

Published in final edited form as:

*J Biomed Inform.* 2008 April ; 41(2): 357–370.

## A Multi-Level Model of Information Seeking in the Clinical Domain

Peter W. Hung, MD, Stephen B. Johnson, PhD, David R. Kaufman, PhD, and Eneida A. Mendonça, MD, PhD

*Department of Biomedical Informatics Columbia University*

### Abstract

**Objective:** Clinicians often have difficulty translating information needs into effective search strategies to find appropriate answers. Information retrieval systems employing an intelligent search agent that generates adaptive search strategies based on human search expertise could be helpful in meeting clinician information needs. A prerequisite for creating such systems is an information seeking model that facilitates the representation of human search expertise. The purpose of developing such a model is to provide guidance to information seeking system development and to shape an empirical research program.

**Design:** The information seeking process was modeled as a complex problem-solving activity. After considering how similarly complex activities had been modeled in other domains, we determined that modeling context-initiated information seeking across multiple problem spaces allows the abstraction of search knowledge into functionally consistent layers. The knowledge layers were identified in the information science literature and validated through our observations of searches performed by health science librarians.

**Results:** A hierarchical multi-level model of context-initiated information seeking is proposed. Each level represents (1) a problem space that is traversed during the online search process, and (2) a distinct layer of knowledge that is required to execute a successful search. Grand strategy determines what information resources will be searched, for what purpose, and in what order. The strategy level represents an overall approach for searching a single resource. Tactics are individual moves made to further a strategy. Operations are mappings of abstract intentions to information resource-specific concrete input. Assessment is the basis of interaction within the strategic hierarchy, influencing the direction of the search.

**Conclusion:** The described multi-level model provides a framework for future research and the foundation for development of an automated information retrieval system that uses an intelligent search agent to bridge clinician information needs and human search expertise.

### Keywords

information retrieval; information seeking; online searching; search strategies; problem solving; user expertise; cognitive model; clinician information needs; intelligent agent

---

Address reprint requests to: Eneida Mendonça Vanderbilt Clinic, 5<sup>th</sup> Floor 622 West 168<sup>th</sup> Street New York, NY 10032. Corresponding author: Peter Hung Vanderbilt Clinic, 5<sup>th</sup> Floor 622 West 168<sup>th</sup> Street New York, NY 10032 Tel: (718) 309-4667 Fax: (212) 342-1647 peter.hung@dbmi.columbia.edu.

**Publisher's Disclaimer:** This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Data Supplement. Full librarian search.

## Introduction

Multiple information needs arise in the patient care setting during the normal course of work (1-4). Gaps in knowledge can adversely affect the ability of a physician to make the best care decisions. Unfortunately, these needs often remain unmet for a variety of reasons, despite the presence of an ever-growing variety of online resources. Some common obstacles that prevent physicians from pursuing or finding answers to their information needs include: doubt that relevant information exists, lack of time to initiate a search, uncertainty about where to find the information, and the absence of relevant topics in the resources searched (5). Additionally, the process of searching an information resource for an answer may be disruptive to natural work flow. A recent survey of more than three thousand physicians showed that while the Internet was perceived as an important source of information, barriers to Internet information use were also high, with 57% of physicians reporting significant navigation and searching difficulties, 49% reporting too much information to scan, and 45% reporting the belief that specific information to answer their question was not available (6). These findings underscore a need to develop information retrieval systems that can help bridge these barriers.

One challenge is that while physicians have expert knowledge in the domain of patient care, they often do not possess the expertise necessary to translate their information needs into search strategies that produce the desired answers (7;8). Beyond the initial barrier of articulating one's information need, many physicians do not know how to choose and search online resources effectively in order to find correct answers to their questions (9;10). Furthermore, online information resources differ significantly in terms of their features and search interfaces, such that an effective search strategy for one resource may not work for another. Because of this, a clinician may rely on a few "favorite" resources while lacking the expertise to effectively search other more appropriate resources for a particular information need.

On the other hand, health science librarians represent a wealth of search knowledge within the biomedical domain. As a result of formal training and practical experience, they are adept at matching questions to information resources that are likely to yield appropriate answers. They also routinely employ systematic techniques for conducting a productive search for information, and often know how to navigate specific information resources in ways that casual searchers do not (11;12).

Often overlooked is the fact that the searcher's choice of strategies and techniques plays a central role in the effective retrieval of information, and that "online information retrieval is a problem-solving activity of a high order, requiring knowledge and understanding for consistently good results" (13). Creating a system capable of such problem solving—one that employs an intelligent agent to make decisions about how to look for an answer by generating flexible search strategies based on human expert search knowledge—could be a particularly useful approach to meeting the information needs of clinicians. Loosely related efforts in biomedicine have focused on constructing hand-coded, pre-determined queries for narrow question types and particular resources (e.g., diagnosis questions in PubMed) (14-17). In other domains, researchers have built systems that incorporate particular algorithmic search strategies based on those used by reference librarians (18), but little has been done to explore and develop technologies that are able to accept a defined information need and autonomously generate complex, adaptive search strategies that will conduct a high-precision search to retrieve an appropriate answer, changing course as necessary during the process. A system that employed an intelligent search agent with such functionality would undoubtedly be a valuable tool for satisfying specific information needs that arise within the context of clinicians' work, particularly if supported by a user interface that accepts questions and their context in a form that clinicians are familiar with, i.e. natural language.

Before one can feasibly build a system that leverages human expert search knowledge, the nature, structure, and process of searching for information, as well as where and how search experts apply their knowledge, must be understood. Therefore, it is important to elucidate an appropriate theoretical model of information seeking that will guide and facilitate the formal representation of search expertise. Towards this purpose, the present paper describes the development of a theoretical model of online information seeking within the biomedical domain.

Before moving further, we more precisely describe the target of our modeling efforts. Consider an information retrieval system that entails the participation of two parties: an information *requester* (e.g., a clinician) and an information *searcher* (e.g., a human librarian or intelligent agent). The question that we are trying to answer is: given a well-defined information need that has been both adequately expressed by the requester and understood by the searcher, what types of search knowledge would the searcher require in order to solve the problem of retrieving an appropriate answer that meets the information need? The focus of the model to be presented is therefore on the generation and execution of a search strategy in the context of a single search session and in response to a single information request (as opposed to a more open-ended information seeking process involving multiple search sessions over time, as described elsewhere in the literature). Additionally, our interest lies particularly in the use of electronic resources and the strategies utilized to identify and seek information from these resources.

From the outset, we emphasize the limited scope and pragmatic nature of our model, as it is ultimately meant to serve as the architectural basis of a strategy-generating module to be used by an intelligent agent within an automated search system. As such, the model that we present does not aspire to be exhaustively descriptive nor explanatory of information seeking (e.g., how information needs emerge, why people conduct a search in the first place, what they do with the answers). The paper does not explicitly present a cognitive model of information seeking, but rather an explanatory vocabulary that can be used to articulate such a model. The goal of the current model is to present the dimensions of search knowledge needed to execute a context-initiated search process.

## Background

In presenting a nested view of information seeking behavior, Wilson differentiates between information behavior and information seeking. *Information behavior* is defined as “those activities a person may engage in when identifying his or her own needs for information, searching for such information in any way, and using or transferring that information.” More finely grained is *information seeking*, which concerns “the variety of methods people employ to discover, and gain access to information resources” (19). The model developed in this paper can best be described as dealing with a subset of information seeking that has been variously characterized as *embedded* (20;21), *task-oriented* (22), or in other words, *context-initiated*, by virtue of the fact that it occurs within the context of work-related tasks, and specifically in our case during the course of patient care.

Much of the modeling research within the information science literature has focused on information behavior or generalized information seeking at a macro level (23-27). Although these models are usually too broad and lack the granularity necessary to be used as the architectural basis for an intelligent search agent, certain constructs are helpful in focusing our scope. In a model describing information seeking on the Web that extends work by Ellis (28), Choo delineates four main modes of online information seeking: undirected viewing, conditioned viewing, informal searching, and formal searching. The reasons for entering each mode are distinctly different. Users with non-specific information needs will engage in undirected viewing (e.g., browsing a general news website) or conditioned viewing (e.g.,

monitoring a bookmarked industry website relevant to one's work for updates). Those with specific information needs will utilize informal search (e.g., submitting a simple, unstructured query to a local search engine) or formal search (e.g., systematically performing a comprehensive search on a particular topic) to extract particular information (29). As pertains to our model, this influences the direction of our modeling efforts towards the representation of searching modes as opposed to viewing modes, since the end goal is to support information seeking strategies for specific clinician information needs rather than non-specific ones.

The information seeking model put forth by Marchionini closely describes a formal online search from a cognitive perspective, which we found useful to articulate the general components of such a search at a high level. The model is meant to provide an overview of the online information seeking process, rather than a detailed representation of the search itself. The information seeking process is described as being iterative in nature and composed of a set of interconnected subprocesses: (1) recognize and accept an information problem, (2) define and understand the problem, (3) choose a search system, (4) formulate a query, (5) execute search, (6) examine results, (7) extract information, and (8) reflect/iterate/stop. The subprocesses can be thought of as activity modules that default into a sequence of serial phases but may also proceed in parallel, depending on the intermediate results of the search (30).

Marchionini's model outlines some of the types of knowledge that are necessary for the information seeking process. For example, the knowledge that needs to be applied to choose an appropriate search system is clearly distinct from the knowledge that one needs to formulate an effective query. Additionally, Bates has made the distinction between search strategy and search tactics, the first of which encompasses knowledge about how to conduct a search at an overall level while the second requires knowledge about short-term moves to advance a search (31). Recent work by Bhavnani and Bates supports the idea that searching for information indeed utilizes several types of search knowledge. Using hierarchical goal decomposition, which is a method to systematically break down goals into component levels of subgoals, the search task can be decomposed into intermediate, resource, and keystroke layers. These layers require knowledge that pertains to search strategy, resource-specific search methods, and motor actions/keystrokes, respectively (32).

The search for information has previously been compared to other complex problem-solving activities. For instance, Harter likens the online information searcher to an investigator engaged in scientific inquiry. Scientific experimentation is a trial-and-error process which involves identifying important variables, formulating hypotheses, gathering data to test hypotheses using pragmatic methods and procedures, and evaluating the results in response to a research question. Similarly, when a librarian performs an online information search, he or she identifies concepts of interest, formulates queries that represent "best guesses" of what will yield the desired information, utilizes standard procedures and methods in the language of the search system to gather data, and evaluates the results in view of a particular request for information. As with scientific inquiry, online searching is an iterative process that requires refinement and revision to move closer to the goal (13).

Given the similarities of information seeking to scientific inquiry, the theoretical framework put forth by Klahr and Dunbar to model scientific discovery is informative. This framework, which is taken from a cognitive scientific perspective, considers scientific discovery as a complex problem-solving task. As defined by Newell and Simon, a *problem* consists of an initial state, a goal state, and a set of operators that can be applied to reach the goal state through a series of intermediate steps. There are also constraints that must be satisfied before a given operator can be applied. The *problem space* is defined as the set of all possible states, operators, goals, and constraints (for instance, the set of all possible moves on a chess board that a player can make during the course of a game), and the *problem-solving process* is the search for a

path through the problem space that connects the initial state to the goal state (33). Initial research on simple laboratory studies described a model consisting of two problem spaces, one of experiments and one of hypotheses. This was based on the observation that participants in these studies appeared to be focused either on the space of possible manipulations or on the space of possible explanations for experimental results (34). In later work, Klahr expanded the framework to more adequately represent a wider range of problem solving related to scientific discovery. This expanded framework views scientific discovery as requiring coordination between several problem spaces (the number varying with the context), including spaces of hypotheses, experiments, paradigms, data representations, strategies, and instruments (35).

## Formulation Process

We applied the information seeking models of Choo and Marchionini described in the previous section to focus our model formulation efforts to the segment of information seeking most applicable to our particular purpose. Thus, we limited our model to searching--as opposed to viewing or browsing--since our goal is to address those questions arising from clinical care that are fairly well-defined and thus potentially answerable through online information sources. In terms of Marchionini's information seeking model, it was apparent the clinician would have perceived a gap in knowledge recognized the nature of the problem as being answerable through an online information search. In formulating our model, we also assumed that the clinician's expressed information need would have been analyzed, adequately defined, and understood by the searcher.

The formulation of our model was a fundamentally iterative process that, following the generation of an initial set of ideas, was characterized by cycles of refining these ideas through the lens of published literature and synthesizing additional ones at weekly meetings between the authors. After an initial review of the biomedical, computing, and information science literature failed to uncover suitable models of contextualized information seeking for our goals, we proceeded to explore information search at a more fundamental level to deepen our understanding of the nature of the problem. Early discussions directed us towards the artificial intelligence literature to explore possible parallels between problem solving and information seeking.

Given this approach, it was natural at first to conceptually equate the process of information seeking to the idea of a formulaic search algorithm. Conception of contextualized information seeking as a prototypical problem solving task initially led us to consider representing information search within a single problem space, as some have suggested (18). In such a representation, online information search is a sequence of individual moves—such as navigating to an online resource, entering a search term, or reformulating a query—that proceeds through a single problem space. Search knowledge is encoded directly above the level of allowable moves (as a set of rules, for example).

However, further exploration identified the need for a different representation due to the complexity of the task. The limitation of the single space approach was apparent, since for any given information need, the number of ways to proceed with the search and the possible results at each step are almost unlimited, creating an exceedingly large problem space. Also unclear was how to adequately represent the different types of search knowledge necessary to effectively traverse this problem space in order to reach the goal.

We began examining how complex tasks in other domains have been modeled for points of applicability. Klahr's work on modeling scientific discovery as a complex problem solving task was found to be particularly relevant, given the apparent similarities with information seeking mentioned earlier. In light of this complexity and the commonalities between information seeking and scientific inquiry, we determined that modeling online contextualized information



seeking across multiple spaces would facilitate the abstraction of search knowledge into functionally consistent layers, or levels. Besides being compositionally elegant, such multi-level abstraction is advantageous in that it reduces the complexity of decisions to be made at each level while improving transparency by explicitly separating the distinct types of knowledge required to solve the problem (36). The distinct layers of knowledge were identified through the information science literature as discussed previously and validated through our observations of searches performed by health science librarians.

## Model Description

### Overview

We propose a hierarchical multi-level model of contextualized information seeking. Each of the four levels in the hierarchy—grand strategy, strategy, tactics, and operations—represents a separate problem space that is traversed during the online search process. The final level, assessment, provides feedback and guides the direction of the search. Every level represents a distinct layer of search knowledge that is required to execute a successful search. The model is an idealized, abstracted representation of information search strategy generation/execution, rather than an explanatory model of observed human information seeking behavior. As a starting point, we assume that the information to be sought arises from the clinical domain, and from within a particular context related to patient care. We also assume that the searcher has ascertained the information requester's information need such that an initial set of relevant concepts can be used as the basis of the search. The model describes the online search for information beginning with the goal of formulating an overall plan to retrieve an appropriate answer.

In general, problem solving is concerned with achieving the ultimate goal while also minimizing the resources used. Solving complex problems often encompasses multiple hierarchical levels of strategy, with each layer characterized by a different scope. Short-, medium-, and long-range objectives are represented by different strategy layers, and the goals of a given layer are accomplished by the one beneath it. Similarly, the overall goal of an information search is to retrieve information relevant to the requester in the most efficient manner possible. The relevant goals and subgoals vary in scope and fall into a natural hierarchical structure, which we refer to as the *strategic hierarchy*. An important characteristic of the strategic hierarchy is that the problem space at each level is constrained by the level(s) above it. Additionally, the search knowledge used to accomplish these goals separates into distinct layers, with the goals of one layer being accomplished by the methods of the one beneath it.

Bates borrows military terms to introduce the concept of search tactics, distinguishing them from search strategies (31). Such terms have been used variously in other domains such as business and game-playing, and are useful to describe and distinguish various aspects of online information seeking as a goal-oriented, problem-solving activity. In the present model, we expand on this naming scheme and apply it more widely to information seeking.

We define an *information resource* (sometimes referred to simply as *resource*) as a source of information that is available online and accessible to be searched for relevant information. The forms that such a resource could take are variable, and include online textbooks, databases, informational websites, and search engines. Also included in the definition of resources are meta-resources, such as a search engine that queries other search engines, or a searchable collection of online journals.

The iterative nature of searching for information as articulated in the model in this paper runs parallel to Norman's Theory of Action, which generally describes human action upon the world

as being a cycle that involves stages of execution and evaluation. During execution, an overall goal is translated into intentions, which in turn are mapped to executable action sequences. The system state is then evaluated in view of goals and intentions, thus completing an action cycle. (37) Similarly, the goal of performing a search for information in the multi-level model is accomplished via choices made within the strategic hierarchy—abstract “intentions” of increasing granularity—which ultimately map to concrete operations that execute the search. The assessment layer completes a search cycle by providing relevance feedback to the searcher in light of current search goals.

## Grand Strategy

In general, grand strategy is the broadest conception of how to integrate available resources to attain the ultimate objective. As it applies to our information seeking model, the *grand strategy* comprises the highest level plan of the online search, and sits atop the strategic hierarchy. The grand strategy determines what resource or resources will be searched, for what purpose, and in what order. Grand strategy is not concerned with the operational details of how to actually search the resources chosen. Appropriate resource selection depends on the type of question (e.g., therapy), the expected answer type (e.g., journal article), and the context that initiated the question (i.e., why the question is being asked). It is also influenced by whether the objective of the search is to retrieve results of high relevance or high precision. An effective grand strategy accurately maps the characteristics of a given question arising from a particular context to the information resources most likely to yield the desired answer. Search experts are more likely to be able to do this than unskilled searchers because there is a substantial difference in knowledge of where to go for answers to particular kinds of questions.

Formulating a grand strategy involves more than resource selection, and should not be equated with simply selecting a search system as described by Marchionini and others (27;30). Rather, grand strategy consists of one or more *phases*, with each phase calling for a single resource (e.g., Harrison's Online), or type of resource (e.g., online medical textbook) to be searched for a particular purpose. Note that searching a meta-resource can also be specified here. The purpose for searching a resource may be to answer a simple question in its entirety, or to answer part of a complex question that contains several concepts, each requiring a different resource. Phases may be executed in series or in parallel, depending on what is specified by the grand strategy. There is no inherent limit on how simple or complex a grand strategy must be. A common grand strategy chosen by search novices involves utilizing a single general search engine, such as Google. Other grand strategies may be much more complex, involving several types of online resources to corroborate or piece together possible answers that are retrieved. Grand strategies are partially the product of experience (e.g., having successfully addressed an information need using a particular resource) and partly based on explicit knowledge (e.g., knowing the scope and limitations of different resources).

Should a phase of the grand strategy fail, a decision must be made about what course to take within the grand strategy problem space, e.g., to continue with the current grand strategy by making changes to the phase, or to choose another valid grand strategy. A grand strategy is successful when the overall goal of the search has been accomplished through the retrieval of an appropriate answer. The overall search goal will not be reached if (1) all grand strategies relevant to the search fail, or (2) constraints on resources (e.g., time) prevent the search from continuing.

## Strategy

Up until this point in the paper, the term strategy has been used loosely to refer to any plan relating to information seeking, regardless of scope. We now define *strategy* to mean an overall approach for searching a single resource to answer an information need. This is more

constrained and specific than Bates' definition of search strategy as “a plan for the whole search” (31).

Narrowing the scope of the term is necessary, because there is an important distinction between grand strategy and strategy. The grand strategy level may, and often does, specify the searching of more than one information resource, while the strategy level dictates how to proceed with the searching of a single resource that was specified by the grand strategy. More precisely, strategies are always instantiated within the context of a particular phase of a grand strategy, and have as their goal the completion of that associated phase.

As with grand strategy, a strategy to search a particular resource may be simple or complex. Search experts routinely employ a variety of systematic strategies, often in combination with each other (30). A naïve information searcher, on the other hand, may utilize a very simple strategy of choosing a few terms, executing the search, and evaluating the results with little iteration.

Examples of strategies that are described in the information science literature include briefsearch, building blocks, successive fractions, and citation pearl-growing. “Briefsearch” as described by Harter (also known as “quick and dirty” or “easy” search) consists of composing a simple query within the chosen resource using a few terms linked by a Boolean AND operator and evaluating the resulting output. There is little to no iteration in this simple strategy (13). Popular among search novices, search experts also use this strategy at times for background or intermediate information (30). The “building blocks” strategy involves breaking the query down into separate concepts, identifying terms for each concept, searching each concept separately, and combining the results of each sub-search in a final search (13). In “successive fractions,” a large initial set of results is iteratively pared down by successively adding concepts that are specific to the problem. The “pearl growing” strategy takes a specific document or citation known to be of high relevance and uses characteristics of this reference item such as index terms, keywords, citations, and publication data to retrieve additional results; the process is repeated until the “pearl” has reached desired size (38).

A strategy consists of one or more *objectives*, which represent conceptually distinct elements of the strategy, or intermediate subgoals. The strategy meets its overall goal of retrieving an answer when all objectives are accomplished. These objectives may or may not have an order in which they must be achieved. The mechanics of how to execute a strategy within a given resource are not addressed at the strategy level, but rather left to layers that are lower in the model hierarchy.

Should an individual objective fail to be achieved, a decision will need to be made whether to choose another objective in an attempt to complete the strategy, or to abandon the strategy for a different one. Obviously, if all alternative objectives have been exhausted, the strategy will fail. In the event that a strategy has failed, the two options are to try another valid strategy within the strategy problem space or to terminate efforts to search the current resource, in which case the current phase of grand strategy will fail.

A key point to recognize is that the strategy problem space, and thus choice of strategy, is constrained by the resource being searched. In other words, the number of appropriate strategies to search a given resource is substantially smaller than the universe of all strategies. Some resources will be searchable using certain strategies while others will not. As an example, PubMed may be amenable to a search strategy that makes use of a citation's references to find additional relevant articles, while a typical informational website may not. Choice of strategy is also influenced by the overall search goal—for instance, a high precision search will likely require a different strategy than a high relevancy search.



## Tactics

From a general problem-solving perspective, the term *tactics* represents a limited plan of action for using available means to achieve a specific limited objective. Tactics are intended for use in implementing a wider strategy, and there is little use in performing a tactic without a guiding strategy. A tactic in the context of an information search can be defined as a localized maneuver made to further a strategy. The information science literature describes numerous tactics available to the searcher depending on the situation at hand (31;39;40).

For the purposes of this model, tactics can be applied at the limit, query, facet, and term level. A *limit* refers to a restriction that can be imposed on a set of results, such as limiting the publication year to 2007. A *query* consists of one or more search concepts (facets) expressed in the language of the information resource. Queries can be connected using Boolean AND, OR, or NOT operators. *Facets* in a query represent distinct concepts of interest (13), and are connected to other facets using Boolean AND or NOT. Each facet is represented by one or more related *terms* that are connected by a Boolean OR.

Generally, tactics are employed to increase either the recall or precision of a search, and/or to increase or decrease the size of the resulting set. Examples of tactics that may be available at each level for a typical Boolean search system are listed in Table 1. Note that this is not meant to be a definitive list, and that other restricting or expanding tactics may be available as well. For instance, proximity operators may be supported which are more restrictive than AND, such as NEAR (requiring terms to be within a set distance of each other) or WITH (requiring terms to be within the same field). Additionally, different resources will support a different range of tactics; thus, not all tactics will apply to all resources. As an example, a resource that only allows simple keyword searches may not support the use of certain Boolean or proximity operators.

At any given point in the search, the tactics problem space is constrained by the information resource being searched (as determined by the grand strategy) as well as the current objective of the chosen strategy. Also important to note is that the descriptions of tactics are inherently generic; at this level, we are not concerned about how tactics are actually implemented within a given resource. This is left instead for the level of operations, described next.

## Operations

In general, an operation refers to “a process or series of acts involved in a particular form of work” (41). Thus far in the current information seeking model, nothing has actually been entered into the search system to change its state. To this point, the levels described have been abstract. However, at the level of *operations*, which is the lowest level in the strategic hierarchy, intentions are mapped to specific, concrete input (commands, keystrokes, mouse clicks, etc.). Operations are defined in terms of the mechanisms available to implement a particular tactic within an information resource, and thus are resource-specific translations that execute tactical decisions. Depending on the resource, there may be more than one way to implement a desired tactic (e.g., either entering a command-line statement or indicating the action through a point-and-click interface). While tactics may be general to several systems and are independent of the actual interface, operations are specific to the kind of interaction supported by the system (e.g., links vs menu selections). For example, PubMed and Ovid support somewhat different kinds of interaction and this would necessitate the use of different operations.

In some instances, a tactic may translate directly into a single system command. For example, removing a term from a particular facet in PubMed can be accomplished by eliminating the term and the operator that precedes it (if there is one) from the query string. In other instances, translating a tactic may require a string of commands or the combined use of more than one

resource-specific tool. For example, replacing a search term with a broader one in PubMed might involve identifying a parent term using the MeSH browser, then replacing the original term in the search string. An operation thus consists of one or more atomic *moves* that when performed, complete the intended tactic. The final move in any operation is to actually execute the search, whereby results are retrieved and the search process turns to the assessment task.

## Assessment

The *assessment* level provides the basis of interaction with the levels of the strategic hierarchy just described. This level differs from the preceding ones in that it is distinct from the strategic hierarchy, being involved with providing relevance feedback that informs the progress and status of the search to this point. The level of assessment can thus be thought of as part of a control structure that influences the direction of the search at all levels of the hierarchy.

Execution of a search cycle yields a set of results. Evaluation of the relevance of these results takes place on the level of assessment, and can be performed at varying degrees of detail. The amount of information that each type of assessment yields about the search results is proportional to the amount of processing time required to complete the assessment (Figure 1). At one end of the spectrum, the searcher can assess whether the present search produced “hits” or not. This type of assessment costs very little processing time, but also yields a small amount of information. At the other end of the spectrum, the searcher can scan the full text of the retrieved results for relevance, which has a high time cost but yields much more information on relevance.

Choice of which level(s) of relevance to assess at any given point in the search will be determined by: (1) what the particular search strategy requires, (2) the stage of the search, and (3) the amount of time available to commit to the search. While some search strategies will use simple measures of relevance to guide the progression of a search, others will depend on more detailed feedback to provide a substrate for further search. Additionally, most searches are highly iterative in nature, such that assessment is influenced by the stage of the search; exploratory stages may only require the assessment of hits to “test the waters” or ensure that one is on the right track, while later stages where the answer is believed to be in the current set of documents may necessitate the scanning of document sections. Finally, time constraints will factor into assessment decisions—a search where time is not a factor will be able to afford more time-costly assessments than one where speed is a priority. Particularly in the context of patient care, speed is of the essence.

From a problem-solving perspective, the assessment level represents the evaluation of the present state within the context of the strategic hierarchy, which has the effect of determining if the last step in the search was successful in meeting the current limited objective. Assessment of the retrieved results feeds back through the hierarchy in a dependent fashion, beginning with the lowest level of operations. Successful accomplishment of the current goal at one level (e.g., operations) leads to an evaluation of the problem state at the level directly above it (e.g., tactics). On the other hand, failure to meet a goal at the level being evaluated necessitates the decision to either explore other states in the current level’s problem space, or to abandon the space and return to the level above it. At the highest level, should a grand strategy fail and all other feasible grand strategies been exhausted, the entire search will fail.

## Integrated Model Summary

We now briefly summarize the various parts of the aforementioned model and explain the general flow of the search process in relation to those parts (Figure 2). At the starting point of our model, the searcher has ascertained the requester’s contextualized information need and determined an initial set of relevant concepts to be used as the basis of the search.

The search proceeds through the multiple levels defined by the strategic hierarchy. Recall that each strategic level is a distinct problem space that contains possible choices at that level. Search begins at the grand strategy level, with the formulation of a grand strategy that is appropriate for the type of information need and anticipated answer type. The rest of the strategic hierarchy is constrained to the particular resource chosen. After a resource has been selected, the search progresses to the level of strategy and the searcher decides upon a strategy to search the resource. Once this is accomplished, attention turns to the tactics level (further constrained by the current objective of the chosen strategy) in order to advance the strategy at this point in the search. The tactic may be at the limit, query, facet, or term level. Generally, tactics are employed to increase either the recall or precision of a search, and/or to increase or decrease the size of the resulting set. Executing a particular tactic in the context of a given resource requires appropriate knowledge at the operations level in order to utilize one or more tools available through the resource's interface.

At this point, the chosen search operation is executed and a results set is generated. At the assessment level, the searcher can perform various types of evaluations to ascertain if the goal state has been reached (answer has been found), and if not, what the next step should be. Among the possibilities would be to choose another tactic to move the current search strategy forward, choose a new search strategy, try a different resource, select a different grand strategy, or conclude a failed search (perhaps after a certain amount of resources has been expended, e.g. time).

### Example of Application of the Model

To illustrate the different levels of the strategic hierarchy, as well as the flow of a search as it unfolds within the model's constraints, we use the example of a search performed by a health sciences librarian on an actual clinician information need arising from patient care. It is important to note that this exercise is designed as a sufficiency test rather than as a comprehensive validation; the question we address here is whether the model and explanatory vocabulary can adequately account for the information searcher's behavior. For this particular search task, the librarian was given a clinical scenario along with the following question to answer using Ovid MEDLINE: what is the cause of GI bleeding related to renal failure in a patient with diabetes? A think-aloud protocol was employed and the librarian's actions were recorded using Morae (TechSmith Corporation, Okemos, MI), a commercially available usability tool that performs screen capturing, keystroke logging, and video recording. A partial listing of the search is shown in Figures 4 and 5, while a complete listing can be found in the accompanying Data Supplement.

The grand strategy of this particular search was pre-defined by the search task, i.e. to search for the answer using Ovid MEDLINE. Within Ovid, the searcher chose a systematic strategy that closely followed a building blocks approach: separately defining terms for each of the three facets representing diabetes, GI bleeding, and renal failure (corresponding to Objectives 1-3, shown in Figure 3); combining the facets to yield an initial answer set (Objective 4), and reformulating as necessary (Objectives 5-9; the last two are shown in Figure 4). To complete each objective of the building blocks strategy, certain tactics of limited scope were employed. Each tactic was mapped to specific operations within Ovid, and each operation was carried out via a series of observable atomic moves. Upon execution of each search cycle, the librarian performed an assessment of the results generated. This assessment ranged from noting the presence of hits to evaluating the number of hits to scanning the titles of the results set, depending on the current stage of the search. From here, the search continued on to the next objective (in the case of successfully completing the current objective) or a new objective was formulated (in the case of failing to complete an objective). Feedback from the assessment

level prompted adaptive decisions at the strategic objective level. Continuing with the initial strategy eventually yielded an appropriate answer set.

We now proceed to step through the search in more detail. As shown in Figure 3, the grand strategy was to search Ovid for the answer. Within this resource, the searcher then elected to approach the search using a building blocks strategy (previously described). The first objective of this strategy was to define the diabetes facet. To accomplish this objective, the tactic of adding broad, structured terms to define the facet was employed. By structured terms, it is meant terms that exist as controlled indexing items within the resource (in contrast to free text or keywords, which would have constituted a different tactic). Operationally, performing this tactic in Ovid involved making moves to enter “diabetes” into the main search box, which yielded a list of possible structured MeSH terms for selection; checking “Diabetes Mellitus” as the structured term of interest; selecting the “Explode” and “All Subheadings” options to broaden the terms used; and executing the search. A large results set of 181,558 hits was retrieved. Assessment of the search results appeared to simply consist of noting the presence of hits, as the librarian did not pause in response to the results but rather continued on with defining the next facet. An identical process was followed for defining facets for gastrointestinal bleeding (Objective 2) and renal failure (Objective 3).

With the three major facets defined, Objective 4 was to combine all facets together in order to formulate an initial answer set. The tactic to accomplish this objective involved combining all three facets with Boolean AND. This mapped to entering “#1 AND #2 AND #3” on the operations level. The retrieved result (Result 4), consisting of 14 hits, led to an assessment that involved scanning the titles of the retrieved citations. No relevant articles were found based on their titles (Assessment 4), which necessitated the formulation of different objectives to increase relevance.

The objectives chosen to continue the search were still consistent with the overall building blocks strategy—objectives consisted of either redefining individual facets, or intersecting facets in various combinations. Next, the librarian tried combining only the gastrointestinal bleeding and renal failure facets (Objective 5) using similar tactics and operations as in the previous objective, which yielded 384 hits (Result 5). The number of hits returned was assessed to be large enough to continue to Objective 6, redefining the diabetes facet to represent the concept of “Diabetes Complications” as opposed to “Diabetes Mellitus.” Tactics, operations, and assessment were similar to those used to complete Objective 1, except that “Diabetes Complications” was selected as the MeSH Subject Heading of interest. For Objective 7, the new diabetes complications facet was combined with the other two facets (Result 5), yielding 6 hits. Assessment involved scanning the titles of citations, which were again found to not be relevant to the answer being sought. The decision was made to redefine the diabetes facet again (Figure 4, Objective 8), keeping the same structured term but using the tactic of narrowing its reach by restricting it to the main focus of the article. Operationally, this was accomplished through a series of moves that included selecting the “Focus” option after choosing “Diabetes Complications” from the MeSH Subject Headings list. Objective 9 was to combine this newly redefined facet with the facet of gastrointestinal bleeding. This was performed in the manner previously described, yielding 10 hits, the titles of which were scanned and judged to be appropriate (Assessment 9). The librarian concluded the search at that point, having successfully met the goal of answering the question.

As this example has shown, the model's explanatory vocabulary is sufficient to categorize and describe the searcher's behavior within a single search session such that key types of search knowledge are capable of being distinguished. In addition, certain aspects of the model only implicitly demonstrated by the validation deserve comment. Although the grand strategy in the example was to search Ovid for the answer, it should be noted that a number of other

possible grand strategies could have yielded the desired information, be it using a different reference database, an online textbook, a generic search engine, or a combination of these, to name a few. Within the model, these different options would be represented as possibilities within the grand strategy problem space. Furthermore, the grand strategy for this search was simple, consisting of a single phase. An example of a more complex grand strategy involving multiple phases could have been to search different sources (perhaps textbooks in endocrinology, gastroenterology, and nephrology) to identify and corroborate the answer.

A key feature of the model is that the choices at each level are constrained by the level(s) above it in the strategic hierarchy. For instance, less sophisticated resources such as some online textbooks may not have supported the building blocks strategy that was pursued; this would have added a constraint such that a different initial strategy would need to be chosen. Similarly, the fact that the types of tactics observed in the example were relatively few is not surprising given that according to our model, tactics are also constrained by each of the higher levels of the hierarchy.

Also of note is that sometimes there was more than one operational possibility that the librarian could have chosen to execute a particular tactic. For example, combining results sets could have been accomplished by entering the results sets of interest separated by the AND connector (e.g., “#2 AND #3”), or by performing a series of mouse clicks through Ovid’s “Combine Searches” tool. In this example, the result would have been the same regardless of which operation was chosen. In other instances, if the various operational possibilities have different mechanisms of action and/or precision, the outcome may not be identical.

Feedback from the assessment level in this particular search prompted adaptive decisions exclusively at the level of strategy such that over the course of the search several strategic objectives were reformulated. In this search, the tactics employed either successfully completed the higher strategic objective or the results led to the formulation of an entirely different objective, i.e., backtracking at the level of strategy. However, a different search may well have led to movement at other levels in the model hierarchy. One can imagine that sometimes it could be necessary to try a second tactic to move the current objective forward because the first tactic did not yield an appropriate result, i.e. backtracking at the level of tactics. Also, should the search have continued to produce unsatisfactory results, the assessment process could have led at some point to the decision of abandoning the building blocks strategy for another strategy, or perhaps of trying an entirely different grand strategy.

## Discussion

The multi-level model of context-initiated information seeking described in this paper serves as an idealized representation of search strategy from a problem-solving perspective. The layered nature of the model facilitates the conceptual separation of different types of search knowledge. For example, the knowledge that permits the formulation of grand strategy is inherently different from that which provides the basis for tactical decisions. Our model provides a framework that distinguishes between these distinct levels of knowledge, and also defines the relationships between them. We anticipate that the model will lend structure to future cognitive studies of information seeking that as a result may yield richer characterizations of the search process. Through repeated observation, one can imagine developing a knowledge base consisting of usable grand strategies, strategies, tactics, operations, and assessments.

The model may inform how context-initiated information searches are evaluated in the future. Comparisons between clinician searchers and librarian searchers, for instance, can be described in terms of distinctions present at any of the particular levels described by the model. By the



same token, the model should prove useful in helping to meaningfully characterize differences between searches that fail and those that are successful. It may also be possible to pinpoint and describe where the inefficiencies are within a particular search, which in turn may lead to an improved understanding of what constitutes an “ideal” search.

Previous research has shown that goals and subgoals often shift during the information seeking process as a search evolves and progresses (42). Our model may provide a useful way to study and understand such shifts. An interesting question is whether the process of shifting between various levels of the model hierarchy in response to feedback is influenced by an attempt to minimize computational or cognitive costs. For example, one might consider the possibility that as the searcher progresses further down the model hierarchy, the cost of changing direction at a higher level increases due to the resources invested to that point. If we accept this for the sake of discussion, reformulating a grand strategy to answer a particular question would be more expensive than executing another low-level operation. Assuming that a searcher has adequate knowledge of the possibilities at each level, it would seem reasonable that decision-making regarding what to do next during a search might proceed from the bottom of the hierarchy upward in an attempt to minimize the amount of additional resource expenditure at any given point. Thus, searchers may be more likely to consider making changes to available operations before tactics, tactics before strategy, and so on. Of course, search experts have superior knowledge at multiple levels compared to novices (43); thus, novices will be forced to move up the strategic hierarchy more quickly than a search expert in order to further a search, resulting in more abandoned searches at the level of strategy and grand strategy. Empirical research is needed to ascertain the validity of these hypotheses.

The model is designed as an explanatory device for information seeking. It is articulated in top-down fashion from the grand strategy level down to the operations level. We do not presuppose that humans conduct searches in a strictly top-down fashion, as much of search emerges in the context of interaction. For example, it is unlikely that a searcher begins the process with a fully mapped out strategy that accounts for every contingency. Nor do we intend to suggest that all of the knowledge is in the head of the searcher. External representations (e.g., displays, lists of search terms) and affordances (e.g., selection choices, dialogue boxes, and specific tools) play an important role in guiding the search process.

Additionally, our model should not be taken as an attempt at representing the entire information seeking process. Rather, the focus of our modeling efforts was to elucidate the types of knowledge necessary to build and execute a search strategy in response to a specific information need initiating from a patient care context within the clinical domain. We did not, for instance, try to model the information need itself. One of the model's assumptions is that the information need is in fact unambiguous and well-articulated, having emerged from within the context of a particular patient care task. In cases when this assumption cannot be met and the information need perhaps requires clarification to resolve ambiguity, an additional explanatory mechanism that reflects such a state may need to be invoked. Also, while the model is predicated on the assumption that the context from which an information need arises plays an important role in determining the appropriate direction of a search (reflected in the choice of grand strategy, for example), context is not explicitly modeled at this point. Future work to extend the current model's explanatory power could include formal representations of the information need as well as its context, while elaborating on their relationships to existing model constructs.

A main objective for developing our model of information seeking was to provide a theoretical construct that could serve as the pragmatic basis for the design of an automated information retrieval system that employs an intelligent search agent to find information in the clinical domain. In general terms, such a system might be expected to: (1) accept some sort of clinician query as input, (2) analyze the query in preparation for the search, (3) perform the search itself,

(4) package the results of the search, and (5) deliver the packaged answer to the clinician. The present model specifically expands on the third step. Table 2 shows representative functionality enabled by various aspects of the model for the system designer's consideration. Note that the model is adaptable to the particular needs of the system, such that certain elements of the model can be removed if the associated functionality will not be needed. For example, if the envisioned system will be used to autonomously generate a search strategy within a particular resource (as opposed to across multiple ones), then the grand strategy layer can be removed. Similarly, if the system is meant to suggest and execute appropriate tactics from within a resource while assuming the framework of user-driven strategy and assessment, then the relevant levels to focus on would be limited to tactics and operations.

Our multi-level model thus provides the system designer with a framework to separate the system's search features in terms of discrete search knowledge constructs. Such separation facilitates the development of appropriate architectural components to support the search knowledge needed in the system. Instead of having a single search module, the designer may alternatively implement separate modules corresponding to grand strategy, strategy, tactics, etc. A potential benefit of modularizing the architecture around the model's knowledge layers is that the developed components will likely be easier to maintain, reuse, and interchange with other components.

Within the field of software engineering, recurring solutions that have been identified and described for use in solving common software design problems are known as design patterns (44). In a similar way, one can think of our model as providing something akin to intelligent search agent "design patterns" that assist in elucidating the important architectural components which support the implementation of systems intending to use such agents. Now that we have successfully laid the relevant theoretical groundwork, our next steps in the development of a system architecture based on the model will involve the formalization of computable data structures and methods at each level. We anticipate that the system architecture, when complete, will involve several modules patterned along the levels of the model's strategic hierarchy, and will include a control module to drive the direction of an ongoing search based on our modeling of assessment tasks.

## Conclusion

We have described a multi-level model of context-initiated information seeking that provides the basis for leveraging the knowledge of human search experts. This work is significant with respect to future research, in that it (1) facilitates the development of intelligent search agents capable of generating adaptive strategies to search for information, (2) informs the creation of search interfaces that support more satisfying user experiences, and (3) deepens a program of empirical cognitive research in information seeking. The ultimate objective of our research is to increase the percentage of clinician information needs that are satisfactorily addressed. By developing a system that bridges the gap between clinician information needs and the search expertise of healthcare librarians, we aim to bring providers closer to the knowledge that can satisfy their information needs at the point of care.

## Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

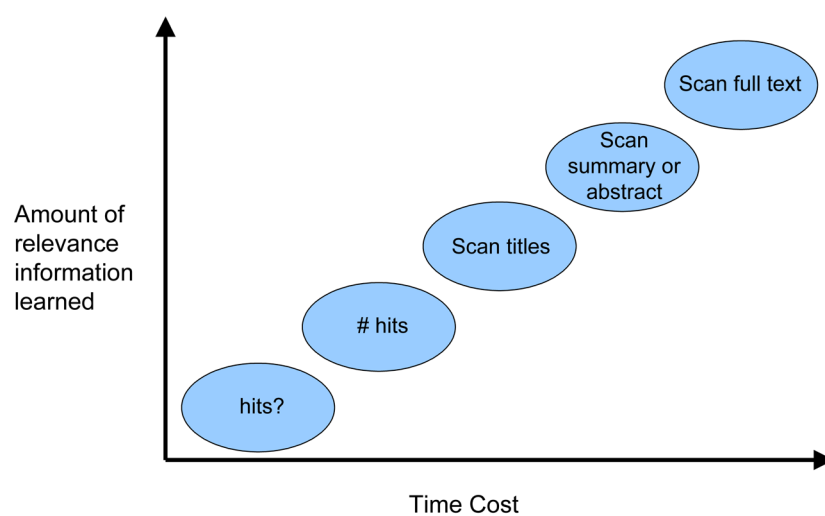
## Acknowledgements

This research is supported by NIH grant 1R01LM008799-01A1 and NLM training grant T15LM007079.

## Reference List

1. Covell DG, Uman GC, Manning PR. Information needs in office practice: are they being met? *Ann Intern Med* Oct;1985 103(4):596–9. [PubMed: 4037559]
2. Ely JW, Osheroff JA, Ebell MH, Bergus GR, Levy BT, Chambliss ML, et al. Analysis of questions asked by family doctors regarding patient care. *BMJ* 1999;319(7206):358–61. [PubMed: 10435959]
3. D'Alessandro DM, Kreiter CD, Peterson MW. An evaluation of information-seeking behaviors of general pediatricians. *Pediatrics* 2004;113(1):64–9. [PubMed: 14702450]
4. Green ML, Ciampi MA, Ellis PJ. Residents' medical information needs in clinic: Are they being met? *Am J Med* 2000;109(3):218–23. [PubMed: 10974185]
5. Ely JW, Osheroff JA, Chambliss ML, Ebell MH, Rosenbaum ME. Answering physicians' clinical questions: obstacles and potential solutions. *J Am Med Inform Assoc* Mar;2005 12(2):217–24. [PubMed: 15561792]
6. Bennett NL, Casebeer LL, Kristofco RE, Strasser SM. Physicians' Internet information-seeking behaviors. *J Contin Educ Health Prof* 2004;24(1):31–8. [PubMed: 15069910]
7. Green ML, Ruff TR. Why do residents fail to answer their clinical questions? A qualitative study of barriers to practicing evidence-based medicine. *Acad Med* 2005;80(2):176–82. [PubMed: 15671325]
8. Haynes RB, McKibbin KA, Walker CJ, Ryan N, Fitzgerald D, Ramsden MF. Online access to MEDLINE in clinical settings. A study of use and usefulness. *Ann Intern Med* Jan 1;1990 112(1):78–84. [PubMed: 2403476]
9. McKibbin KA, Haynes RB, Walker Dilks CJ, Ramsden MF, Ryan NC, Baker L, et al. How good are clinical MEDLINE searches? A comparative study of clinical end-user and librarian searches. *Comput Biomed Res* Dec;1990 23(6):583–93. [PubMed: 2276266]
10. McKibbin KA, Fridsma DB. Effectiveness of clinician-selected electronic information resources for answering primary care physicians' information needs. *J Am Med Inform Assoc* Nov 1;2006 13(6):653–9. [PubMed: 16929042]
11. Hsieh-Yee I. Effects of search experience and subject knowledge on the search tactics of novice and experienced searchers. *J Am Soc Inf Sci* 1993;44(3):161–74.
12. Debowski S. Wrong way: go back! An exploration of novice search behaviours while conducting an information search. *Electron Libr* Nov 28;2001 19(6):371–82.
13. Harter, SP. Online information retrieval: concepts, principles, and techniques. Academic Press; Orlando: 1986.
14. Haynes RB, Wilczynski NL. Optimal search strategies for retrieving scientifically strong studies of diagnosis from Medline: analytical survey. *BMJ* May 1;2004 328(7447):1040. [PubMed: 15073027]
15. Haynes RB, McKibbin KA, Wilczynski NL, Walter SD, Werre SR, for the Hedges Team. Optimal search strategies for retrieving scientifically strong studies of treatment from Medline: analytical survey. *BMJ* May 21;2005 330(7501):1179. [PubMed: 15894554]
16. Montori VM, Wilczynski NL, Morgan D, Haynes RB, for the Hedges Team. Optimal search strategies for retrieving systematic reviews from Medline: analytical survey. *BMJ* Jan 8;2005 330(7482):68. [PubMed: 15619601]
17. Wilczynski NL, Haynes RB. Developing optimal search strategies for detecting clinically sound prognostic studies in MEDLINE: an analytic survey. *BMC Med* Jun 9;2004 2:23. [PubMed: 15189561]
18. Chen HC, Dhar V. Cognitive process as a basis for intelligent retrieval-systems design. *Inf Process Manage* 1991;27(5):405–32.
19. Wilson TD. Models in information behaviour research. *J Doc* 1999;55(3):249–70.
20. Hoppe, U.; Schiele, F. Towards task models for embedded information retrieval. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems; ACM Press; Monterey, California, United States. 1992. p. 173-80.
21. Sutcliffe A, Ennis M. Towards a cognitive theory of information retrieval. *Interact Comput* Jun;1998 10(3):321–51.
22. Vakkari P. Task-based information searching. *Annu Rev Inform Sci* 2003;v37:p413–64.2003

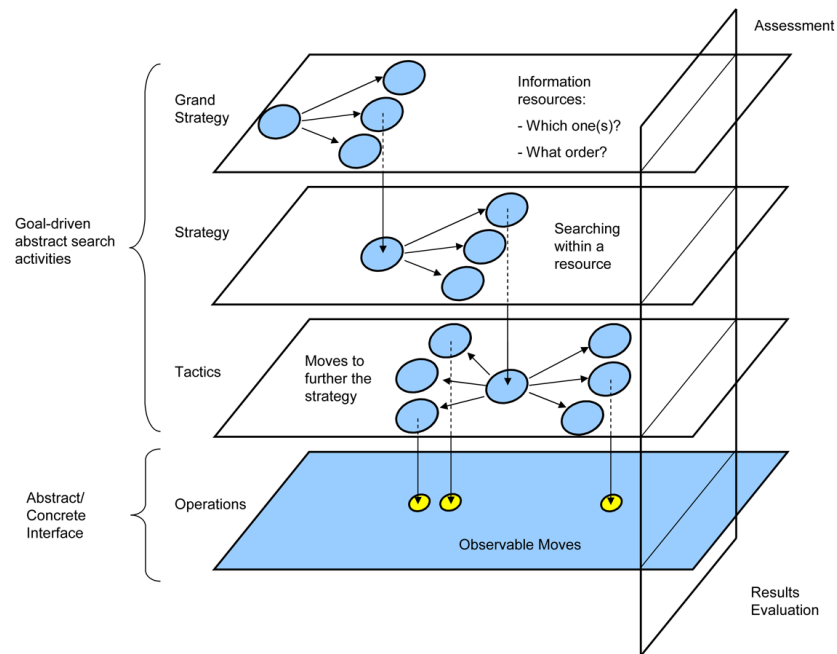
23. Kuhlthau CC. Inside the search process - information seeking from the users perspective. *J Am Soc Inf Sci* 1991;42(5):361–71.
24. Wilson, TD.; Walsh, C. Information behaviour: an interdisciplinary perspective. University of Sheffield Department of Information Studies; Sheffield: 1996.
25. Leckie GJ, Pettigrew KE, Sylvain C. Modeling the information seeking of professionals: a general model derived from research on engineers, health care professionals, and lawyers. *Libr Quart* 1996;66(2):161–93.
26. Vakkari P. A theory of the task-based information retrieval process: a summary and generalisation of a longitudinal study. *J Doc* Jan 1;2001 57:44–60.
27. Ingwersen P. Cognitive perspectives of information retrieval interaction: elements of a cognitive IR theory. *J Doc* Mar;1996 (1):3–50.
28. Ellis D. A behavioural approach to information retrieval system design. *J Doc* Sep;1989 45(3):171–212.
29. Choo CW, Detlor B, Turnbull D. Information seeking on the Web: an integrated model of browsing and searching. *First Monday* Feb;2000 5(2)
30. Marchionini, GM. Information seeking in electronic environments. Cambridge University Press; Cambridge, UK: 1995.
31. Bates MJ. Information search tactics. *J Am Soc Inf Sci* 1979;30(4):205–14.
32. Bhavnani, SK.; Bates, MJ. Separating the knowledge layers: cognitive analysis of search knowledge through hierarchical goal decompositions; Proceedings of the 65th ASIST Annual Meeting; 2002. p. 204–13.
33. Newell, A.; Simon, HA. Human problem solving. Prentice-Hall; Englewood Cliffs, NJ: 1972.
34. Klahr D, Dunbar K. Dual-space search during scientific reasoning. *Cognitive Sci* 1988;12(1):1–48.
35. Klahr D, Simon HA. Studies of scientific discovery: complementary approaches and convergent findings. *Psychol Bull* 1999;125(5):524–43.
36. Stefik M. Planning and meta-planning (MOLGEN: Part 2). *Artif Intell* May;1981 16(2):141–69.
37. Norman, DA.; Draper, SW. L. Erlbaum Associates; Hillsdale, N.J.: 1986. User centered system design : new perspectives on human-computer interaction.
38. Hawkins DT, Wagers R. Online bibliographic search strategy development. *Online* 1982;6(3):12–9.
39. Fidel R. Moves in Online Searching. *Online Rev* 1985;9(1):61–74.
40. Harter SP, Peters AR. Heuristics for online information-retrieval - a typology and preliminary listing. *Online Rev* 1985;9(5):407–24.
41. The American Heritage Dictionary of the English Language. 4th ed.. Houghton Mifflin Company; Boston: 2000. Operation.
42. Xie H. Shifts of interactive intentions and information-seeking strategies in interactive information retrieval. *J Am Soc Inf Sci* 2000;51(9):841–57.
43. Sutcliffe AG, Ennis M, Watkinson SJ. Empirical studies of end-user information searching. *J Am Soc Inf Sci* 2000;51(13):1211–31.
44. Gamma, E.; Helm, R.; Johnson, R.; Vlissides, J. Design patterns: elements of reusable object-oriented software. Addison-Wesley Longman Publishing Co., Inc.; 1995.



**Figure 1.**

The amount of relevance information learned is proportional to the time it costs to perform the assessment. The nature of the relationship depicted is not necessarily linear.





**Figure 2.**

A multi-level model of contextualized information seeking. Each level from Grand Strategy to Operations represents a distinct problem space within the model's strategic hierarchy. The goals of a given level are accomplished by the level beneath it, until the lowest level of the hierarchy is reached (Operations). Grand Strategy, Strategy, and Tactics are levels that are abstract in nature (white planes), while the Operations level represents a mapping of intention onto concrete search actions (shaded plane). Orthogonal to the other levels of the strategic hierarchy, the Assessment level provides evaluative feedback of search results and serves as a control structure.

Grand Strategy: search Ovid to retrieve answer

Strategy: building blocks (separately define each facet, then combine facets)

Objective 1: define facet for diabetes

Tactic: add structured terms to broadly define diabetes facet

Operation: add MeSH terms for diabetes in Ovid

Move: enter “diabetes” to identify MeSH Subject Heading

Move: choose “Diabetes Mellitus” from list of MeSH Subject Headings

Move: select “Explode”

Move: choose “All Subheadings”

Move: execute. Result 1: 181558 hits.

Assessment: presence of hits→yes

- operational, tactical, and strategic objective success
- go to next objective

Objective 2: define facet for gastrointestinal bleeding

Objective 3: define facet for renal failure

Objective 4: combine all facets to produce answer set

Tactic: combine facets using Boolean AND

Operation: intersect Result 1, Result 2, and Result 3 in Ovid

Move: enter “#1 AND #2 AND #3”

Move: execute. Result 4: 14 hits

Assessment: scan titles→not relevant

- operational and tactical success, strategic objective failure
- try different objective

**Figure 3.**

Initial stages of librarian search. Note that Objectives 2 and 3 are left unexpanded for space considerations, but unfolded in similar fashion to Objective 1. The search continues on with Objective 5 (not shown).

Objective 8: redefine diabetes facet

Tactic: restrict structured terms in diabetes facet to the major heading

Operation: restrict current MeSH heading for diabetes complications facet to main focus in Ovid

Move: enter “diabetes complications” to identify MeSH Subject Heading

Move: choose “Diabetes Complications” from MeSH Subject Headings

Move: select “Focus”

Move: choose “All Subheadings”

Move: execute. Result 8: 9384 hits.

Assessment: presence of hits→yes

- operational, tactical, and strategic objective success
- go to next objective

Objective 9: combine two facets to produce answer set—GI bleeding and diabetes complications

Tactic: combine two facets using Boolean AND

Operation: intersect Result 2 (GI bleeding) and Result 8 (diabetes complications) in Ovid

Move: enter “#2 and #8”

Move: execute. Result 9: 10 hits.

Assessment: scan titles→results relevant

- operational, tactical, strategic, and grand strategic success
- end search

**Figure 4.**

End of librarian search. The final stages of the search introduced in Figure 3 are shown.

**Table 1**  
Example tactics within a resource supporting Boolean search

---

**Limit tactics**

- Add limit
- Remove limit
- Broaden existing limit
- Narrow existing limit

**Query tactics**

- Join two queries using AND operator
- Join two queries using OR operator
- Join two queries using NOT operator
- Add facet using AND operator
- Add facet using NOT operator
- Remove facet

**Facet tactics**

- Add term to facet using OR operator
- Remove term from facet

**Term tactics**

- Select more general term (broaden term)
- Select more specific term (narrow term)
- Increase truncation of a term
- Decrease truncation of a term
- Replace existing term with synonym
- Restrict field that term appears in

---

**Table 2**

Knowledge layers and associated examples of intelligent search agent functionality supported by the model

<b>Model Level</b>	<b>Knowledge Represented</b>	<b>Illustrative Intelligent Search Agent Functionality</b>
Grand Strategy	Information resources that can be accessed to meet a particular information need. Search approaches for handling complex information needs requiring multiple information resources.	Choose appropriate information resources that will most likely answer the information requester's question <i>Example: risk factors of thromboembolism in atrial fibrillation→select UpToDate</i> Provide an overarching plan of what resources to search, in what order, and for what purpose in response to the question being asked <i>Example: What oral antifungal agents for the treatment of onychomycosis do not interact with benzodiazepines?</i> <i>Phase 1: treatment of onychomycosis→select Harrison's</i> <i>Phase 2: benzodiazepine interactions with Phase 1 results→ select Micromedex</i>
Strategy	Strategies for how to systematically (and often iteratively) search a resource to meet an information need.	Choose and carry out different strategies for searching a particular resource depending on the goal of the search (e.g., high precision, high recall) and the selected resource. <i>Example: high precision search in PubMed→citation pearl growing strategy</i>
Tactics	Localized search maneuvers to meet current strategic objectives, by either increasing recall or precision at the level of search terms, facets, or queries.	Select and execute appropriate, finely grained search actions to improve the current search. <i>Example: current results set too large→restrict existing search terms to title</i>
Operations	Resource-specific methods/tools that actualize supported search tactics.	Choose or suggest resource-specific interface(s) that can be utilized by the user/agent to execute the desired tactic. <i>Example: add more specific indexed terms to current facet in Ovid→Use Explode function in MeSH tree tool</i>
Assessment	Evaluation of results at each cycle in the search process.  What to try next in the case of local failure to meet success criteria based on results evaluation	Decide how to evaluate current search results in the manner most appropriate to the progress and stage of the search. <i>Example: initial stage of building blocks search strategy→ assess presence/absence of search hits</i> Decide how to backtrack/adjust search, i.e., whether to pursue a new operation, tactic, strategy, or grand strategy when evaluation of search results indicates failure to find relevant results. <i>Example: tactic to increase specificity of search terms fails to produce relevant hits→try tactic to add synonyms</i>