In-Order:
Enhancing Google via Stigmergic Query Refinement

Garrett Camp and Mihaela Ulieru*

Emergent Information Systems Laboratory
The University of Calgary
http://eis.enel.ucalgary.ca/

*Canada Research Chair
Industrial Research Chair
Adaptive Risk Management Laboratory
The University of New Brunswick
http://www.cs.unb.ca/~ulieru/

Abstract

InOrder is a query refinement tool that works on top of Goolge and helps individual users to collaboratively participate in best Web query formulations. The incremental refinement works via an indirect communication process facilitated by a visual interface which adapts significantly to reflect user contributions. The interface visually guides users in an implicit manner via a "what you click is what you mean" approach which assists semantic visualization of past interactions by highlighting relevant search terms in brighter colors. Users simply click on anything that seems relevant while being 'attracted' to the terms already clicked by other users. This elicits conceptual refinements and reduces the rating effort required by many collaborative systems. InOrder functions as a "visual search WIKI", which represents search *intent* rather than search of formal articles. Since it takes less effort to click than type, the system increases search usability by reducing the interactive effort required to discover documents when search goals are unclear.

Keywords: Query Formulation, Google, Interactive Search Guidance, Relevance Feedback, Stigmergic Collaboration, Visual Search WIKI, Semantic Self-Organization, Information Scent, Memetics,

1. Introduction

1.1. Search Relevance in Today's Information Flooding

The concept of search is central to the human condition. Individuals must continually acquire new information and resources to ensure their survival, and the process of search helps one meet such goals. Within early societies search focused on the location of food or shelter, and increasing division-of-labor and technological progress have now progressively focused search activity towards the acquisition of textual knowledge. First within libraries and other physical collections of knowledge, and now within electronic networks “search" has become a primary mechanism of information retrieval. The rate of growth of online data has been exponential, with over 100 billion documents now estimated to be accessible through the web [24].

This vast collection of electronic knowledge is leading us towards a greater understanding of human behavior, and removing physical constraints to communication and learning. But with such growth’ so have difficulties involving organization and access of this knowledge. Plenty of data is available, and the fundamental issue has now become the acquisition of relevant information without excessive human effort. On the World-Wide-Web search-engines have emerged as the primary mechanism for information retrieval, and are now able to query billions of documents within a fraction of a second. Yet within such large document collections users are often not fully aware of available contents or useful strategies for retrieval. Users experience information overload when presented with thousands of search results, and observations from natural knowledge sharing systems are not typically employed to assist the search process. Document-centric techniques also limit resource awareness because typically only a few pages are examined before requirements change or a task is abandoned. Studies show that users usually examine less than 10 pages before stopping or modifying search, and often before satisfactory results have been found [34].

Query-based web search also presents significant challenges when users are unsure of the best keywords to use. Formation of accurate queries relies on prior domain knowledge, and without use of appropriate terminology search results often bear little relevance to the task at hand. Guidance is not offered to users when information needs are vague or unclear, and elicitation of such requirements is typically not integrated into the process of web search. Combined with the navigational challenges caused by complex hypertext networks, this means use of standard search-engines is often only practical when one knows exactly what to search for, and has a clear understanding of how to go about it. Such issues suggest that new search tools that elicit search requirements and assist hypertext document navigation have potential to improve retrieval efficiency as well as the search experience. Bill Gates even notes that natural language and contextual semantic approaches are the next step in search engine development [1], which will reduce the cognitive effort required of searchers and improve information access.

1.2. System Goals

This work investigates an interactive search guidance system that uses such semantic and contextual techniques to support the web retrieval process of relevant information. InOrder aims to provide users with an improved search experience when requirements are unclear by achieving the following milestones:

1. Improve search usability for people with vague or complex search goals.
2. Guide the search process, and reduce time required to accurately specify search requirements.

3. Provide a consistent interface to refine and expand queries, which will improve search result relevance.
4. Recommend useful terminology to increase awareness within a domain as search occurs. This leads to longer and more specifc queries which produce more relevant Google results.
5. Help users visualize domain knowledge, so relevant terms may be recognized and utilized more rapidly. Visual states should reflect prior feedback from like-minded searchers within the context (ie. search group), so searchers may effectively ‘see’ the insights of others.

1.3. Framework for Evaluation of Information Relevance

The first principle to consider when designing an information retrieval (IR) interface is that of relevance. The notion of relevance is the central criterion by which information retrieval systems are judged, and for the past 30 years there has been no practical substitute for the concept of “degree of relevance" when evaluating IR tools. Definitions of relevance generally relate to the importance or significance of information delivered by a retrieval tool. Information relevance judgments are subjective measures describing how well a document satisfies a user's information need, where useful information is facts and data organized and expressed in a coherent and meaningful form. Definitions of relevance include: “Pertinence to the matter at hand" and “Applicable or useful information".

Within the context of search, the best definition of relevance is likely “the capability of an information tool to retrieve data appropriate to the users need" given by Greisdorf in his excellent overview of interdisciplinary research into relevance [18].

Information Retrieval Theory also includes the philosophical branches of epistemology (what can be known) and ontology (what exists) as well as to the concepts of language, context and feedback. Work by Croft examines such concepts, discussing how language models can be used to represent context and support context-based IR technique [7]. Croft states that the goal of personalization [29] is to improve the effectiveness of information retrieval systems, and concludes that consideration of context is vital when eliciting relevance feedback during query disambiguation. Croft also notes that, for the most part, relevance feedback has not been successfully integrated into web search systems because it is difficult to get users to supply ratings.
Users are typically not motivated to provide relevance feedback unless immediate personal benefit occurs and the interface adapts significantly to reflect user contributions.

Related work by Mizzaro entitled “How many relevances in information retrieval" [27] finds that there are many aspects to relevance which may be formally represented within a multi-dimensional model. Mizzaro claims inconsistent concepts are frequently used to discuss the notion of relevance, mixing terminology such as utility, usefulness, pertinence, interest, topicality, association, authority and popularity. Because such judgments are subjective and dependent upon specific tasks and goals, Mizzaro offers a more objective framework for modeling relevance using the following four dimensions:

A. **Information Resources**, consisting of three entities: a. Document - the physical entity which the user will obtain; b. Surrogate - a representation of the document (ie. metadata) and c. Information - the non-physical entity a user receives after examining the document. Surrogates commonly examined by users are usually the textual snippets displayed on search results pages - information that helps users to evaluate relevance before actually visiting a page.

B. **Representation of the Users' Problem**, dealing with entities which represent user need: a. Real Information Need (RIN) - the actual information needed by a user; b. Perceived Information Need (PIN) - a user's mental model of their requirements (operation: Perception); c. Request - the users expression of their PIN, typically in natural language based on known terms (operation: Expression) and d. Query - a representation of the user request, usually as a Boolean query (operation: Formalization). These take into account user interaction with a search system, realizing that progression from A to D creates inaccuracy since each is an imperfect model of the prior representation.

C. **Time**, which considers that the context or situation determines if something is presently relevant. Mizzaro states that as users dynamically interact with an IR system a set of time points is created representing each interaction, starting from when the RIN initially arises (a search task is conceived) to satisfaction or abandonment of the search task.

D. **Components**, into which the first dimension may be decomposed. These include a. Topic, that is the subject area which is interesting to the user; b. Task, that is the activity which motivates the retrieval process and c. Context, everything not pertaining to the topic or task which influences how search is performed.

Mizzaro also notes that the three operations that involve representation of the user problem frequently limit retrieval capabilities. The Perception operation is limited when a user does not know what to search for. Based upon the ASK paradigm (Anomalous States of Knowledge) this occurs when a user does not know what they want to know. The Expression operation is hindered by the vocabulary problem within unfamiliar domains, resulting in use of vague or ambiguous terms when a user does not possess terminology which will accurately expresses their needs. Finally the Formalization operation is limited because users often do not understand Boolean logic or how search systems function, and thus do not accurately formalize questions or statements into appropriate system language (a suitable query). Within this four-dimensional model Mizzaro suggests several methods of improving an information retrieval system by assisting the operations of Perception, Expression and Formalization. One suggestion is to equip an IR system with a set of stereotypical tasks from users, from which users may select the most appropriate one. He also notes that designing an IR system that combines direct manipulation with intuitive visualization and relevance feedback could improve information retrieval systems. The next section will examine how InOrder uses these suggestions within Mizzarro's four-dimensional model to assist relevance refinement and search requirements specification, presenting associated semantic knowledge that helps users realize, construct and adjust their queries.

1.4. In-Order's Stigmergic User-Centric Interface: a Use Case

InOrder, Fig. 1, enables communities of like-minded users to collaborate via a visual interface which enables a natural 'what you click is what you mean' indirect communication (stigmergic) process [5], [10]. The layout of the InOrder interface (Fig. 1) reveals 17 user-centric interface features, described in the sequel on a "Carribean Travel" example: [1]

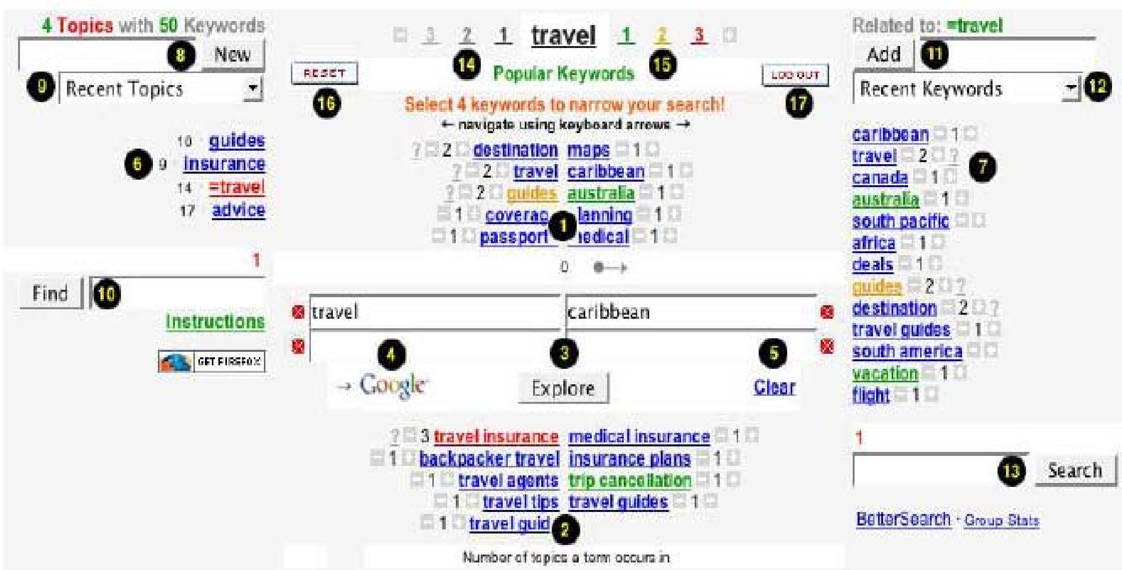

Fig. 1: In-Order's stigmergic visual interface

1. Suggested Words: any of these words may be added to the search query. Terms with higher counts and brighter colors have been used more frequently by searchers in the past. In Fig. 1 the terms 'travel' and 'Caribbean' have been selected, and added to the list of recently used keywords.

2. Suggested Word-pairs: same as at point 1 - any of these may be selected to be added to the query.

[1] This description matches the "instructions" link appearing at the bottom left-hand side of the interface.

3. Explore Button: Once some terms are selected, if more suggestions are needed the user may click the explore button to further narrow a search. In this example, clicking "Explore" would suggest keywords related to traveling in the Caribbean.

4. Google Button: Once 4 keywords (or less) have been selected, top Google results are shown for these terms. Keywords may be deleted using the red X buttons.

5. Clear Button: This removes all currently selected keywords, for a fresh restart.

6. List of Topics: These are topics within a group that one may browse by selecting the most appropriate one.

7. Selected keywords: These are keywords that have been recently used within the selected topic (in this case the default "travel" topic). Essentially this is a log of what you or others have clicked in the past because it seemed useful.

8. New Topics: If no topic of interests is displayed, one can be created here by clicking "New".

9. Topic Sorter: One may sort the list of topics by recent-cy (default), popularity, or alphabetically.

10. Finding Topics: By clicking on a question mark next to a recently used keyword (near 7), one will be shown all the topics in the group which contain similar keywords. This helps with finding related topics.

11. Adding Keywords: If one knows of a keyword that is useful but not present, one can feel free to add it.

12. Keyword Sorter: Sets of keywords in a topic may be sorted by recent-cy (default), sequence or alphabetically.

13. Searching for Keywords: At any time one may search for other keywords within the group that are related to what one is looking for. Simply typing a word in here, or clicking on a question mark next to any word/word-pair in the center section of the screen (1 or 2) this will show all similar matching keywords.

14. Useless Terms: When noticing a term that isn’t helpful by clicking the ‘minus’ button next to it one can contribute to ‘vote it down’ – so it become less and less noticeable within the emerging group of like-minded searchers.

15. Interesting Terms: These are helpful terms which have been voted up 1, 2 or 3 times. Clicking the plus button next to any term which was interesting or helpful one can vote them ‘up’ within the search group so these will then be noticed more often.

16. Reset Button: Starts a fresh search by returning one to the default group view for the particular topic, an overview of all popular keywords within the group.

17. Logout Button: Takes one back to InOrder.org, where one can browse/create other groups.
As shown in Figure 2, the interface encourages visual markup of semantic content and represents the relevance of potential refinements using a discrete set of colored levels. Similar to the function of scent when foraging for food, this “information scent" [4] provides perceptual semantic cues that draw user attention in proportion to utility. This robust, visual approach helps users recognize the most appropriate keywords within a given context and assists decision-making during search.

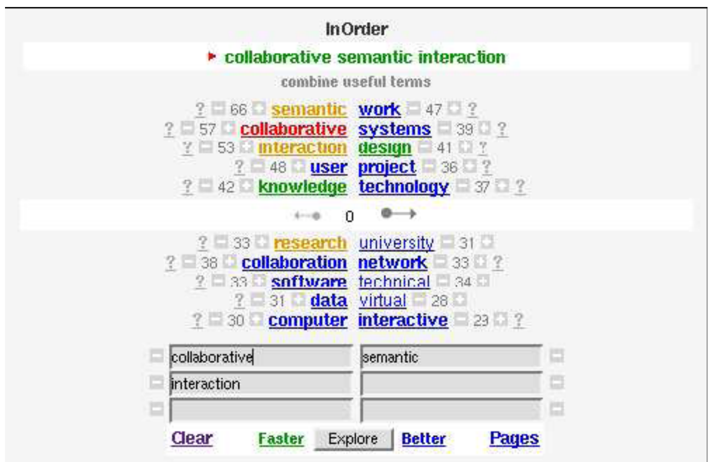

Fig. 2: Visual semantic blackboard

All ‘group click activity’ is collected into a WIKI-like hypertext interface that allows users to explore associated semantic knowledge prior to document evaluation and reduce effort required to execute unclear search tasks. No concrete relationships between suggested keywords are given (only their relative conceptual prevalence within a context) yet such conceptual information is still very useful for supporting semantic exploration and query reformulation (Fig. 3). In situations where a user does not possess sufficient domain knowledge to manually perform effective query refinement, InOrder enables users to quickly survey relevant terminology and offers search decision support. Through this process of keyword discovery and guided query reformulation InOrder helps users utilize PageRank [2] more efficiently.

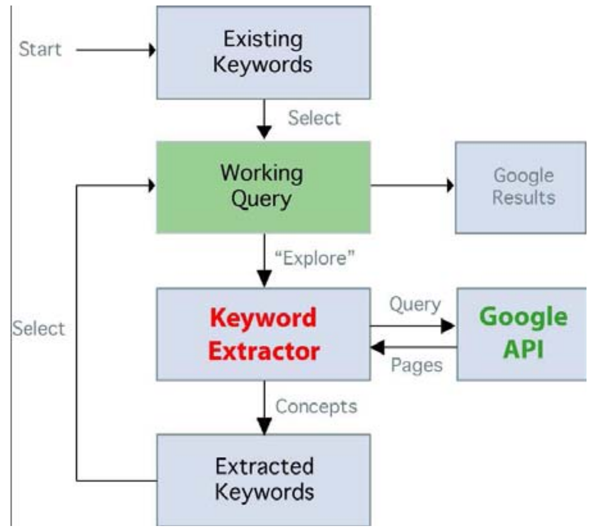

Fig 3: Interaction during query refinement

1.5. How In-Order Improves Search Relevance

Using the intuitive 'what you click is what you mean' search interface previously described InOrder seeks to increase the quality and quantity of relevance feedback during search. To do so semantic information is elicited in an interactive fashion, using incremental guidance to help users focus on relevant web resources. Semantic relevance feedback is provided automatically as interaction occurs, creating collections of relevant terminology which subsequent users may examine to assist decision making during search. Figure 2 shows how InOrder's *relevance feedback approach* compares with Mizzaro's four dimensional framework for relevance modeling presented in Section 2.

A. **Information resources**: Within InOrder resources consist of pages indexed by Google, and keywords which frequently appear within fetched pages. The domain-centric search groups contain terminology commonly used by authors within a given domain - common contextual information that helps users formulate better queries.

B. **Representation of the problem**: InOrder assists the three operations of Perception, Expression and Formalization within the second dimension to assist relevance optimization. The Perception operation is assisted by presenting summaries of associated terminology which illustrate domain concepts. Such observation of prior use of search terms assists the process of "sense making" during search. The Expression operation is facilitated through suggestion of query refinements. Users may simply click on any keyword they recognize to be relevant to the task at hand, an activity that is often easier than pondering and manually reformulating queries. Finally, the Formalization operation is aided by enabling easy expansion and reformulation of a query. This transpires in users incrementally constructing more accurate representations of search requirements based on terminology others found useful in the past.

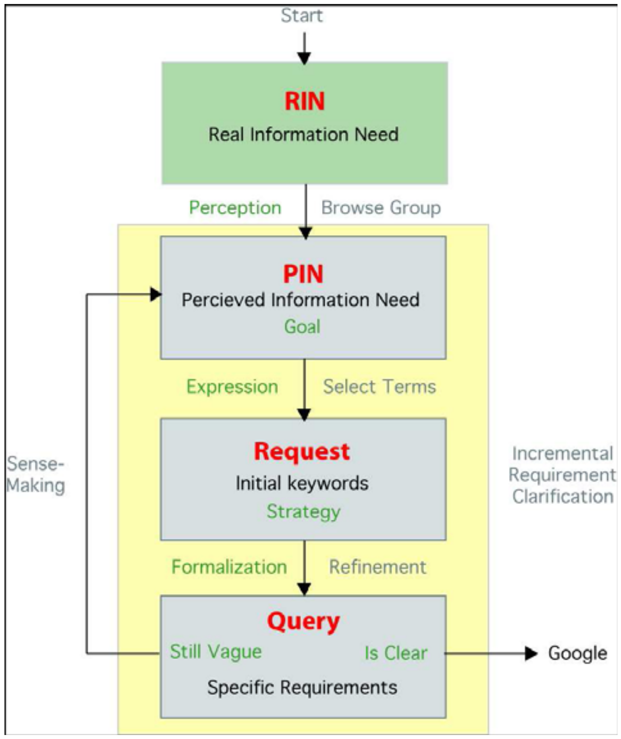

Fig. 4 The Information seeking process

C. **Time**: Each click act is a new time point within InOrder, a semantic refinement explicitly chosen within the current context. These semantic selections reflect better representations of users needs since they are instructed to “select terms to narrow their search". These keyword selections are sequentially logged and presented by default in reverse temporal order to encourage exploration of timely domain associations.

D. **Components**: Within InOrder a Topic is represented by the search group (group-name. inorder.org) containing keywords extracted from relevant web resources.

A ‘Task’ is represented by the topic which displays sets of refinements representing a particular sub-domain. Finally, ‘Context’ is represented by all the surrounding terminology within a search group, including other topics, popular keywords and recent semantic interactions. Such context is established not only via each ‘personal refinement’ activity, but also via an implicit interaction with other searchers who faced similar search tasks.

Using this 4-dimensional model as a basis for InOrders interaction design helps users retrieve more useful search results. The stigmergic interaction (by indirect communication) within a group guides users from *perceived* to *real* information needs and reduces the time required to locate relevant resources. This stigmergic ‘what-you-click-is-what-you-mean’ interaction model also facilitates collaborative query reformulation among groups of people who are interested in a particular search domain.

2. InOrder: Advancing the State-of-the-Art in Query Refinement

Within the state-of-the art work in query refinement systems, InOrder is positioned at the confluence of research in Search Usability, Data Mining and the Semantic Web (Fig. 5)

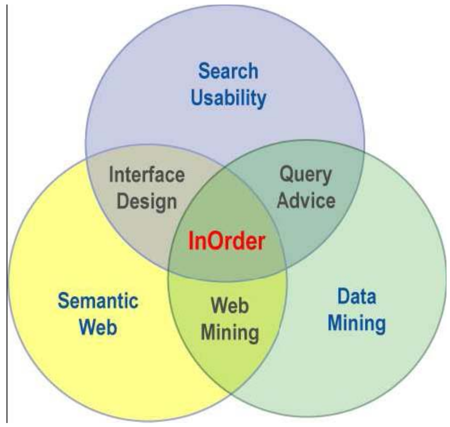

Fig. 5: In-Order within the State-of-the-art in query refinement

2.1. InOrder as a Query Advice System

2.1.1. Search Usability

Semantic theory suggests that language usage within a context reflects its importance, and that lexical co-occurrence is a meaningful measure of contextual utility. Applying this theory to InOrder's interactive web-mining process enables incremental semantic interaction to build *collective domain models* (Fig. 6).

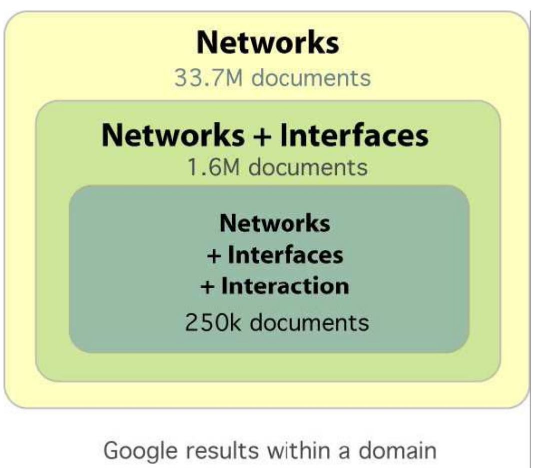

Fig. 6: Domain modeling using contents of search results

Oxford defines a "domain" to be a sphere of activity or knowledge, a field or subject of interest. Domain Modeling refers to the process by which computerized models of such activity or knowledge may be constructed and maintained. Within InOrder a domain is represented by a search group, a collective history of query refinements that characterize the needs of users who chose to interact with a group. InOrder utilizes such domain models to assist keyword discovery, using the visual semantic blackboards (Fig. 2) to present users with concise summaries of interaction history. This assists users when they are unsure of search requirements, as they are able to examine stereotypical selections made by others. As illustrated in Figure 6, InOrder constructs domain models by extracting terminology from pages returned by Google, and using explicit human selection to incrementally associate keywords. This may also be thought of as *search task modeling*, in which search requirements are elicited from anonymous users and aggregated into contextual repositories. These models (search maps) are concise and meaningful because they are formed through explicit user validation of extracted suggestions. These search maps are extracted from search results in response to a user request using the Google API (Fig. 3). Documents are fetched for a concatenation of the group label + working query terms, forming a HTML corpus which may have concepts extracted from it. For example, if a “principles" topic were created within the “interface design" group, top results for “principles interface design" would form the set of HTML from which frequent keywords would be parsed. As different topics are created within a group and keyword extracted relating to various sub-domains, InOrder aggregates the semantic selections from these sub-domains into a default summary that future users will see when they first visit the group.

This domain modeling approach enables gradual construction of knowledge bases as exploration occurs. Query-by-query, potentially useful semantic information is aggregated into groups to form more accurate models of domain search interest. Anyone may create new groups and topics within them, capturing innovative search strategies within a domain. This collaborative, bottom-up approach to domain modeling reinforces relevant terms, and semantic trails may be re-explored later by like-minded searchers, gradually constructing search groups that guide future query refinement activities. From this perspective, InOrder behaves much like a “domain zeitgeist" representing the current state of search intent within a search group. In contrast to Google's Zeitgeist [15] which demonstrates global query frequency trends within a given time frame, InOrder presents users with cumulative interaction statistics within a group. Since specific domains are much less dynamic than large global systems such as Google, this offers users a concise overview of search intent relating to a subject of interest.

2.1.2. Evolutionary Data Revision Mechanism

InOrder's knowledge base is continuously revised as interaction occurs. Such revision occurs with every click act within a group, and any time a user creates a new topic. Terminology which is not present within a topic may also be added to a group using the “Add" input field. This allows a user to add specific terms which they know would be useful within the given context, in the process recommending other searchers utilize it. However, since new terms and topics are automatically integrated into the existing knowledge base an adaptive mechanism is needed to ensure that such new useful information will replace older less useful ideas, and groups remain a reasonable size. To achieve this InOrder uses fixed-size groups with a few simple replacement policies to revise semantic maps as domain interests shift.

Each InOrder group may contain a maximum of 250 topics, which themselves may each contain 250 keywords. Once these limits have been reached old terminology is replaced by new semantic refinements. In a similar fashion, when additional topics are added the smallest topic is replaced, as such topics did not contain as many useful refinements as other topics, so they are the first to be replaced. This oldest-first replacement policy ensures terminology which is replaced is likely no longer useful.

Terminology may also be deleted from specific topics within a group. If a term is voted down to the m2 level it may be deleted using the red X button, which removes it from the topic and decreases the count/position of the keyword in the default list. This allows useless terminology to be pruned from a group as participants see fit, but raises barriers to deletion by requiring 3 clicks to delete. Voting activity guides moderation of keywords present, causing useful terms to be seen more often and irrelevant ones to be replaced by new search strategies. As a group gradually grows to 62,500 terms and overlap occurs between topics, the default summary of popular clicks converges towards a meaningful representation of a domain (Fig. 7). This evolutionary contextual approach also helps avoid positive feedback within a domain model.

Since order of presentation biases selection and thus reinforcement, steps should be taken to ensure selections are not reused simply because they are appear higher within a list of terms. Care must be taken to avoid such inappropriate reinforcement which would cause initial selections to influence the direction of a group to a great extent, and lead to the "Winners take all" behaviour. To avoid such self-reinforcement InOrder uses a conditional approach to keyword reinforcement. Default lists of recommended keywords are calculated using a “keyword frequency across topics" approach, so the number of topics a keyword appears in influence rank rather than straight click counts. If a term is used within a variety of contexts (it was clicked on within several topics) then it's rank will rise, but initial terms within the default list will not rise if they are repeated clicked within the default topic. Only use within a different group context will reinforce a term, which helps InOrder avoid inappropriate positive reinforcement of initial group concepts.

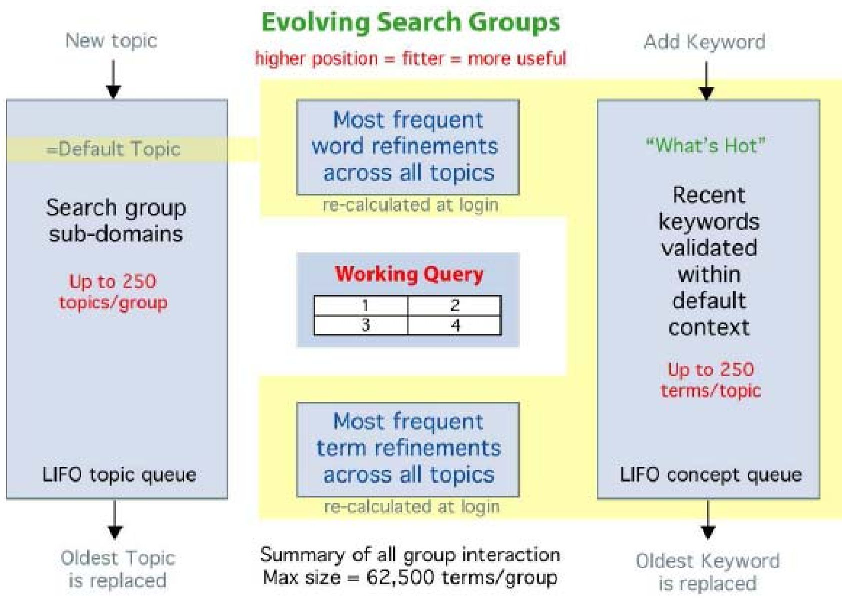

Figure 7: Evolving interaction summaries of up to 62,500 refinements in 250 topics

2.3. In-Order and the Semantic Web

Oxford dictionary defines Semantics to be the branch of linguistics and logic concerned with meaning. Within computer science semantics refers more specifically to the meaning associated with a string of characters, as opposed to syntax that governs their combination. And within linguistics, semantics often concerns either "formal semantics" or "lexical semantics" which delves into the logical and relational aspects of meaning.

From the perspective of communications research, Principa Cybernetica Web (PCW 1996) states that "semantics" concerns the study of how and what a sign, symbol, message or system means to an observer. Within this framework InOrder is a “semantic interface" because it captures search intent by eliciting conceptual feedback. Within InOrder, these symbols are keywords presented to users; either extracted or declarative knowledge which helps users construct or refine search queries. These informal yet meaningful representations of search intent are not concerned with detailed specification of logical or relational information, and concentrate instead on the capture of useful concepts explored during strategic search.

While the current vision of a Semantic Web built using XML is a valuable one, most research currently focuses on knowledge representation issues rather than interface development. And since most of the web currently consists of unstructured HTML content, this suggests that interfaces that facilitate cooperative translation of unstructured content into structured formats could be of considerable value. InOrder focuses on this interface-centric aspect of Semantic Web design, using an interactive web mining approach to extract meaningful (i.e. “semantic") data from unstructured HTML into InOrder's semantic framework. This web-mining approach enables InOrder’s keyword recommendation service, and in turn assists the refinement of broad or complex search tasks. Explicit keyword selection validates keyword utility within a given context, so search groups may converge towards more useful semantic structures as interaction occurs. This approach to semantic collaboration helps users explore relevant semantic knowledge when clarifying queries, and builds structured knowledge bases that may support other Semantic Web tools.

InOrder is a "semantic interface" in a conceptual sense rather than formal or lexical one, as it seeks to model "conceptual semantics" which illustrate the cognitive structure of meaning. Such structure within InOrder consists of associated keywords used during contextual query refinement.

From a software design perspective InOrder may be considered an agent for exploring, manipulating and visualizing semantic information. The interface displays collections of “search tips" within a domain along with relevance indicated by term color. This helps users quickly scan large quantities of textual content for relevant ideas, focusing on the brightest and most useful concepts first. Collaborative voting activity aggregates individual opinions of semantic relevance, and establishes a shared visual representation of semantic interaction within a group. Combined with interaction design that serves both personal and collective goals, such visualization improves the usability of semantic foraging and helps coordinate exploration among like-minded searchers.

Figure 8 illustrates how the stigmergic process of ‘clickthrough’ via which users interact within InOrder progressively organizes and validates semantic knowledge. An initial click on an extracted term (plain font) includes it within a topic. Subsequent clicks, if performed in different topical contexts, increment click counts to increase term prominence within the default view. This promotes terminology which is useful within a variety of subcontexts (different topics) and creates an index of decreasing domain generality. This leads to general or central domain concepts being viewed more frequently since they have greater potential utility.

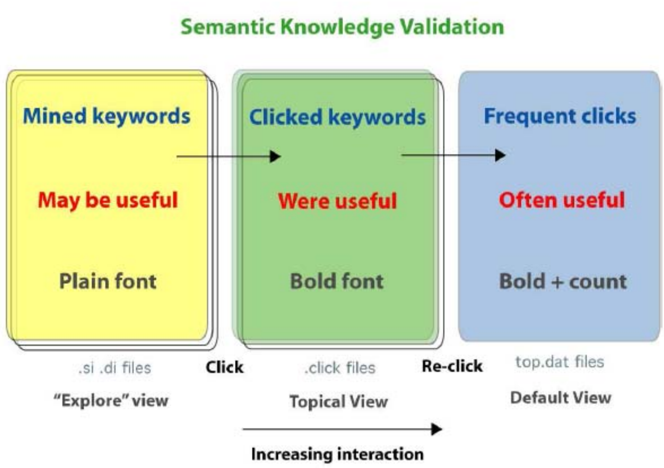

Figure 8: Knowledge validation through repeated interaction

This approach to domain model validation using explicit interaction is supported by InOrders' web-mining process in several ways. First, the web-mining functionality significantly reduces levels of user interaction required to contribute knowledge. Relevant terminology is mined from search results, such that selective validation requires a single click, and specification of search intent becomes a simpler and easier process. Second, clickthrough generates meaningful repositories because such acts are explicit and unambiguous. As mentioned earlier, in constrast to document clickthrough data these acts of semantic clickthrough are inherently meaningful because the purpose and outcome is clear. Click-through within search results may have been prompted by a misleading or outdated title or summary, but this does not occur within InOrder because a user knows clicking a term adds it to their query and the group. Selective acts offer both personal and collective benefits, so the process of query refinement becomes one of conceptual design. Third, InOrder provides motivation for interaction to occur. Unlike some collaborative recommendation systems, users gain a direct and immediate benefit from interaction, developing collective semantic maps as personal search goals are pursued.

Query refinement within InOrder is supported using a shared semantic blackboard. InOrder applies this framework to the exchange of search strategies within a domain. The blackboard functions as a central repository for data consisting of prior semantic selections performed during specification of search requirements. This architecture, Fig. 9, allows users to select appropriate refinements with minimal effort, and integrates semantic endorsements into the collective semantic model. Both extracted and validated keyword lists are organized by topic - that is sub-domains which are created within a search group. Following instructions to “select terms to narrow your search" opinions on term relevance are integrated into a collective log of search behavior. This semantic blackboard informs users of search strategies employed by like-minded searchers, leading to more productive search sessions.

InOrder's blackboard consists of three panes which control the presentation of topics and terms within them:

1. The Topic Pane: Selection of concepts within this pane loads a new set of extracted terminology and matching validated terms, presenting users with associate concepts within a subdomain. These correspond to extracted and validated terminology for the topical label which are presented to the left of this pane.

2. Extracted Terminology: The central blackboard pane presents keywords which frequently occur within top search results. Single keywords (.si) are shown above the working query, and keyword pairs (di-words, in short ‘di’s) shown below. Selection of any word or word pair adds it to the list of recently validated keywords.

3. Validated Terminology: This pane presents keywords which users have clicked on in the past within the central pane, or added manually using the “add" input at the top of the pane. These keywords inform users which concepts have been recently utilized within a topic, and may be useful. Selection of a term from this pane moves it from its prior position to the end to preserve the temporal selective sequence.

Modification of the semantic blackboard (Fig. 9) occurs with every selective act, where clicking within the Topic Pane changes the present topic, and clicking elsewhere validates terminology within the selected topic. Selection of either extracted or validated keywords transfers them from the repository into the first free input available in the working query, replacing the oldest term if four have already been selected. The unique blackboard for each search group may contain up to 250 topics of the 250 most recently utilized keywords, with selection replacing oldest terms first. This causes keywords lists to be continually updated as refinement occurs. Since recent is the default sort mode, this visually corresponds to the selected term being added to the top left-hand side of the pane.

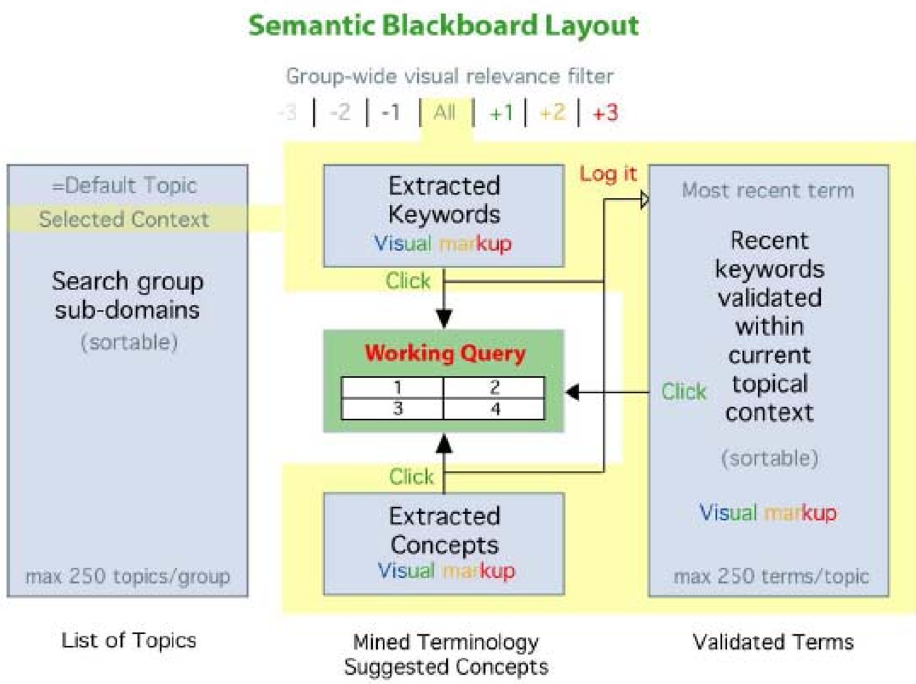

Figure 9: InOrder's semantic blackboard layout

When a topic reaches the limit of 250 terms, the newest terms replace the oldest ones, those that have not been utilized within the last 250 clicks within a topic.
This “use it or lose it" semantic retention policy within the shared semantic space helps ensure only interesting contextual refinements persist. The simplicity of this visual feedback approach makes InOrder relatively simple for users to understand to encourage participation. Following the principle of Stigmergy [22], InOrders' semantic blackboard, Fig. 9) is a completely open knowledge base which may be altered by anyone searching within a group. Direct manipulation of the blackboard structure causes a groups' content to evolve in an incremental fashion. Human agents (searchers) cooperatively evolve the semantic indices which comprise the blackboard, and refine domain model through individual participatory action. Interaction indicating a term was useful (clicking or voting up) increases visual activation and prominence of a term, whereas lack of use or negative feedback (voting a term down) decrease future likelihood of use. This helps ensure promising search strategies get noticed quickly, whereas those poor are not explored due to lower perceptual priority. Since all participants have read/write access,

interaction is designed such that a significant amount of intentional effort is required to alter blackboard composition. The cognitive effects of such relevance visualization are increased neural activation for terminology which previous searchers deemed useful, and this stimulates recognition and sense-making according to participant consensus.

2.3. InOrder and Interface Design

InOrder is a collaborative interface that facilitates web search through conceptual exploration and requirements specification. It uses interactive and collaborative interface design techniques to overcome drawbacks associated with document-centric search systems, and provide a flexible, visual interface for guided query refinement.
This creates an integrated tool which lets users explore and reuse past search terms, to find relevant web documents while surveying domain terminology.
Within this collaborative search system interaction automatically moderates query repositories, improving information access by querying keyword suggestions that help clarify search requirements. This interface and system design draws upon insights from research into knowledge engineering and human factors as well as social and biological systems, and creates an intuitive semantic search tool.

The success of an information retrieval tool depends greatly on its interface design, and careful observation of human factors is essential when designing such tools. Several studies have been conducted to understand such interaction with search engines. Jansen has complied an extensive overview of recent search behavior research [21], focusing on usage metrics such as query length, number of pages viewed and probability of use of search syntax. Findings from this study are shown in Table 1.

TABLE 1: Typical search engine use

Study	Fireball (1998)	Altavista (1999)	Excite (2000)
Time Period	31 days	43 days	1 day
Number of Documents	3 million	100 million	40 million
Number of Queries	16 million	1 billion	54 thousand
Query Length	1.66 terms	2.35 terms	2.21 terms
Use of Boolean Operators	2.5%	-	8.5%
Viewed 10 results or less	59%	85%	58%

Examining the results of these studies we note 3 common trends:

- Queries are typically quite short at approx. 2 terms per search query.
- A small percentage of users take advantage of Boolean syntax (e.g. AND, OR, NOT) to improve search results.
- The majority of users examine fewer than 10 documents for a given query, indicating a very small fraction of results are viewed.

These trends indicate that users seldom use advanced techniques and do not methodically survey search results, due to lack of time and/or motivation. Long queries are quite rare, and only a few pages are typically examined. Since query length and the relevance of results is closely tied to the effort users invest in requirements specification, an interface which assists the expansion and reformulation of queries will enable more effcient discovery based search. InOrder fulfills this purpose, creating an interface that quickly elicits clarifications to initial search requirements. The keyword recommendation service not only suggests refinements and enables like-minded users to share search advice, but also reduces the cognitive effort required to make well-informed search decisions.

Collaborative interface design is a fundamental mechanism InOrder employs to assist information exploration and retrieval. A collaborative interface is an effective method of implementing a guided search application because it encourages participation constructing the semantic knowledge bases, which in turn helps users clarify search goals. InOrder seeks to have all searchers interact and contribute knowledge, so it should be designed to support a wide range of users and platforms. This means the interface must be platform independent, and employ a server-side solution that is easy to use so that individuals who are experts within a domain but have limited web experience are able to share search strategies. By making the usability of semantic collaboration and refinement central to all design decisions, InOrder is able to establish collaborative repositories of such strategies that assist decision-making during search.

InOrder's collaborative interface shares many functional characteristics with a WIKI [39] but is focused on the elicitation and sharing of search strategies rather than formal authoring. InOrder behaves as an “associated keywords WIKI" that aggregates query refinements into a shared semantic repository. Rather than groups of authors cooperating to create comprehensive articles, groups of casual searchers simply click on keywords they believe will improve their queries and this shares opinions of semantic relevance with future searchers. In contrast to a WIKI that may be modified heavily by one user and thus requires sophisticated revision control, within InOrder the process of stigmergic interaction incrementally adapts the composition of a search group. The interactive e®ort required to significantly alter a group is fairly extensive which dissuades abuse, and user interaction is rewarded with improved search queries.

Collaborative interfaces must consider as well motivational factors, and observations from peer-to-peer (P2P) systems offer insight in this regard. Vassileva's examination of usage within P2P communities [38] notes that participation is essential for a system to achieve critical mass. She argues that participation in P2P systems falls into 5 distinct classes of users, and that an imbalance among them will prevent the positive feedback and critical mass achieved by services such as Napster [28]. The 5 classes of users corresponding to decreasing cooperative participation are: 1. Create service, 2. Allow service, 3. Facilitate search, 4. Allow communication, and 5. Uncooperative free-rider (eg. not allowing uploads). She suggests that the service offered should be tied to levels of participation, to ensure a significant proportion of the users operate within the first few levels and the majority of users don't monopolize the resources of a few (on Gnutella, one study found only 5 percent of users provide over 70% of content [13]. Although such observations focus on P2P communities they are equally applicable to participatory interfaces such as InOrder.

Vassileva's arguments suggest a collaborative interface should encourage interaction that contributes new information while consuming it. This indicates that the cost-benefit ratio user natural usage should be designed to discourage free-ride behavior so the majority of users contribute value to the system. InOrder utilizes such principles to overcome participatory barriers and provide a simple, one-click interface with direct visual feedback. Users interact with collections of keywords, to discover useful terminology, and such exploratory efforts immediately benefit other users.

To ensure interaction compels participation InOrder also considers the relationships between *user perception, contribution and reward*.
More specifically, these 3 fundamental elements of the user experience must be balanced so that users feel interaction is worth their time and effort. *Interface perception*, concerns how quickly a user understands the purpose and benefits of interaction. *Contribution*, requires that a sense of contribution is present, which InOrder achieves through the Direct (Visual) Manipulation approach which will be detailed in Section 5.

3. Comparison with other Query Advice Systems

InOrder has similar goals to services such as Google Suggest [16], but differs in it's exact approach and application. Whereas Google Suggest offers query completion based on the first few characters, InOrder encourages interactive construction of queries from associated domain terminology. Such queries are incrementally constructed by choosing frequently occurring keywords extracted from top-ranked Google search results. Associated terminology is mined from authoritative web search results obtained using the Google API [14], Fig. 3, so that frequently co-occurring concepts may be incorporated into a search query. This effectively allows selective conceptual reuse of ideas contained within authoritative documents, reducing the effort required to perform query refinement tasks. To casual users seeking information on general topics this process may be viewed as suggestions of things to search for, whereas more advanced users may use the interface as a power search tool for surveying a domain and recombining semantic information into precise queries. Integration of such requirements into a search group facilitates reuse past search strategies, and helps users identify more relevant conceptual paths to search more effectively.

InOrder extends the approach of other query refinement systems such as the OBIWAN developed by Cooper and Byrd [6]. OBIWAN (One Button Interface With Associated Network) is a visual interface which prompts query refinement during search, providing cues which help users formulate better queries. Cooper et al. observed that when users translate information requirements into queries the “vocabulary problem" -described by [11] - is a significant barrier to query construction. Users need assistance increasing precision, reducing ambiguity and avoiding synonymy when constructing queries. OBIWAN was designed to meet this need. The OBIWAN interface is backed by a suite of tools called TALENT (Text Analysis and Langauage ENgingeering tools) which builds collections of lexical resources to overcome the vocabulary problem. TALENT extracts domain specific vocabularies (also referred to as context thesauri) to inform users of relevant conceptual knowledge. The interface uses a “lexical network" in which terms are the nodes and relations the links, to present useful associated information and elicit query modifications. InOrder provides similar functionality within a more flexible, usable framework. InOrder uses a server-side HTML interface instead of a client-side UNIX GUI used by OBIWAN. InOrder also uses Google's API (Fig. 3) and web-mining to acquire lexical resources from over 8 billion WWW pages, rather than IBM's NetQuestion (IBM, 2005) which queries smaller, more specific document collections. This endows InOrder with better usability and greater flexibility. InOrder also displays complete document summaries on result pop-ups, rather than just document titles displayed by OBIWAN. InOrder also alternates between query refinement and document browsing activities, providing one interface which does both simultaneously. This process of interaction guidance (pop-up Google results after 4 clicks) encourages selection of a greater number of terms because interaction improves the quality of results from a system most users are already familiar with (Google). However, the biggest differences between InOrder and OBIWAN are it's collaborative aspects. InOrder encourages collaborative, visual evaluation of semantic relevance which further assists comprehension and scanning of vocabulary. Visual relevance moderation is possible after keyword extraction, allowing participants to validate and reuse only the most useful keywords within the open knowledge bases. This enables the quality of InOrder groups to improve as participation increases, as collective and synergistic organization emerges from personal query refinement activity. This collaborative mechanism for relevance determination extends OBIWAN's approach, and enables incremental convergence towards more useful search groups.

The VQUERY [23] query refinement tool also has similarities to InOrder. Just as Cooper noted, a primary difficulty gathering information from large document collections is the necessity of reading large volumes of text, Jones believes a major bottleneck in conventional information retrieval is that results returned do not assist query reformulation. The VQUERY tool was designed to support refinement of initial queries and help users explore related ideas without scanning large volumes of text. VQuery shows “query previews" and “dynamic queries" to update results when queries are changed. It presents relevant keywords which may narrow search and guides the process of web search. Like OBIWAN, VQUERY was constructed using UNIX GUI toolkits to query small specialized collections, and usability and flexibility were not primary considerations. InOrder performs both of VQuery's "query preview" and "dynamic query" functions but within a simplified HTML interface connected to Google. The “query preview" functionality corresponds to InOrder's “New Group", “New Topic" or “Explore" functions, and Google Pop-up windows to the “dynamic queries" feature as 4th, 5th and subsequent keyword selections are made. Such improvements allow InOrder to serve a wider range of users and retrieval contexts, and provide a consistent visual interface which is more usable than OBIWAN or VQuery.

InOrder's query refinement approach is also in many ways an extension of the “Also try" feature on Yahoo [40] or the “gigabits" offered by the Gigablast search engine [12], the primary difference being that they are fewer in number and integrated with the document results, rather than using InOrder's alternation approach. The visual conceptual interface lets users build open repositories of query refinements that guide collective search activity. In contrast to OpenMind that has the ambitious goal of modeling common sense, InOrder simply elicits semantic associations within a search context. Rather than creating formal knowledge repositories, InOrder builds simple, open models that informally represent search interests of like-minded users. This process of collaborative modeling of search goals within a domain forms public collections of semantic knowledge and helps users reuse successful search strategies.

The design of InOrder's approach to interaction, and how semantic foraging prior to document exploration assists search in many information retrieval contexts is based on interaction models. Two in particular which have inspired InOrder's interaction model are Information Foraging [30] and the StumbleUpon web discovery service [36]. These two approaches improve the Standard Interaction Model by applying principles from social and biological systems to the process of information retrieval. InOrder uses a new interaction model combining aspects of Berrypicking, Sensemaking and Information Foraging with a hybrid rating approach found within StumbleUpon. This interaction model focuses query formation activities by recommending useful associated keywords as potential refinements.

InOrder's semantic query refinement process suggests terms commonly used by web authors as potential refinements. The acquisition of these concepts guides a users search process, and integrates relevant concepts into the collective knowledge base. Moderation of these knowledge bases occurs in an incremental and evolutionary fashion, in which potential query refinements compete for the attention of users. Selective recombination of such keywords increases conceptual awareness and facilitates query construction. There are primarily three differences between InOrder's approach and the standard interaction model, Fig. 10.

First, users examine keywords rather than pages during the initial stages of search. The reasoning being, in situations when search goals are not clear it is more efficient to browse ideas than pages to narrow a search. Second, within InOrder users judge conceptual relevance rather than page relevance. This creates meaningful models of query relevance and avoids ambiguous or inaccurate judgements caused by poor page design or aesthetics. Third, users receive guided query formulation, rather than having to browse several documents to identify helpful terms. Reducing the effort to refine queries, increases query length and search relevance [32]. This unique interactive model is further supported by web-mining which generates these query suggestions.

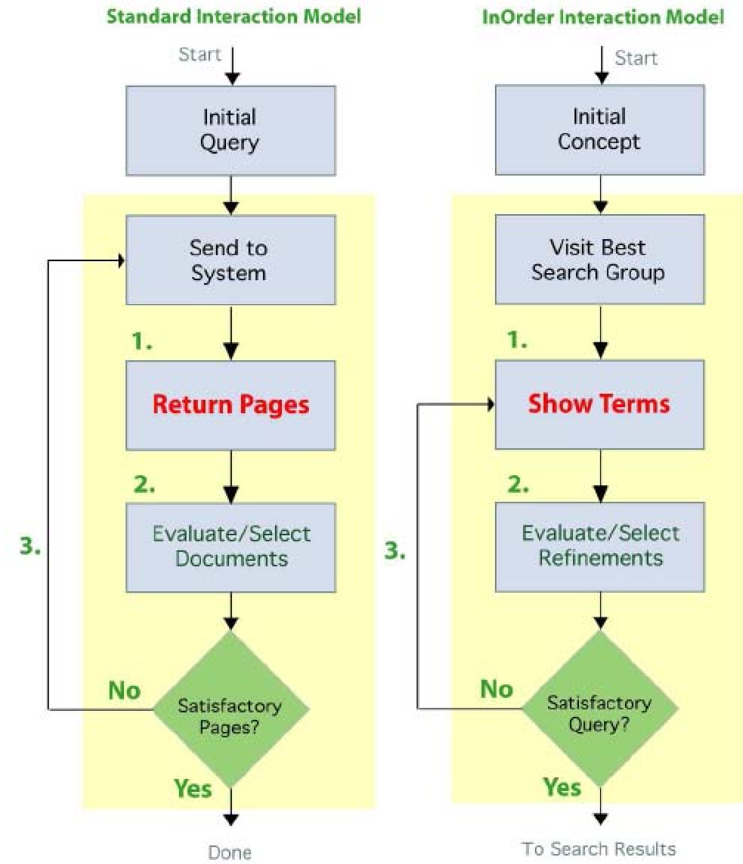

Fig. 10: Standard Interaction Model vs. In Order Interaction Model

From an ontological perspective, InOrder constructs informal associative models of a domain. Interaction with the semantic blackboard validates keywords relevant to a search group, and encourages re-examination of successful search strategies. These simple selective acts create useful semantic models, and with greater levels of interaction groups converge towards more objective representations of conceptual interest. This knowledge acquisition approach is less ambitious than projects such as OpenMind [42] or CYC [43]. InOrder seeks to capture only knowledge of semantic relevance rather than formal conceptual structures to support automated reasoning. While the simplicity of such associative models does reduce InOrder's flexibility compared to the mentioned projects, it also has potential to collect greater amounts of data because it is collected through the byproduct of searching (a common task) rather than a non-essential stand-alone process.

4. In-Order's Foundational paradigms

In-Order operates on the following paradigms of complex adaptive systems functionality:

4.1. Direct Visual Manipulation

As mentioned above, one technique employed by InOrder to increase usability of the blackboard interface is an interaction style called direct manipulation. [33] points out that

this approach to user interface interaction usually leads to more usable interfaces, and InOrder allows direct, visual manipulation of all semantic knowledge within a search group. Shneiderman also notes that one of the reasons for the success of direct manipulation interfaces is that visibility of objects of interest enable interaction close to the high-level task domain. Users don't need to decompose tasks and figure out what is going on, they simply interact and receive immediate feedback on the outcome of each action. Each act produces a comprehensible result in the task domain whose effect is immediately visible, reducing user anxiety because the system is comprehensible and actions are easily reversed. This interactive approach elicits meaningful semantic knowledge while helping users understand the purpose and
benefits of participation. Direct manipulation supports stigmergic coordination of InOrders' semantic collections. Each click or vote creates immediate visual notification confirming interactive intentions were met. Combined with visual presentation which reflects relevance, this facilitates dynamic visualization of search domain concepts and offers a simple and meaningful way to explore the search advice of others. Ideas are explored before pages, establishing visual semantic trails which adapt to reflect search intent. This technique also preserves a users sense of control and contribution, to enable participation without personalization.

4.2. Orienting Response and Information Scent

In the context of human experience, the Orienting Response (recognized by the Russian physiologist Sechenov in the 1850s, and studied extensively by Ivan Pavlov in 1927) explains why moving objects attract our attention. Early humans took advantage of this reflex not only to hunt and avoid predators, but also to detect fire, lightning and other environmental threats. Colors such as blue or green which are common in nature evoke the least response, whereas rare ones such as yellow or red draw greater levels of attention. Within modern society such behavioral reflexes now also affect activities such as driving automobiles or watching television. Consideration of the Orienting Response has guided the design of red traffic lights, yellow highlighters, terrorist alert systems and television ads, each in an attempt to provide novel stimuli which gain and hold human attention. InOrder instead applies this principle to the exploration and markup of keyword repositories which assist query refinement. Much like ants modify their behavior in response to chemical scents, coloration within InOrder guides the eyes and thus use of suggested semantic information. By combining the visual Orienting Response and the principle of Stigmergy within a simple interface, InOrder faciliates visualization of relevant conceptual information and perceptually guides information foraging activities. An example of such markup invoking the Orienting Response is illustrated in Figure 2 Brighter terms such as green or yellow ones draw more attention because they have been voted up most often, suggesting they were useful in the past. Such cues direct attention towards promising semantic information within a given context, and support the selective process during query reformulation. This approach is based upon Chi and Pirolli's work on Visual Information Scent and Searching [4]. Visual markup informs users of the Information Scent of potential refinements, with 7 colors and 7 relevance states providing attention guidance. Terms with greater utility within a context influence future exploratory efforts simply because they are noticed more often.

In short, the color draws attention to most relevant information so time is spent exploring the most promising options. Similar to the cognitive maps formed via pheromone trails in ant colonies, the perceptual semantic interface guides conceptual exploration, and explicit semantic use/voting (i.e. explicit moderation) maintaining the relevance of shared semantic structures. This indirect communicative approach enables self-organizing semantic repositories to emerge from query interaction as per Section 6.

4.3 InOrder as a Memetic Algorithm

InOrder behaves much like a memetic algorithm, one which adapts according to human choice rather than natural selection. InOrder achieves this by treating query terms as memes, pieces of knowledge which reside within host groups and may be utilized for search. Yet unlike societies where information is spread through selective oral and written discourse and stored within human memory, InOrder uses the open semantic blackboard to share concepts or memes in an indirect and collaborative fashion. Within this framework search interaction deposits keywords into a shared environment where they are subject to public examination and review. Contextual click frequency influences position, color and thus adequacy of a concept, and as domain characterization occurs memes with the greatest contextual utility rise to the top.

Table 2 shows the representational analogies between genetic and memetic [8] systems and InOrder's components.

Genetic Algorithm	Memetic System	InOrder
Gene	Meme	Keyword
Chromosome	Paradigm	Topic
Population	Society	Group
Fitness	Urge to repeat	Urge to click
Initialization	Cultural Norms	Popular Terms
Selection	Most Compelling	Most Useful
Crossover	Human Discourse	Term Combination
Mutation	New Ideas	New Terms

Table 2: Representational analogies between genetic/memetic systems and InOrder

As a composable memetic interface InOrder facilitates the collection, management and reuse of search memes. The interaction design promotes incremental acquisition of useful ideas, and re-evaluation or re-combination of frequently used terms. This process becomes one of intelligent crossover and mutation as human searchers intentionally explore new conceptual terrain, and “memetic foraging" prior to document evaluation allows the keywords to emerge via cooperative selection. This creates collaborative knowledge communities, in which individuals with similar interests can easily share opinions, attitudes and ideas within a conceptual ‘commmons’. Potential applications of such modes of use include collective brainstorming and strategy formation, in which the shared environment becomes a sketchpad for the exploration of related ideas. In the next section we will examine how this evolutionary/memetic framework has the potential to organize conceptual information without any central guidance.

4.4. Semantic Self-Organization

As illustrated in Table 3, InOrder's components also have direct analogies to neural and ant systems which both exhibit properties of self-organization [37]. Like these systems, InOrder employs a swarm strategy to coordinate information exchange among system entities. Inspired by the parallels drawn between natural and information systems in [22], InOrder is designed such that actions of many independent agents lead to increasing system organization. Directed human choice [25] (click on preferred keywords) - as opposed to pheromone secretion - creates persistent but adaptable semantic communities, which evolve and self-organize via collaborative keyword selection.

Brain	Colony	InOrder
Neurons	Ants	People
Hebbian process	ACO algorithm	HBGA
Cognitive Map	Pheromonal Field	Semantic Blackboard
Learning	Foraging	Query Reformulation
Neural path	Foraging path	Semantic Search path
Activation level	Pheromone intensity	Keyword rank
Reinforcement	Deposit Pheromone	Positive interaction
Inhibition	Evaporation	Negative interaction
Perception	Chemical detection	Color draws attention
Recruitment	Move along gradient	Notice relevant terms
Associative priming	Past food discovery	Past keyword utility
Minimizes required energy	Path distance	Human evaluative effort
Efficient neural structures	Efficient food collection	Efficient search

Table 3: Analogy InOrder and other self-organizing systems

This indicates InOrder's conceptual framework has potential to form self-organizing search maps. InOrder may be viewed as a Human-Based Genetic Algorithm (HBGA) which constructs domain-specific search groups [25]. The query refinement and moderation actions of participants are human-driven selective behavior, and the semantic blackboard represents chromosomes evolving through direct manipulation by cooperating human agents. This facilitates semantic discovery as conceptual mutation/recombination occurs and selections are integrated into a knowledge base.

Fitness within these semantic structures is reflected by both position and color of the keywords, and collaborative interactions adapt present structures towards more useful search maps.

5. Architectural Design

5.1 Design Principles

InOrder's collaborative query refinement approach helps users build better queries. By helping users focus their searches and clarify what they are looking for, this helps them utilize Google more efficiently (Fig. 11). Although current search interfaces do provide spelling corrections and related directories, they do not guide the process of query construction. InOrder goes one step further, providing conceptual guidance that assists this process.

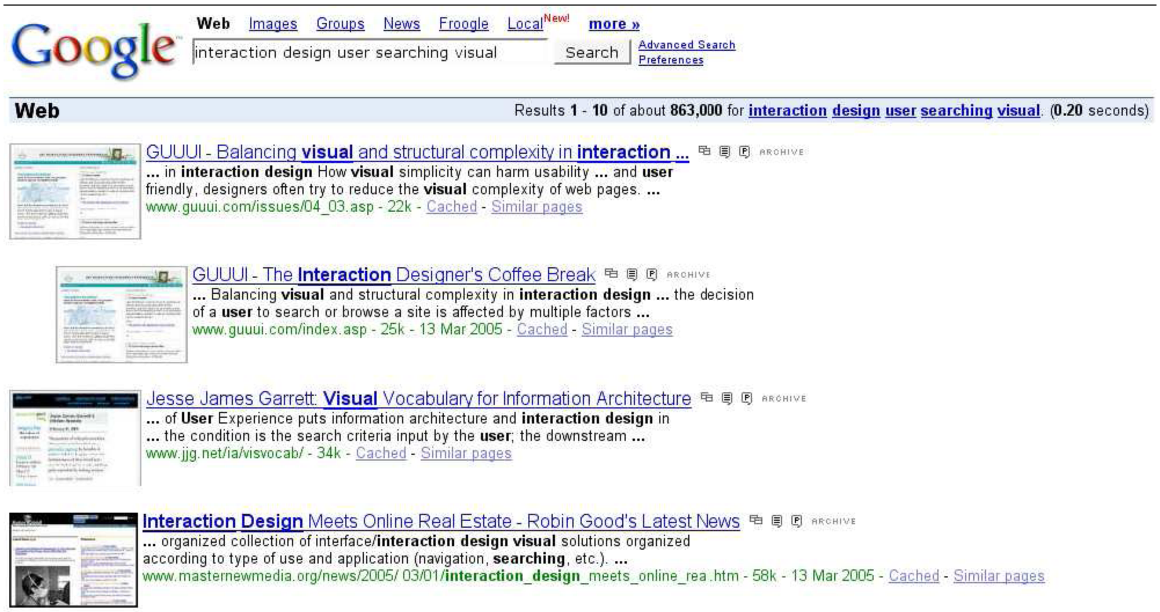

Fig. 11: Google pop-up triggered by selection of keywords ‘searching, usability, navigation and interaction design’

To achieve this InOrder was designed following these principles:

- The interface should be platform independent to improve accessibility. This prompted use of server-side PHP solution which supports any forms-capable browser. Barriers to entry are reduced by avoiding signup, login or submission of any personal information.
- Use adaptive hypertext to structure the process of query reformulation. The system should log search tasks within the appropriate context and aggregate relevance opinions from individual users into a shared repository. This semantic blackboard with visual markup should behave in a stigmergic manner.
- Encourage explicit semantic interaction to ensure contributions are consistent and meaningful. A single click should perform actions which offer both personal and collective benefits, such that users participate in a process of conceptual design as they refine personal queries. This will allow casual participation to construct meaningful search maps.

- Prompt selection and evaluation of concepts before pages to assist accurate requirements specification. This is achieved through interaction design which conducts semantic search in parallel with document search. Web mining also supports this procedure, by incrementally extracting relevant keywords from relevant search results.
- Employ the principle of *direct manipulation* to increase system comprehension and motivate participation. All information should also be presented in a manner that assists visualization of recent interactions and fosters a sense of contribution.

5.2. Architecture

InOrder's backend is designed such that casual interaction naturally validates and organizes semantic information into meaningful data structures. The architecture, Fig. 12, helps users incrementally focus search tasks and easily build and manage groups, which assists exploration by like-minded searchers. Within InOrder a domain is represented by a search group, a collective history of query refinements which characterize the needs of users who chose to interact with a group. InOrder utilizes such domain models to assist keyword discovery, using the visual semantic blackboards to present users with concise summaries of interaction history. This assists users when they are unsure of search requirements, as they are able to examine stereotypical selections made by others. InOrder constructs domain models by extracting terminology from pages returned by Google, while using explicit human selection to incrementally associate keywords. This may also be thought of as search task modeling, in which search requirements are elicited from anonymous users and aggregated into contextual repositories. These models are concise and meaningful because they are formed through explicit user validation of extracted suggestions. These search ontologies do not formally specify relationships between entities, simply recording adjacent keyword selections within a given context.

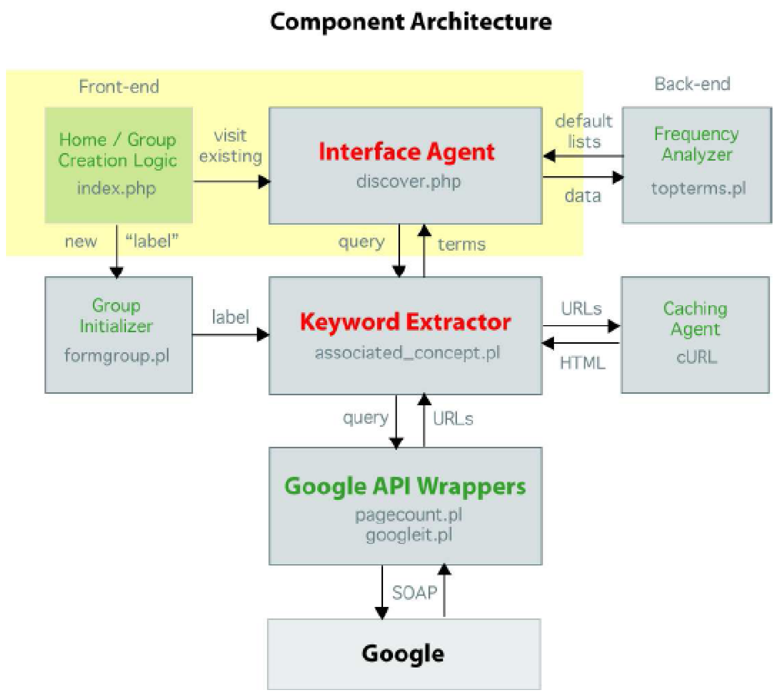

Fig. 12: InOrder System Architecture

6. InOrder Interaction With Google

6.1. Group Creation Strategy

InOrder's group creation strategy (Fig. 13) determines whether a new or existing group should be presented to a user. This is based on the belief that users should reuse existing groups if it is similar to the concept requested to avoid creation of very similar groups (e.g. "car care", when "cars" already exists containing a lot of related terms).

Currently In-Order's knowledge base consists of approx. 1700 groups so the majority of queries will create new groups, and only if subjects related to collaborative systems or interface design are queried will the custom Google search likely present a list of similar groups. As the knowledge base grows to hundreds of thousands of groups, it is anticipated that the majority of queries will take users directly to a matching group, or direct them to very similar groups. This will result in faster average group access times, as smaller percentages of queries will trigger the keyword extraction process which takes the majority of time during group formation.

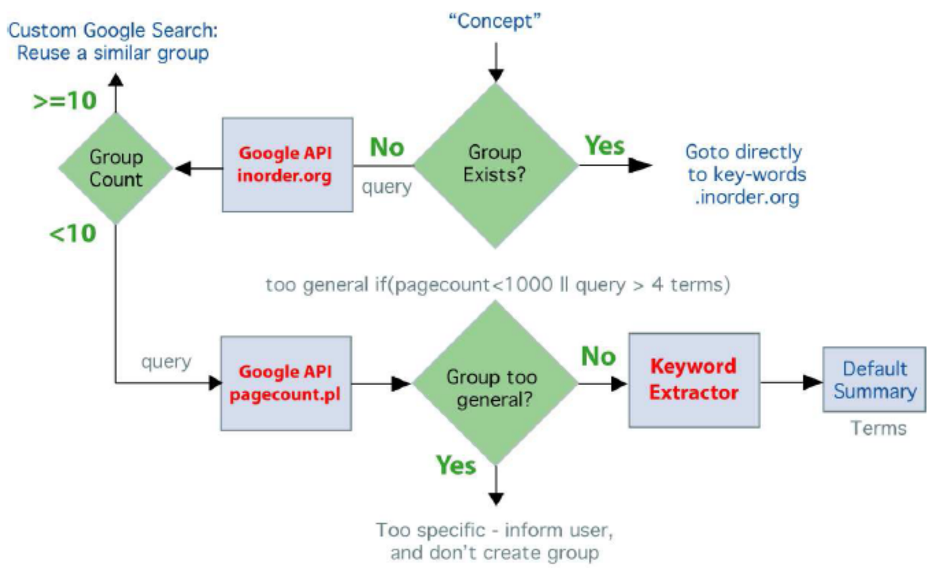

Figure 13: Conditional logic for group creation or reuse

The logic behind this decision-making process is invoked within the default index.php users see when visiting InOrder.org. It prompts search for a "concept" relating to their search needs. The following conditional logic is then applied as illustrated in Figure 13 above:

1. If the query matches that of an existing group, take user to group.

2. If the query is too specific (more than 3 terms, or less than 1000 Google results) instruct user to submit a new concept.

3. If several groups already exist which contain the queried terms (on custom Google 10+ groups already exist) then direct to Google results.

4. If the query isn't too specific and a group of that label doesn't exist, create a group of terms associated with the query (call extractor). If a specific group is desired which didn't meet the 3rd criteria users may visit the "create" link on the main page in which case this restriction is removed, allowing dedicated users to create similar groups if they are willing to put in the extra effort. This group creation strategy allows unique new queries to automatically create new search groups, yet encourages collaboration among users searching for similar subjects.

Step 2 and 3 of this decision logic utilizes a script that returns the number of pages in Google's index matching a query. This script (pagecount.pl) accesses the Google API using SOAP (simple object access protocol) and reads the estimatedTotalResultsCount which variable it returns. This variable represents an approximate count of the total number of pages containing the given query terms within Google's index of over 8 billion pages. The index.php file invokes this script when a search is performed from the default page, which returns the pagecount. This count is returned for both pages in Google's entire index, as well as pages indexed within the inorder.org domain, so custom Google search is only invoked if several similar InOrder groups exist, and new groups aren't created for very similar or misspelled subjects. The following call performs this functionality:

```
result = googleSearch -> doGoogleSearch(key, modquery, 0, 10,
"false", "", "false", "lang_en", "latin1", "latin1");
pagecount = result->{'estimatedTotalResultsCount'};
```

Using pagecount to guide interaction is based upon the idea of Googleshare [17]. Based on the concept of ‘mindshare’ that seeks to understand how prevalent an idea is among the populaction, Googleshare is
defined as the percent of Google's index which contains a given term. Thus a term such as 'interfaces' has a Googleshare of 0.59%, compared with 0.042% for 'collaborative interface design', and this metric indicates relative conceptual prevalence among web authors. InOrder applies this approach but uses a hard limit rather than a percentage requirement. Empirical experimentation found that when groups were created for very specific subjects, meaningful terms could not be easily extracted, so a requirement of 1000 pages on Google was imposed to create a group. Combined with a limit of 3 words to describe a search group (so very specific concepts are not entered) this avoids creation of groups within domains in which InOrder likely wouldn't be useful, since within such a specific domain PageRank would likely have already placed the most authoritative and relevant documents on the first page of results.

If several InOrder groups containing the query terms supplied at the default entry page already exist, custom Google search recommends examination of these similar groups. This capability is provided by restricting search within Google's index to pages within the inorder.org top level domain. This is performed by supplying two additional hidden form inputs when querying Googles index. This group content is obtained by Googlebot via the group index listed at http://inorder.org/index.html.

```
<FORM method=GET action=http://www.google.com/custom>
<input type=hidden name=domains value=\"inorder.org\">
<input type=hidden name=sitesearch value=\"inorder.org\">
```

InOrder also uses the rel=nofollow attribute within certain hyperlinks to avoid redundant indexing of search groups. The nofollow parameter is one obeyed by Googlebot during crawling that instructs the crawler not to follow a particular link. This prevents pages downstream from that link from being indexed, and helps webmasters control how PageRank is allocated within local domains. Rel=nofollow is used on all InOrder interface links which perform coordinative functions. These links (of the form inorder.org/?function=parameter) should not be followed because they control adaptation of the interface rather than leading to new content. Such links were added after observing that Googlebot had indexed many redundant pages within a search group, and had manipulated the interface in an undesired manner.

Current Google has indexed 2300 pages on InOrder, of which 1700 are valid groups. Over time as traffic increases and the index is re-spidered this number should converge towards an accurate count of InOrder groups.

InOrder's domain-centric search groups use a contextual, frequency-based semantic model. When examining potential refinements users should be presented with terms most likely to meet their needs, so InOrder takes a utilitarian approach to construction of these models. It assumes terminology used most frequently within a specific context (a search group) is most useful, and that presentation order should reflect estimated utility. Following this principle, InOrder presents all semantic data within the default central pane according to contextual frequency of use. More frequently occurring terms are presented first, for either extracted terminology or explicitly validated refinements. This frequency-based approach ensures suggested refinements reflect the opinions of authoritative web authors. Viewing publishing as one-way discourse, it is assumed authors use terminologies that efficiently express their ideas and that more important concepts are included more frequently. To generate these suggested refinements InOrder's Keyword Extractor parses out frequently occurring terminology from the top 30-50 Google results for a given query or subject, and looks for consensus among authors using these terms. This process depicted in Figure 14 provides a conceptual set from which searchers may forage for relevant ideas. This information extraction and presentation strategy is chosen based on observations of language utilization, in particular Zipf's Law [41].

Some believe the power-law frequency-rank distribution seen in large textual collections is a consequence of human desire for communicative efficiency. Based on Zipf's Principle of Least Effort, such thinking assumes individuals utilize vocabulary that maximizes understanding and minimize communicative effort; indicating frequently used terms within a context are the most meaningful. Intuitively this makes sense

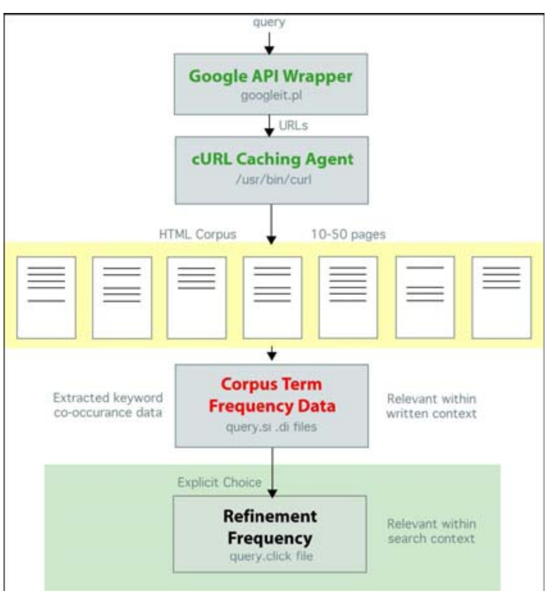

Figure 14: Using term frequency within search results to bootstrap search groups

The Keyword Extractor associatedconcepts.pl parses term usage patterns among authors of top Google results, to produce related suggestions within a domain of interest. This improves search usability since users may now simply scan conceptual summaries rather than clicking back and forth when query refinement is necessary.

Whereas the Pagecount script uses the Google API to supply only pagecount information, the Keyword Extractor uses the URL's which the API returns to dynamically create HTML corpora [14]. When an InOrder query is made, associated concepts are extracted from content linked to by search engine results. These top Google results for “new topic" and “explore" requests within a group forms a set of documents from which frequently occurring keywords are extracted. A Perl script called googleit.pl creates this list of URL's which is then read by the keyword extractor. Using the Google API to support such dynamic web mining enables InOrder to reduce the effort required to formulate queries within unfamiliar domains. Users are able to incorporate new concepts into queries with less effort than manual reformulation because they may simply select refinements from suggested concepts that reputable authors felt were important.

6.2. Keyword Extraction via Datamining

InOrder's Keyword Extractor is a vital component supporting interface usability. It caches HTML for top Google results and extracts semantic data from these corpora, so potentially useful keywords may be suggested to assist query refinement. This significantly increases the usability and flexibility of the InOrder interface, as it allows users to utilize concise semantic summaries of authoritative web resources. As new groups, topics and \explore" maps are created, the Keyword Extractor constructs a conceptual overview of the domain, which searchers may then collaboratively forage for the most useful terms. AssociatedConcepts.pl implements the Keyword Extractor, consisting of 554 lines of PERL code which cache's and extracts semantic information. This data is loaded by discover.php, processed and presented to the user.
Following the frequency-of-use based strategy for domain modeling outlined above, this section describes InOrder's data model which is created incrementally as interaction occurs. The 7 data types are manipulated through four computer-initiated and two human-initiated actions, such that explicit selective actions extract meaningful data in a recursive fashion. Starting from the HTML cache, all the way to keyword relevance states, knowledge bases are constructed by interactively eliciting of search requirements from users. Since no additional effort (beyond that which offers personal benefit) are necessary to construct these search maps, this helps InOrder overcome the Knowledge

Acquisition Bottleneck, and quickly bootstrap search groups. Figure 15 illustrates actions

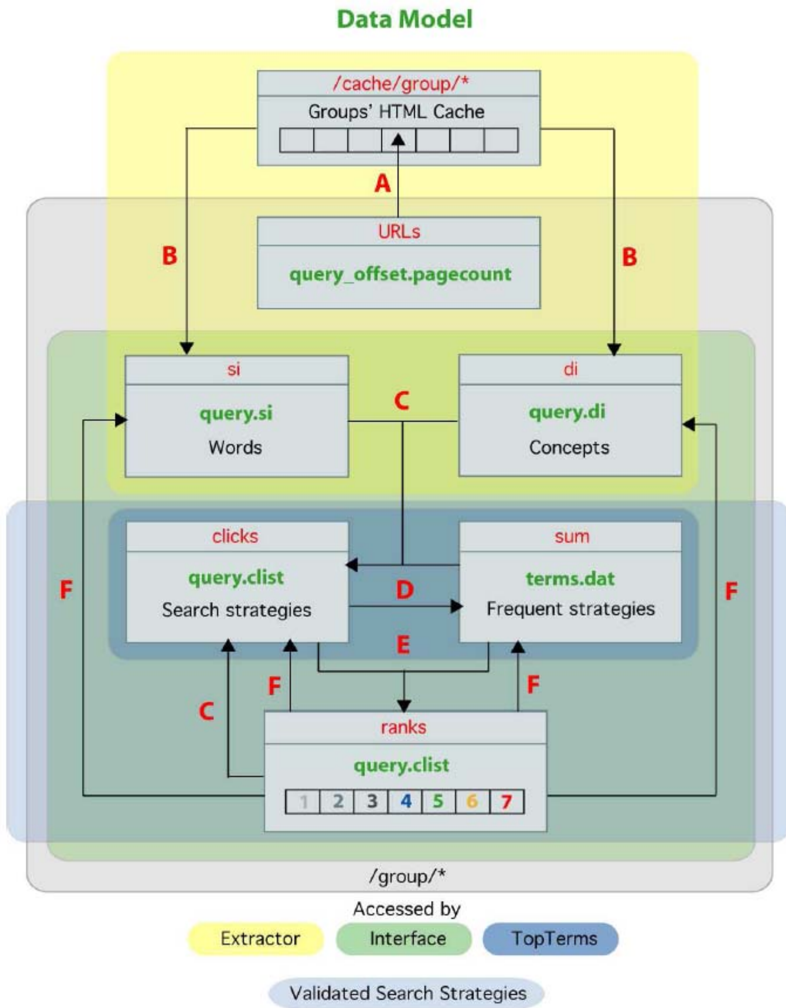

and data

Fig. 15: InOrder Data Model

InOrders' seven data types (urls, cache, si, di, clicks, sums, ranks) reside with folders of the same names, with extensions of the .pagecount, .html, .si, .di, .clist, .dat and .dat respectively. The current implementation stores this data in flat text, and a future improvement would be to a relational database such as MySQL implementation with XML-based export functionality. The three system components which manipulate search group data are:

1. The Keyword Extractor, which invokes cURL [3] to produce cache which matches the contents of url returned by the Google API. Cache is then processed to produce si and di outputs sets of suggested terms and term pairs.

2. The Interface, which coordinates selective transfer of keywords from .si and .di-s into .click -ables, and movement of keywords within ranks.

3. The TopTerms script, which processes all present click -ables into the .dat residing inside the sum folder. This aggregates click behavior within each topic into a summary which becomes the groups default view.

The 6 actions which control data flow and crystallization of HTML into more concise semantic representations are:

A. cURL fetches html page of each URL returned by Google API,

B. The Keyword Extractor gathers frequently co-occurring terms from HTML corpus,

C. User selects refinement which validates the term, and moves it from si/di -bles into .clist for current topic,

D. Topterms.pl performs frequency analysis on .clist -bles (all topics) to create default group summaries residing in the sum folder,

E. Users vote on validated keywords, moving them between 7 various keyword relevance state -bles (p1-3.dat, m1-3.dat) and

F. Interface utilizes keyword rank status to color potential refinements. As a result, higher ranked terms become noticed more often than lower ranked ones.

This data model allows the three simple user actions of *request extraction*, *select keyword*, and *vote on keyword* to incrementally develop semantic blackboards and support 1-click query refinement.

As shown in Figure 16, InOrders' visual stigmergic interface coordinates this incremental specification process.

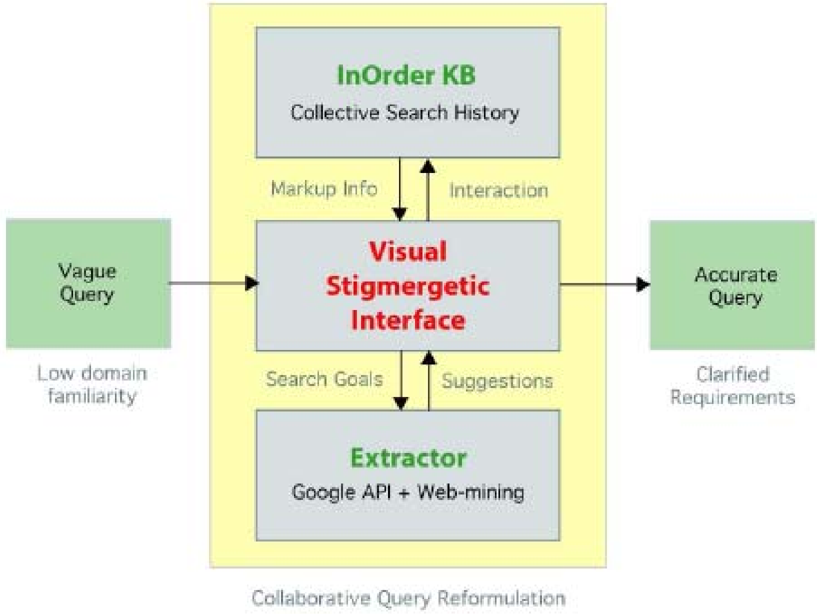

Figure 16: Conceptual Search approach

7. System Evaluation

To gauge the potential value of InOrder's query support system we will now examine how InOrder usage compares to regular Google search sessions.

7.1. Case Study: InOrder vs. Google

A comparison of search behavior between InOrder and Google seeks to answer several questions. First, how does InOrder affect the manner in which users formulate their queries? Second, on which types of search tasks is InOrder most helpful? And third, is InOrder better for searches within familiar or unfamilar domains? To answer these questions the search session behavior of a group of volunteer beta testers was analyzed within an experiment designed to measure search session productivity. Testers were asked to perform similar search tasks within various levels of domain expertise on both Google and InOrder, allowing session characteristics to be compared. This enables understanding of exactly how InOrder affects information seeking during search, and by how much it may improve the search experience.

7.1.1 Expected Use Cases

Since InOrder is a flexible tool for semantic collaboration it may be used in several different ways. Depending on a users' knowledge of computing systems and level of motivation during search the following use-cases were anticipated:

a. Query refinement - The most common use-case was expected to be one of refining search: users click suggested terms to construct better queries, and thus ¯nd better pages. This offers the highest personal benefit, in particular when search is being performed within a subject a user has little knowledge of.

b. Brainstorming - A user browses/clicks primarily to explore and learn about a domain. This was expected to be less common, since usage would be primarily for gathering ideas and concepts to assist broad or complex searches. While this still provides personal benefit, it would likely be useful only for more involved search tasks.

c. Collective Research - Users engage in purposeful conceptual design and cooperative search activity. This was expected to be the least common use-case and require the most time and effort.
Usage would be driven by the goal of passing ideas on to anonymous search peers, and contribute ones knowledge towards an index of “collective wisdom" within a subject. Since keywords would be clicked without necessarily examining any pages, the personal benefit derived from such use would often be lower, with interaction performed for the purpose of characterizing a domain.

7.1.2. Results

Here we compare Google and InOrder session metrics for 10 volunteers recruited in June 2005. These sessions represent roughly 3 hours of search interaction on each system. Once testers had completed the experiment and emailed questionnaire responses, the Search History logs were analyzed to examine the differences in search behavior. The interaction data was tabulated into an Excel spreadsheet as follows: Average query length for each query during a given task was calculated by taking an average length in words.

Time in minutes from first page visit to task-end indicator was used to measure the time to complete a task. If recorded minutes were the same 0 was used, indicating the task likely took less than 30s to complete. The number of visited pages was obtained using a simple count, and search success rate determined by simply comparing the ratio of successful to unsuccessful search tasks.

Compiled results indicate that InOrder does improve search efficiency by providing a visual semantic interface for query refinement. This supports the hypothesis that one-click query refinement improves web search in certain instances. The effects on search session behavior illustrated in Figure 17 are:

- Increased search success rate. This indicates that users successfully completed search tasks more frequently when using InOrder than when using Google alone.

- Fewer page visits required to complete a task. This means less effort was required when using InOrder vs. Google alone, with 42% fewer pages being viewed.

- Less time spent searching to complete a task. Less time was required, thus search sessions were more efficient. On overage, 32% less time was required to complete a task.

- Longer queries. The guided reformulation resulted in more detailed queries, which improves the relevance of Google results.

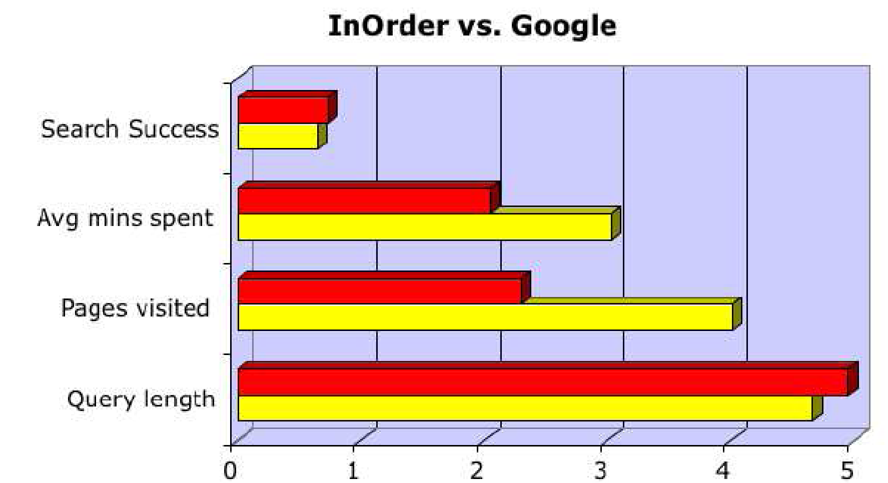

	Query length	Pages visited	Avg mins spent	Search Success
InOrder	4.948666667	2.3	2.05	0.733333333
Google	4.6695	4.016666667	3.033333333	0.65

Figure 17: Comparison of InOrder and Google for 4 search metrics

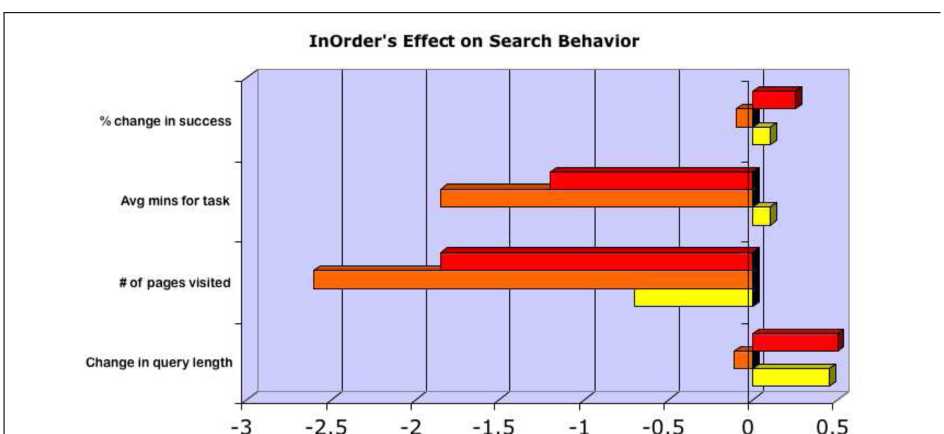

	Change in query length	# of pages visited	Avg mins for task	% change in success
No background	0.5	-1.85	-1.2	0.25
Some background	-0.1125	-2.6	-1.85	-0.1
Good background	0.45	-0.7	0.1	0.1

Figure 18: Behavioral effects, broken down by domain familiarity

Further examination of Figure 18 indicates that InOrder is:

- Most useful for unfamiliar domains (promotes learning, concept formation, sensemaking). All metrics increased in desired directions by a prominent amount: better search success rate, fewer pages and minutes required to complete a task, and longer queries. If you didn't know what to search for InOrder was found to be quite useful.

- Somewhat useful for somewhat familiar domains (exhaustive keyword search, brainstorming and research). Slightly more pages were visited on average, but success rate was better and query length/num mins showed slight improvement. So for a research tool InOrder does assist brainstorming, but requires a bit more effort to conduct a successful search.

- Not very useful for very familiar domains. Queries were longer and slightly fewer pages/minutes were required, but search success rate was significantly lower. So in this instance InOrder just got in the way. Users generally already knew what keywords were needed, so InOrder just encouraged unnecessary refinements which took search off track.

7.2. Discussion

During the course of development several changes were made to improve the usability and functionality of InOrder. One such change was group creation permissions. Transparency is one of the primary reasons open source development has been so successful; bad code gets fixed more quickly and good code spreads faster. Yet “gate-keepers" are also vital within open source development, so it was expected that each search group should be approved by a human moderator to reduce spam.

Initially this approach was taken but it stiffed participation. While it was intended to keep InOrder spam free, it proved to be a significant stumbling block since users wanted immediate gratification, and were dissuaded by a group approval queue. An additional problem was replication of very similar groups. For example, if a popular “evolution" group already existed and a user searched for “evolving", a unique group would still be created. This resulted in duplication of effort, and inadequate reuse of prior conceptual development. To solve these two issues the group creation logic described in Fig. 13, Section 6.1, was added. This meant groups were created immediately upon request, yet the filtering logic avoided creation of low-quality or very similar groups. The result of this change was more group creation and better collaboration since the custom Google search logic directs users to existing groups more often.

Other changes included interface description and layout changes to improve comprehension and usability. Since InOrder was initially developed for use by motivated researchers seeking powerful reformulation capabilities, clear instructions or guidelines were not provided. Changes included a switch from the instruction of “click on helpful terms" to “select keywords to narrow search", and a better description of what InOrder does on the default page. The label “group" was changed to “concept" on the front page, which reduced submission of very specific search queries rather than research subjects.

Another important change made was addition of group terms to all queries submitted to Google pop-up and Google API. This was required because users were often clicking refinements but not keeping the original terms within the working query to establish context. For example, after 4 clicks “motorcycle" specific terms would no longer be present within the query block, and results would no longer be relevant. Adding the group terms within all queries to provide context solved this issue.

Upon completion of the beta testing it was also realized that several areas of the experimental design could use improvement. Although 34 beta testers were initially recruited, only 14 fully completed the experiment due to its complexity or time requirements, and of those several did not complete all necessary actions to make appropriate comparisons. Only 10 beta testers completed the entire experiment properly, so further testing would be helpful in gathering a wider range of feedback. Simplification of the experiment would assist this, in particular by addressing the common mistakes testers made when carrying out the 9 steps. Common mistakes were:

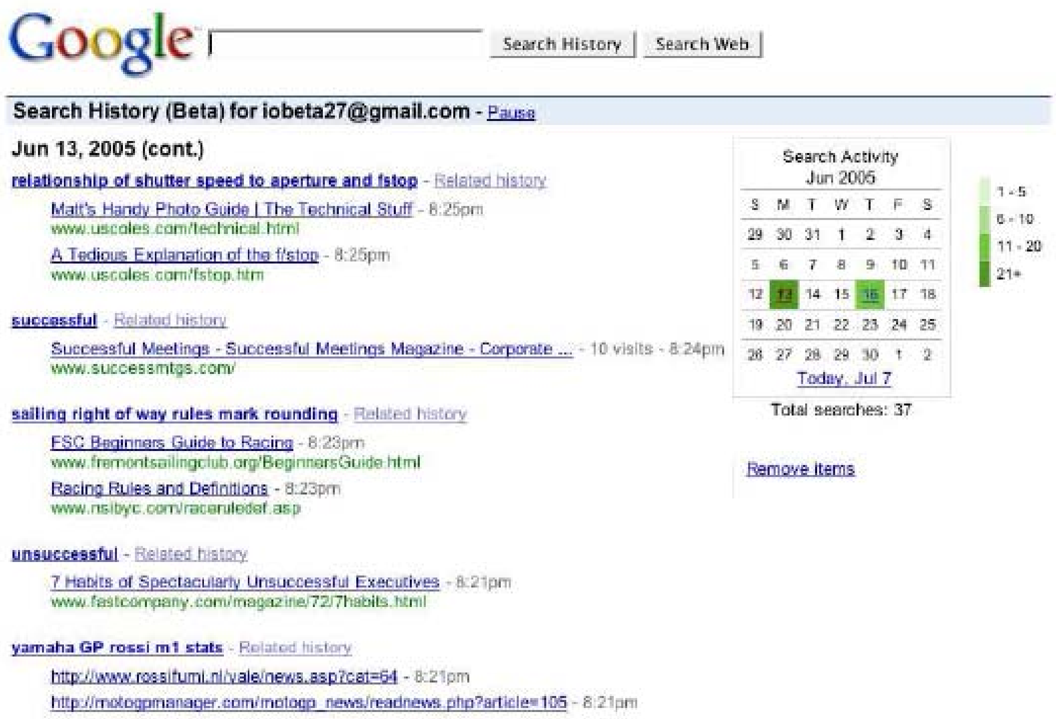

Fig. 19: Google “Search History” function

(1) users forgot to say if a search task was successful or not - instead they simply proceeded to the next search task. A special frame to guide interaction and elicit such feedback would avoid this issue.
(2) Comparisons made with Google Search History (Fig. 19) were accurate only to the minute level as they did not display second markers. While these differences even out over 60 datapoints, tracing the time between initiation and completion or abandonment of

a search task to the second would produce more accurate estimates of the time reduction InOrder offers.
(3) Users were asked if it was “successful" or “unsuccessful" but this is somewhat of a subjective measure. Additional feedback which said why the search was successful or not would enable better understanding of why a search was not successful, so that such issues could be addressed. Replication of this experiment with simpler instructions and the aforementioned improvements for 50 additional volunteers would validate these findings.

Feedback on InOrder also shows that the interface is still too complex for average web users. Although the interface has been developed with usability in mind, feedback suggests that comprehension of the purpose of value of InOrder is still limited by interface complexity. While this is not surprising considering the simplicity of conventional search interfaces, this indicates the interface should be simplified to accomodate a wider range of users. The number of interface elements needs to be reduced, such that new terminology is gradually presented alongside progressively more relevant results. Such simplification should enable InOrder to provide the same search support, yet open up participation to individuals who are used to interfaces as simple as Google and don't wish to become accustomed to a new search interaction model.

8. Future Directions

8.1 AJAX/MySQL Implementation

The current InOrder prototype uses PHP to implement the interface. Due to re-rendering after each interaction this requires transmission of 24kB of data with each click. For users without broadband this makes interaction very slow. Use of Asynchronous Javascript and XML (AJAX) would avoid this issue, reducing bandwith requirements and improving interface responsiveness. Another useful improvement would be to implement the relational semantic structures using MySQL tables rather than flat text. This would enable faster retrieval, less complex concurrency handling and simplification of the PHP code to avoid complex session variable adjustment.

8.2 Improved Web-Mining and Classification

InOrder's web-mining technique could also be improved significantly. First, the Keyword Extractor could be specified to use cached content (such as Google's Cache) rather than fetching documents each time. This would reduce the time required for corpus formation from over 5 seconds to a fraction of a second. The Keyword Extractor itself could also be optimized, and re-written in C++ instead of Perl. Combined with more sophisticated parsing techniques (adaptive window lengths for co-occurence calculation rather a global document approach) this would enable the Keyword Extraction component to return more useful refinement suggestions with much less delay. InOrder's semantic data models also have potential applications in dynamic classification of web content. Within Recommender Systems such as StumbleUpon automated content classification is performed to improve system usability. Since InOrder's search groups contains concise associative data relating to a concept or subject of interest, such datasets could prove very useful when automatically classifying web content. InOrders' dynamic semantic repositories could enable robust classification of textual content based on explicitly confirmed semantic associations collected during query refinement activities.

8.3 Social Semantic Clustering

InOrder's extraction process could also benefit from more sophisticated clustering techniques such as Levenshtein distance to calculate morphological difference among textual strings. This would allow terms such as “evolution" and “revolution" to be associated without requiring specific stemming rules which may not work in all instances. Semantic clustering among sets of individual semantic interaction would also be useful after privacy issues are considered and resolved. Clustering on personal semantic interaction profiles could establish semantic similarity networks similar to ReferralWeb, and enable discovery of like-minded searchers and relevant resources they have endorsed (matching tagged content). This is a promising area of future investigation. InOrder could also be integrated with social bookmarking systems such as del.icio.us [9] or StumbleUpon [36] to facilitate discovery of better information providers. The aforementioned semantic clustering techniques could be used to identify one's “search peers" and provide access to relevant content they have endorsed. Such integration would facilitate creation of more effective personalized search and recommendation engines, which would cluster participants together and re-rank search results based not only on tags, interests or binary ratings but also keywords utilized during prior semantic exploration.

9. Conclusions

InOrder assists the execution of vague or complex search tasks by providing a composable visual interface for semantic knowledge sharing. The collaborative search interface recommends keywords to assist query formation by extracting associated textual knowledge from authoritative search results. This allows users to improve their queries through selective use of associated terms which better specify search requirements, and helps users realize “what to search for" as recommendations are examined. Use of terminology which search peers found helpful is also possible, by examining bold keywords which peers endorsed when examining prior extracted keywords. This service of semantic recommendation prior to document retrieval causes search within InOrder to become a two stage process, in which InOrder displays aggregate semantic interaction guiding term selection behavior. Groups of searchers may use and reuse these search maps to clarify search requirements and construct better queries as they alternate between conceptual and document browsing modes. This creates open semantic indices with open reread/rewrite access in which all participants are moderators, and simple selective acts automatically improve index relevance.

From a Recommender Systems perspective, InOrder focuses on conceptual knowledge sharing rather than document discovery. Retrieval systems such as Google return authoritative pages matching a query but do not assist requirements specification. InOrder offers such support, assisting the formation of more precise queries when goals require clarification. This keyword recommendation approach enables InOrder to improve the capabilities of existing search engines and deliver more relevant results to users.

Beyond it's primary purpose of query refinement InOrder also functions as a collaborative interface assisting brainstorming and collaborative inquiry. The open, visual interface helps users not only gather and evaluate search strategies but also survey ideas associated with a domain. Participants may easily forage for ideas and make their opinions known by simply clicking on terms they find interesting. Cognitive and biological mechanisms observed within social insect systems are used to guide attention towards relevant domain concepts, and this supports conceptual learning and sensemaking even if queries aren't constructed. Incremental interaction automatically creates maps of associated keywords, and this helps InOrder overcome the “knowledge acquisition bottleneck". These semantic maps have several potential applications beyond keyword recommendation including classification, personal semantic modelling, and information filtering for e-commerce applications.

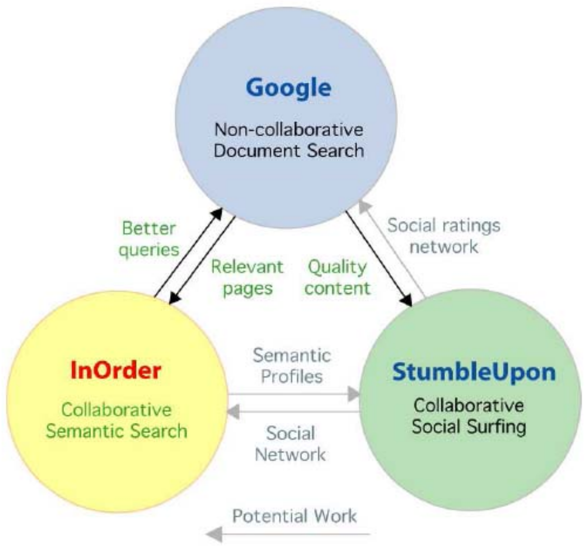

Fig. 20: InOrder in the context of today’s search engines

As shown in Figure 20, InOrders' anonymous, collaborative and semantic approach gives users the benefits of collaboration seen in systems such as StumbleUpon [36], and improves conventional search engines by encouraging exploration of concepts before pages. This knowledge acquisition strategy and semantic recommendation interface allows InOrder to integrate easily with other retrieval systems. The semantic collections guide query formation and direct users to more relevant search results, which may then be shared with others using social bookmarking services such as StumbleUpon or del.icio.us [9].

REFERENCES

[1] Battelle, J., “Bill Gates speaking at Wall Street Journal's Digital Conference", Website, 2004, http://battellemedia.com/archives/000693.php.

[2] S. Brin and L. Page, 1998 The Anatomy of a Large-scale Hypertextual Web Search Engine", Computer Networks and ISDN Systems, vol. 30, no. 1{7, pp. 107{117, 1998.

[3] Cache URL utility, Website, 2005, http://curl.haxx.se/.

[4] Chi, E.H., P. Pirolli, K. Chen, and J. E. Pitkow, \Using information scent to model user
information needs and actions and the web", in Proceedings of CHI 2001, 2001, pp. 490-497.

[5] D. R. Chialvo and M. M. Millonas, How Swarms Build Cognitive Maps", in The Biology and Technology of Intelligent Autonomous Agents, L. Steel, Ed., vol. 144, pp. 439{450. Nato ASI Series, 1995.

[6] Cooper, J. W. and R. J. Byrd, \OBIWAN - a Visual Interface for Prompted Query Refinement", 1998.

[7] Croft, W. B., S. Cronen-Townsend, and V. Larvrenko, \Relevance Feedback and Personalization: A Language Modeling Perspective", in DELOS Workshop: Personalization and Recommender Systems in Digital Libraries, 2001.

[8] Dawkins, Richard, The Selfish Gene, Oxford University Press, 1989, ISBN 0192860925

[9] Del.icio.us - Social Bookmarking, Website, 2003, http://del.icio.us/.

[10] M. Dorigo and G. Di Caro, \The Ant Colony Optimization Meta-Heuristic", in New Ideas in Optimization, D. Corne, M. Dorigo, and F. Glover, Eds., pp. 11{32. McGraw-Hill, London, 1999.

[11] Furnas, G. W., T. K. Landauer, L. M. Gomez, and S. T. Dumais, \The Vocabulary Problem in Human-System Communication", Communications of the ACM, vol. 30, no. 11, pp. 964{971, 1987.

[12] Gigablast Search Engine, Website, 2001, http://www.gigablast.com/.

[13] (Gnutella, 2000) Free-Riding on Gnutella", Website, 2000, http://www.firstmonday.dk/issues/issue5-10/adar/.

[14] (Google 2004) The Google API", Website, 2004, http://www.google.com/apis/.

[15] (Google 2005a) \The Google Zeitgeist", Website, 2005, http://google.com/press/zeitgeist.html/.

[16] (Google 2005b) “Google Suggest Beta", Website, 2005,
http://www.google.com/webhp?
complete=1.

[17] (GoogleShare, 2001)
http://www.answers.com/topic/googleshare.

[18] Greisdorf, H., “Relevance: An Interdisciplinary and Information Science Perspective", 2000, Informing Science (3:2), 2000, pp. 67-71.

[19] Ibm NetQuestion, Website, 2005,
http://www-306.ibm.com/software/data/textsearch390/about.html.

[20] Ito, Yoichi, “Emergent democracy",
http://joi.ito.com/static/emergentdemocracy.html.

[21] (Jansen, 2001) B. J. Jansen and U. W. Pooch, \A Review of Web Searching Studies and a Framework for Future Research", Journal of the American Society of Information Science, vol. 52, no. 3, pp. 235{246, 2001.

[22] S. Johnson, \Emergence: The Connected Lives of Ants, Brains, Cities, and Software", 2002, Scribner.

[23] Jones, S., “Vquery: a Graphical User Interface for Boolean Query Specification and Dynamic Result Preview", 1998.

[24] Kautz, H. A., B. Selman, and M. A. Shah, \The Hidden Web", AI Magazine, vol. 18, no. 2, pp. 27{36, 1997.

[25] A. Kosorukoff, \The Human-Based Genetic Algorithm", Website, 2000, http: //citeseer.ist.psu.edu/kosorukoff01human.html.

[26] Mizzaro, S. How many relevances in information retrieval? - Interacting with computers, 10(3):305--322, 1998.

[27] Mizzaro, S., “How many Relevances in Information Retrieval", Proceedings of the Workshop `Information Retrieval and Human Computer Interaction', GIST Technical Report GR96-2, Glasgow University, The British Computer Society, Glasgow, UK, pp. 57{60.

[28] (Napster, 1999) „Napster", Website, 1999, http://www.napster.com/.

[29] Personalization Article, Website, 2001, http://en.wikipedia.org/wiki/Personalization.

[30] (Pirolli, 1999) Pirolli , P. and S. Card, “Information Foraging", 1999, Psychological Review.

[31] (PCW 1996) “Principa Cybernetica Web", Website, 1996, http://pespmc1.vub.ac.be/.

[32] (Rememberance, 1996) \The Rememberance Agent", 1996, http://remem.org

[33] Schneiderman, B. Direct Manipulation: A step beyond programming languages , IEEE Computer, 16(8), August 1983, pp.57-69.

[34] A. Spink, D. Wolfram, B. Jansen, and T. Saracevic, Searching the Web: The Public and their Queries", 1988.

[35] Steve, J. (2001) Emergence: The Connected Lives of Ants, Brains, Cities, and Software, Schribner 2001, ISBN 068486875X

[36] (Stumbleupon, 2001) Stumbleupon, Website, 2001, http://www.stumbleupon.com/.

[37] (Ulieru and Este, 2004) Mihaela Ulieru and Robert Este, “The Holonic Enterprise and Theory of Emergence”, *International Journal Cybernetics and Human Knowing* (Imprint Academic), Vol. 11, No. 1, 2004, pp. 79-99

[38] (Vassileva 2002) J. Vassileva, \Motivating Participation in Peer to Peer Communities", 2002, ESAW 2002: 141-155.

[39] “Wikipedia", Website, 2001, http://www.wikipedia.org.

[40] Yahoo! Search, Website, 2005, http://search.yahoo.com/.

[41] (Zipf, 2002) http://en.wikipedia.org/wiki/Zipfs-law.

[42] P. Singh, \The Public Acquisition of Commonsense Knowledge", 2001, http: //www.openmind.org/.

[43] D. B. Lenat, \CYC: A Large-Scale Investment in Knowledge Infrastructure", Communications of the ACM, vol. 38, no. 11, pp. 33{38, 1995.