' q_{ij} should not drop. This can only be achieved if the user spends more effort on the decision, which increases e_{ij} .
3. Maximizing q_{ij} by avoiding erroneous positive decisions (but keeping p_{ij} high): Again, this will increase the user's effort for deciding about a choice.

Overall, we can see that the system has to find a good compromise between these three strategies in order to maximize the expected benefit.

As a simple example, assume that the system proposes some terms for query expansion. As one possibility, only the terms themselves are listed. Alternatively, for each term, the system could show a few example term occurrences in their context, thus giving the user some information about the usage of the term. The user effort per choice is lower in the first case, but the decisions will also be more error-prone.

5 Optimum Ranking for IIR

So far, we have regarded single choices, and discussed ways for optimizing the expected benefit of a choice. Now we consider the complete set of choices to be presented in a situation. As mentioned above, we assume that these choices are presented in linear order. So we have the problem of arranging the set of choices in an optimum order — which will ultimately lead us to the PRP for IIR.

Symbol	Meaning
S	set of situations
s_i	situation
C_i	set of choices in situation s_i
n_i	$= C_i $ (number of choices)
c_{ij}	single choice
p_{ij}	probability that choice c_{ij} is accepted
q_{ij}	probability that acceptance of c_{ij} is correct
e_{ij}	user effort for evaluating c_{ij}
b_{ij}	benefit from accepting c_{ij}
g_{ij}	correction effort if c_{ij} was chosen erroneously
a_{ij}	'average benefit' if c_{ij} is accepted

Table 2: List of symbols

In order to simplify the following discussion, let

$$a_{ij} = q_{ij}b_{ij} + (1 - q_{ij})g_{ij}$$

denote the 'average benefit' of a choice, thus simplifying the formula for the expected benefit to

$$E(c_{ij}) = e_{ij} + p_{ij}a_{ij}$$

5.1 Expected benefit of a selection list

Now we assume that the set of choices C_i of a situation s_i is ordered in a linear list $r_i = \langle c_{i1}, c_{i2}, \dots, c_{i,n_i} \rangle$.

For computing the expected benefit for this list, we assume that the user considers the choices in linear order, and the first positive decision will move the user to a new situation.

$$\begin{aligned}
E(r_i) &= e_{i1} + p_{i1}a_{i1} + \\
&\quad (1 - p_{i1})(e_{i2} + p_{i2}a_{i2} + \\
&\quad (1 - p_{i2})(e_{i3} + p_{i3}a_{i3} + \\
&\quad \dots \\
&\quad (1 - p_{i,n-1})(e_{in} + p_{in}a_{in}))
\end{aligned} \tag{2}$$

$$= \sum_{j=1}^n \left(\prod_{k=1}^{j-1} (1 - p_{ik}) \right) (e_{ij} + p_{ij}a_{ij}) \tag{3}$$

(Here we assume that the iterative product yields 1 for the case of an empty range.)

5.2 Optimum ranking of selections

For discussing the optimum ranking of selections, we are regarding an arbitrary pair of choices c_{il} and $c_{i,l+1}$ which appear in adjacent order at positions $l, l+1$ (with $1 \leq l < n_i$) in the list of choices. Then we can rewrite the expected benefit $E(r_i)$ as follows

$$E(r_i) = \sum_{\substack{j=1 \\ l \neq j \neq l+1}}^n \left(\prod_{k=1}^{j-1} (1 - p_{ik}) \right) (e_{ij} + p_{ij}a_{ij}) + t_i^{l,l+1} \quad (4)$$

where

$$\begin{aligned} t_i^{l,l+1} &= (e_{il} + p_{il}a_{il}) \prod_{k=1}^{l-1} (1 - p_{ik}) + \\ &\quad (e_{i,l+1} + p_{i,l+1}a_{i,l+1}) \prod_{k=1}^l (1 - p_{ik}) \end{aligned}$$

In the following, we only regard the case where $p_{ij} < 1$ for $j = 1, \dots, l-1$; otherwise, choices c_{il} and $c_{i,l+1}$ would never be reached, and their sequence would not matter. Now we assume that we would change the order of these two choices; in this case, only the term $t_i^{l,l+1}$ in (4) changes, and let us call the corresponding term $t_i^{l+1,l}$. So the difference between the expected benefits of these two lists is $t_i^{l,l+1} - t_i^{l+1,l}$. In order to simplify the derivation, we divide this difference by the probability that the user did not select any of the choices before, i.e. the product of the corresponding counter-probabilities. This simplified difference can be transformed as follows:

$$\begin{aligned} d_i^{l,l+1} &= \frac{t_i^{l,l+1} - t_i^{l+1,l}}{\prod_{k=1}^{l-1} (1 - p_{ik})} \\ &= e_{il} + p_{il}a_{il} + (1 - p_{il})(e_{i,l+1} + p_{i,l+1}a_{i,l+1}) - \\ &\quad (e_{i,l+1} + p_{i,l+1}a_{i,l+1} + (1 - p_{i,l+1})(e_{il} + p_{il}a_{il})) \end{aligned} \quad (5)$$

$$= p_{i,l+1}(e_{il} + p_{il}a_{il}) - p_{il}(e_{i,l+1} + p_{i,l+1}a_{i,l+1}) \quad (6)$$

Since $\prod_{k=1}^{l-1} (1 - p_{ik})$ is positive, the expected benefit of the original list is not less than that of the modified list iff $d_i^{l,l+1} \geq 0$.

Now let us first regard the special cases where $p_{i,l} = 0$ or $p_{i,l+1} = 0$. If $p_{i,l} = 0 \neq p_{i,l+1}$, then the difference is negative and the two choices should be reordered in for increasing the benefit. Otherwise, if $p_{i,l+1} = 0$, then the difference will be nonnegative, and the choices should remain in the current order. So these two conditions would lead to the effect that all choices with zero selection probability would be moved to the end of the choice list (and thus, they better should not be included in this list).

In the following, we assume that $p_{i,j} > 0$ for $1 \leq j \leq n_i$. Then the combination of the condition $d_i^{l,l+1} \geq 0$ with eqn (6) yields the following criterion:

$$a_{il} + \frac{e_{il}}{p_{il}} \geq a_{i,l+1} + \frac{e_{i,l+1}}{p_{i,l+1}} \quad (7)$$

So we have a condition for bringing two adjacent choices into the right order, for increasing the expected benefit of the complete choice list. By applying this condition iteratively, and reordering two adjacent choices in case the condition is not satisfied (similar to bubble sort), we can bring the whole list into an order where the expected benefit is maximized.

So we can formulate our *probability ranking principle for interactive information retrieval (IIR-PRP)*: rank choices c_{ij} by decreasing values of

$$\varrho(c_{il}) = a_{il} + \frac{e_{il}}{p_{il}}. \quad (8)$$

5.3 Analysis

The first interesting observation is the fact that our ranking criterion ϱ for a choice c_{ij} is different from its expected benefit. As a simple example, assume that we have two choices c_{i1} and c_{i2} where $p_{i1} = 0.5$, $a_{i1} = 10$, $e_{i1} = -1$ and $p_{i2} = 0.25$, $a_{i2} = 16$, $e_{i2} = -1$. Then we have $E(c_{i1}) = 4$ vs. $E(c_{i2}) = 3$. However, $\varrho(c_{i1}) = 8$ vs. $\varrho(c_{i2}) = 12$. It can be checked that our ranking criterion indeed maximizes the expected benefit of the list: $E(\langle c_{i1}, c_{i2} \rangle) = 4 + 0.5 \cdot 3 = 5.5$ vs. $E(\langle c_{i2}, c_{i1} \rangle) = 3 + 0.75 \cdot 4 = 6$. The reason for this difference lies in the fact that the expected benefit of a list is not just the sum of the expected benefits of the single choices (as is the case with the classic PRP — see below), as shown in eqn (3).

Another important issue is the comparison with the classical PRP. We can show easily that our IIR-PRP is a generalization of the classical PRP. Let $e_{ij} = \bar{C} < 0$ (the cost for reading a document) and $a_{il} = C$ (the benefit of a relevant document); substituting these values in eqn (7), we get:

$$\begin{aligned}
C + \frac{\bar{C}}{p_{il}} &\geq C + \frac{\bar{C}}{p_{i,l+1}} \Rightarrow \\
p_{il} &\geq p_{i,l+1}
\end{aligned}$$

So we have the classical PRP, where documents are ranked by decreasing values of their probability of relevance.

Although the 'probability of relevance' p_{ij} still plays a major role in the IIR-PRP, we see that the major extension of our new model is the consideration of varying values for the effort e_{ij} and the average benefit a_{ij} , as well as the tradeoff between these two parameters. The classical PRP minimizes the cost of a search by ranking documents according to increasing values of expected cost (or decreasing benefit), which is estimated as $p_{il}C + (1 - p_{il})\bar{C}$. In the example from above, we have shown that a ranking according to decreasing benefit in general is not optimum in our case. The reason for this difference is the variability of the effort and benefit values in our model. In case these values are constants, our model reduces to the classical PRP. In fact, in this case our model is equivalent to the PRP for finding one relevant document (since the first positive decision brings up a new situation), and for this problem, the PRP is known to yield the optimum solution. For finding more relevant documents, the PRP assumes that the information need remains unchanged and that the relevance judgments of documents are independent of each other; since our approach abandons these assumptions, we are not able to make predictions about further relevant documents (this task is left to other, more specific models which may use certain additional assumptions — e.g dropping only the second of these assumptions, so that the user wants to see more relevant, but substantially different documents)

[Bookstein 83] describes a generalization of the classical PRP to multi-valued relevance scales, where different relevance values are associated with different cost factors. Then it is shown that optimum retrieval is achieved when document are ranked according to increasing costs. However, in terms of our model, Bookstein regards the term $p_{ij}(a_{ij} + e_{ij})$, whereas we separate the effort for a decision from its potential benefit in case of acceptance. So the two models are not directly comparable. Only for the binary case with constant effort and benefit values, our model corresponds to the PRP.

Finally, readers familiar with Markov models may notice that our approach describes in fact such a model: situations correspond to states, and a choice c_{il} is a transition with probability $p_{il} \prod_{k=1}^{l-1} (1 - p_{ik})$. In our approach, we pose no restrictions on the number of situations/states: In general, each sequence of choices $c_{1i}, c_{2j}, \dots, c_{mk}$ may lead to a unique situation that can only be reached via this sequence. Thus, we have a Markov model where the

number of possible states is infinite, but countable. Since we assume that the transition probability is always positive ($0 < p_{ij} < 1$), the Markov chain is irreducible. However, the problem is more complicated due to creative actions of the user (like adding a term to the query that was not explicitly proposed by the system). In terms of our model, we would represent such a decision by estimating the corresponding effort and its expected benefit. Indeed, the actual benefit depends very much on the term chosen. In the followup situation, the system knows this term, and can react appropriately. Due to this problem, there is no straightforward way for applying more elaborated methods from Markov models, which aim at analyzing paths through the model. Further research is needed in this area.

6 Towards application

As mentioned before, the work presented here forms a framework similar to the classical PRP. Thus, it describes *which* parameters should be considered, but does not specify *how* these parameter can be estimated. Nevertheless, we want to outline here some directions of further research that we deem useful for accomplishing this task.

With regard to the kind of research required, we can distinguish three groups of parameters:

1. Selection probability p_{ij} : Many IR models are addressing this problem. Besides the probability of relevance of documents, there are also many approaches for computing query expansion terms, or for generating document summaries. Thus, for most kinds of selection lists, there is already a substantial amount of research which provides useful solutions (or at least starting points) for estimating this parameter. However, since most of these approaches are still based on the assumption of a static information need, more work is required to make these models more dynamic.
2. Effort parameters e_{ij} , g_{ij} and the success probability q_{ij} : In this area, most research is needed. Here empirical studies with real users should be performed, closely monitoring the users' actions (involving eye-tracking) — see e.g. [Joachims et al. 07]. As an additional problem, visualization aspects may affect results heavily (see e.g. [Malik et al. 06]). Thus, this kind of research should also develop some 'best practice' methods that serve as reference points for the parameter values derived.

3. Benefit b_{ij} : Of course, there is the general problem of information value, which is heavily application-dependent. Another possible definition is that of saved effort, relative to some baseline. Below, we outline an estimation method following the latter approach. There may be other — even better — methods for estimating these parameters, but we want to demonstrate that there are already solutions to this problem.

By defining benefit as saved effort, we of course depend on the problem of estimating the effort of certain actions. However, for quantifying benefit, here we consider a single type of action only, namely that of scanning through a ranked list of documents. So we assume a unit effort per document in the rank list — the problem of scaling is to be solved in connection with the methods estimating user effort.

The basic assumption of our method is the following: for the current situation, the system has constructed (explicitly or implicitly) the sub-optimum query q' , and the user's choice will now lead to the optimum query q , where he only will have to scan the ranked list of documents. In addition, we assume that the user wants one relevant document only (or only one more, in case he has found some already). There may be many other user standpoints, but here we regard the most simple case only.

For a given query q , [Nottelmann & Fuhr 03b, Nottelmann & Fuhr 03a] describe methods for estimating the number r_q of relevant documents in the database as well as their proportion among the top k documents.

For the latter, we need an assumption about the retrieval performance of the system. As a simple model, we use a linear recall-precision curve of the form

$$P(R) := P^0 \cdot (1 - R) \tag{9}$$

where P denotes precision, R stands for recall, and the parameter P^0 is the initial precision to be chosen. Let n_q be the position of the first relevant document in the ranked list. For this point, we have $P = 1/n_q$ and $R = 1/r_q$. Substituting these values in eqn (9), we get as approximation of the position of the first relevant document

$$n_q = \frac{r_q}{P^0(r_q - 1)} \tag{10}$$

So we know the effort for locating the first relevant document in the ranking list of the optimum query q . In the current situation, however, we are still dealing with the query q' , and we want to know how many documents the user would have to scan in the corresponding result list until he finds a relevant document. For that, we define the probability $P(q|q')$ that a random document from the result list of q also occurs in the result list of

q' (of course, here we would have to limit the length of the result lists in a reasonable way — one possible approach would be the assumption of Boolean retrieval). Based on the data available in the IR system, this parameter can be computed easily (e.g. by retrieving the top k documents for q , and then determining their position in the output of the current query q'). The probability $P(q|q')$ obviously has a multiplicative effect on the precision, so that we get as modification of eqn (10) the position of the first relevant document in the ranked list of q' to

$$n_{q'} = \frac{r_q}{P(q|q')P^0(r_q - 1)} \quad (11)$$

Based on these results, the benefit for moving from q' to q can be estimated as $n_{q'} - n_q$.

For illustration, let us return to our Java example in Table 1. Assuming that half of the documents in which all query terms occur are also relevant, and with an initial precision of $P^0 = 0.5$, we would get $n_q \approx 2$ in all cases. If we estimate the values $P(q|q')$ based on Boolean retrieval, they are identical to the corresponding p_{ij} value shown in the table. Using these estimates, we would arrive at the $n_{q'}$ values as shown in the second to last column, from which we would have to subtract $n_q = 2$ for computing the final benefit. Assuming further that the effort for selecting a term is $e_{ij} = 1$, we would arrive at the values for the ranking criterion as shown in the last column. Obviously, this would lead to a reverse ranking of the choices; moreover, the expected benefit $e_{ij} + p_{ij}(n_{q'} - n_q)$ for the term 'program' would be negative (-0.33), so this choice should not be shown. So it turns out that our initial ranking for this example may not have been correct — it all depends on the actual effort and benefit parameters.

7 Related work

The shortcomings of the classical PRP have been noticed already in [Stirling 75], where a theoretical model for considering dependencies between documents is presented. On a more practical side, [Carbonell & Goldstein 98] describe experiments where the similarity of the top-k documents is used for re-ranking, in order to present the most dissimilar, but potentially relevant documents to the user. In [Chen & Karger 06], different metrics considering dependencies between retrieved documents are regarded, and corresponding methods for optimizing result ranking are presented.

The dynamic nature of information needs has been emphasized by several authors following the cognitive approach to IR [Belkin et al. 82], [Borlund &

Ingwersen 98], [Ingwersen 96]; e.g. the latter asks 'to view relevance in IR as situational, relative, partial, differentiated and non-linear'.

However, the only actual IR system following these ideas is the implementation of the ostensive model [Campbell 00], which uses a kind of 'aging' mechanism for relevance feedback data in order to determine the next documents to be presented to the user. Moreover, this system is highly dynamic (partly due the task of image retrieval studied here), and each selected choice creates a new situation (according to our terminology).

From the area of human-computer interaction, [Williamson & Murray-Smith 04, Williamson 06] present interfaces for displaying time-varying information; they use probabilistic predictions of user behavior and their potential goals for arranging the information displayed to the user.

Our model is remotely related to the Page Rank model [Page et al. 98] which also regards transition probabilities between interaction states (i.e. page views); however, in our approach, we would consider the order in which the different links are encountered by the user when looking through a page, whereas the Page Rank model ignores this factor.

[White & Drucker 07] describe query trails of Web searches and their analysis; however, this approach monitors only the positive decisions made by the user, but not the choices they were faced with in each situation. In contrast, the work presented in [Joachims et al. 07] uses eyetracking for observing users during Web searches, thus registering e.g. the items from the result list users were looking at. Along with their time measurements, this kind of research could be a good starting point for implementing the generic model presented here.

On a more general level, interactive IR systems as proposed here can be seen as an instantiation of interactive computing [Goldin et al. 06] where systems interact continuously with the user and/or their environment. Thus, for the actual design of IIR systems, this new computing paradigm may provide a fruitful basis.

8 Conclusion and Outlook

In this paper, we have presented a framework for extending probabilistic IR to interactive retrieval. Based on the notions of situations and decision making, we first have shown how the expected benefit of a single choice can be maximized. The most important result of our paper is the derivation of the optimum ordering of choices — the probability ranking principle for interactive IR. We also have shown that the classical PRP is a special case of our new model.

Similar to the classical PRP, our model uses certain parameters, but does not specify how the values of these parameters can be estimated. This is the subject of more specialized models (similar to the broad variety of probabilistic models that are all founded on the PRP).

On the other hand, with the IIR-PRP as described here, there is a point of reference for the development of IIR models and systems. IIR systems are a commodity nowadays, but the functional design of these systems lacks an underlying theory. The work presented in this paper is a first step towards the development of such a theory.

Acknowledgment

I wish to thank the Glasgow IR group, especially Keith van Rijsbergen, for their hospitality and fruitful discussions when staying with them in August 2007, while I was writing this article. The suggestions by the three anonymous reviewers were very helpful in improving the initial version of this paper.

A Definition of the event space

Let $S = \{s_0, s_1, s_2 \dots\}$ denote a (possibly infinite) set of situations. In each situation $s_i \in S$, we have a set of choices $C_i = \{c_{i1}, c_{i2}, \dots, c_{i, n_i}\}$ with $c_{ij} \in S$, i.e. choices are a partial mapping $c : S \times \mathcal{N} \rightarrow S$. Our event space is situation-specific, and we make no assumptions how the event space changes when the user moves to a new situation. Let U denote all uses of our system, and $U_i \subseteq U$ is the set of all these uses which arrive at situation s_i (e.g. in a Web search engine, all uses starting with the same query — provided that no additional information is available — will lead to the same situation). Now our event space is $\Omega = C_i \times U_i$. Unfortunately, we have only judgments about a subset $J \subseteq C_i \times U_i$ of the elements of the event space — due to the fact that the user leaves the situation as soon as he accepts a choice. Associated with each element J , we have the acceptance decision of the user, which can be modelled as a relation $A \subseteq J \subseteq C_i \times U_i$. Furthermore, not all of the accepted choices will turn out to be right from the user’s point of view (so that he will return to situation s_i — which we would model as another use). Thus, right decisions are a subset $R \subseteq A$ of the accepted ones. The probabilistic parameters we define now are all independent of the actual uses — the system knowledge about the use is implicitly represented by the actual situation. Let X denote a random variable ranging over Ω , and Z a variable ranging over U_i . Then we define $p_{ij} = P(X \in A | X = (c_{ij}, Z) \wedge X \in J)$ as

the probability that a use in situation s_i will accept choice c_{ij} .

The only independence assumption we now have to make is the following: the probability of a user accepting a choice c_{ij} is independent of the choices he rejected before. In most cases, this supposition will be fairly valid (e.g. ranked list of documents, or list of expansion terms). Please note that this assumption is much weaker than that of the classical PRP, where independence of both positive and negative relevance judgments is assumed. With this presupposition, we exclude any sequence effects, i.e. changing the order of the choices being presented does not affect their probability of being accepted. More formally, the independence assumption can be written as follows:

$$P \left(X \in A | X = (c_{ij}, Z) \wedge X \in J \wedge \bigwedge_{k:(c_{ik}, Z) \in (J-A)} Y_k = (c_{ik}, Z) \right) = P(X \in A | X = (c_{ij}, Z) \wedge X \in J)$$

(Here the Y_k 's are random variables ranging over all the choices of the same use Z .) Furthermore, let $q_{ij} = P(X \in R | X = (c_{ij}, u) \wedge X \in A)$ denote the probability that acceptance of this choice is not revised later.

References

- Bates, M. J.** (1989). The design of browsing and berrypicking techniques for the online search interface. *Online Review* 13(5), pages 407–424. <http://www.gseis.ucla.edu/faculty/bates/berrypicking.html>.
- Belkin, N.; Oddy, R.; Brooks, H.** (1982). ASK for Information Retrieval: Part I. Background and Theory. *The Journal of Documentation* 38(2), pages pp. 61–71.
- Bookstein, A.** (1983). Outline of a General Probabilistic Retrieval Model. *Journal of Documentation* 39(2), pages 63–72.
- Borlund, P.; Ingwersen, P.** (1998). Measures of relative relevance and ranked half-life: performance indicators for interactive IR. In: *SIGIR '98: Proceedings of the 21st annual international ACM SIGIR conference on Research and development in information retrieval*, pages 324–331. ACM Press, New York, NY, USA.
- Campbell, I.** (2000). Interactive Evaluation of the Ostensive Model Using a New Test Collection of Images with Multiple Relevance Assessments. *Information Retrieval* 2(1), pages 89–114.

- Carbonell, J.; Goldstein, J.** (1998). The use of MMR, diversity-based reranking for reordering documents and producing summaries. In: Croft, W. B.; Moffat, A.; van Rijsbergen, C. J.; Wilkinson, R.; Zobel, J. (eds.): *Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 335–336. ACM, New York.
- Chen, H.; Karger, D. R.** (2006). Less is more: probabilistic models for retrieving fewer relevant documents. In [Eftimiadis et al. 06], pages 429–436.
- Eftimiadis, E. N.; Dumais, S. T.; Hawking, D.; Järvelin, K. (eds.)** (2006). *SIGIR 2006: Proceedings of the 29th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, Seattle, Washington, USA, August 6-11, 2006*. ACM.
- Fuhr, N.** (1992). Probabilistic Models in Information Retrieval. *The Computer Journal* 35(3), pages 243–255.
- Goldin, D. Q.; Smolka, S. A.; Wegner, P.** (2006). *Interactive Computation. The New Paradigm*. Springer.
- Ingwersen, P.** (1996). Cognitive perspectives of information retrieval. *The Journal of Documentation* 52(1), pages 3–50.
- Joachims, T.; Granka, L.; Pan, B.; Hembrooke, H.; Radlinski, F.; Gay, G.** (2007). Evaluating the accuracy of implicit feedback from clicks and query reformulations in Web search. *ACM Trans. Inf. Syst.* 25(2).
- Malik, S.; Klas, C.-P.; Fuhr, N.; Larsen, B.; Tombros, A.** (2006). Designing a User Interface for Interactive Retrieval of Structured Documents — Lessons Learned from the INEX Interactive Track. In: *Proc. European Conference on Digital Libraries*.
- Nottelmann, H.; Fuhr, N.** (2003a). Evaluating different methods of estimating retrieval quality for resource selection. In: Callan, J.; Cormack, G.; Clarke, C.; Hawking, D.; Smeaton, A. (eds.): *Proceedings of the 26th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, New York.
- Nottelmann, H.; Fuhr, N.** (2003b). From Retrieval Status Values to Probabilities of Relevance for Advanced IR Applications. *Information Retrieval* 6(4).

- O’Day, V. L.; Jeffries, R.** (1993). Orienting in an Information Landscape: How Information Seekers Get From Here to There. In: *Proc. of the INTERCHI ’93*, pages 438–445. IOS Press.
- Page, L.; Brin, S.; Motwani, R.; Winograd, T.** (1998). *The PageRank Citation Ranking: Bringing Order to the Web*. Technical report, Stanford Digital Library Technologies Project.
- Robertson, S. E.** (1977). The Probability Ranking Principle in IR. *Journal of Documentation* 33, pages 294–304.
- Stirling, K. H.** (1975). The Effect of Document Ranking on Retrieval System Performance: A Search for an Optimal Ranking Rule. In: *Proceedings of the American Society for Information Science* 12, pages 105–106.
- Turpin, A. H.; Hersh, W.** (2001). Why batch and user evaluations do not give the same results. In: Croft, W. B.; Harper, D.; Kraft, D. H.; Zobel, J. (eds.): *Proceedings of the 24th Annual International Conference on Research and development in Information Retrieval*, pages 225–231. ACM Press, New York.
- Turpin, A.; Scholer, F.** (2006). User performance versus precision measures for simple search tasks. In [Eftimiadis et al. 06], pages 11–18.
- Voorhees, E.; Harman, D.** (2000). Overview of the Eighth Text REtrieval Conference (TREC-8). In: *The Eighth Text REtrieval Conference (TREC-8)*, pages 1–24. NIST, Gaithersburg, MD, USA.
- White, R.; Drucker, S.** (2007). Investigating behavioral variability in Web search. In: *Proc. WWW*, pages 21–30.
- White, R. W.; Jose, J. M.; Ruthven, I.; van Risjbergen, C. J.** (2005). Evaluating Implicit Feedback Models Using Searcher Simulations. *ACM Transactions on Information Systems* 23(3), pages 325–361.
- Williamson, J.; Murray-Smith, R.** (2004). Granular Synthesis for the Display of Time-Varying Probability Densities. In: *Int. Workshop on Interactive Sonification, Bielefeld*. <http://www.dcs.gla.ac.uk/jhw/ison.pdf>.
- Williamson, J.** (2006). *Continuous Uncertain Interaction*. PhD thesis, University of Glasgow, Computer Science.