An Analysis of Theories of Search and Search Behavior

Leif Azzopardi
School of Computing Science
University of Glasgow
Glasgow, United Kingdom
Leif.Azzopardi@glasgow.ac.uk

Guido Zuccon
Information Systems School
Queensland University of Technology (QUT)
Brisbane, Australia
g.zuccon@qut.edu.au

ABSTRACT

Theories of search and search behavior can be used to glean insights and generate hypotheses about how people interact with retrieval systems. This paper examines three such theories, the long standing Information Foraging Theory, along with the more recently proposed Search Economic Theory and the Interactive Probability Ranking Principle. Our goal is to develop a model for ad-hoc topic retrieval using each approach, all within a common framework, in order to (1) determine what predictions each approach makes about search behavior, and (2) show the relationships, equivalences and differences between the approaches. While each approach takes a different perspective on modeling searcher interactions, we show that under certain assumptions, they lead to similar hypotheses regarding search behavior. Moreover, we show that the models are complementary to each other, but operate at different levels (i.e., sessions, patches and situations). We further show how the differences between the approaches lead to new insights into the theories and new models. This contribution will not only lead to further theoretical developments, but also enables practitioners to employ one of the three equivalent models depending on the data available.

Categories and Subject Descriptors

H.3.3 [**Information Storage and Retrieval**]: Information Search and Retrieval: Search Process; H.3.4 [**Information Storage and Retrieval**]: Systems and Software: Performance Evaluation

General Terms

Theory, Experimentation, Economics, Human Factors

Keywords

Information Foraging Theory, Search Economic Theory, Interactive Probability Ranking Principle

1. INTRODUCTION

The field of Information Seeking and Retrieval (ISR) seeks to understand, predict and explain how people interact with search systems. However, the interaction between a person and search system is non-trivial. It is affected by a host of factors including variables from the person's context (e.g. background, experience, expertise), the interface and the system's configuration [16]. During the course of an ad-hoc topic search session, information seekers pose a number of queries, browse snippets, and assess documents in order to fulfil their information need. This requires them to make many implicit and explicit decisions regarding: what queries to pose, what documents to view, what facets/features to try, whether to continue inspecting snippets for the current query, when to issue a new query, and when to stop searching [32]. Consequently, understanding how information seekers behave and interact with search systems has been a long standing and challenging area of research [8, 17].

While most researchers have focused on cataloguing search behavior based on empirical and observational evidence [19], a number of attempts to formalize the interactions have been proposed [2, 14, 23]. Such formal models use a mathematical framework in which the most salient variables between an information seeker and a system are represented. Given such models, it is possible to describe, predict and crucially explain how and why seekers behave the way they do. Such models also enable the generation of hypotheses regarding search behavior which can be subsequently empirically tested and validated (e.g. [5, 24]). However, these models also have numerous limitations ranging from the low level assumptions engaged by the different models, the variables that they include/exclude and the difficulties arising from the complexities of human behavior. While more sophisticated models are being developed in order to address such limitations (e.g. [3, 4]), our focus is on understanding how the different theories and ensuing models relate to each other. Consequently, in this paper, we will analyze, compare and develop three ISR theories: Information Foraging Theory (IFT) [23], Interactive Probability Ranking Principle (iPRP) [14] and Search Economic Theory (SET) [2]. In this paper, we model the task of ad-hoc topic search using each theory within a common framework. This will enable us to explore how they are similar and different, and what we can learn from each of them. Specifically, we focus on whether these theories make similar predictions about search behavior and how we can develop, refine and extend the different models.

2. BACKGROUND

Numerous models of Information Seeking and Retrieval have been proposed in the literature [6, 7, 8, 9, 13, 16, 17, 18, 20, 36]. Such models fall into two main categories: (1) conceptual and descriptive, and (2) formal and mathematical. Conceptual and descriptive models typically depict the interactions and variables at a high level - they describe the process a searcher takes, and the stages that they may go through. These are particularly useful for understanding the various components involved in the process and which factors are likely to have an influence. Such models are called pre-theoretical in [17] because they provide a picture of the problem domain, which can be used to build more formal models. On the other hand, formal and mathematical models in particular are more precise, enumerating the phenomena as a list of variables and showing how each variable functionally relates to each other[1].

A well known conceptual model of information seeking is the Berry Picker model proposed by Bates [6]. The model draws an analogy between a searcher and a forager (in this case a berry picker). As shown in Figure 1, the berry picker moves from patch to patch collecting the juiciest and ripest berries, before moving onto the next patch. Similarly, the searcher goes from one patch of results to another patch of results selecting the most relevant documents, and moving on to another patch of results. During the course of their search, the searcher's information need evolves and so the type of information they find valuable at any given point also changes (e.g. the berry picker might first collect strawberries, then blueberries, then raspberries, and so on).

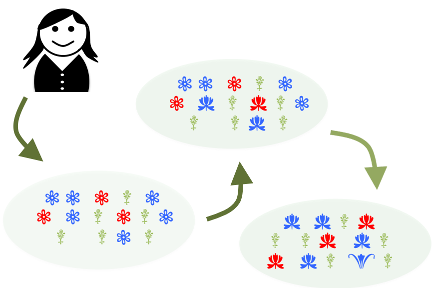

Figure 1: The berry picker foraging patches models a searcher with an evolving information need.

This conceptual model of information seeking is intuitive and most people can easily relate to the idea. However, the model does not provide an indication of how long the searcher will stay in a particular patch or how the time to reach a patch will affect their behavior. Bates suggested that searchers would weigh up the costs and benefits in order to decide what action to take next [7]. Other researchers also considered how searchers are like foragers and initial attempts [27, 29] suggested that Optimal Foraging Theory could be used to model the searching process. Subsequently, this lead to the development of Information Foraging Theory [23] (see subsection below).

The idea of using cost-benefit and decision making tools is not new and a number of other formal models have been developed using such a framework. Indeed, early Information Retrieval (IR) research exploited such tools to examine IR systems in a number of ways ranging from purchasing decisions [1, 26] to ranking [14, 25, 35] to user behavior [2, 10, 12]. Initial attempts focused on the trade-off between the cost of an IR system and its effectiveness. In [1, 26], Axelrod and Rotheberg compare different mechanized IR systems available during the late 1960s and early 1970s by performing a cost-benefit analysis in order to decide which system to purchase[2]. In [12], Cooper took a more user-oriented perspective. He modeled the trade-off between the amount of time a user should spend searching versus how much time the system should spend searching. In the same period, Robertson [25] examined the problem of ranking in terms of the costs and benefits of ranking one document above another. This led to the formulation of the Probability Ranking Principle (PRP) which essentially applies decision theory to the ranking problem [25]. More recently, Fuhr revised and extended the PRP to consider a series of interactions in the interactive Probability Ranking Principle (iPRP) [14]. This generalized model accounts for the different costs and benefits associated with particular choices when ranking documents (see subsection below).

In [34], Varian outlined three directions in which economics could be useful for search: (1) to obtain better estimates of the probability of relevance, (2) to apply Stigler's theory of Optimal Search Behavior to IR [31], and (3) to examine the economic value of information using consumer theory, "where a consumer is making a choice to maximize expected utility or minimize expected cost" [34]. A number of different works have begun to examine these directions. For example, in [35], Wang and Zhu used Portfolio Theory to obtain better estimates of relevance by accounting for the uncertainty associated with the probability estimates when ranking. While in [10], Birchler and Butler explain how Stigler's theory can be applied to search in order to predict when a user should stop examining results in a ranked list. In a variation on Varian's third suggestion, Azzopardi showed how Production Theory [33] could be used to model the search process [2]. This led to the development of Search Economic Theory (SET) which has been specifically developed to model ad-hoc topic retrieval (see subsection below).

Now that we have provided the context for the different models, we will provide an overview of IFT, iPRP and SET, before developing and comparing the different approaches.

2.1 Information Foraging Theory

Information Foraging Theory is composed of three types of models: *Information Scent model*, *Information Patch model* and *Information Diet model* [23]. Of relevance to this work is the patch model, which describes how long foragers will stay in a patch before moving to a new patch. Under the patch model, the analogy with an information seeker is as follows. Moving between patches is like expressing a new query (and thus incurs a moving/querying cost), while staying within a patch is akin to assessing documents (where each document takes a certain amount of time to process). The Information Patch model predicts how long a forager should stay in a patch before moving on to the next patch.

Under IFT, it is assumed that the forager is rational in that (i) they will visit the patch with the highest yield first, and (ii) they wish to maximize their gain per unit of time. To instantiate the model a gain function parameterized by

[1]Note that a formal model need not be mathematical, it could be expressed in some other formal language or construct.

[2]Note the system, while mechanized, also included librarians and technicians as part of the search process.

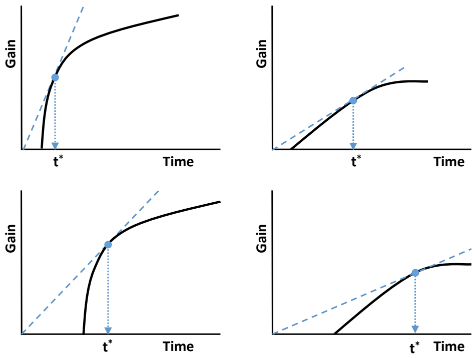

Figure 2: Information Foraging Theory: Top plots are when between-patch times are low, bottom plots are when between-patch times are high. For longer between-patch time. the model predicts foragers will stay longer. For higher yields (left plots), it predicts foragers will move sooner. $t^\star$ indicates optimal time in patch.

time, i.e., $g(t)$ is required. The point where a forager should move to the next patch is when the maximum gain per unit of time is achieved. This depends on the time it takes to get to a patch, the cost of processing documents, and the distribution of relevant information (as specified by the gain function). In Figure 2, the time spent in each patch is shown graphically for patches of different yields. A searcher moves to a patch, they receive no gain for traveling to the patch. They then start to assess and thus extract gain from documents. By drawing a line from the origin to the gain curve, it is possible to determine when the forager should stop, because the tangent is when the forager has achieved the maximum gain per unit of time, and should now switch to a new patch. If the forager stays longer then they will obtain less and less gain (i.e., diminishing returns).

2.2 Interactive PRP

The Interactive Probability Ranking Principle (iPRP) [14] forms an extension to the well known Probability Ranking Principle [25]. However, the iPRP relaxes a number of modeling assumptions made by the PRP: (i) the notion of a fixed information need, and (ii) the relevance of a document is independent of previously seen documents. While these assumptions are reasonable approximations, they have been shown to break down under certain circumstances [15]. The main requirements, when developing the iPRP, were specified as: (i) consider the complete interaction process (i.e., not just document ranking, but other activities), (ii) allow for different costs and benefits of different activities (e.g., a longer document takes longer to process than a shorter document), and (iii) allow for the information need to change through the course of interaction. The motivation of this later point was to incorporate the idea or notion of the Berry Picker model [6] where as the searcher moves from patch to patch their information need changes. Consequently, under the iPRP, when the searcher encounters information, this may change their information need. To instantiate the principle, a number of further assumptions are also made:

1. focus on the function level of interaction (i.e., the different activities a searcher can take and the cost/gain associated with each interaction),
2. decisions form the basis of interaction (i.e., what the searcher does next is based on a decision from the possible set of activities available at that point in time),
3. the searcher evaluates these choices in a linear order (this ordering is either explicit or implicit), and,
4. only decisions which are positive and correct are of benefit to the user.

A key concept in the iPRP is the notion of situations. Each situation essentially represents the current state of the system, and the choices it offers to the user at a particular point during the search. When a user takes a choice, then the system moves to a new state, and the user enters another situation. Note that the system may present entirely new choices, or present the same set of choices during any particular situation. However, the benefit and probability of accepting choices that have been previously presented are likely to change. A situation could contain, for example, the choice to enter a new query, along with a series of choices to examine the documents.

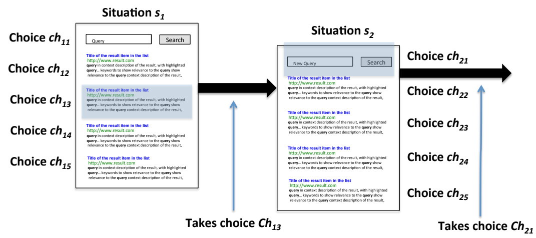

Figure 3: Each choice results in a new situation.

More formally, for a given situation s_i, the user is faced with a number of choices $Ch_i = \{ch_{i1}, \ldots, ch_{in}\}$. For each choice, there is a corresponding probability p_{ij} that the user will accept choice ch_{ij} in situation s_i (see Figure 3). For each situation, it is assumed that the choices are considered in a linear order.

When a user takes a particular choice, ch_{ij}, they will accrue some benefit if they made a good decision (or lose some benefit if they made a bad decision). On average, though, a particular choice ch_{ij}, will result in the average benefit[3], a_{ij}. Under the iPRP, benefit is used to refer to negative costs, and so costs are expressed as negative benefit. This means that the framework uses the same units (such as time) to denote both costs and benefits. This implies that a good decision will result in the user saving time, while a bad decision will result in wasting time. Every choice ch_{ij} also requires a certain amount of effort e_{ij} to be expended, which is also expressed in the same units (i.e., time).

The expected benefit that is accrued from a particular choice can then be formulated as follows:

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

and the expected benefit of taking n choices is:

$$E(ch_{i1}, \ldots, ch_{in}) = \sum_{j=1}^{n} \Big(\prod_{k=1}^{j-1} (1 - p_{ik}) \Big) (e_{ij} + p_{ij}a_{ij})$$

[3] $a_{ij} = q_{ij}b_{ij} + (1 - q_{ij})r_{ij}$, where q_{ij} is the probability that the choice was correct and yields b_{ij} benefit. Otherwise the decision incurs a cost r_{ij} with probability $1 - q_{ij}$ to backtrack.

It is shown in [14] that the expected benefit is maximized when the following criterion is met:

$$a_{ij} + \frac{e_{ij}}{p_{ij}} \geq a_{i,j+1} + \frac{e_{i,j+1}}{p_{i,j+1}} \quad (1)$$

2.3 Search Economic Theory

The initial model proposed in [2] draws upon an analogy with Production Theory [33]. In production theory, a firm produces an *output* (such as goods or services), and to do so requires *inputs* to the process (usually termed, capital and labor) [33]. The firm will utilize some form of *technology* to then produce the output given the inputs. The process of production is similar to the search process. The inputs to the search process are:

Q the number of queries that the user will issue, and,

A the number of documents the user will assess per query.

The output of the search process is a certain amount of utility or gain that the searcher extracts from the relevant documents found during the process. The technology engaged by the user to produce (i.e., find) relevant documents is a retrieval system. This abstraction reduces the search process down to the core variables which directly influence how much utility a user receives through the course of interaction with the system. Consequently, the total amount of gain is proportional to the number of queries and number of documents assessed per query, i.e., the total cumulative gain G is equal to the function $g(Q, A)$. Depending on the particular retrieval system employed, different **technological constraints** are imposed upon the search process such that only certain combinations of inputs will produce a given or specified amount of gain. Under the model a range of different search strategies are potentially possible. For example, for a particular level of gain G, a searcher may pose many queries, and examine a few documents per query, or pose few queries and examine many documents per query. Figure 4 shows an example of a search production function. Each point along the curve represents a combination of inputs that yields the same gain. However, each combination comes at a price. So a cost function $c(Q, A)$ was also defined in [2]. Now, given this model of the search process (i.e., $g()$ and $c()$), a rational searcher would either minimize $c(Q, A)$ in order to obtain a particular level of gain G, or maximize $g(Q, A)$ for a given cost C.

In [4], this model was revised and extended to include a number of other parameters such as the number of snippets examined, the number of result pages viewed, and the probability of examining a document. In this paper, to facilitate the mapping between theories we will focus on building models based on the initial economic model of search and leave such extensions for future work.

3. THEORETICAL DEVELOPMENT

Given the three theories outlined above, the goals of this paper are as follows:

- to develop a model of search based on each theory using the same notion and a common set of assumptions,
- to derive hypotheses about search and search behavior,

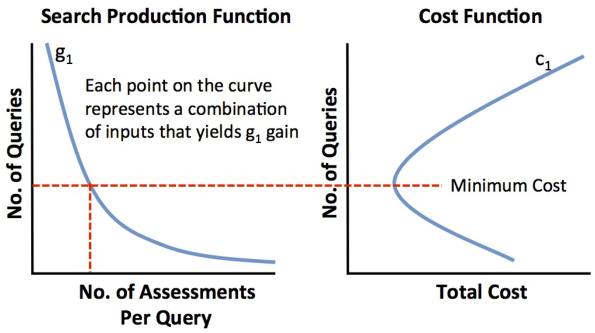

Figure 4: The right plot shows the gain curve that yields g_1 gain, and the left plot shows the corresponding cost function c_1.

- to determine the similarities and differences between the approaches, and,
- to learn how the differences can be used to extend the other models.

We will undertake the theoretical development in the context of ad-hoc topic retrieval where the user would like to find a number of documents relevant to the topic (and have a limited amount of time to perform this task). Gain will be measured as cumulative gain, and the effort / cost will be measured in time (seconds). First, we need to establish a common notation. Let $g()$ denote a cumulative gain function, which can be parameterized either by time or actions. Thus, let $g(t)$ be the cumulative gain at time t, and let $g(Q, A)$ be the cumulative gain given the number of queries Q and the number of assessments per query A. Let G denote a particular level of cumulative gain, and let T denote a particular length in time in seconds, where over a session of length T a searcher receives G cumulative gain. For each action there is an associated cost, which we will represent as time: the time it takes to issue a query t_q and the time it takes to assess a document t_d. Since effort is negative benefit under the iPRP then $e = -c = -t$.

Given this notation, we will first develop the initial SET model proposed in [2], as it already provides a model of search for ad-hoc retrieval over a session. From this model, we will draw a number of hypotheses. We will then develop the other theories to model the same scenario.

4. ECONOMIC MODEL OF SEARCH

As previously mentioned, the basic economic search model [2] is composed of a gain $g()$ and cost $c()$ function, parameterized by Q and A:

$$g(Q, A) = kQ^{\alpha}A^{\beta} \quad (2)$$

and:

$$c(Q, A) = t_qQ + t_dQA \quad (3)$$

First, note that we have revised the cost function to be expressed in terms of time, where the total time taken, ($T = c(Q, A)$), is the sum of the amount of time spent querying (t_qQ) and the time spent assessing (t_dQA). The parameters in the gain function represent (and summarize) a number of different elements over the search session. α represents how independent each query is from the next. If

α is set to one, then all queries are independent (i.e., no overlap of the result lists); if α is set to zero, then all the queries return the same result list. While, k and β represent the quality of the result list, where β is typically less than one and suggests that, as the searcher moves down the ranked list, they receive less and less gain (i.e., diminishing returns).

Given these functions, we want to determine the optimal number of queries $Q^\star$ and the optimal number of assessments per query $A^\star$ that maximize the gain for a fixed amount of time T. To solve this, we first create a Lagrangian Multiplier:

$$\mathbf{\Delta = kQ^\alpha A^\beta - \lambda\Big(t_q Q + t_d A Q - T\Big)} \quad (4)$$

Then take the partial derivatives.

$$\begin{aligned}\frac{\partial\Delta}{\partial A} &= kQ^\alpha\beta A^{\beta-1} - \lambda\Big(t_d Q\Big)\\ \frac{\partial\Delta}{\partial Q} &= k\alpha Q^{\alpha-1}A^\beta - \lambda\Big(t_q + t_d A\Big)\end{aligned}$$

By setting the partial derivates to zero, we can re-arrange each equation to equal lambda, and then substitute to remove the lambdas, to obtain the optimal number of assessments per query:

$$A^\star = \frac{\beta t_q}{(\alpha-\beta)t_d} \quad (5)$$

The corresponding number of queries then can be found by substituting $A^\star$ into Equation 2, where we assume that $g(Q, A)$ is equal to G:

$$Q^\star = \Big[\frac{G}{k}\Big]^{\frac{1}{\alpha}}\Big[\frac{(\alpha-\beta)t_d}{\beta t_q}\Big]^{\frac{\beta}{\alpha}} \quad (6)$$

Note that from Equation 5, α must be greater than β because $A^\star$ is required to be zero or greater.

4.1 Search and Search Behavior Hypotheses

Using a method called comparative statics [33], where all variables are held constant, except the one in question, it is possible to generate a number of hypotheses regarding search and search behavior[4]. Assuming that the searcher has a fixed amount of time T, and that they seek to maximize G then the model predicts how the search behavior, characterized by Q and A, would change in response to changes in other variables. For example, according to the model, if k, which relates to the amount of gain in documents, increases, then the model predicts that a user will issue fewer queries. The five hypotheses are summarized below:

k-hypothesis : as k increases, then $Q^\star$ will decrease, but $A^\star$ will stay constant.

α-hypothesis : as α increases, then $A^\star$ will decrease and $Q^\star$ will increase.

β-hypothesis : as β increases, then $A^\star$ will increase and $Q^\star$ will decrease.

t_q-hypothesis : as t_q increases, then $A^\star$ will increase and $Q^\star$ will decrease.

[4]Note this simpler economic model of search provides a subset of the hypotheses presented in [4].

t_d-hypothesis : as t_d increases, then $A^\star$ will decrease and $Q^\star$ will increase.

Given this economic model of ad-hoc topic search and the ensuing hypotheses regarding search behavior, an open question is whether the other approaches make similar claims about search behavior. In the following subsections, we will develop models of ad-hoc topic search using IFT and iPRP and determine whether they make similar predictions.

5. IFT MODEL OF SEARCH

In this section we apply IFT to model the ad-hoc search task over a session. Essentially, we wish to determine how many patches will be visited (which corresponds to the number of queries), and how long a forager should stay in a patch (which translates into how many documents they should examine per patch/query) given a fixed amount of time T.

First, we need to formulate IFT in the same terms of SET. To do this, we can take the cost and gain functions from SET (which are expressed in terms of Q and A) and re-express them in terms of gain given time. To simplify this process we shall consider the case when Q is one. By re-arranging Equation 3, the number of documents that are assessed is equal to the total patch time spent t minus the query time t_q, divided by the time per document t_d:

$$A = \frac{t - t_q}{t_d} \quad (7)$$

We can then substitute A into Equation 2 to arrive at:

$$g(t) = k\Big[\frac{t-t_q}{t_d}\Big]^\beta \quad (8)$$

which is the gain given time for one query, where if $t < t_q$ then $g(t) = 0$. To find the optimal time t in a patch, we need to determine when the rate of gain, i.e., $\frac{g(t)}{t}$, is maximized. As shown graphically in Figure 2, this happens to be the tangent to the curve $g(t)$ from the origin. To work this out algebraically, we first need to obtain the slope of the line given the curve. To do this we first take the derivative of the gain function:

$$g'(t) = \frac{k}{t_d}\beta\Big[\frac{t-t_q}{t_d}\Big]^{\beta-1} \quad (9)$$

Since we know that the line passes through two points $(0, 0)$ and $(g(t), t)$ and has a gradient given by Equation 9 it is possible to determine the equation for the line. In general, the slope (or gradient) m of a line is given by $m = (y_1 - y_0)/(x_1 - x_0)$; thus equating m to Equation 9 and solving for t, we obtain:

$$\begin{aligned}\frac{k}{t_d}\beta\Big[\frac{t-t_q}{t_d}\Big]^{\beta-1} &= \frac{k\Big[\frac{t-t_q}{t_d}\Big]^\beta - 0}{t-0}\\ \frac{\beta}{t-t_q} &= \frac{1}{t}\\ \beta t &= t - t_q\\ t^\star &= \frac{t_q}{1-\beta}\end{aligned} \quad (10)$$

By solving the equation above we arrive at the optimal amount of time a forager should spend in a patch, $t^\star$. Substituting $t^\star$ into Equation 7, we arrive at:

$$A^\star = \frac{\beta t_q}{(1-\beta)t_d} \tag{11}$$

To obtain a total of G gain then the forager would have to visit a number of patches (by issuing queries). So the number of queries issued would be equal to G divided by the gain obtained per patch. We can determine this as follows.

$$\begin{aligned} Q^\star &= \frac{G}{g(t^\star)} \\ &= \frac{G}{k}\Big[\frac{(1-\beta)t_d}{\beta t_q}\Big]^\beta \end{aligned} \tag{12}$$

In this case, the number of patches visited (i.e. number of queries issued) in a particular period of time T would be equal to $\frac{T}{t^\star}$.

Key Result: The IFT model results in a similar set of equations for $A^\star$ and $Q^\star$ to those of SET. In fact, the formulations are equivalent when $\alpha = 1$, and thus IFT and SET would make the same predictions regarding search behavior (as outlined in subsection 4.1).

α-difference: However, it also shows that there is a clear difference between the models. The IFT model does not immediately cater for the situation that patches (sets of results) may overlap, whereas in SET the α parameter expresses how much overlap there is between the result lists (where $\alpha = 1$ denotes no overlap, and $\alpha = 0$ denotes complete overlap).

Common Assumption: For the equivalence to hold it is assumed that in both IFT and SET the patch/result quality is the same across all patches/result lists. Of course, in practice, result quality varies from query to query. IFT addresses this limitation by employing the following theorem.

5.1 Charnov's Maximum Marginal Theorem

Charvnov's Maximum Marginal Theorem (CMMT) [11] states:

> "that a forager should remain in a patch so long as the slope of the gain function is greater than the average rate of gain in the environment." [23]

The theorem implies that if the forager is within a patch that yields less than the average rate of gain, then the forager should move to another patch, but if they are in a patch that has a higher than average yield, then they should stay in the patch. In Figure 5 the left plot shows the average patch distribution, where the tangent represents the average rate of gain. Let's now assume the forager moves to a patch, such as the one in the right plot, where the patch distribution is lower than the average. Instead of staying until time $t_2^\star$, which would be optimal if all patches were similarly distributed, now the forager would stay only until $t_1^\star$. This is because after this point the rate of gain would be less than the average rate of gain. Note, the theorem assumes that a forager has some idea of the average distribution of yields in patches.

More formally, we can mathematically determine how long an optimal forager would stay in a given patch as follows. Let's assume the forager visits a patch with k_i and β_i and the average patch is k and β. The optimal amount of time to spend in a patch is when the rate of gain in patch i equals the average rate of gain. Consequently, under CMMT, the

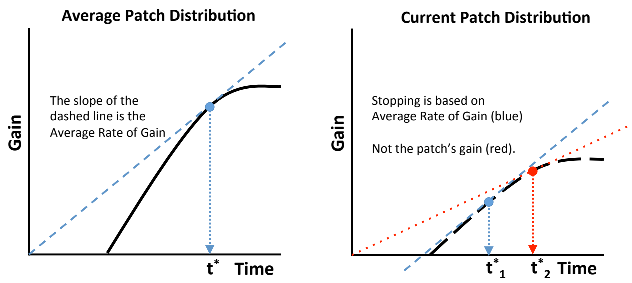

Figure 5: Applying Charnov's Maximum Marginal Theorem. Left plot shows the average patch distribution where the dashed blue line denotes the average rate of gain. The right plot shows that based on the average rate of gain the forager would spend less time in the patch leaving at $t_1^\star$ and not at $t_2^\star$.

model predicts that a forager's behavior changes depending on the patch distribution.

Essentially we would like to know when the rate of gain over time in the average patch (denoted by the slope of the line $(0,0)$ to $(t^\star, g(t^\star,\beta))$) equals the rate of gain over time in the new patch. Algebraically this becomes:

$$g'(t^\star,\beta) = g'(t_i,\beta_i) \tag{13}$$

since we know $t^\star$ for the average patch, we can solve for t_i, the total amount of time spent on patch i.

$$\begin{aligned} \frac{k.\beta}{t_d}\Big[\frac{\beta t_q}{(1-\beta)t_d}\Big]^{\beta-1} &= \frac{k_i\beta_i}{t_d}\Big[\frac{t_i-t_q}{t_d}\Big]^{\beta_i-1} \\ \Big[\frac{k\beta}{k_i\beta_i}\Big]^{\frac{1}{\beta_i-1}}\Big[\frac{\beta t_q}{(1-\beta)t_d}\Big]^{\frac{\beta-1}{\beta_i-1}} &= \frac{t_i-t_q}{t_d} \\ t_i = t_q + t_d\Big[\frac{k\beta}{k_i\beta_i}\Big]^{\frac{1}{\beta_i-1}}&\Big[\frac{\beta t_q}{(1-\beta)t_d}\Big]^{\frac{\beta-1}{\beta_i-1}} \end{aligned} \tag{14}$$

It is easy to check that if $k = k_i$ and $\beta = \beta_i$ then t_i will equal $t_\star$. If $\beta > \beta_i$ then the amount of time t_i decreases (thus less documents are assessed), and conversely so, if $\beta < \beta_i$. This result is consistent with the β-hypothesis.

k-difference: However, if $k > k_i$ then the amount of time t_i decreases and is less than $t_\star$ (thus less documents are assessed), and conversely so, if $k < k_i$. This result is inconsistent with the k-hypothesis and suggests that k will also influence the time/number of assessments per query. This suggests that the k-hypothesis should be revised, and the SET model should be revised.

6. SET REVISION

Before developing a model of search based on the iPRP, we will first consider how the insights from IFT can be incorporated into the SET model to make it more realistic. From the IFT model, we know that if k_i and β_i in the current patch are different from the average k and/or β then the foragers will change their behavior. However, under the current SET model, changes in k only influence Q because all the patches are assumed to be identically distributed. Below, we remove this limitation from the SET model.

First, assume that the searcher will issue Q queries. For a query q, the gain is characterized by k_q and β_q, and the searcher will examine A_q documents. The total cumulative gain then is the sum over all queries:

$$g(A_1 \dots A_Q) = \sum_{q=1}^{Q} k_q \alpha(q) A_q^{\beta_q} \tag{15}$$

and the corresponding cost function is:

$$c(A_1 \dots A_Q) = Qt_q + \sum_{q=1}^{Q} A_q t_d \tag{16}$$

where $\alpha(q) = q^\alpha - (q-1)^\alpha$, i.e., this breaks down the Q^α for each query q. Note that if all $k_q = k$, and all $\beta_q = \beta$ then all A_q will equal A, and we return back to the SET model described in Equations 2 and 3.

Now, let us consider the scenario where we have two queries, which have different k values. If we plot the marginal gain curves then we observe that as more documents are assessed the searcher will receive less and less gain (see Figure 6 where the marginal gain curve on the left has a lower k value than the one on the right). According to the CMMT, the depth the searcher ought to go to is when the marginal gain is equal across queries (see the dashed blue line in Figure 6). This shows that for the query with a higher k, an optimal searcher would examine more documents.

To formalize this intuition, lets consider the case where we have the average distribution k and β and the new query q yields k_q and β_q, and $\alpha = 1$; the gain function becomes:

$$g(A, A_q) = kA^\beta + k_q A_q^{\beta_q} \tag{17}$$

and the corresponding cost function is:

$$c(A, A_q) = 2t_q + t_d A + t_d A_q \tag{18}$$

Using the Lagrangian Multiplier method, we can solve the equation for A_q to obtain the following:

$$A_q = \left[\frac{k\beta}{k_q \beta_q}\right]^{\frac{1}{\beta_q - 1}} \left[A\right]^{\frac{\beta-1}{\beta_q-1}} \tag{19}$$

The optimal depth to go to when query q yields $k_q = k$ and $\beta_q = \beta$ is $A^\star$ (see Equation 5). However, if k_q or β_q decreases, then A_q decreases, while if k_q or β_q increases, then A_q increases. In relation to the IFT model, this revised SET model now makes the same prediction, i.e., if A_q in Equation 19 was expressed as time, then it would equal t_i from Equation 14.

Revised k-hypothesis: Thus, as k_i decreases w.r.t. the average patch distribution k, then A will decrease, and Q will increase. Under the revised SET model, the hypotheses between IFT and SET are now consistent with each other.

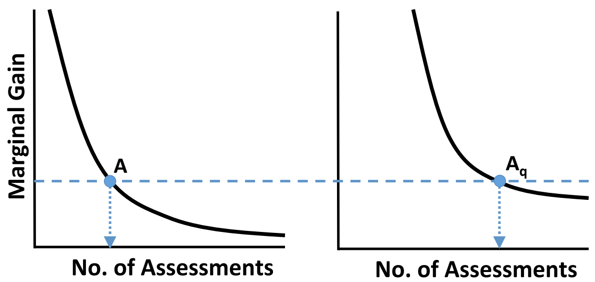

Figure 6: Plot of the Marginal Gain functions.

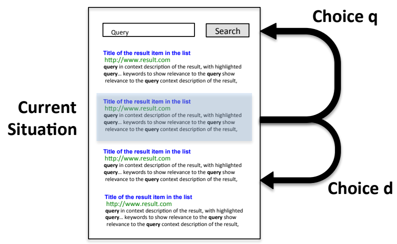

Figure 7: iPRP situations: Choice q: issue a query or choice d: examine the next document.

7. IPRP MODEL OF SEARCH

The iPRP provides a general framework for making low level decisions during the search process. To investigate how the iPRP relates to the other models, we need to frame the problem of ad-hoc topic retrieval in a similar manner. To do this, we consider that at any one point in time during the search process, the user is presented with a decision problem, which comprises of two choices: (i) issue a query or (ii) examine the next document.

More formally, we consider the user being in situation j, where they have just examined document $i - 1$. The user can take one of the following two choices:

Choice d: move to situation $j + 1$ by examining document i, with probability p_d which requires effort e_d; or

Choice q: move to situation $j + 1$ by issuing the query q with probability $p_q = 1 - p_d$ which requires effort e_q.

Since there are only two possibilities then $p_d + p_q = 1$. When the user starts the search they are in situation $j = 0$. As there are no documents to examine then $p_d = 0$ when $j = 0$, and so the only choice they have is to query.

If *choice d* is made, the user acquires a benefit of a_d, i.e. the benefit yielded by document d at rank i; if *choice q* is made, we assume they acquire a benefit of a_q. For the purposes of this analysis, we shall assume that when a user issues a query it also implies that they examine the first document in the ranked list and so the benefit that the user receives comes from the first document in the ranked list[5].

In SET and IFT, querying and assessing efforts are expressed in terms of the time required to form and issue a query, and assess a document. Under iPRP, effort is expressed as negative benefit, so $e_q = -t_q$ and $e_d = -t_d$. The benefit acquired from either of the two choices also needs to be expressed in terms of time. If we assume the same gain function as in SET (and IFT), i.e. $g(A) = kA^\beta$, then the gain of a document at rank i is the difference between position i and $i - 1$, i.e., $g(i) - g(i-1) = k(i^\beta - (i-1)^\beta)$. To calculate the benefit w.r.t. time, we substitute the gain of a document at rank i into Equation 8 and solve for t, where t equals the benefit a_d. Thus, the benefit of a document (at rank i) is:

$$a_d = t_q + t_d(i^\beta - (i-1)^\beta)^{1/\beta} = t_q + \gamma t_d$$

For simplicity of notation, we have set $\gamma = (i^\beta - (i - 1)^\beta)^{1/\beta}$. Following the assumption that the benefit of a query is provided by the benefit of the first ranked document

[5]Note we performed the same derivation but assuming the benefit of the query was zero, and came to similar findings.

(i.e., $i = 1$), we have:

$$a_q = t_q + t_d(i^\beta - (i-1)^\beta)^{1/\beta} = t_q + t_d$$

Using these values of efforts and benefits in Inequality 1, we obtain:

$$(t_d + t_q) - \frac{t_q}{1-p_d} \geq (\gamma t_d + t_q) - \frac{t_d}{p_d}$$

and multiplying each side of the inequality by $p_d(1-p_d)$:

$$p_d(1-p_d)(t_d+t_q) - p_d t_q \geq p_d(1-p_d)(\gamma t_d + t_q) - (1-p_d)t_d$$

This inequality can be further developed to derive a condition on the relationship between t_q, t_d and the gain γ. Such a condition is expressed as a function of the probability of examining a document, which is informed by a user model (see below). Thus, the above Inequality can be rewritten as:

$$-\frac{t_q}{t_d} \geq p_d(1-\gamma) + \gamma - \frac{1}{p_d} \tag{20}$$

When the above condition is satisfied, the iPRP predicts that the user is better off issuing a new query rather than examining the next document.

Differently from SET and IFT, iPRP considers a stochastic user, which assesses a document with probability p_d (or vice versa, issues a query with $p_q = 1 - p_d$). Therefore, we need some way to estimate or model p_d. The user model $\mathcal{U}$ assumed will determine how the user behaves. If empirical data was available it would be possible to estimate the probabilities. For the purposes of this paper, we shall utilize the same user model prescribed by the Rank Biased Precision [22] measure to approximate p_d at each rank.

7.1 RBP User Model Based

A user model that is often used for ad-hoc retrieval is that underlying Rank Biased Precision (RBP) [22], where a user examines the document at rank i with a probability $p_d^{\mathcal{U}_{RBP}}(i) = \rho^{i-1}$. Here, $0.5 < \rho < 1$ is a parameter that indicates the persistence/patience of the user, with $\rho = 1$ representing a persistent user that examines every rank position. Under $\mathcal{U}_{RBP}$, the iPRP predicts assessing is abandoned in favor of querying when the following inequality is satisfied (from Inequality 20):

$$-\frac{t_q}{t_d} \geq \rho^{i-1}(1-\gamma) + \gamma - \frac{1}{\rho^{i-1}} \tag{21}$$

Given these Inequalities, the iPRP model can be used to predict the search behavior. To illustrate the Inequalities in action, we have plotted the left (LHS) and right hand side (RHS) in the plots in Figure 8.

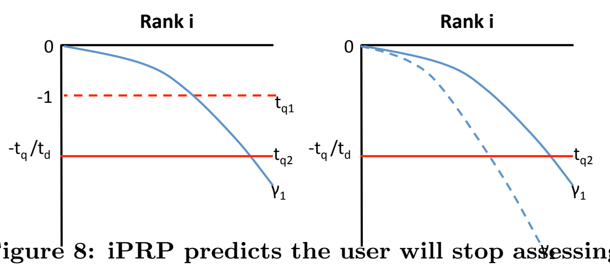

Figure 8: iPRP predicts the user will stop assessing documents and issue the next query when Inequality 21 is satisfied. Graphically, the inequality is satisfied when the red line (LHS) is above the blue line (RHS).

Key Results: First, if t_q increases (see Figure 8, left plot, where $t_{q2} > t_{q1}$), then the LHS becomes more negative, suggesting that a user should examine more documents before issuing a new query. If t_d increases, then the LHS becomes less negative, suggesting that a user would examine less documents before issuing a new query. These predictions are consistent with the t_q and t_d hypotheses.

If the gain increases (either via k or β) then γ increases (see Figure 8, right plot, where $\gamma_1 > \gamma_2$) and the RHS does not curve down as fast, suggesting that a user would examine more documents before issuing a new query. This is also consistent with the k and β hypotheses[6]. These findings mean that the SET, IFT and iPRP models all make similar predictions about search behavior.

***p*-difference**: Differently from the other theories, iPRP models the user as a stochastic agent (rather than fully rational) and thus introduces the probability of accepting choices. In the instantiation presented here this probability is modeled using the user model underneath RBP and thus Inequality 21 is also characterized by the persistence ρ that is attributed to the user. When $\rho = 1$, RHS becomes parallel (and greater) to LHS, meaning that the iPRP predicts the user will keep assessing documents without querying. As ρ becomes smaller, the RHS shape curves down faster and thus meets the LHS earlier. This suggests that as the user becomes more impatient, then they will switch from assessing to querying sooner.

8. EMPIRICAL EXAMPLE

To show the differences between the different models, we have calculated the number of assessments per query for a number of different conditions: (i) when the cost of a query is 30 seconds versus 15 seconds and (ii) when β (quality of result lists) is varied from 0.2 to 0.4. The estimates of k, β, t_q and t_d are based on, and are similar to, the empirically grounded values reported in [4] to provide a realistic set of values for ad-hoc topic retrieval. Table 1 shows the $A^\star$ for the SET model (and through equivalence the IFT model), along with the i value at which the user should switch from assessing to querying for different ρ values for iPRP (thus i and A both represent the number of assessments per query). The first thing to note is that as β increases, A and i increase. Under SET/IFT, the doubling of the query cost, also doubles the depth of assessment A. However, under the iPRP model, the increase in i is much smaller. For the iPRP model, as ρ increases i also increases, where the less patient the searcher the closer they become to the optimal stopping point with respect to the SET/IFT models. Intuitively, though, the predictions regarding the stopping point based on the iPRP seem to be more in line with how actual searchers tend to behave. In terms of the differences, it appears that the user's probability of accepting a choice is more influential in how far down the ranked list they will go than the change in performance (via β) and to a lesser degree the cost of a query (i.e. t_q). These observations motivate a number of lines of future investigation: (1) which model most closely reflects actual behavior, (2) how such a probability can be encoded within the IFT/SET models, and (3) which parameter(s) are the most important or influential in determining search behavior.

[6] Note that if we included $\alpha(i)$ into the model, then we could also derive the α hypothesis.

k	β	t_q	t_d	$A^\star$	$i, \rho 0.1$	$i, \rho 0.5$	$i, \rho 0.9$
When $t_q = t_d$							
0.25	0.20	15	15	0.25	1.21	1.70	5.57
	0.30			0.43	1.22	1.70	5.57
	0.40			0.67	1.23	1.72	5.59
When $2t_q = t_d$							
0.25	0.20	30	15	0.50	1.39	2.28	9.37
	0.30			0.86	1.39	2.28	9.37
	0.40			1.33	1.40	2.29	9.38

Table 1: Stopping points for different conditions. k is held constant across patches. The number of assessments per query is reported for each condition for $A^\star$ and i, ρ. As ρ increases so does the depth, but the change is rather invariant to changes in performance i.e., β increases.

9. DISCUSSION AND CONCLUSION

In this paper, we have taken three theories of Information Seeking and Retrieval and applied them to model ad-hoc topic retrieval. We created models of ad-hoc topic retrieval using the same notation to show how SET, IFT and iPRP are related. We enumerated a list of hypotheses about search behavior stemming from these models, and showed that each model makes similar predictions.

However, our analysis revealed a number of differences between models stemming from three parameters: k (result list quality), α (the result list overlap) and p (probability of accepting a choice).

k-difference: This difference arose because IFT predicted that foragers would stay longer in patches when k_i was increased w.r.t the patch average, whereas SET did not. This is because the original SET model was invariant to k, assuming it was constant across all patches/result lists. In the revised SET model, each result list was parameterized with an individual k_q making the model more realistic and generalizable. As a result, the initial k-hypothesis was revised, such that: as k_i increases w.r.t the average k, then the number of documents assessed (time spent in patch) also increases, thereby reducing the overall number of queries across the session.

α-difference: The α-difference came about because the SET model included an exponent that denotes how related subsequent queries were (essentially how much patch overlap there is). While we did not revise the IFT model, it would be possible to include another parameter in the gain function to denote this when visiting subsequent patches, i.e., the gain at patch q would be $g(t,q) = \alpha(q)g(t)$ where $\alpha(q) = q^{\alpha} - (q-1)^{\alpha}$ and q is the qth query issued. Similarly the gain (benefit) from a new query in the iPRP could also be discounted accordingly.

p-difference: The third difference stemmed from the probability of accepting a choice p that is within the iPRP, but not within the other models. The probability of accepting a choice could be incorporated into the other models, where instead of considering the gain and cost functions, they are modified to be expected gain and expected cost functions. The probability of accepting a choice is essentially the probability of taking an action, and so can be used to compute the expected gain and expected costs. This would mean we could frame the problem in the same way for IFT and SET. For example, the gain function for a single query q would need to be updated such that: $g(A_q) = k_q \sum_i p_i (i^{\beta_q} - (i-1)^{\beta_q})$ where i represents the ith document in the ranked list. Further work will be needed to perform the complete derivation to determine whether the updated gain function results in IFT and SET making predictions similar to iPRP (in terms of direction and magnitude). A key implication of the inclusion of p, under the RBP user model, is that the patience/persistence parameter has a greater influence on the depth a searcher goes to than the performance of the system.

In conclusion, this work represents the first major attempt to compare and integrate the different models/theories - and as such this theory-based paper formally shows under what conditions these models are equivalent, and what we can learn by exploring and developing such theories. This work, therefore, provides a bridge between these theories, paving the way forward for further theoretical developments and innovations in an area that is central to the field of Information Seeking and Retrieval.

This bridge lets us understand the user and system interactions at different levels through the different models, at the session level through SET, the patch level through IFT and the choice level with the iPRP. Thus, testing and validating the predictions of one model will essentially provide evidence to support the others. Also, depending on the specific data available, it is possible to instantiate one (or all) of the models to form common predictions about how changes to the system/interface will affect search behavior. The equivalence means observations and concepts in one model can be transferred to the other models further improving the models and refining the predictions regarding search behavior.

However, there are many future challenges and open questions that remain. Firstly, we need to understand more precisely how the different parameters impact and influence search behavior and search performance. From the examples we have provided we have seen that certain parameters have a greater impact on search behavior and performance than others, e.g., the probability of accepting a choice plays a major role in shaping the searcher's behavior and this appears to dictate interactions more so than a change in performance or cost. Further work is required here to explore a range of user stopping models for p taken from other evaluation metrics (other than RBP) and determine which is the most appropriate/accurate (i.e., which most closely resembles actual user stopping behavior). Secondly, we have only examined one possible gain function, however there are many other possible functions to be explored. Selecting and/or estimating an appropriate and realistic gain function poses a significant challenge. A possible direction here lies in drawing upon estimates/functions used in new measures such as the Time Biased Gain [30] and the U-measure [28] which encode in different ways how searchers extract gain over time from their interaction. This leads to two further points: (i) creating more realistic and accurate models of the gain/cost/interaction, and (ii) the integration between measures and models. Regarding (i) we focused on the most basic model of search, however search is more complex and so more sophisticated models need to be developed in order to better capture how people interact with search systems. Already some work has been done in this area [3, 4] extending the basic model by including interactions with search result page and snippets. The next step is to show how these additional variables and parameters can be added to

IFT and iPRP. Furthermore, we have also seen that through the iPRP model of search, the patience/persistence strongly influences the prediction on how far a searcher will go down in the ranked list. It seems that such model is perhaps more realistic and encoding such probabilities into the models provides a way to relax the rationality assumption. Thus, extending the IFT and SET models in this regard would provide a novel and valuable extension. With respect to (ii) we have shown that the user models employed by evaluation measures can be injected within these models of search. The obvious extension of this work is to explore the range of user models derived from evaluation measures [21], encode them within these models of search, derive the different predictions each one makes, and then empirically explore which user model best fits observed data. Less obvious is that we can look to develop more sophisticated measures of search performance by building these models of search into evaluation measures. Finally, but not exhaustively, is the need to test and validate the predictions and the assumptions given these models of search, and to build a body of evidence to support (or not) these models. Consequently, more empirically based studies that test and examine these hypotheses are required to show when they hold, and in what instances they breakdown. This is a vital step in the model building process as it enables further developments and refinements. The goal of this work was to show how these theories and their corresponding models related, while we have made significant headway in this direction, it is clear that much more research is needed.

Acknowledgments We would like to thank Professor Norbert Fuhr for his help and advice regarding the iPRP.

10. REFERENCES

[1] C. W. Axelrod. The economic evaluation of information storage and retrieval systems. *Information Processing & Management*, 13(2):117–124, 1977.

[2] L. Azzopardi. The economics in interactive information retrieval. In *Proc. of the 34th ACM SIGIR conference*, pages 15–24, 2011.

[3] L. Azzopardi. Economic models of search. In *Proceedings of the 18th Australasian Document Computing Symposium*, ADCS '13, pages 1–1, 2013.

[4] L. Azzopardi. Modelling interaction with economic models of search. In *Proc. of the 37th ACM SIGIR conference*, SIGIR '14, pages 3–12, 2014.

[5] L. Azzopardi, D. Kelly, and K. Brennan. How query cost affects search behavior. In *Proc. of the 36th ACM SIGIR conference*, pages 23–32, 2013.

[6] M. J. Bates. The design of browsing and berrypicking techniques for the online search interface. *Online Information Review*, 13(5):407–424, 1989.

[7] M. J. Bates. Training and education for online. chapter Information search tactics, pages 96–105. Taylor Graham Publishing, London, UK, 1989.

[8] N. J. Belkin. Some(what) grand challenges for information retrieval. *SIGIR Forum*, 42(1):47–54, 2008.

[9] N. J. Belkin, R. N. Oddy, and H. M. Brooks. Ask for information retrieval: part I: background and theory; part II: results of a design study. *Journal of Documentation*, 38(2) 61-71 and 38(3) 145-164, 1982.

[10] U. Birchler and M. Butler. *Information Economics*. Routledge, 1st edition edition, 2007.

[11] E. L. Charnov. Optimal foraging, the marginal value theorem. *Theoretical population biology*, 9(2):129–136, 1976.

[12] M. D. Cooper. A cost model for evaluating information retrieval systems. *Journal of the American Society for Information Science*, pages 306–312, 1972.

[13] S. Erdelez. Information encountering: a conceptual framework for accidental information discovery. In *Proc. of the ISIC conference*, pages 412–421, 1997.

[14] N. Fuhr. A probability ranking principle for interactive information retrieval. *Information Retrieval*, 11(3):251–265, 2008.

[15] M. D. Gordon and P. Lenk. A utility theoretic examination of the probability ranking principle in information retrieval. *JASIS*, 42(10):703–714, 1991.

[16] P. Ingwersen and K. Järvelin. *The Turn: Integration of Information Seeking and Retrieval in Context*. Springer-Verlag New York, Inc., 2005.

[17] K. Järvelin. IR research: systems, interaction, evaluation and theories. *ACM SIGIR Forum*, 45(2):17–31, 2012.

[18] K. Järvelin and T. D. Wilson. On conceptual models for information seeking and retrieval research. *Information Research*, 9(1):9–1, 2003.

[19] D. Kelly and C. Sugimoto. A systematic review of interactive information retrieval evaluation studies, 1967-2006. *Journal of the ASIST*, 64(4):745–770, 2013.

[20] C. C. Kuhlthau. Developing a model of the library search process: Cognitive and affective aspects. *RQ*, pages 232–242, 1988.

[21] A. Moffat, P. Thomas, and F. Scholer. Users versus models: What observation tells us about effectiveness metrics. In *Proc. of the 22nd ACM CIKM conference*, pages 659–668, 2013.

[22] A. Moffat and J. Zobel. Rank-biased precision for measurement of retrieval effectiveness. *ACM Trans. on Information Systems*, 27(1):2:1–2:27, 2008.

[23] P. Pirolli and S. Card. Information foraging. *Psychological Review*, 106:643–675, 1999.

[24] P. Pirolli, P. Schank, M. Hearst, and C. Diehl. Scatter/gather browsing communicates the topic structure of a very large text collection. In *Proceedings of the ACM SIGCHI conference*, pages 213–220, 1996.

[25] S. E. Robertson. The probability ranking principle in ir. *Journal of documentation*, 33(4):294–304, 1977.

[26] D. H. Rothenberg. An efficiency model and a performance function for an information retrieval system. *Information Storage and Retrieval*, 5(3):109 – 122, 1969.

[27] D. M. Russell, M. J. Stefik, P. Pirolli, and S. K. Card. The cost structure of sensemaking. In *Proceedings of the INTERACT/SIGCHI*, pages 269–276, 1993.

[28] T. Sakai and Z. Dou. Summaries, ranked retrieval and sessions: A unified framework for information access evaluation. In *Proc. of the 36th ACM SIGIR conference*, pages 473–482, 2013.

[29] P. E. Sandstrom. An optimal foraging approach to information seeking and use. *The library quarterly*, pages 414–449, 1994.

[30] M. D. Smucker and C. L. Clarke. Time-based calibration of effectiveness measures. In *Proc. of the 35th ACM SIGIR conference*, pages 95–104, 2012.

[31] G. J. Stigler. The economics of information. *The Journal of Political Economy*, 69(3):213–225, 1961.

[32] P. Thomas, A. Moffat, P. Bailey, and F. Scholer. Modeling decision points in user search behavior. In *Proc, of the 5th IIiX conference*, pages 239–242, 2014.

[33] H. R. Varian. *Intermediate microeconomics: A modern approach*. W.W. Norton, New York:, 1987.

[34] H. R. Varian. Economics and search. *SIGIR Forum*, 33(1):1–5, 1999.

[35] J. Wang and J. Zhu. Portfolio theory of information retrieval. In *Proc. of the 32nd ACM SIGIR conference*, pages 115–122, 2009.

[36] T. D. Wilson. Human information behavior. *Informing science*, 3(2):49–56, 2000.