Proceedings of the Estonian Academy of Sciences,
2012, **61**, 2, 89–106
doi: 10.3176/proc.2012.2.02
Available online at www.eap.ee/proceedings

INFORMATICS

Complex search: aggregation, discovery, and synthesis

Georg Singer, Dmitri Danilov*, and Ulrich Norbisrath

Faculty of Mathematics and Computer Science, University of Tartu, J. Liivi 2, 50409 Tartu, Estonia

Received 8 June 2011, revised 18 October 2011, accepted 18 November 2011, available online 21 May 2012

Abstract. Search engines such as Google, Bing, and Yahoo, supplemented by other information search portals like Wikipedia, have become the means for searching information on the Internet. Along with the increasing popularity of search engines, the academic interest in search has shifted from analysing simple look-up query and response patterns to analysing rather complex information seeking needs. Current search tools seem to support complex search not as well as they do in the case of look-up. Especially support for aggregating search results from multiple search-queries, taking into account discoveries made during a complex search task, and synthesizing them to some newly compiled document of information is only at the beginning and motivates researchers to develop new tools for supporting those information seeking techniques. We focus in this work on the exploratory search concepts aggregation, discovery, and synthesis. Our idea is that these are today the most time consuming activities, especially when fulfilling a complex information need. We will use these three concepts throughout this paper to evaluate different approaches in exploratory search and give an overview of the state of the art and current ongoing research in respect to these concepts.

Key words: exploratory search, complex search, aggregation, discovery, synthesis, graph visualization, data exploration.

1. INTRODUCTION

According to Maslow's hierarchy of needs [1], information seeking is a fundamental human activity. However, information overload and information duplication impose a growing problem upon our knowledge societies. The quest for efficient information search tools is more relevant than ever. Today's search systems are designed to follow the "query–response" or shortly look-up concept. Users with an information need enter queries into search systems and those search systems produce ranked lists of search results. Ideally those search results are relevant for the queries used [69]. They are among the most basic types of search tasks, usually happening in the context of question answering and fact finding. Queries can be classified into mainly navigational (looking for a web address), informational (looking for a piece of information) and transactional (looking to carry out a transaction on the web like buying a certain pair of Nike shoes queries [12]. We intentionally do not focus on a specific type of query in this paper. As this paper deals with discovery, aggregation, and synthesis of information, the majority of the queries in such information search processes would be classified as "informational". A query-based approach is usually insufficient for the tasks in focus: *complex search tasks*. Especially aggregating search results from multiple search-queries, taking into account discoveries made in a complex search task, and synthesizing them to some newly compiled document of information or as a mental idea are poorly supported.

We will now first define what complex search and complex search tasks are. We will give a scenario of complex search, which will serve as a base for imagination and clarification. Then, we will present a selection of state of the art methods and concepts addressing exploratory search. We will use the concepts of aggregation, discovery, and synthesis to compare these. The next section will discuss the ongoing research and possible future directions of exploratory search, again with a focus on aggregation, discovery, and synthesis. We will finish this paper with a conclusion discussing the exploratory search evaluation based on these aforementioned concepts.

* Corresponding author, dmitri.danilov@ut.ee

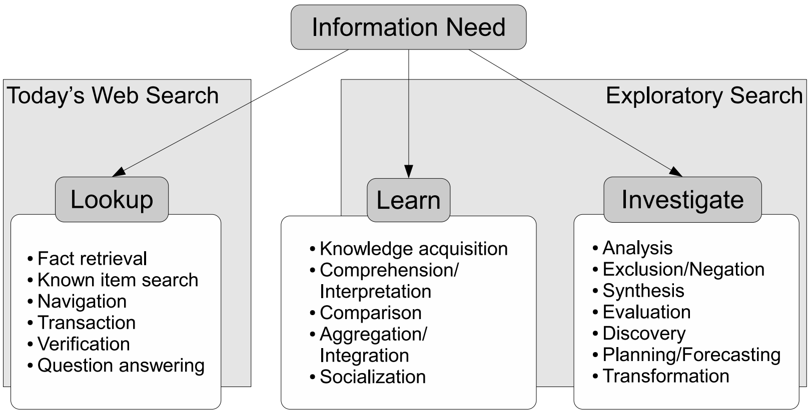

Fig. 1. Types of information needs [44].

The area dealing with *complex search tasks* is coined by Marchionini in *Exploratory Search* (*ES*) [44]. A problem with the definition of Exploratory Search is that it is mainly defined as encompassing all aspects of search not covered by current search. As depicted in Fig. 1, in 2006 Marchionini defines Exploratory Search basically as searching in a way which is not supported by "Today's Web Search". As web search is of course a very fast developing area, this is a rather unstable definition. Therefore we want to focus in this work on the exploratory search concepts aggregation, discovery, and synthesis instead. Our idea is that these are the most time consuming activities when fulfilling an information need. We will use these three concepts throughout this paper to evaluate different approaches in exploratory search and give an overview of established tools to support exploratory search and current ongoing research in respect to these concepts.

It is important to mention that the field of interactive information retrieval (IIR) [55], which also focuses on a longer term search process and tries to especially focus on breaking the separation of a search into a user and a system, exists already longer than Exploratory Search. IIR is very close to Exploratory Search, however Exploratory Search focuses stronger on the discovery aspect in Search. Synthesis and aggregation are not explicitly covered in IIR.

For further discussion, we have to define the terms *search task*, *search*, *complex search*, *complex search task*, *aggregation*, *discovery*, and *synthesis*. A *search task* expresses a goal and a starting point for the fulfilment of a specific information need. An example could be "the search for Wolfgang Amadeus Mozart's birthday". Of course, typing in "Mozart's birthday" on any current search engine will yield the user the right answer in the first listed hit. This is definitely not a complex task. *Complex search* on the contrary is a multi-step and time consuming process that requires multiple queries, scanning through many documents, and extracting and compiling information from multiple sources. An example here would be the search for all information relevant for assessing the safety situation in Afghanistan or an open task like finding the difference between the two composers Mozart and Bach. We will refer to the actual execution or execution process of the search task, independent of fulfilling the goal to satisfy the stated information need or not, as *search*. A *complex search task* is the corresponding search task that leads to a *complex search* activity. A (complex) search task is therefore the description in contrast to (complex) search itself, which is an interactive process. The search task is the motivator of the search (process). This does not mean that the actual search task cannot change during a search. An example for a complex search task would be this description: "Find the best universities for your child wanting to study either architecture or political science, assuming you live in Germany and are able to support your child with 1500 EUR per month" under an assumption that there is not yet a search engine that just takes the social profile of a parent and child and their preferences and monetary possibilities as input and generates the corresponding result for this. Nowadays, a search for such a task will still involve a lot of *aggregation* of different information sources, *discovering* lots of new facts of what is important in such a search (like understanding that accommodation is a very important cost factor to account for apart from tuition in our university example), and the need to *synthesize* all this information in a manner allowing making a decision. A more detailed definition of aggregation, discovery, and synthesis will follow below.

For studying search tasks, we used the following tasks in our experiments:

1. Please find the date when Wolfgang Amadeus Mozart was born.
2. What is the corresponding value of 20 000 Estonian kroon in USD?
3. When was penicillin invented and by whom?
4. How do you plan a budget but still nice wedding; write down the 20 most important points to consider.

5. Your daughter Anna is to graduate from high school and wants to study abroad in the field of either political sciences or architecture. You can support her with 1500 euro per month. Which universities can you recommend that are affordable and will offer the best career perspectives after graduating? Please compile a list of 10 universities.
6. Mr. Johnson and his family are planning a trip to Paris. Please write a summary of the following: What cultural events are there (for example in August 2010)? What sightseeing to do?
7. Barbara has been offered a very well paid position in Kabul, Afghanistan. She wants to know how risky it is to live in Kabul at the moment. Compile enough information to have a good indication for the risk.

Cases 1 and 5 are the examples that we used to motivate complexity of search tasks. Tasks 1–3 are clearly easy to solve with existing search engines. Also task 4 might be very simple to solve with the right query. Tasks 5–7 are definitely still very complex and time consuming. We have proven this complexity in [59]. Why do we think that tasks 1–3 are easy and tasks 5–7 are difficult? This has to do mainly with the time spent on these tasks. The first tasks took none of the contestants longer than four minutes. However, tasks 5 to 7 took all the contestants significantly more time. Especially the time for task 5 summed often up to several hours. As to tasks 4–7, all require the contestants to compile down some kind of result. It seems that this compiling of results is very time consuming and not so well supported. This is the reason why we are especially interested in examining the support for this compiling. When taking a look at Fig. 1, we can see that only aggregation, discovery, and synthesis are relevant for such a compilation.

We will understand *aggregation* as the support for selecting, storing, and accessing the relevant documents for a certain aspect of a search need. It also covers the possibility of accessing documents from a set of relevant documents interactively with for example bookmarking techniques. Aggregation finds documents in known categories. If we take another look at our best university search example, aggregation here means bookmarking different university ranking sites preferably for architecture and political science and corresponding sites with information about tuition cost. A quite related field of study here is Personal Information Management (PIM) [32], which also refers to aggregation activities: acquiring, creating, storing, organizing, maintaining, retrieving, using, and distributing the information. In spite of covering similar activities there are several differences. Aggregation also covers collaboration aspects and does not stick to only the personal information space. Aggregation focuses only on the document set that is relevant for a certain aspect of a search need.

Discovery means the support for finding new, relevant documents and at the same time extending the searcher's view of the topic towards previously unknown, related categories and concepts. The discovery step points the user towards interesting additional facts relevant for an information need based on the search already undertaken or selections made. In the university case discovery would offer study reports of other students in the cities selected, information about cost of living, and maybe also about security and public transportation.

Synthesis is the support for compiling multiple documents into one and extracting relevant information – this means also just parts – from the already found or selected documents. A key part of synthesis is building summaries, which comprises information fusion and compression [6]. We also will take a look at visualization of documents and information in general and its impact on these three. In the university search example, this would be at first a good overview of necessary ranking, tuition, and living information on the results page of the used search engine, but also support for compiling your own university ranking list, easily gathering snippets from other web pages and some automatic sort function with adaptable filters for such a personal ranking.

For further reference and imagination, we introduce one more scenario illustrating a complex search task. Imagine Brian and Sarah are journalists in the popular newspaper *FooTimes*. They are currently working on a political analysis article related to the conflict between North Korea and South Korea. Sarah and Brian decide to divide the work. Brian will search for related information and events that happened in North Korea during the last three-month period, while Sarah will search for events and facts which took place in South Korea. The two journalists want to dig out all conflict related events, the relationships among found events, and people that were involved in the conflicts, possible reasons of conflicts, and other related facts. In the resulting article, they want to describe and prove relationships between events and give an overview of the main influencing factors. Brian and Sarah have a week for this research task. The journalists have to read a lot of material in detail in order to find subtle relationships. They intend to bring out hypotheses on various relationships and show evidence to support them or show the rationale behind their hypotheses.

In the given case, the journalists need to make many search sessions, store the relevant articles and information sources, combine them, and research relations between different events. This complex search goes far beyond a simple look-up task and includes discovery, aggregation, and synthesis tasks. In addition to this, the journalists need to collaborate helping each other to focus on the important facts.

As a first step, most journalists will try to find an expert in the given area and have them interviewed. Let us assume here that Brian and Sarah do not have any contact and decide to make a simple overview of possible reasons of a conflict between the two countries. First, the journalists search for reliable information sources. They start from searching for articles on major news web portals like BBC, CNN, or Financial Times. Brian and Sarah also use search engines and services like Delicious to search for local newspapers and other public information sources like information on embassy web sites. During the search Brian and Sarah quickly examine information on the web pages and create bookmarks for the relevant ones. During that activity they make notes trying to remember the most important facts to later find relations between the information in the relevant documents. Every day Brian and Sarah arrange a small meeting where they share their findings, bookmarks, and most important – discuss the direction of the search for the next day. After several days, the journalists decide that they have found enough information. Their findings can now cover the major aspects of the future article. Starting from this point, the journalists mostly research their collected documents and notes trying to find the most influencing factors. Brian and Sarah print out documents and work together making notes in the text and discussing the influence of the found facts. They try to combine information by drawing schemes and discussing connections between facts relying on their cognitive abilities. After working out and agreeing on a set of ideas Brian and Sarah write their respective parts of the article.

Aggregation, discovery, and synthesis aspects are evident in this case. The work of the journalists starts with iterations of aggregation and discovery activities. Relevant articles and discoveries made during the process can be used to adjust the research direction. The complex search process then continues with collaborative work on an overview, which is aimed at discovering unknown relations between events. Current tools are not focusing on discovery of unknown facts based on a relevant document collections. This case demonstrates the need for such tools to facilitate discovery and support the user in hypothesis generation. There are two challenging tasks in that case. The first one covers the collaboration of journalists in the meetings. During the meetings, they discover important and poorly covered information and adjust the direction of the search for the next day. The second is the try to discover and prove ideas for their final article compilation.

After illustrating the flow of a complex search task with the help of a journalist research example, we will in the next section analyse established search tools regarding their support for aggregation, discovery, and synthesis.

2. ESTABLISHED TOOLS AND CONCEPTS TO SUPPORT EXPLORATORY SEARCH

In this section we will show that when trying to improve the support for Exploratory Search it makes sense to look at the three aspects: aggregation, discovery, and synthesis. We will also investigate how far current search systems support aggregation, discovery, and synthesis aspects of Exploratory Search. We will give examples and relate these to the tools covered by White and Roth [68].

White and Roth [68] highlight the following eight features that a search system must have in order to be used for exploratory search purposes: (1) enable understanding and learning, (2) querying and real time query refinement, (3) result filtering, (4) take advantage of the search context, (5) offer advanced visualizations, (6) new forms of collaboration, (7) histories and progress updates, (8) task management. In order to improve the support for Exploratory Search, one would try to improve the individual support for each of those features. A closer look at them reveals that those features are quite softly and vaguely defined. For example, both "enable learning" and "offer advanced visualizations" can mean a lot. They are difficult to measure. However, this would be an important prerequisite for improvement.

In order to get an understanding of the relationship between White and Roth's features and the three aspects we suggest, we juxtaposed the eight features of Exploratory Search systems with our three main search process steps aggregation, discovery, and synthesis and analysed in how far the support for aggregation, discovery, and synthesis can be derived from those eight features. For example, if a search system enables learning, this also means that it supports discovery (as defined in the introduction) because discovery is a major part of learning. Learning does not support aggregation and synthesis. Real time query refinement does support finding more documents for a known category and it might also point the user towards new, unknown categories. Results filtering functionality is clearly supporting the aggregation of documents for a known category, but offers little support for discovery and synthesis. If a system takes advantage of the search context, it will also use this additional information to help the user in finding new categories so far not known to them. Advanced visualizations that for example make neighbourhood relationships in information visible, also allow the user to discover new information categories. The last three features around collaboration, progress updates, and task management seem to mainly improve the efficiency of the search process in general. Yet it is thinkable that new forms of collaboration facilitate synthesis.

Table 1 summarizes our findings and shows that each of White and Roth's eight features to some extent supports either aggregation, or discovery, or synthesis or even all features at the same time. Yet, as already mentioned, when trying to improve the support for complex search tasks, starting with White and Roth's features does not seem to be the best choice as many of the features have an overlap, as in the case of "enable learning" and "offer advanced visualizations". Of course, offering better visualizations also improves learning. Considering that aggregation, discovery, and synthesis, as we have defined them, can be mapped to White and Roth's eight features of Exploratory Search systems (as outlined in Table 1), we hypothesize that improving them would automatically also enhance White and Roth's eight features. As the three concepts have

Table 1. Relevance of exploratory search features for aggregation, discovery, and synthesis

Exploratory Search Features	Support for		
	Aggregation	Discovery	Synthesis
Enable learning	no	yes	no
Offer real time query refinement	yes	yes	no
Offer result filtering	yes	no	no
Take advantage of search context	no	yes	no
Offer advanced visualizations	no	yes	no
Facilitate new forms of collaboration	yes	yes	yes
Offer histories and progress updates	no	yes	no
Task management	no	no	no

no mutual overlap and are measurable, focusing on them seems to be a more promising approach for any complex search improvement initiative.

After justifying our suggestion to focus on aggregation, discovery, and synthesis instead of White and Roth's eight features, we took current search tools and analysed them regarding their support for these three concepts. Some of the following examples were also partly used by White and Roth [68] in their analysis of current exploratory search systems. We updated the ones that are still available and supplemented them with additional examples such as the universal search paradigm or social search features, both recently adopted by commercial search engines as follows.

Most of the current commercial search engines offer dynamic query interfaces as depicted in Fig. 2. Once the user starts entering a query into the search window, the search engine automatically extends this query with additional terms based on internal statistical analysis of frequently entered search terms [33]. For instance, typing "north korea" into Google immediately offers to extend the query with "news, nuclear weapons, documentary, wiki" etc. as outlined in Fig. 2. Entering a query into a dynamic query interface helps users to discover more than they had known before. Users would

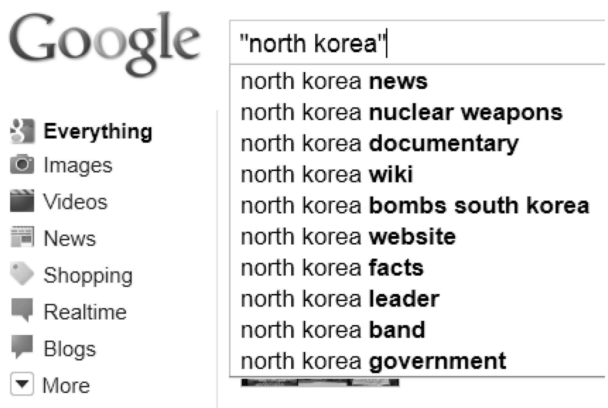

Fig. 2. Screenshot of the dynamic query extension feature in Google.

enter keywords that are already existing in their knowledge space based on prior experience and those would automatically be extended by the search engine based on the search statistics. Those query completion techniques have been proven to be useful for the user in early stages of their search [67]. Apart from aggregation, the query extension feature of current search engines supports the discovery aspect and can therefore be used to aggregate relevant documents for a certain aspect. This feature does not support synthesis.

Faceted browsing (also sometimes called *faceted search* or *faceted navigation*) allows users to explore a collection of documents based on a predefined classification. Each document is classified according to certain aspects (*facets*) that all objects in the collection have in common like the date of publication, author, or topic. Those facets can be used to rearrange the set of documents in multiple ways, enabling the user to discover unknown aspects in a certain domain. Wilson et al. [72] published a multicolumn-faceted browser called mSpace Explorer as depicted in Fig. 3. It is a client for exploring large data sets. The mSpace Explorer demo (accessible at http://research.mspace.fm) allows the user to browse a huge set of articles by choosing the date and the year an article was published and the theme and the subject as outlined in Fig. 3. We looked for all articles in the data set that were published in 2000 in the health area, and obtained four results. We might then also be interested in articles published in the 1920s and also deal with this topic. We can easily filter them out by simply changing the "decade" slider to 1920. The data set contains 43 articles dealing with health published during this time. The mSpace Explorer demo illustrates the advantage for the users. The result space categories are given to the user. All documents in the set are classified according to those categories. The user can systematically explore the data set.

The mSpace Explorer supports aggregation of documents in known categories and also discovery in unknown categories and unknown documents in those categories. It does not support the synthesis of a set of collected documents into a new entity. Taking our example again, we are given all the dimensions that we can use to explore the data set, like health, lifestyle,

Fig. 3. Screenshot of the mSpace Explorer.

politics, and science and technology. This means that all necessary aspects are in a way preloaded and the effort for discovery is minimized. However, when we look at the support for synthesis, the mSpace Explorer does not offer any automatic functionality. After having found our relevant set of documents we would then have to compile, e.g. a summary, or find relations by hand.

Web searching is usually done by information seekers alone. However, many of the search tasks that users are confronted with could be more efficiently carried out by working together with others. Collaborative search has recently been investigated by a number of researchers [3,45,47,60]. Their kind of search approach tries to take advantage of the collaborative nature of search by offering various forms of collaboration support tools. Searching in a team can be more effective than if done alone as stated by Golovchinsky et al. [23]. This is mainly due to the fact that people in a team bring in a variety of views, perspectives, ideas, experiences, and finally keywords to look for. Especially at the beginning of an exploratory search task the combined experience of many users helps to cover a search topic better and will lead to better retrieval results [5].

Microsoft's SearchTogether tool [46] as outlined in Fig. 4 is an Internet Explorer plug-in that allows groups to carry out web searches in a collaborative manner. Once people decide they will carry out a search together they can sign in to SearchTogether. The tool provides two views, the contacts view and the summary view. In the contacts view users can see each other and the search queries they are using during searches. In the summary view they can share relevant documents that they have found and also add comments.

The SearchTogether tool supports discovery of new content by its team work functionality and also assists in aggregating relevant documents in known categories. Other team members benefit from what their colleagues find. The SearchTogether does not support synthesis.

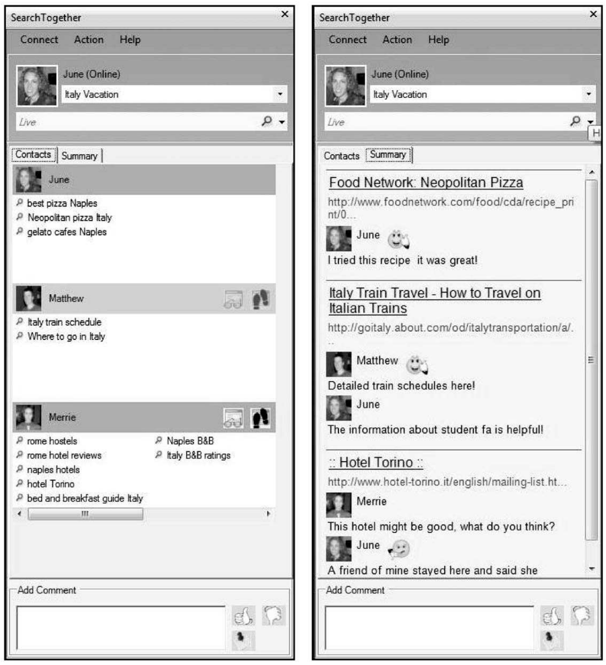

Fig. 4. Screenshot of Microsoft's SearchTogether tool.

Despite the fact that search is usually considered to be an activity people carry out on an individual basis, the individual can benefit from the wisdom of the crowd. Evans and Chi [21] define social search as follows:

> "Social search" is an umbrella term used to describe search acts that make use of social interactions with others. These interactions may be explicit or implicit co-located or remote, synchronous or asynchronous.

They present a canonical model of user activities during search and show where shared information can be beneficial to individual searchers. Also Freyne et al. [22] demonstrate that an integration of social search and social navigation allows users to combine their search and browsing behaviours and provides a significant advantage for users. Carmel et al. [13] have managed to prove that social search based on users' social relations outperforms non-personalized social search. Recently Google started to provide additional information from social connections as outlined in Fig. 5. Based on a query, relevant search results are retrieved by Google. If a certain organic search result was already shared on Twitter by somebody in a person's Twitter network, this person is added to the corresponding search result [60]. As illustrated in Fig. 5, the results page for the query "cloud computing economics" shows that this result was also posted on Twitter at the end.

This social search feature supports finding additional content to a known category. Integrating social aspects into organic search results does not support synthesis.

Universal search is commonly referred as "contextually enriching generic web search results with results from other data sources" [54]. Big search engines like Google or Bing have started to add more exploratory search features to their technology. The universal search model, first revealed in 2007, integrates document surrogates, videos, and images in the first search result pages and therefore offers a more complete and comprehensive user experience [62]; for example, the query for "Barack Obama" in Bing as outlined in Fig. 6 yields a results page showing the Wikipedia entry as the first result, followed by some images of Barack Obama, then the first organic search results, accompanied by suggestions of related searches for Obama speeches and other topics.

The support of universal search for aggregation is high as it shows the user even different media in addition to organic search results. It does not support users in discovering new categories or help in synthesizing new information.

After having given an overview of the support of the concepts of aggregation, discovery, and synthesis in current established search tools, we summarize our findings in Table 2 and compare these to the functionality of standard web search interfaces. Our analysis shows that aggregation of information is supported to some extent by all systems that we investigated. Aggregation means finding relevant documents in a known category and was already supported by the classic search interfaces of Google and Bing. Also dynamic query interfaces, faceted browsing, collaborative search tools, social search, and universal search interfaces do support aggregation. Yet the question that remains is: How good is the support and what could be improved?

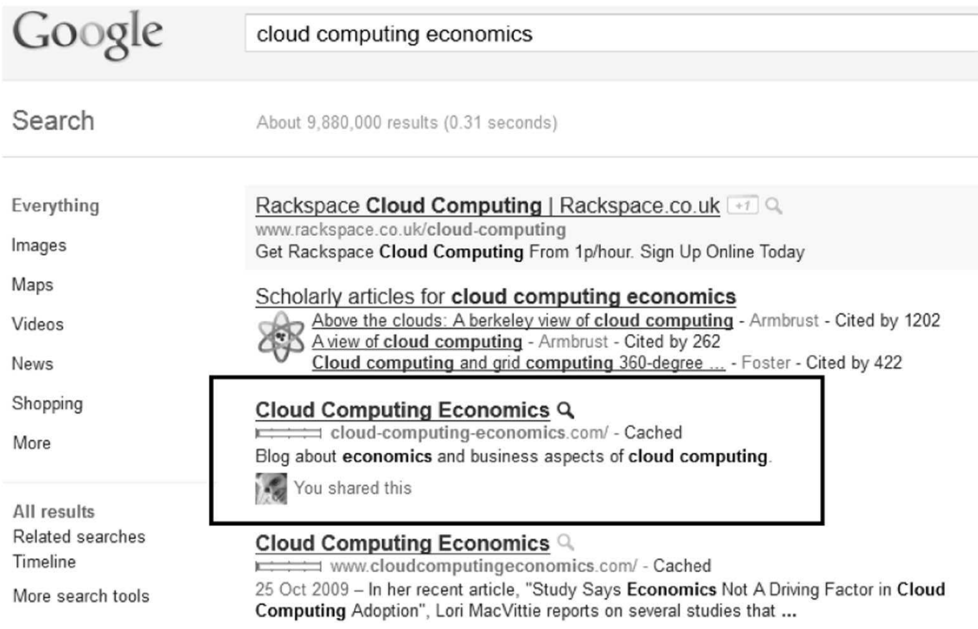

Fig. 5. Screenshot of the social search feature in Google.

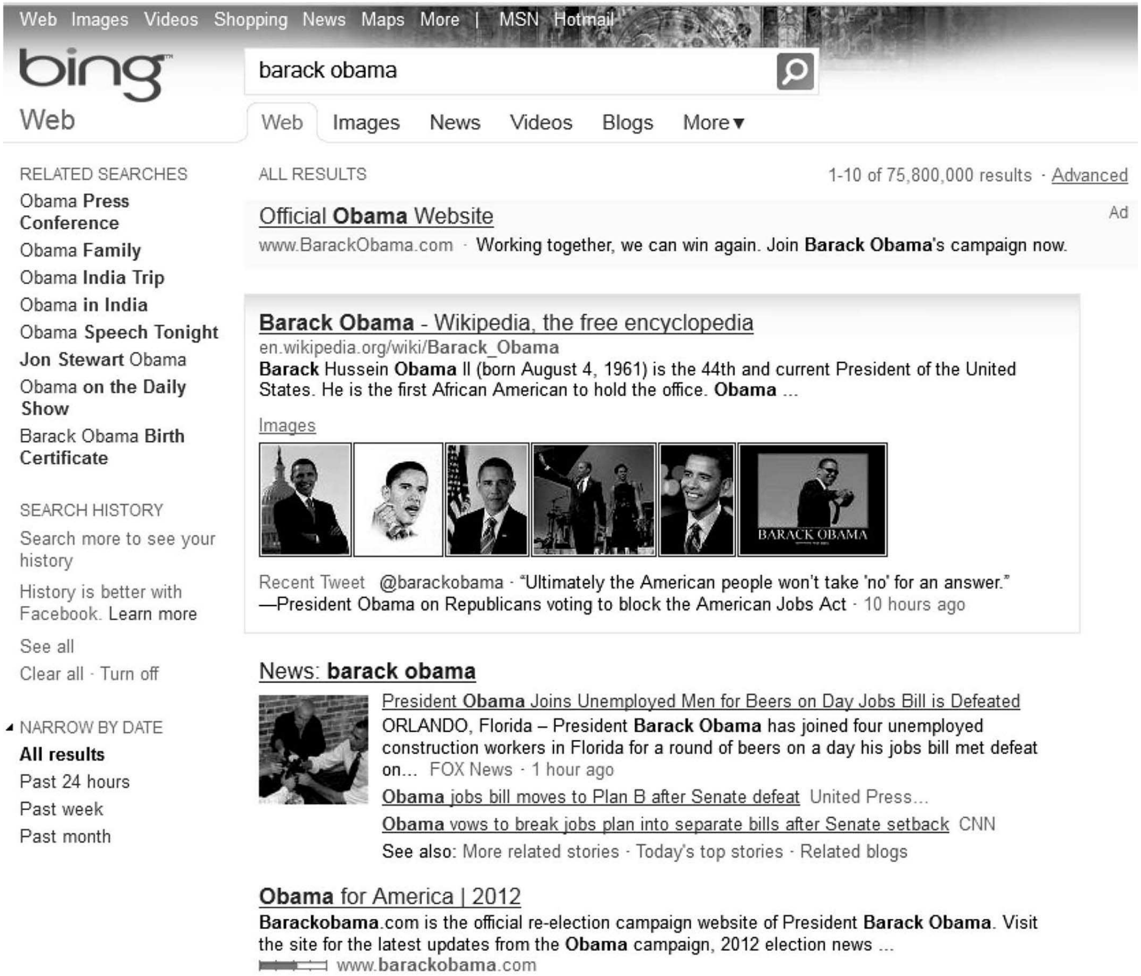

Fig. 6. Screenshot of a universal search feature in Bing.

Table 2. Support of aggregation, discovery, and synthesis in current search systems

Search system	Aggregation	Discovery	Synthesis
Standard web search interfaces	yes	no	no
Dynamic query interfaces	yes	yes	no
Faceted browsing	yes	yes	no
Collaborative search tools	yes	yes	no
Social search	yes	no	no
Universal search interfaces	yes	yes	no

Discovery is not supported in standard web search interfaces, as those only allow users to enter keywords that they already know. First support for information discovery is given by dynamic query interfaces, collaborative search tools, faceted browsing, and universal search interfaces. Discovery means learning unknown aspects in a new domain. Faceted browsing supports discovery as all the documents in a repository are classified by hand along various relevant dimensions. Hence the important aspects and concepts of a certain domain are made evident and can be explored accordingly. Of course faceted browsing requires a significant amount of human input in advance in order to make a set of documents searchable this way. Discovery can either be done automatically (where the dynamic query interfaces use query statistics to extend queries) or it can be achieved by enabling people to collaborate. As people have different knowledge about a domain, "short circuiting" their brains will help the involved information seekers to extend their knowledge spaces. Universal search interfaces enable discovery insofar as they present a selection of different media (document surrogates, images, and videos) to the reader and hence increase the probability of discovering new aspects.

When it comes to synthesis, our survey showed that this aspect of exploratory search is very poorly supported in present search systems. This task also has not received enough attention from the scientific community.

It needs to be stated at this point that the support for aggregation and discovery of current established search systems is just taking root and a lot of room for improvement is still evident. After having investigated the support for aggregation, discovery, and synthesis in established search systems, in the next section we want to present a selection of ongoing research that is being carried out in this area.

3. ONGOING RESEARCH

This section will give an overview of the research that is being carried out in the area of exploratory search and the respective support of aggregation, discovery, and synthesis.

3.1. Aggregation

In this section, we will give an overview of the aggregation aspects in current research related to web search. The term *information aggregation* is mainly used in the economics domain for the idea that markets have the ability to aggregate and reveal dispersed information [51]. Aggregation of information in the context of search can be described by activities like selecting, storing, and maintaining data objects. However, when we describe complex search task evaluation in the Exploratory Search domain we need to add an important property for the retrieved and stored document collection: documents should be relevant to the information need. We want to especially highlight that aggregated documents can be irrelevant to the initial task, but can still be relevant to the information need as the information need is static and the initial task can change during the search because it is highly dependent on new information discoveries. It is important to keep the relevance property in mind when we look at applications in the area of social bookmarking. One of the best examples is social bookmarking services like del.icio.us [28,37], which provide users with already aggregated collections relevant to a certain topic. These collections can become a great help if the topic corresponds to the search task of the search session [73]. The idea behind the social bookmarking services is simple and relies on the fact that an information need can be shared among a group of users and thus support implicit collaboration [64]. Recent research concerning the collaborative aspect of search [52] states that collaborative search can outperform post hoc merging of similarly instrumented single user runs. Thereby, aggregation of the relevant information objects and further activities related to the data maintenance are likely to have collaboration support during the whole process, and social bookmarking is one of the example realizations of this concept.

Another successful information aggregation application is the web project Twitter [31] where participants can receive textual information from different sources including already aggregated information of other participants. Recent research [50] on the Exploratory Search application TweetMotif is a good example of the presented classification. TweetMotif can be considered as a try to utilize Twitter as an information aggregation platform and perform discovery tasks based on it.

Social networks like Facebook undoubtedly pose a growing source of information. For example, looking at the news feed section of Facebook, all the posts and posted media can be seen as an aggregation process. However, just taking Facebook's news stream is not aggregation as we understand it as the posts and posted media lack the relevance for a certain topic. Hong et al. [30] published a tool called FeedWinnower, which is an enhanced feed aggregator to help people filter feeds as the aforementioned one. Users can select topics, people, sources, and time. Use of FeedWinnower would then turn a feed of potentially relevant documents into relevant ones and hence can be seen as implicit collaboration in our narrower sense.

A quite related model that should be mentioned here is called berrypicking [10]. It describes techniques where users move through a variety of sources, and each new piece of information gives them new directions to follow as well as ideas for their next queries. The term berrypicking is used as an analogy to picking blueberries in the forest where berries are scattered on the bushes. This model can be used to describe the iterations of aggregation and discovery activities for a specific information need and corresponding evolving search task. Each such iteration consists of relevant data aggregation and discovery tasks performed using the aggregated relevant set. The search task can change after performing the discovery tasks and documents can become irrelevant. However, the influence of discoveries to the search task evolution does not necessarily mean rejection from all previously aggregated documents. Some documents can still be relevant and the aggregated relevant set can include documents from different iterations. Thus, berrypicking demonstrates that the aggregation activity is driven by discoveries and search task evolution.

3.2. Discovery

In this subsection we will take a close look at current research supporting information discovery. "Information discovery" is referred to generating new ideas while browsing relevant information [35]. Information discovery is one of the key elements of information search, on and off the web. Due to the commercial success of current search engines it is difficult for their users to judge what search needs they should be used for and where their limits are. Schraefel [57] gives a list of approaches that support more complex search and knowledge building, which contains amongst others "knowledge building by association" (connecting formerly disconnected domains), "annotations and notes" (context dependent additional information for documents and document collections), "history review" (allows following the paths an information seeker has taken), and "collaborative knowledge building" (sharing progress in brainstorming and search sessions).

All these approaches do not follow the current search paradigm, which is tuned towards optimizing precision, offering a set of most relevant documents for each query. As Crimmins et al. [16] already diagnosed in 1999, "it is impossible to achieve both high precision and high recall in a Web search because of the size and the diversity of the Web's contents". Since then the size of the web has multiplied. A very specific aspect of discovering information relates to information seekers entering a domain new to them. With no or little domain knowledge such an information seeker is at risk of not taking very relevant aspects of a search problem into consideration [57]. There are different approaches to support the discovery of formerly unknown aspects, some are automatic approaches, others are human-supported ones. El-Arini and Guestrin [20] offer an automated approach for discovering additional aspects to a specific information seeking problem. The information seeker can take a set of relevant papers and use this set as a "query". The system will then suggest additional papers that "the information seeker should also read". Koh et al. [36] developed a system to support creativity in learning by enabling information discovery and exploratory search. With their system called "combinFormation" users are able to make clippings of found documents, arrange them on the screen, and also add conceptual relationships.

Social search is a human-supported way of discovering unknown aspects for a search problem. There is one group of approaches that tries to use social annotations [7] such as bookmarks from bookmarking sites like del.icio.us to enhance the ranking algorithm. Those approaches fall into the classical search engine paradigm that tries to find "the" optimal document for a single query and hence does not help much on the discovery side.

The second group of approaches tries to take advantage of the collaborative nature of search [3,45,60] by offering various forms of collaboration support tools. One such approach is a commercial offering called Haystack.com, which enables users to easily share documents that they have found during their searches, re-find documents that they had already found before, and organize their searches better [60]. Another, a more academic approach, is called CoSearch [3]. This approach uses common standard hardware and enables users to work together at a single computer with multiple mice attached. CoSearch comes with a special colour coded user interface. Each user can click search results which are then put into a page queue. SocialSearchBrowser [14] is a tool that tries to turn mobile devices into social search and information discovery tools. It enables users to share queries and interactions with peers and ask them questions, filters what queries are displayed based on the level of friendship, and maps questions and answers on a Google map interface.

There is one more group of approaches based on the advanced information visualizations that can be personal and collaborative. One of the approaches is visualization of document relations using graphs. The main feature of graphs is their focus on the information relationship. Sometimes it becomes an important shift in the data representation to spot unknown facts. It can be seen that certain search tasks are solved better with graph-based representation using TouchGraph [19] (see Fig. 7). Exploratory search information visualization focuses on the visual representation of large data collections to help people understand and analyse data [68]. In graph-based representation this property can be implemented by graph clustering and 3D graph visualization [71]. In many cases 3D graphs can outperform 2D graphs [38]. The 3D Graph Explorer (3DGE, accessible at http://graphexplorer.ulno.net) represents a 3D application that visualizes the directed graph with a 2.5D layout (see [29] for a 2.5D layout example) and provides a possibility of exploring the graph data (see Fig. 8). This application offers 3D graph-based visualization to explore data relations to support discovery tasks as well as hypothesis generation.

According to the list of exploratory search system features (see Section 2 above and Chapter 4 in [68]), the software should offer visualizations to support insight and decision making: systems should present customiz-

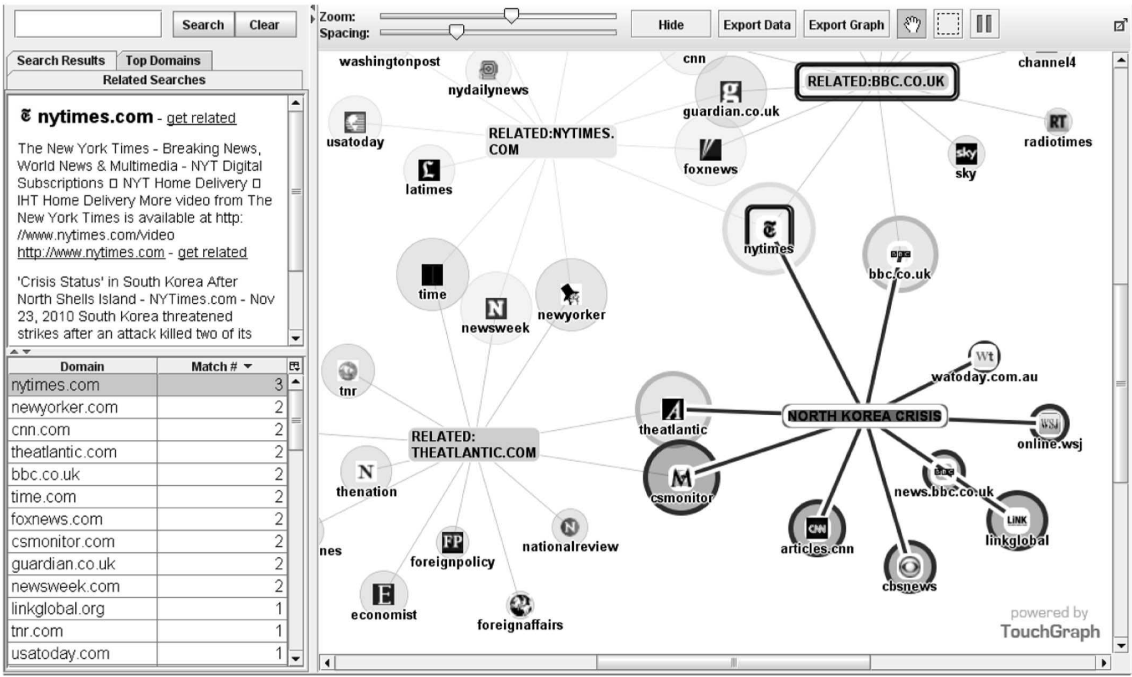

Fig. 7. Screenshot of TouchGraph.

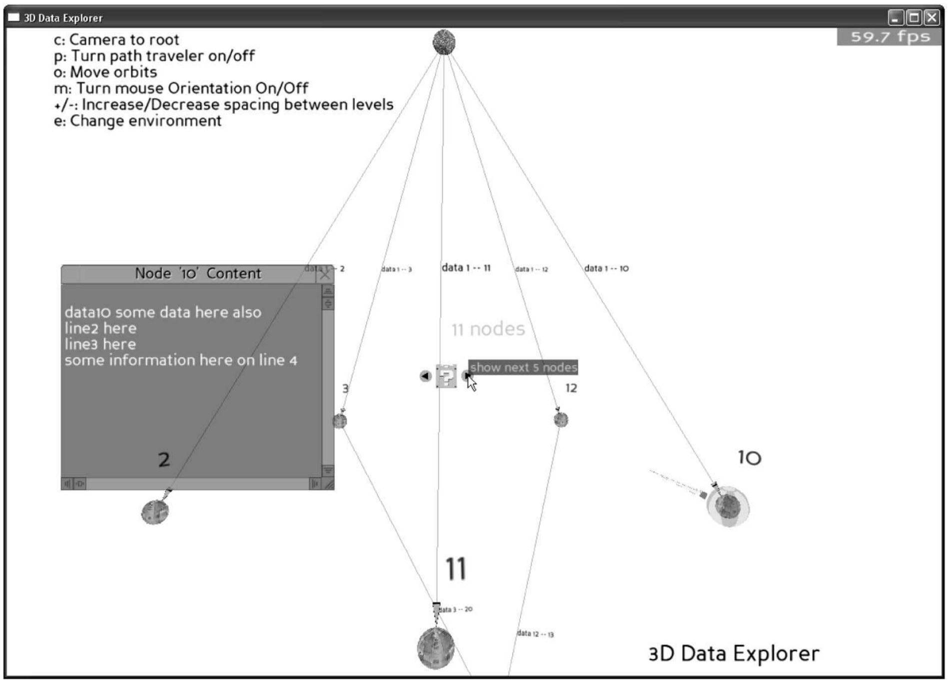

Fig. 8. Screenshot of the 3D Graph Explorer prototype.

able visual representations of the collection being explored to support hypotheses generation and trend spotting. Drawing graphs is an efficient technique to explore complex object relations and thus helps users to generate hypotheses based on the explored relations. Thus, graph-based representation is helpful for discovery tasks.

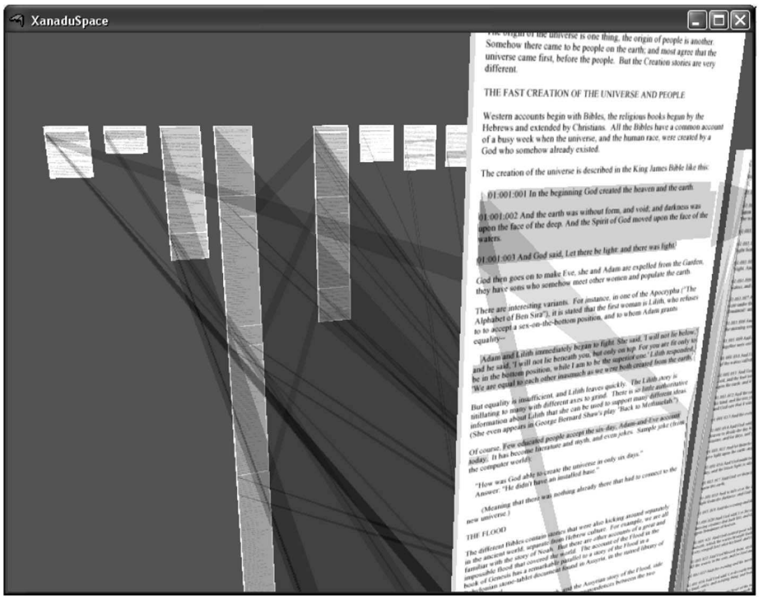

Fig. 9. Screenshot of XanaduSpace application.

Graphs can also be used as a data structure. One of the most famous projects that employs graph-based structuring for document storing is the Xanadu project, which was founded in 1960 by Ted Nelson (project info is accessible at http://www.xanadu.com/). The Xanadu project proposes a system with a way for people to store information not as individual "files" but as a connected literature with an ownership identification for every piece of created data (see screenshot of demo application in Fig. 9). However, the project was not as successful as hoped and lost competition with the World Wide Web[1].

3.3. Synthesis

Synthesis is commonly referred as combining parts of separate items into one single and new entity. A definition which fits our view is given by Amigó et al. [4]. They define synthesis as "the process of (given a complex information need) extracting, organizing and inter-relating the pieces of information contained in a set of relevant documents, in order to obtain a comprehensive, non redundant report that satisfies the information need." In their report they compare summaries of several documents created by humans with automatic summary generation. They prove that the information synthesis task is not a standard multi-document summarization problem.

However, summarization is currently the most tackled synthesis technique in research. Nenkova et al. [49] state three main tasks in a summarization process, which are (1) content selection; (2) information ordering; and (3) automatic editing, information fusion, and compression [6]. Regarding content selection, a key challenge is to find an answer to the question: What is a document about? In order to find the topic of a content, researchers mainly apply mathematical models, amongst others word frequency models [42,48], lexical chains [8,56], latent semantic analysis [24,26,61], and content models [27]. Qazvinian and Radev [53] published a novel approach to summarize scientific articles. Their automatic clustering method allows generating common themes among a subset of research papers that are co-cited.

As far as information ordering is concerned, it has not received enough attention from the scientific community. The challenge is to organize the selected information so that text summaries are still coherently readable. Barzilay and Elhadad [9] investigated a multi-document summarization scenario in the news genre. They tested the chronological ordering algorithm against a majority ordering algorithm (keep the order of facts as it is in the majority of the documents) and found that majority ordering is the more promising strategy.

Research papers do not currently cover other aspects of synthesis like conclusions drawn or prior knowledge of the candidates integrated into a synthesis report. Using different visualization techniques like graphs or diagrams for synthesizing information is not well studied as yet either.

3.4. Metrics and evaluation

This section (based on [59]) gives an overview of established evaluation methods in the areas of classical search engine performance measurement, user search

[1] Or to be more precise, to developments taking place in computer history from the 1960s to today.

experience research, and Exploratory Search. It highlights important usability metrics for aggregation, discovery, and synthesis. It also discusses evaluation suggestions for complex search tasks.

The study methods used in information retrieval (IR) can be divided into two main groups: juror based studies and click-through data based studies. In juror based study methods to measure the retrieval effectiveness of IR systems, a static set of documents plus a set of queries are taken and the respective results are evaluated by jurors (see [39]). This so-called Cranfield methodology [15] is widely used to evaluate IR systems. The Text Retrieval Conference (TREC) chose it as the main paradigm for its activities of large-scale evaluation of retrieval technologies [66]. Many of the measures used in established methods to evaluate search engines and their performance on a technical level, like relevance and recall (see [40]), leave out important user-related aspects.

There are different established ways to capture user experience. Conducting studies by applying user survey techniques is a relatively easy way to gather user feedback – demographics about the users are available and asking detailed questions is possible. Laboratory experiments are commonly used to capture user experience to the very detail. The experiments are carried out in a laboratory setting in the presence of an interviewer. Although laboratory studies and also user surveys have the advantage of comprehensively capturing the user experience and additional information about the users is available (like demographics), both impose the risk of a user trying to look better in a laboratory environment (depending on the formulation of the questions and the interviewer). In addition, as laboratory experiments are costly, sample sizes are often small. Hence they are mostly not sufficiently representative. Complementary approaches to laboratory experiments are log file analysis and live ticker inquiry analysis. Search engines provide huge amounts of data out of their everyday operations and hence sample sizes are sufficiently large. However, other problems arise: log files are anonymous data files that lack any kind of user feedback or additional demographic information about the users like gender, age, or profession, or the beginning, end, or nature of a search [25,58].

New web-specific measures to better reflect current search trends have been developed. Yet those are still limited in giving a comprehensive quality indication of present web search services [17,43,65]. A search engine quality measurement framework that reflects the system-centric approach and also the user-centric approach by taking into consideration index quality, quality of the results, quality of search features, and search engine usability was proposed by Lewandowski and Höchstötter [41]. Their main argument is that for measuring the user experience empirical studies are indispensable. While IR mainly follows the query–response paradigm (considering one single query and its resulting documents at a time), interactive information retrieval (IIR) [55] focuses on the search process, and exploratory search [68] comprises more complex activities such as investigating, evaluating, comparing, and synthesizing. Especially for such activities researchers are striving to integrate the user into the measuring process as can be seen in the TREC Interactive Track [18] or TREC High Accuracy Retrieval of Documents Track [2]. The main issue so far is the repeatability and comparability of experiments on different sites. In an exploratory search context, covering as many aspects of a topic as possible is equally important to optimizing for relevance in classical IR, where precision is the main performance measure [63]. In order to carry out evaluation experiments, researchers have suggested developing special measures for exploratory search systems [11]. The following measures were identified as being appropriate [70]: (1) engagement and enjoyment, (2) information novelty, (3) task success, (4) task time, and (5) learning and cognition.

Finally it is important to mention that time is an especially important aspect not to be neglected when trying to evaluate exploratory search systems. Because exploratory search sessions dealing with exploratory search tasks can span over days and weeks, long-term studies are essential in order to get realistic study results [34]. As measuring exploratory search tasks is a more complicated endeavour than measuring search at the query level [34], Singer et al. [59] present a new evaluation method and a corresponding tool for evaluating exploratory search tasks. Their user-centred approach tries to narrow the gap between evaluation methods purely focused on technical aspects and expensive laboratory methods. The Search-Logger tool can be used to carry out task-based user studies independent of a laboratory environment. It collects implicit user information by logging a number of significant user events. Explicit information is gathered via user feedback in the form of questionnaires before and after each search task. The method tries to answer the following research questions (amongst others): What are characteristics of exploratory search tasks? Is it possible to measure the complexity of an exploratory search task and if yes, how? How much do current search tools support exploratory search? What is the user experience when carrying out exploratory search tasks with current means? Can information needs be classified in terms of search tool support?

When taking into account the aforementioned performance evaluation methods and reviewing their applicability for establishing metrics for aggregation, discovery, and synthesis, we suggest that a complex task evaluation be divided into three sub-tasks and propose that separate measurements be performed on aggregation, discovery, and synthesis aspects of the application. Powerful aggregation tools do not necessarily lead users to discoveries just like hundreds of relevant bookmarks do not automatically give an overview of the complex problem. Similarly, the discovered facts are not automatically ordered and compressed in a meaningful way, which is performed during the synthesis task. With the classification we present in this article, we propose that these concepts be separated during the evaluation of applications supporting complex search tasks to achieve more consistent results. As for aggregation tasks document surrogates are essential, graphs usually do not present enough information to be a helpful tool for aggregating relevant documents. The same does not hold true for discovery tasks. Here graphs enable the searcher to find unknown relations between relevant documents already found.

For all three aspects (aggregation, discovery, and synthesis) we consider using formative usability studies to focus on application design issues and to generate recommendations for improving their support in search tools. The overall task effort, time, and performance are important for the aggregation tasks. The same usability metrics can be used for the discovery tasks, but the most important measure would be the efficiency of cognition and learning of unknown facts in the relevant information set. Another important usability metric is user satisfaction in terms of how well the discovery tools were able to inform the user about the aggregated documents. The first two sub-tasks, aggregation and discovery, are iterative, meaning that the user may return to the aggregation activity after discovering a new fact. Thus, we need to consider this behaviour as a standard work-flow during the task progression. In this case it could be important to measure how fast and how many relevant documents the user was able to cover with discovery tools during one iteration. Synthesis is usually the final task, when users unite their findings and compile the resulting document. It can be omitted and depends on the search task, because the information need does not necessarily require a resulting document compilation. Thus, the synthesis task evaluation is more straightforward, but depends on the application's aggregation features, that is how well the application supports manipulating aggregated document parts and discovered (e.g. images from the visual discovery tools) information. In the case of synthesizing texts, we can again use the overall effort, task time, and quality metrics to evaluate the efficiency of this summary step.

4. CONCLUSION

In this paper we present a focused view on the area of Exploratory Search. Exploratory Search combines all research on search tasks which are weakly covered by the classical query–response paradigm. It claims to especially cover open-ended complex search processes, discovery, and learning. As this definition encompasses a very vast and quickly changing field due to the fast progress in classical information retrieval, we are focusing in this paper on measurable and time consuming aspects of complex search: aggregation, discovery, and synthesis. The roots of the presented aspects go back to the work by Marchionini, who coined the term Exploratory Search in 2006 [44].

The present paper takes the three concepts aggregation, discovery, and synthesis, defines them, and argues that actually they are the main steps in the process that Marchionini calls Exploratory Search. It does so by proving that those three concepts can be mapped to Marchionini's features of Exploratory Search systems, by analysing how well those aforementioned three concepts are supported in established search tools, and by covering ongoing research in the respective areas. It uses a show case example (from the journalism domain) and clearly points out the aggregation, discovery, and synthesis aspects to highlight the lack of their support. Further on, it tries to motivate the importance of the concept of Exploratory Search, but as its definition is vague, sets the focus on aggregation, discovery, and synthesis as the main pillars of complex search. This gives a clearer view on the challenges in this area of search. When analysing current search tools with regard to their support for aggregation, discovery, and synthesis, our main findings were that the concept of information aggregation is supported to some extent by all systems that we investigated. Aggregation means finding relevant documents in a known category and was already supported by the classical search interfaces of Google and Bing. Also dynamic query interfaces, faceted browsing, collaborative search tools, social search, and universal search interfaces do support aggregation. However, we believe that even the support for aggregation is just in its initial stage and could theoretically serve users better (like offering users ways to remember and interrelate what documents they have aggregated in a search process). As to discovery, we found that it is nearly not supported in standard web search interfaces and that first support for information discovery is given by dynamic query interfaces, collaborative search tools, faceted browsing, and universal search interfaces. Discovery is only supported on a very low level, and we believe that for example graph visualization techniques will tremendously impact the field of information discovery in the middle and long term. When it comes to synthesis, our survey showed that this aspect of Exploratory Search is not supported at all in the present search systems. Information synthesis comprises amongst others the steps of information ordering, automatic editing, information fusion, and information compression. All these aspects require heavy use of artificial intelligence and research in those areas has not progressed far enough to come up with commercially usable solutions.

Our review of the ongoing research in the respective areas shows that regarding improving the support for aggregation, promising approaches are coming from the areas of social search and collaborative search. Also on the discovery side, domain-specific solutions (e.g. for repositories of research papers) and collaborative as well as graph-based approaches have produced interesting prototype solutions. Synthesis, as the least supported feature, is tackled by researchers by applying mathematical models (e.g. frequency models, lexical chains, and content models), and some interesting articles on automatic summary building seem to point into the right direction.

We hope that the analysis of aggregation, discovery, and synthesis in complex search made in this paper will focus researchers on those concepts more when trying to analyse complex search tasks and can offer them a more concrete, better measurable, and less vaguely defined research framework than the concept of Exploratory Search does.

ACKNOWLEDGEMENTS

This research was supported by the Estonian Information Technology Foundation (EITSA) and the Tiger University programme.

Georg Singer is a researcher and lecturer at the University of Tartu in Estonia. He is doing his PhD in Informatics, working on improving the support for complex search tasks in commercial Web search engines. Prior to coming to Estonia he held various management positions in Austria, Germany, and the UK. Georg is also active as an entrepreneur. He has an MBA in Entrepreneurship from the Danube University, Austria, and an MSc in Theoretical Physics from the University of Vienna, Austria.

Dmitri Danilov is a PhD student of Informatics at the University of Tartu, Estonia. He is doing his PhD in the area of web search, where he researches different discovery techniques as well as visualization solutions to improve support of complex search tasks. During his MSc research in the web search he created a 3D Graph Explorer data visualization tool to examine 3D graph based approach to complex data visualizations. He has a 4-year experience in the US electronical medical record software (EMR software provided by Raintree Systems, Inc.), especially in the HL7 interface solutions.

Ulrich Norbisrath is a senior research fellow at the University of Tartu, Estonia. He has been working in Estonia since the end of 2006. He teaches courses on agile and visual software development and systems administration. He is involved in research projects from the areas of Cloud Computing, Distributed Computing, and Knowledge Management. He invented the concept of Friend-to-Friend (F2F) Computing (http://f2f.ulno.net) and is one of the founders of the Search-Logger project (http://www.search-logger. com). He is currently working on cloudsourcing cost prediction, advancing his F2F Computing technology, and researching new ways of exploratory search. In his PhD (2006) at the RWTH Aachen University (Aachen, Germany) he researched a way to enable an easy deployment of smart home services in random physical environments.

REFERENCES

1. Abraham, M. *Motivation and Personality*. Harper & Row, New York, 1954.
2. Allan, J. HARD track overview in TREC 2005: high accuracy retrieval from documents. In *Proceedings of the 14th Text Retrieval Conference*. Massachusetts Univ. Amherst Center for Intelligent Information Retrieval, 2005.
3. Amershi, S. and Morris, M. R. CoSearch: a system for co-located collaborative web search. In *CHI '08 Proceedings of the Twenty-sixth Annual SIGCHI Conference on Human Factors in Computing Systems*. ACM Press, 2008, 1647–1656.
4. Amigó, E., Gonzalo, J., Peinado, V., Peñas, A., and Verdejo, F. An empirical study of information synthesis tasks. In *ACL '04 Proceedings of the 42nd Annual Meeting on Association for Computational Linguistics*. Association for Computational Linguistics, Barcelona, 2004, 207–214.
5. Baeza-Yates, R. and Pino, J. A. A first step to formally evaluate collaborative work. In *Proceedings of the International ACM SIGGROUP Conference on Supporting Group Work: The Integration Challenge*. ACM New York, NY, USA, 1997, 56–60.
6. Balinsky, H., Balinsky, A., and Simske, S. J. Automatic text summarization and small-world networks. In *DocEng '11 Proceedings of the 11th ACM Symposium on Document Engineering.* ACM Press, 2011, 175.

7. Bao, S., Xue, G., Wu, X., Yu, Y., Fei, B., and Su, Z. Optimizing web search using social annotations. In *Proceedings of the 16th International Conference on World Wide Web*. ACM New York, NY, USA, 2007, 501–510.
8. Barzilay, R. and Elhadad, M. Using lexical chains for text summarization. In *Proceedings of the ACL Workshop on Intelligent Scalable Text Summarization*. 1997, Vol. 17, 10–17.
9. Barzilay, R. and Elhadad, M. Inferring strategies for sentence ordering in multidocument news summarization. *J. Artifl. Intell. Res.*, 2002, **17**, 35–55.
10. Bates, M. The design of browsing and berrypicking techniques for the online search interface. *Online Inform. Rev.*, 1989, 407–424.
11. Borlund, P. and Ingwersen, P. Measures of relative relevance and ranked half-life: performance indicators for interactive IR. In *Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM New York, NY, USA, 1998, 324–331.
12. Broder, A. A taxonomy of web search. *ACM SIGIR Forum*, 2002, **36**, 3.
13. Carmel, D., Zwerdling, N., Guy, I., Ofek-Koifman, S., Har'el, N., Ronen, I., Uziel, E., Yogev, S., and Chernov, S. Personalized social search based on the user's social network. In *Proceedings of the 18th ACM Conference of Information and Knowledge Management – CIKM '09*. ACM, Hong Kong, China, 2009, 1227–1236.
14. Church, K., Neumann, J., Cherubini, M., and Oliver, N. SocialSearchBrowser: a novel mobile search and information discovery tool. In *Proceedings of the 15th International Conference on Intelligent User Interfaces*. ACM, Hongkong, China, 2010, 101–110.
15. Cleverdon, C. W., Mills, J., and Keen, E. M. *An Inquiry in Testing of Information Retrieval Systems*. Aslib Cranfield Research Project, College of Aeronautics, Cranfield, UK, 1966.
16. Crimmins, F., Smeaton, A. F., Dkaki, T., and Mothe, J. TetraFusion: information discovery on the internet. In *IEEE Intelligent Systems and Their Applications*. 1999, Vol. 14, 55–62.
17. Ding, W. and Marchionini, G. A comparative study of web search service performance. In *Proceedings of the ASIS Annual Meeting*. 1996, Vol. 33, 136–142.
18. Dumais, S. T. and Belkin, N. J. *TREC Experiment and Evaluation in Information Retrieval*. Chapter "The TREC Interactive Track: Putting the User Into Search". MIT Press, 2005.
19. Efe, K. and Ozerturk, S. Evaluating user effectiveness in exploratory search with TouchGraph Google Interface. In *Human–Computer Interaction. New Trends*. Springer, San Diego, 2009, 405–412.
20. El-Arini, K. and Guestrin, C. Beyond keyword search. In *Proceedings of the 17th ACM SIGKDD International Conference of Knowledge Discovery and Data Mining – KDD '11*. San Diego, California, USA, 2011, 439–474.
21. Evans, B. M. and Chi, E. H. Towards a model of understanding social search. In *Proceedings of the 2008 ACM Conference on Computer Supported Cooperative Work*. ACM, San Diego, CA, USA, 2008, 485–494.
22. Freyne, J., Farzan, R., Brusilovsky, P., Smyth, B., and Coyle, M. Collecting community wisdom. In *Proceedings of the 12th International Conference on Intelligent User Interfaces – IUI '07*. Honolulu, Hawaii, USA, 2007, 52–61.
23. Golovchinsky, G., Qvarfordt, P., and Pickens, J. Collaborative information seeking. *Computer*, 2009, **42**, 47–51.
24. Gong, Y. and Liu, X. Generic text summarization using relevance measure and latent semantic analysis. In *Proceedings of the 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2001, 19–25.
25. Grimes, C., Tang, D., and Russell, D. M. Query logs alone are not enough. In *Workshop on Query Log Analysis at WWW*. Citeseer, 2007.
26. Hachey, B., Murray, G., and Reitter, D. The embra system at duc 2005: query-oriented multi-document summarization with a very large latent semantic space. In *Proceedings of the Document Understanding Conference (DUC) 2005, Vancouver, BC, Canada*. 2005.
27. Haghighi, A. and Vanderwende, L. Exploring content models for multi-document summarization. In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics, Stroudsburg, PA, USA, 2009, NAACL '09, 362–370.
28. Heymann, P., Koutrika, G., and Garcia-Molina, H. Can social bookmarking improve web search? In *Proceedings of the International Conference on Web Search and Web Data Mining*. ACM Press, 2008, 195–206.
29. Hong, S. H., Nikolov, N. S., and Tarassov, A. A 2.5D hierarchical drawing of directed graphs. *J. Graph Algorithms Appl.*, 2007, **11**, 371–396.
30. Hong, L., Convertino, G., Suh, B., Chi, E. H., and Kairam, S. FeedWinnower: layering structures over collections of information streams. In *CHI '10 Proceedings of the 28th International Conference on Human Factors in Computing Systems*. ACM Press, 2010, 947–950.
31. Huberman, B. A., Romero, D. M., and Wu, F. Social networks that matter: Twitter under the microscope. *First Monday*, 2009, **14**, 8.
32. Jones, W. and Teevan, J. *Personal Information Management (PIM)*. University of Washington, 2007.
33. Jones, R., Rey, B., Madani, O., and Greiner, W. Generating query substitutions. In *Proceedings of the 15th International Conference on World Wide Web*. ACM New York, NY, USA, 2006, 387–396.
34. Kelly, D., Dumais, S., and Pedersen, J. Evaluation challenges and directions for information seeking support systems. *IEEE Computer*, 2009, **42**.

35. Kerne, A. and Smith, S. M. The information discovery framework. In *Proceedings of the 2004 Conference on Designing Interactive Systems Processes, Practices, Methods, and Techniques – DIS '04*. Cambridge, MA, USA, 2004, 357–360.
36. Koh, E., Kerne, A., and Hill, R. Creativity support: information discovery and exploratory search. In *Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval – SIGIR '07*. Amsterdam, The Netherlands, 2007, 895–896.
37. Krause, B. Hotho, A., and Stumme, G. A comparison of social bookmarking with traditional search. In *Advances in Information Retrieval* (Macdonald, C., Ounis, I., Plachouras, V., Ruthven, I., and White, R. W., eds). Springer, Berlin, Heidelberg, 2008, Vol. 4956, 101–113.
38. Kumar, N. and Benbasat, I. The effect of relationship encoding, task type, and complexity on information representation: an empirical evaluation of 2D and 3D line graphs. *MIS Quarterly*, 2004, **28**, 255–281.
39. Lewandowski, D. The retrieval effectiveness of web search engines: considering results descriptions. *J. Doc.*, 2008, **64**, 915–937.
40. Lewandowski, D. *Web Search Engine Research*. Emerald Group Publishing, 2012.
41. Lewandowski, D. and Höchstötter, N. Web searching: a quality measurement perspective. *Web Search*, 2008, **14**, 309–340.
42. Luhn, H. P. The automatic creation of literature abstracts. *IBM J. Res. Dev.*, 1958, **2**, 159–165.
43. MacCall, S. L. and Cleveland, A. D. A relevance-based quantitative measure for internet information retrieval evaluation. In *Proceedings of the ASIS Annual Meeting*. 1999, Vol. 36, 763–770.
44. Marchionini, G. Exploratory search: from finding to understanding. *Commun. ACM*, 2006, **49**, 41–46.
45. Morris, M. R. A survey of collaborative web search practices. In *Proceedings of the Twenty-sixth Annual SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 2008, CHI '08, 1657–1660.
46. Morris, M. R. and Horvitz, E. SearchTogether: an interface for collaborative web search. In *Proceedings of the 20th Annual ACM Symposium on User Interface Software and Technology*. ACM, New York, NY, USA, 2007, UIST '07, 3–12.
47. Morris, D., Ringel Morris, M., and Venolia, G. SearchBar: a search-centric web history for task resumption and information re-finding. In *Proceedings of the Twenty-sixth Annual SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 2008, CHI '08, 1207–1216.
48. Nenkova, A., Passonneau, R., and McKeown, K. The pyramid method: incorporating human content selection variation in summarization evaluation. *ACM Transactions on Speech and Language Processing (TSLP)*, 2007, **4**(2).
49. Nenkova, A., Maskey, S., and Liu, Y. Automatic summarization. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Tutorial Abstracts of ACL 2011*. 2011.
50. O'Connor, B., Krieger, M., Ahn, D., Cohen, W., and Gosling, S. TweetMotif: exploratory search and topic summarization for twitter. In *Proceedings of ICWSM*. The AAAI Press, 2010, 2–3.
51. Ostrovsky, M. Information aggregation in dynamic markets with strategic traders. In *Proceedings of the 10th ACM Conference on Electronic Commerce*. ACM Press, 2009, 253–254.
52. Pickens, J., Golovchinsky, G., Shah, C., Qvarfordt, P., and Back, M. Algorithmic mediation for collaborative exploratory search. In *SIGIR '08 Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM Press, 2008, 315.
53. Qazvinian, V. and Radev, D. R. Scientific paper summarization using citation summary networks. In *Proceedings of the 22nd International Conference on Computational Linguistics. Vol. 1*. Association for Computational Linguistics, Stroudsburg, PA, USA, 2008, COLING '08, 689–696.
54. Quirmbach, S. Universal search: Kontextuelle Einbindung von Ergebnissen unterschiedlicher Quellen und Auswirkungen auf das User Interface. In *Handbuch Internet-Suchmaschinen* (Lewandowski, D., ed.). Akademische Verlagsgesellschaft AKA, 2009, 220–248 (in German).
55. Robins, D. Interactive information retrieval: context and basic notions. *Informing Science*, 2000, **3**, 57–62.
56. Schiffman, B., Nenkova, A., and McKeown, K. Experiments in multidocument summarization. In *Proceedings of the Second International Conference on Human Language Technology Research*. 2002, 52–58.
57. Schraefel, M. Building knowledge: What's beyond keyword search? *IEEE Computer*, 2009, **42**.
58. Silverstein, C., Marais, H., Henzinger, M., and Moricz, M. Analysis of a very large web search engine query log. In *ACM SIGIR Forum*. 1999, Vol. 33, 6–12.
59. Singer, G., Norbisrath, U., Vainikko, E., Kikkas, H., and Lewandowski, D. Search-logger analyzing exploratory search tasks. In *Proceedings of the 2011 ACM Symposium on Applied Computing*. ACM, New York, NY, USA, 2011, SAC '11, 751–756.
60. Smyth, B., Briggs, P., Coyle, M., and O'Mahony, M. *Google Shared. A Case-study in Social Search*. Springer. 2009, 283–294.
61. Steinberger, J. and Jezek, K. Using latent semantic analysis in text summarization and summary evaluation. In *Proceedings of 7th International Conference ISIM'04 Information Systems Implementation and Modelling*. 2004, 93–100.
62. Taylor, B., Mayer, M., and Buyukkokten, O. Interface for a universal search engine, 2008. US Patent 7,447,678.
63. Tunkelang, D. Precision AND Recall. *IEEE Computer*, 2009, **42**.
64. Twidale, M. B. and Nichols, D. M. Collaborative information retrieval. In *Encyclopedia of Library and*

Information Sciences. Taylor & Francis, New York, 2010, Vol. 3, 1080–1087.
65. Vaughan, L. New measurements for search engine evaluation proposed and tested. *Inform. Process. Manag.*, 2004, **40**, 677–691.
66. Voorhees, E. M. and Harman, D. K. *TREC: Experiment and Evaluation in Information Retrieval*. MIT Press, 2005.
67. White, R. W. and Marchionini, G. Examining the effectiveness of real-time query expansion. *Inform. Process. Manag.*, 2007, **43**, 685–704.
68. White, R. and Roth, R. *Exploratory Search*. Morgan & Claypool Publishers, 2009.
69. White, R. W. and Roth, R. A. Exploratory search: beyond the query–response paradigm. *Synthesis Lectures on Information Concepts, Retrieval, and Services*, 2009, 1–98.
70. White, R. W., Muresan, G., and Marchionini, G. Report on ACM SIGIR 2006 workshop on evaluating exploratory search systems. In *ACM SIGIR Forum*. ACM New York, NY, USA, 2006, Vol. 40, 52–60.
71. Widjaja, Y. Y., Pang, C. N. I., Li, S. S., Wilkins, M. R., and Lambert, T. D. The interactorium: visualising proteins, complexes and interaction networks in a virtual 3D cell. *PROTEOMICS*, 2009, **9**, 5309–5315.
72. Wilson, M., Russell, A., and Smith, D. A. mSpace: improving information access to multimedia domains with multimodal exploratory search. *Commun. ACM*, 2006, **49**, 47–49.
73. Xu, S., Bao, S., Fei, B., Su, Z., and Yu, Y. Exploring folksonomy for personalized search. In *Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM Press, 2008, 155–162.

Kompleksotsing: agregatsioon, avastamine ja süntees

Georg Singer, Dmitri Danilov ja Ulrich Norbisrath

Otsingumootorid Google, Bing ja Yahoo ning erinevad veebiprojektid, näiteks Wikipedia, on muutunud peamisteks allikateks internetist otsingute tegemisel. Koos otsingumootorite kasvava populaarsusega on akadeemiline huvi nihkunud lihtsate *look-up*-tüüpi päringute analüüsimisest keerulisemate otsimisvajaduste analüüsimisele. Praegused otsingusüsteemid toetavad kompleksotsingut tunduvalt nõrgemalt kui *look-up*-otsingu mudelit. Eriti halvasti on toetatud mitme päringutulemuse koondamine, selle protsessi jooksul tehtud avastuste arvessevõtt ja leitud info sünteesimine uueks tulemuseks. Otsingumootorite võimetus anda keerulistele probleemidele adekvaatseid vastuseid motiveerib teadlasi arendama uusi infootsimistehnikaid.

Me keskendume selles artiklis kolmele uuriva otsingu mõistele – agregatsioon, uue info avastamine ja süntees –, kuna me leiame, et tänapäeval on just need kõige aeganõudvamad otsingutegevused, eriti kui täidetakse mitmekülgse info vajadust. Me kasutame neid kolme mõistet kogu artikli jooksul erinevate uuriva otsingu lähenemisviiside klassifitseerimiseks, olemasolevate lahenduste hindamiseks ja otsinguvaldkonna uuringutest ülevaate andmiseks.