

Searching in the Context of a Task: A Review of Methods and Tools

Ana Maguitman

Departamento de Ciencias e Ingeniería de la Computación,
Universidad Nacional del Sur,
ICIC CONICET-UNS,
Bahía Blanca, Argentina, 8000,
agm@cs.uns.edu.ar

Abstract

Contextual information extracted from the user task can help to better target retrieval to task-relevant content. In particular, topical context can be exploited to identify the subject of the information needs, contributing to reduce the information overload problem. A great number of methods exist to extract raw context data and contextual interaction patterns from the user task and to model this information using higher-level representations. Context can then be used as a source for automatic query generation, or as a means to refine or disambiguate user-generated queries. It can also be used to filter and rank results as well as to select domain-specific search engines with better capabilities to satisfy specific information requests. This article reviews methods that have been applied to deal with the problem of reflecting the current and long-term interests of a user in the search process. It discusses major difficulties encountered in the research area of context-based information retrieval and presents an overview of tools proposed since the mid-nineties to deal with the problem of context-based search.

Keywords: Context-Based Search, Task-Based Search, Topical Search, Information Retrieval, Personalization.

1 Introduction

Context-based search is the process of seeking material based on the user situation. A context-aware information retrieval system should be able to monitor the activity in which users are involved, identify their search preferences, infer their state and goals, learn what information will help to fulfill these goals, find this information and decide when and how this information should be presented. The user context can contribute to reducing the information overload problem since it can be usefully exploited by a search or suggestion system to define or predict a user information need more precisely. Such a system may take advantage of the user context to enrich an explicit information request, or it may be able to anticipate what the user is looking for without the explicit formulation of a user query.

Typically, context refers to information that is related to the user current state and physical environment. It is usually associated with ephemeral data and may include heterogeneous elements such as the current task, information state, emotional state, date, time, season, location, orientation, temperature, lighting level, noise, resources around, people around and global mood, among others. This information is sometimes supplemented with more persistent data associated with the user profile, such as identity, declared interests, capabilities and Social Media communities to which the user belongs. Monitoring the user activity for extended periods of time allows obtaining additional data from query logs, browsing history, and Social Media interaction. All the above elements help to identify user information needs, and several approaches have been proposed to design context-aware systems that take some of these elements into account to assist users in various situations. Some surveys review a broad-spectrum of context-aware systems [1, 2, 3, 4]. Rather than presenting an overview of context-aware systems from a general perspective, this review focuses on the problem of context-based search, where context refers to the topical context extracted from the user task. It presents a revision of techniques and existing research proposals aimed at dealing with the problem of reflecting the topical context and contextual interaction patterns in the search process. This survey also

traces the evolution of task-based context-aware tools from the mid-nineties to the present days and discusses new progress and research directions.

2 Topical context of a task

As mentioned in [5] the notion of context is ill-defined as a user context may include a great variety of components extracted from the user profile, state and physical environment. Context may result from information acquired manually or automatically. Depending on the application, the explicit and implicit input could give rise to various entities defining the user context. This review focuses on a particular class of context to which we refer to as the *topical context of a task*.

A task is a piece of work required to achieve an objective. Tasks are typically associated with the need to access information to solve problems, evaluate content, construct meaning, create knowledge and make decisions. In an information seeking process, the topical context defines the subject or focus of the information need. For instance, consider a journalist writing an article about “Electronic Voting”. Suppose the journalist has collected a small set of documents related to this topic and would like to retrieve additional material from different sources (the Web, Social Media, digital newspapers, etc.). A context-based search tool that has the purpose of collecting relevant material from these sources can assist the journalist in the task of writing the article. The context-based search tool will monitor the journalist’s work, generate an initial set of queries and incrementally refine these queries to reflect the topic of interest in a better way. The initial queries could be formulated directly from the journalist’s initial context (e.g., the document that is being edited), from a short description provided explicitly by the journalist, or from the set of relevant documents that the journalist has collected so far. The initial queries can be incrementally refined based on the collected material that is relevant to the topic “Electronic Voting”. Feedback on the relevance of the retrieved material can be explicitly indicated by the journalist or can be deduced by the context-based search tool by monitoring the journalist interaction patterns with documents being read or edited, web pages being visited, emails being read or written as well as with the context-based search tool itself. The refined queries may include vocabulary that is incrementally learned from the available information sources. For instance, it is possible to collect new hashtags such as #electoralreform, #cybersecurity or #voterfraud by recognizing them as frequently occurring in tweets that are related to the topic at hand. This new vocabulary can be used to automatically build subsequent queries helping to retrieve novel but topic-related content.

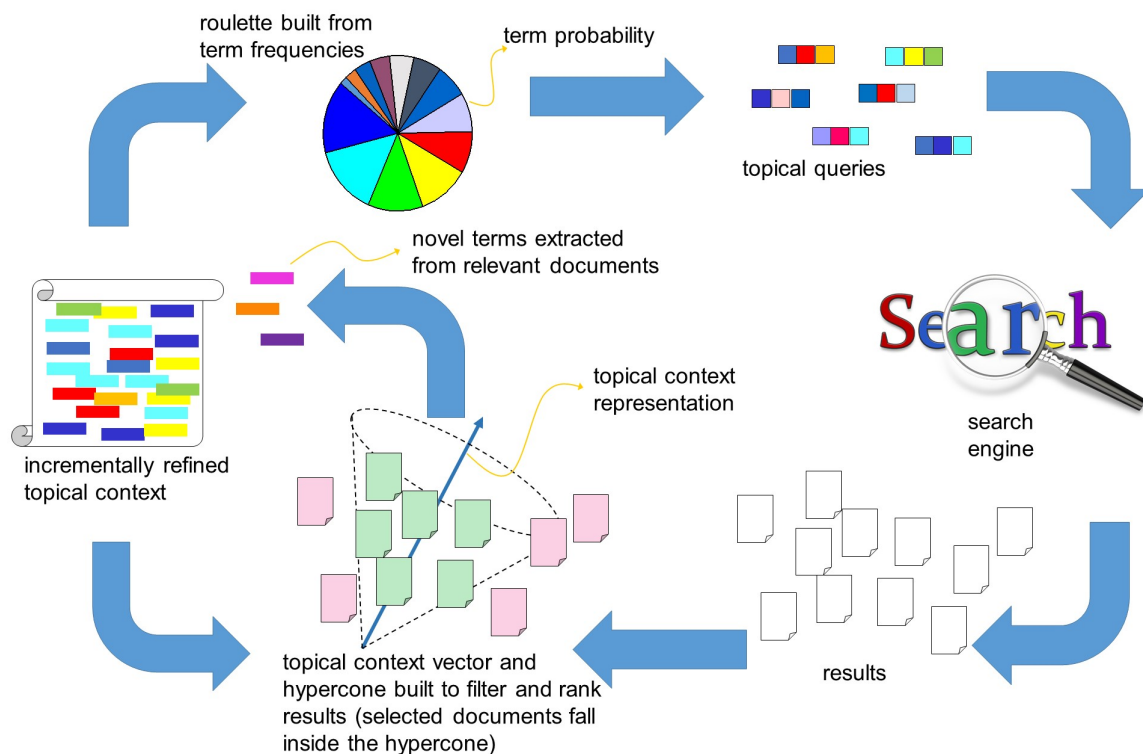


Figure 1: A simplistic approach for incremental context-based retrieval.

It is possible to implement a simplistic approach for incremental context-based search by taking some terms from the given topical context to form queries that are submitted to a conventional search engine (e.g.,

Web search engine, Social Media search engine, news search engine, etc.). Diverse techniques can be applied to select terms for query formation. A basic method for term selection from a text is the *roulette wheel selection* technique [6], where the probability of choosing a particular term to form a query is proportional to the number of times the term appears in the given text. Queries can be submitted to a search engine and a set of documents similar to the topical context can be selected. A hypercone is computed with the axis corresponding to the topical context vector and angle defined according to the given similarity threshold to determine which results are chosen for further analysis. The results whose vector representations fall inside this hypercone are considered relevant results. In the meantime, the terms that often occur in relevant results are considered for inclusion in an augmented context representation. Fig. 1 illustrates this approach to retrieve results based on a topical context.

Topical contexts may be obtained from the user long-term or current interests and can be defined at different levels of granularity. For instance, a topical context may be defined as a collection of documents related to a particular subject of interest. Alternatively, a user may highlight a few words in a document with the purpose of requesting information associated with the highlighted text. This approach gives rise to a more focused view of topical context, which may include the sentence containing the highlighted text, the paragraph where the highlighted text resides, or the whole document the user is reading or editing. The system may also be able to automatically identify the user focus of interest in a document and proactively seek for information based on the inferred context. This could be done by a system monitoring the user interaction with a software application being used. As a result, the user focus of interest may be inferred from various interaction patterns by tracking clicks, dwell time, cursor movement, scrolling, interaction in the area of interest, touch interaction or eye movement. Besides defining the user focus of interest, these contextual interaction patterns can be usefully exploited to recognize user plans, activities and intents.

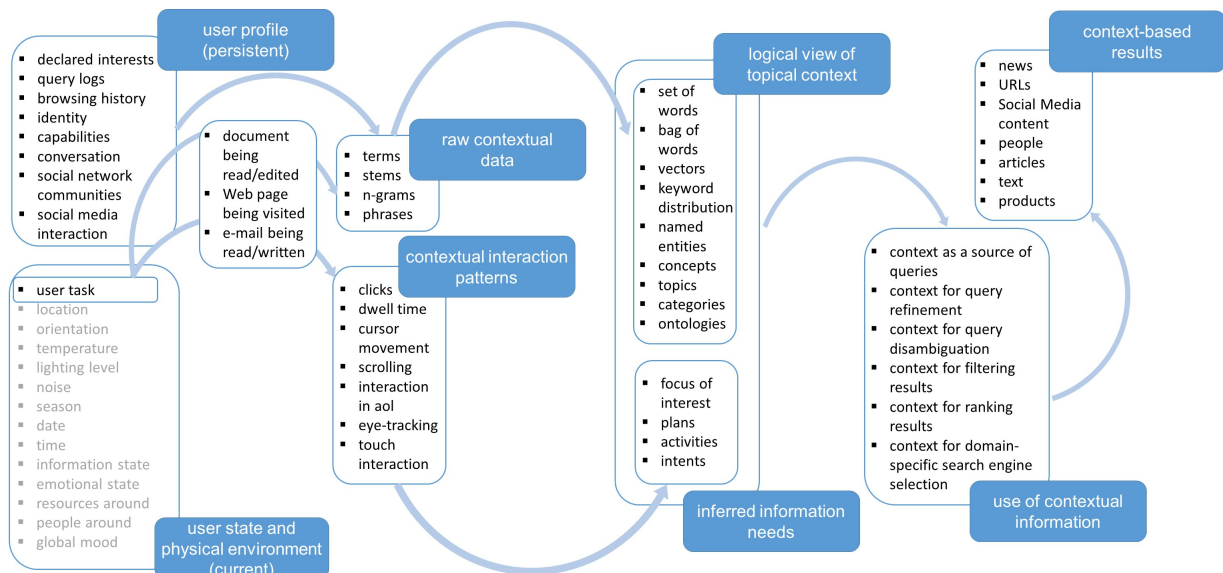


Figure 2: Components of a search system based on the topical context of a task.

A topical context is initially captured as raw contextual data and may include terms, stems, n-grams or phrases. A fundamental problem is how to transform these primary data into more useful context representations. Higher-level topical context representations may take the form of keyword distributions, named entities, concepts, topics, categories, formulas in a logical language, or ontologies, among other high-level representations. A major dimension of analysis in the problem of defining context-based search tools is how the captured context will be reflected in the search process. The captured topical context can be used in various ways, such as a source for automatically generating queries, refining queries and disambiguating queries. It can also be used for filtering results, ranking result or selecting a domain-specific search engine. Another important aspect of a context-based search system is the content of the search results (e.g., news, URLs, Social Media content, people, articles, text, products). Fig. 2 outlines the components and mechanisms typically found in different approaches to context-based search, with focus on techniques for searching based on the topical context of a task.

Additional dimensions of analysis are the way the system learns the user's interests (monitor the user's behavior, receive explicit or implicit feedback, or engage in conversation with the user), the way the system refines the context representation (supervised, semi-supervised, or unsupervised) and the level of intrusiveness. Finally, another element of analysis is determined by the criteria used for assessing the effectiveness

of a context-based method. While precision is the most commonly used measure for determining retrieval evaluation, other classical and ad-hoc measures have also been considered for evaluation purposes, such as recall, diversity, and novelty. However, as it will be discussed in section 3.6, there is no widely accepted methodology for evaluating context-based retrieval.

3 Reflecting a task topical context in search

Diverse objectives can guide the process of searching for online data. There are essential differences between searching for information to fulfill a *consultation need* and searching for material to support the process of collecting *topical resources*. Usually, the purpose of a consultation is to find specific answers to specific questions. In this case, it is crucial to attaining high precision. On the other hand, when collecting topical resources, rather than a specific question, there is typically a topic of interest that can be used to initialize and refine the relevant search space. In this case, while precision is important, other criteria should be considered, such as high recall, novelty, and diversity. This section discusses various aspects involved in the problem of task-based contextual search and reviews different solutions proposed in the literature to deal with this problem. In particular, it addresses the issue of context extraction and modeling and reviews different approaches used to exploit topical context in the search process. It finally discusses criteria that can be used to evaluate a topic-based search process.

3.1 Context extraction and modeling

Context extraction and modeling are typically carried out according to the application domain. Different context dimensions have been identified to characterize general context models [7, 8]. These dimensions include personal context, temporal and spatial context (absolute and relative), task context, social context, historical context, psychological context, cognitive context and event context, among others.

In the scenario of searching in the context of a task, it becomes necessary to identify and model the different aspects of the user information needs with particular focus on the subject of interest. Once these aspects are extracted from the user context, they should be modeled and incorporated into the information seeking system. While the user query provides the most direct evidence to infer a user information need, other explicit or implicit indicators can be exploited to provide additional evidence. For instance, topical context information could be added explicitly by the user in the form of a category. Other approaches attempt to infer context information automatically, which can be done by extracting contextual data from the content of documents or emails being edited or viewed, Web pages being visited, among other sources of contextual data. Some systems extract keywords from the context surrounding the user focus of attention. For instance, by considering the context surrounding the user-selected text, it is possible to collect additional terms. These terms can be used to construct augmented queries, which are then used to guide the search process.

The context extracted from textual data can be modeled using straightforward representations such as sets or bags of words, where all words from the document are considered, disregarding any word ordering or structure in the text. Since not all terms are equally useful for describing a topical context, some modeling approaches have been developed to weigh terms according to their importance. The most basic form of weighting scheme is to assign to each term a weight that is proportional to its frequency [9]. Further initial research in term weighting has demonstrated that when document collections are available, it is useful to account for a global weight that reflects term specificity, such as the term's inverse document frequency [10]. The term-frequency and term-frequency-inverse-document-frequency weighting schemes are the most widely adopted weighting schemes and are usually combined with the vector space model [11] to represent an information need. Additional information available in some textual contexts could be used, such as word position, font size, sentence structure, etc., allowing a more detailed representation of the importance of words and their relations in text.

Central to the concept of topical context is the notion of "semantic context", which refers to the subject or topic where a task, such as searching, is embedded. Representing this notion of context requires going beyond the terms that explicitly occur in the user context. The notion of semantic context is closely linked to that of semantic similarity. Two documents are semantically similar if they belong to the same topic or related topics. Likewise, two words are semantically similar if they represent similar concepts. The study of semantic similarity between documents and terms has long been an integral part of information retrieval and machine learning, and there is extensive literature on measuring the semantic similarity between entities [12, 13, 14, 15, 16, 17, 18, 19]. These methods can be roughly classified into two broad categories: corpus-based and knowledge-based approaches. Corpus-based approaches are based on information obtained from large corpora. Useful methods for deriving semantic contexts rely on latent variable models, which have proved to be successful in revealing the inter- or intra-document statistical structure. Latent Semantic Analysis (LSA)

is a classical approach for capturing the contextual-usage meaning of words by taking advantage of implicit higher-order structure in the association of terms with documents [12]. LSA relies on the vector space model to represent documents as vectors in term space and on the use of singular-value-decomposition to project documents and words to a lower-dimensional semantic space. This dimensionality reduction allows material with similar topical content to be close to one another in the resulting space. Topic modeling techniques, such as Latent Dirichlet Allocation (LDA) [16, 19] or Hidden Topic Markov Model [18], are commonly used mechanisms for discovering and characterizing the topic of a document or a piece of text. These techniques allow to search and explore documents based on their themes. They rely on statistical methods that analyze the words of massive collections of documents to discover their topics and to show how they connect to each other. A topic is defined as a distribution over a fixed vocabulary. Once topic distributions are available, a given topical context can be associated with the most likely topic based on the given distributions. While most topic modeling approaches analyze each word independently of other words to represent topics, some proposals have attempted to capture textual order by considering both individual words and n-gram phrases [20].

Less computationally expensive techniques are based on mapping documents to a kernel space where documents that do not share any term can still be close to each other [21, 22]. Another corpus-based technique that has been applied to estimate semantic similarity is PMI-IR [15]. This information retrieval method is based on pointwise mutual information, which measures the strength of association between two elements (e.g., terms) by contrasting their observed frequency against their expected frequency. Another method for learning topical context representations relies on incrementally collecting topic descriptors and topic discriminators. A topic descriptor is defined as a term that often occurs in documents relevant to the topic while a topic discriminator is a term that occurs only in relevant documents. Topic descriptors and discriminators can be learned based on the user topical context and incrementally collected documents [23, 24].

Knowledge-based approaches for extracting and modeling semantic contexts rely on the use of predefined directories or ontologies to identify semantic relations. Measures of semantic similarity between entities take as a starting point the structure of a directory or ontology where the entities have been previously classified. Examples of such ontologies include WordNet (<http://wordnet.princeton.edu/>) for the case of terms and DMOZ (<http://dmoz.org>) for the case of Web pages. Ontologies are a particular kind of network, and early proposals to estimate semantic similarity have used path distances between the nodes in the network representation [25]. Some successors have looked into the notion of information content [13, 14] or have combined both distance and information content [26] to assess semantic similarity. In an information-theoretic approach, the semantic similarity between two entities is related to their commonality and their differences. Given a set of entities in a hierarchical taxonomy, the commonality of two entities can be estimated by the extent to which they share information, indicated by the most specific class in the hierarchy that subsumes both. Once this common classification is identified, the meaning shared by two entities can be measured by the amount of information needed to state the commonality of the two objects. Generalizations of the information-theoretic notion of semantic similarity for the case of general ontologies (i.e., taxonomies that include both hierarchical and non-hierarchical components) have been proposed by [17].

The framework proposed in [27] combines corpus-based and knowledge-based approaches to build a semantic network to characterize different domains (such as computers, business, and entertainment). The semantic network is constructed by representing each domain as an axis in a multi-dimensional space and each word as a vector in the resulting domain space. Words are obtained from large sets of documents in these domains. The semantic network is further enhanced using linguistic data available through WordNet. By taking advantage of the constructed semantic network, keywords are identified and extracted from the user context (text surrounding the marked query). Another approach that relies on external semantic information sources is based on associating the user context with predefined topic categories [28], which may be obtained from topic taxonomies or ontologies such as DMOZ. Similarly, a topic category can be used to represent the user search profile. Categories can be identified explicitly by users [29], or they can be inferred based on the user past activity through a categorization method [30]. Some approaches attempt to build specific ontologies to represent user interests [31]. Others use ontologies to conceptualize different aspects of the context, such as the themes of the user's information need and the specific data the user is looking for to achieve the task that has motivated the search. This solution is adopted in [32], where a theme domain ontology is used to specify and structure the objects of the themes treated in the documents and their relations. Also, a task domain ontology is used to model the type of task the user is engaged in and to specify the data of interest.

3.2 Context as a source of queries

The availability of powerful search interfaces makes it possible to develop a plethora of applications for information access in context. Existing search services provider commonly establish particular matching

and ranking mechanisms. The criteria considered for retrieval change from service to service (e.g., content relevance, document popularity, document freshness, meta-information) and are typically obscure to those using the interface. In many circumstances, as is, for instance, the case of the Deep Web, the only access point to relevant material is through the submission of queries. As a consequence, a significant research problem in the area of context-based search is learning to formulate effective topical queries automatically.

The effectiveness of a query depends on what the user is trying to achieve. If the goal is to attain good coverage, a broad query will be more useful than a specific one. In this case, the user will try to retrieve most of the relevant material, albeit diluted with irrelevant answers. On the other hand, if the goal is to attain good precision, specific queries will be preferred over broad ones.

Since a topical context may contain numerous terms, selecting useful terms is crucial. Several studies have shown the benefits of having tools that assist in query formulation (and reformulation). Typically, these systems provide a browsing interface where the user intervention must be explicit. However, the burden implied by the need to formulate search requests explicitly can be alleviated if queries are produced automatically based on a topical context [33, 34]. The problem of searching based on a topical context has been studied as an optimization problem [35, 36, 37] where the objective function to be maximized quantifies the effectiveness of a query when presented to a particular search engine. While not all the information contained in a topical context can be summarized in a query, effective mechanisms can be designed to extract small sets of representative terms to construct queries. Possible queries define the search space and solutions can be constructed incrementally by combining terms in different ways, taking the objective function as the guiding criterion. A major difficulty in dealing with query optimization is that this problem does not have optimal substructure, which means that optimal solutions cannot be constructed efficiently from optimal solutions to its subproblems. Consequently, the combinations of effective small queries may not result in an effective longer query. Depending on the system goals, a measure of query effectiveness can be defined using traditional information retrieval notions, such as precision and recall, or other customized performance evaluation metrics. Some distinct features of dealing with topical search as a query optimization problem are:

- High-dimensionality of the search space: query space is a high-dimensional space, where each term that can be part of a query accounts for a new dimension.
- Changeability of the search space: finding suitable combinations of query terms requires exploring the different direction of the search space. In the query optimization problem, the search space is not static and may change as new terms are incrementally learned.
- Suboptimal solutions are acceptable: topical search involves the formulation of high-quality queries even if the formulated queries are not optimal.
- Multiple solutions are expected: topical search requires constructing many high-quality queries rather than a single one.

While the automatic generation of queries offers a convenient mechanism to reflect a topical context in the search process, the use of context as a source of queries only provides a partial solution to the problem of context-based search. If queries are built as a conjunction of terms, short queries will only partially represent the topic at hand, resulting in reduced precision, while long queries will become too specific, resulting in a poor recall. On the other hand, disjunctive queries will typically result in a bad precision for both short and long queries. Learning to formulate queries with more complex syntaxes, such as general Boolean queries, may result in better specifications of the user requirements [38]. Alternatively, search effectiveness can be improved if the context is not only used as a source of queries but also to filter and rank results, as is discussed next.

3.3 Context for filtering and ranking results

Some studies have shown that context-based filtering reduces the user's current workload [39]. Therefore, topical contexts are useful not only as a source of potential terms for query generation and optimization but also as a means for filtering retrieval, by helping to eliminate noise and to distinguish relevant from irrelevant results [40, 41]. For instance, retrieved results can be contrasted against the topical context to preserve only those results that are sufficiently relevant (e.g., based on a minimal relevance threshold between the topical context and the result).

A critical problem is how to determine what results are sufficiently relevant to the topic of interest. Typically, relevance is established using similarity measures, and many algorithms have been used to compute these similarity measures, combining several methods coming from the information retrieval and machine

learning areas. Commonly applied techniques are based on cluster analysis [42], cosine similarity [43], k-nearest neighbors [44], latent semantic analysis [12], topic modeling [16] and Bayesian classifiers [45], among many others.

Some context-based search systems submit queries to several general and domain-specific search engines, requiring an extra step for combining the results of different engines and returning a single list of ranked results. Ranking algorithms may take several criteria. For instance, the ranking algorithm proposed in [27] uses a semantic network to estimate the relatedness of search results to the query context. Alternatively, the context can be explicitly indicated by the user. For instance, a category can be specified by the user and then used by a meta-search engine to select a domain-specific search engine to send queries, to modify queries and to define a ranking on search results. An early system that has adopted this strategy is *Inquirus 2* and its successors [46, 47, 48].

In [49] a spectrum of algorithms to perform automatic contextual searches is provided. One of the proposed algorithms referred to as *Rank-Biasing* uses queries made up of two parts, the selection part, and the optional ranking terms. The selection part of a query is composed of a list of terms and is aimed at retrieving results. The optional ranking terms is a list of weighted terms, where the weight associated with each term determines the term influence on the overall ranking of the retrieved results. A natural approach for implementing contextual search using the Rank-Biasing approach is to use the selection part for representing the regular query and the optional ranking terms to represent extracted terms from the context.

The work presented in [50] addresses the problem of how different types of contexts can be considered and integrated into a ranking model. In particular, context is obtained from query sessions, where queries are classified with respect to previous queries as reformulation, specialization, generalization and general associations. Different context-aware ranking principles based on various types of queries are developed and evaluated.

In [51] a family of context-sensitive retrieval and ranking algorithms based on statistical language models is proposed. The algorithms combine the current query with previous queries and click-through information for better ranking search results. A recent proposal adapts these algorithms to support collaborative retrieval for team members [52]. Users' behavior histories are converted to context-aware statistical language models and then used to re-rank search results. A distinctive aspect of this proposal lies in the fact that search contexts are derived from team members' search histories at the individual and team level, as well as from their explicit collaboration (e.g., chatting).

Some recent approaches that reflect users' topical task in the process of searching the Web for filtering and ranking results include [53, 54, 55]. A review of Web search methods that exploit the search and browse logs to enhance retrieval is presented in [56].

3.4 Context for query refinement

Topical context can be used for query refinement by applying a semi-supervised version of the classical relevance feedback method [24]. Relevance feedback, proposed by Rocchio [57], is a query refinement mechanism used to tune queries based on the relevance assessments of the query's results. A driving hypothesis for relevance feedback methods is that it may be difficult to formulate a good query when the collection of documents is not known in advance, but it is easy to judge particular documents, and so it makes sense to engage in an iterative query refinement process. A typical relevance feedback scenario will involve the following steps:

Step 1: A query is formulated.

Step 2: The system returns an initial set of results.

Step 3: A relevance assessment on the returned results is issued (relevance feedback).

Step 4: The system computes a better representation of the information needs based on this feedback.

Step 5: The system returns a revised set of results.

Depending on the level of automation of step 3 it is possible to distinguish three forms of feedback:

- **Supervised Feedback:** requires explicit feedback, which is typically obtained from users who indicate the relevance of each of the retrieved documents. This kind of feedback is hard to get and as a consequence is rarely available.
- **Unsupervised Feedback:** it applies blind relevance feedback and typically assumes that the top k documents returned by a search process are relevant.

- **Semi-supervised Feedback:** the relevance of a document is inferred by the system based on contextual information.

According to Rocchio’s relevance feedback method, given an initial query vector \vec{q} a modified query \vec{q}_m is computed as follows:

$$\vec{q}_m = \alpha \vec{q} + \beta \sum_{\vec{d}_j \in D_r} \vec{d}_j - \gamma \sum_{\vec{d}_j \in D_n} \vec{d}_j,$$

where D_r and D_n are the sets of relevant and non-relevant documents respectively and α , β and γ are tuning parameters. A common strategy is to set α and β to a value greater than 0 and γ to 0, which yields a positive feedback strategy. Since the information seeking process is commonly performed within a topical context, the relevance of a document to the user’s context can be estimated as the similarity of the document to the user’s current context. This approach allows to automatically identify D_r and D_n based on the given topical context, resulting in a context-based mechanism to query refinement.

Contextual terms can also be used for query disambiguation, as discussed by [58]. In particular, context-sensitive query auto-completion algorithms output the completions of the user’s input that are most similar to the user context or search history [59, 60, 61]. Other methods support query expansion and refinement processes through a query or browsing interface requiring explicit user intervention [62, 63]. In [64] a method is proposed for the automatic refinement of ambiguous queries based on a context list that stores potential context keywords for a large number of topic relevant query n-grams. Then, topical search engines are used to obtain results focused on specific topics. A survey of methods for automatic query expansion is presented in [65].

3.5 Contextual interaction patterns as indicators of plans, activities, and intents

The theme or subject of the user context can help determine the topic of the sought information. However, it may be possible to go beyond the textual information that becomes available through the analysis of a user task to identify further indicators of the user’s information needs. Several forms of interaction between users and applications can provide useful insights into the user plans, activities, and intents. Some approaches have been developed for the recognition of a variety of interaction patterns to provide context-sensitive support [66]. *Microsoft Office Assistant* (typically personified by *Clippy*) is certainly one of the best-known context-sensitive computer aides. It was developed in the framework of the *Lumière* Project [67] and first distributed with *Microsoft Office’97* product suite. The purpose of the *Office Assistant* is to provide support to Microsoft software users. The *Lumière* project [67] is a seminal work on the analysis and modeling of user interaction with software tools to infer user goals and needs. The *Lumière* project identifies *goals* with target tasks or subtasks at the focus of a user’s attention to representing various aspects of user modeling and automated assistance. In the meantime, *needs* are defined as information or automated actions that will reduce the time or effort required to achieve goals. Studies with human subjects have helped identify classes of evidence associated with the goals and needs of software users, such as distinct search patterns, the focus of attention, introspection, etc. The *Lumière* project exploits this evidence for the construction of Bayesian user models that can help to infer the likelihood that a user needs assistance and a particular user need.

Implicit relevance feedback approaches attempt to identify user interest in a particular topic. For instance, the user interaction with a Web browser naturally generates a large amount of implicit feedback information that can be exploited to model the user information needs [68]. This information includes clicks [69], dwell time [70] as well as additional browsing behavior features such as cursor movement, vertical scrolling and interactions in the areas of interest [71]. Eye-tracking analysis has also been used to more precisely interpret implicit user feedback [72], while the interaction of gaze and cursor position has proved to be useful to determine user intent and relevant parts of Web pages [73]. Touch interaction data on mobile devices has also been exploited as a source of relevance feedback [74, 75, 76].

Many works have specifically focused on the user interaction with search engines and have demonstrated the benefit of exploiting query logs [77, 78], query dynamics [79] and immediate query history [80] to better model the user task and to improve retrieval accuracy. Some approaches attempt to identify sequences in data and to connect them to a meaningful process model [81]. This is done with the purpose of learning a conditional probability model that helps to implement predictors with the capability of suggesting the most likely information resource for a given moment. Some approaches apply sequence data mining techniques to learn predictive models [82]. Also, specialized context representations have been proposed to enable capturing temporal information, such as the *Text-Network-Time* platform [83]. A survey of various techniques for detecting query intents using click-through and implicit relevance feedback data is presented in [84].

Table 1 presents a classification of the methods reviewed in this section based on how context is mainly exploited (context as a source of queries, context for filtering and ranking results, context for query refinement and context as an indicator of user intents).

	SOURCE OF QUERIES	FILTERING AND RANKING	QUERY REFINEMENT	INDICATOR OF INTENTS
Interest assessment from browsing activity [54, 55, 70, 68, 85, 86, 87, 88, 89, 90, 91, 92]		✓		✓
Analysis of query history and/or click-through information [51, 53, 69, 77, 78, 79, 80]				✓
Prediction from sequence data mining and temporal process modeling [81, 82, 83]				✓
Monitoring general user activity for context-based retrieval [33, 34, 93, 94, 95, 96, 97, 98, 99, 100, 101, 102, 103, 104, 105]	✓	✓		✓
Accessing digital libraries based on topical context [106, 107, 108]	✓	✓		
Implicit feedback from user interaction with software tools [67, 66, 109, 110]	✓			✓
Query expansion and refinement with explicit user intervention [62, 63]			✓	
Query generation, augmentation and/or refinement from context without user intervention [24, 27, 36, 37, 49, 111, 112]	✓		✓	
Context sensitive query autocompletion [59, 60, 61]			✓	✓
Query generation and rank-biasing based on context [49, 50, 113]	✓	✓	✓	
Query understanding or disambiguation based on context [58, 64, 114]				✓
Implicit feedback from cursor movement, vertical scrolling, interactions in the areas of interest and/or eyetracking [71, 72, 73]				✓
Touch interaction data on mobile devices [74, 75, 76]				✓

Table 1: Context-sensitive methods classified in terms of how context is exploited.

3.6 Assessing context-based retrieval

Traditionally, the performance of topical search methods is evaluated using effectiveness metrics taken directly or adapted from the information retrieval standard evaluation methodology, which requires the availability of a test collection. Given a topic t , the test collection for evaluating a topic-based search system is divided into two sets, the set of documents relevant to t and the set of non-relevant documents. Precision is probably the most commonly used metric to measure the effectiveness of a method for retrieving results relevant to a topic. A generalization of precision is *semantic precision*, which can be naturally computed when other topics are semantically similar (and therefore partially relevant) to the topic at hand. It can be measured based on the notion of semantic similarity, as described in [17]. The traditional performance measure of recall is also a common criterion for evaluating topical retrieval effectiveness. Precision and recall can be combined by their weighted harmonic mean resulting in a single aggregated measure that simultaneously reflects retrieval accuracy and coverage. A semantic framework for evaluating topical search methods is proposed in [115]. The framework relies on the use of topic ontologies and semantic similarity data derived from these ontologies to associate each topic description with a large set of locally indexed relevant and partially relevant web pages. This approach allows assessing the performance of any topic-based retrieval method when tested on the given set of topics and web pages. The framework was successfully applied in the evaluation of different topic-based search methods in terms of precision, semantic precision and recall [24, 37, 112].

While precision and recall are important goals in topical search, some other priorities should also be considered depending on the task at hand. A characteristic of classical information retrieval systems is that they attempt to match requests with the most similar documents. A few approaches take a different position [116, 117] and postulate that in certain circumstances conventional notions of similarity may not be the best criteria for retrieval. In the particular case of searching based on a topical context, the criteria for determining usefulness should not be restricted to similarity. Since topical search is usually intended to include a rich variety of related subtopics, attaining novelty and diversity may be as important, or even more important, than achieving similarity.

The problem of diversification of search results has been extensively studied in the information retrieval community. A study presented in [118] discusses the importance of optimizing the relevance of results at the document set level instead of at the single document level. This approach typically involves paying particular attention to the notion of *subtopic* within a general topic. Existing methods for query diversification sometimes require an initial description defining a topic of interest that allows identifying diverse aspects (or subtopic). In this case, the goal is to generate different queries that can cover multiple aspects associated with a topic or task. In [119] query diversification is attained by taking a document from which multiple aspects can be identified. Since documents are typically long, they can be used to generate several queries covering their various aspects or subtopics. These queries are then filtered by selecting those that are relevant to the query document, as well as new relative to the current query suggestion list. A comprehensive survey of search result diversification approaches and diversity evaluation benchmarks and frameworks is presented in [120].

While test collections containing documents and explicit relevance judgments are usually adopted for evaluating topic-based retrieval, other forms of implicit relevance assessment can be used in the contextual search to capture user activity. Logs reflecting different types of interaction data play a major role in providing an implicit assessment of relevance and other context-related factors. Unfortunately, while documents and relevance assessments are usually publicly available, other forms of contextual information that may help in the evaluation process, such as logs, are not. This limitation on data availability is a consequence of the commercial origin and the legal restriction to invading user's privacy of logs. For these reasons, except for a few exceptions (e.g., [121, 122]), the design of traditional test collections is typically independent of any context. This problem has been identified by some studies that analyze approaches adopted for evaluating context-based retrieval [123, 84]. These studies argue that most of the adopted evaluation frameworks are not realistic as they rely on the laboratory-based evaluation model, which requires predefined test collections and relevance judgments. The use of simulated work task situations in interactive information retrieval evaluation offers a different perspective on the design of context-centered evaluation techniques [124, 125]. An experimental setting for assessing the effectiveness of a context-based search method consists of a simulated work task, where some subjects in an experimental group search the WWW for a given topic during a certain amount of time using a context-based search method under evaluation. On the other hand, the same number of subjects in a control group search the WWW separately for the same topic during the same amount of time as the experimental group. An assessment of the effectiveness of the evaluated context-based method results from comparing the number of noise documents vs. the number of relevant documents accessed by each group. However, these evaluations rely on a user study in a very constrained environment, which detracts from the realism of the search task. As a consequence, in spite of these alternative proposals, context-based

search still lacks a well-recognized evaluation methodology and is typically narrowed to the assessment of topical relevancy.

4 An overview of context-based search systems

Most context-based search systems have been built around content and resources available on the Web. Several of these systems use explicit or implicit context information from users and offer unique functionality to capture long- or short-term user topical interests. *WebWatcher* [85] is an early attempt to assist users locating information on the Web by highlighting hyperlinks in a page based on the declared preferences and browsing history of a user as well as information gathered from other users with similar interests. *Personal WebWatcher* [86] is a successor of *WebWatcher* that learns individual users' interests by observing their browsing behavior. *Letizia* [87] is a user interface agent that unobtrusively assists Web browsing. As the user navigates Web pages, *Letizia* performs a breadth-first search augmented by several heuristics to anticipate what items may be of interest to the user. *Syskill & Webert* [88] uses information retrieval techniques to process the content of a page rated by a user and machine learning to acquire a model, that is utilized to predict which links on a Web page a user will find useful. *SenseMaker* [89] is an interface that facilitates the navigation of information spaces by providing task-specific support for consulting heterogeneous search services. The system helps users to examine their present context, move to new contexts or return to previous ones. *SenseMaker* presents the collection of suggested documents in bundles (their term for clusters), which can be progressively expanded. *Fab* [90] is a hybrid content-based, collaborative Web page recommender system that learns to browse the Web on behalf of a user. *Fab* generates recommendations by the use of a set of collection agents (that find pages for a particular topic) and selection agents (that find pages for a particular user). Users' explicit ratings of the recommended pages combined with several heuristics are used to update personal-agents' profiles, remove unsuccessful agents and duplicate successful ones. *Broadway* [91] is a case-based reasoning system that monitors a user's browsing activity and provides advice by reusing navigational cases extracted from past browsing experiences of a group of users. Another early Web navigation assistant is *SiteSeer* [92], which recommends pages collaboratively by looking at users' bookmarks. Other early Web suggester systems include *LIRA* [126], *BASAR* [127], *ifWeb* [128], *SOAP* [129], *Let's Browse* [130], *SurfLen* [131], *Margin Notes* [132] and *Quickstep* [133], among many others.

The *Watson* system [34, 93, 94] is a context-aware suggester that attempts to find relevant online resources. *Watson* is part of a family of programs known as *Information Management Assistants (IMAs)* developed at the InfoLab of Northwestern University. The purpose of the *IMAs* is to anticipate the user's needs and to provide proactive and on-demand support for the user's current activity. To achieve this goal, the *IMAs* generate a model of the user's task, access information retrieval systems on the user's behalf and unobtrusively deliver useful material. The *IMAs* provide an environment in which resources are retrieved proactively as well as mechanisms that augment users' explicit queries with keywords extracted from the current task. *Point/Counterpoint* [116] is another *IMA* built on top of *Watson*. Instead of retrieving general information, *Point/Counterpoint* brings opposing arguments.

A family of *Just-In-Time Information Retrieval (JITIR)* systems serving as context-sensitive memory augmentation aides has been implemented at the MIT Media Lab. *JITIR* systems are characterized by being proactive, unobtrusive and aware of the user's local context. An early version of a *JITIR* system, *Remembrance Agent* [33], is designed to run in the background of a computer, observing what the user types and reads on a text editor. *Remembrance Agent* uses that information to retrieve related documents and user's old emails, which become available through an unobtrusive interface. *Wearable Remembrance Agent* [95, 96] is a portable, continuously running agent that uses the physical context to find information relevant to the user's situation. Another memory augmentation device developed at MIT Media Lab is *Memory Glasses* [97], a wearable aid that utilizes a context-awareness system based on sensors for vision and listening. Interaction is performed through buttons for user input into a light wearable computing core, while headphones and a clip-on display are used for information delivery. These aids were preliminary versions of what would later become *Google Glass*, a wearable device that is worn like conventional glasses, but that combines a voice recognition software, a touchpad, a display screen, a high-definition camera, microphone and wireless connectivity.

The *Context Toolkit* [134] is a framework for developing reusable solutions for handling context information in interactive applications. It relies on the concept of context widget, which is a software component that provides applications with access to context information from their operating environment. An example of memory augmentation aid built using the *Context Toolkit* is illustrated by the *CybreMinder* system [98]. *CybreMinder* is a context-aware reminder application with an ample view of context, which includes location, time, activity, identity, physical/environmental conditions and potential co-located collaborators. An important characteristic of *CybreMinder* is its support for customizing the delivery of reminders. Based

on this customization the system employs the user's context to choose among different ways of delivering the reminders, including SMS on a mobile phone, e-mail, printing on a local printer or using a nearby display from a wearable, handheld, or static CRT.

A general-purpose and extensible framework for constructing context-aware assistants is provided by *Suitor* [99, 100]. *Suitor* is a collection of "attentive agents" that gather information from the users and the world and post that information on a centralized blackboard. A class of agents called investigators determine users' information needs by monitoring users' behavior and context, including eye gaze, keyword input, mouse movements, visited URLs and software applications on focus. On the meantime, investigator agents watch Web pages and databases for updates. The second group of agents, the reflector agents, interact with the blackboard by prioritizing posted information and matching it with users' needs. Finally, delivery agents display relevant information to users.

CALVIN [101, 102] is a case-based reasoning context-aware system that monitors the user's Web browsing activity to generate a model of the user's task. Also, it provides capabilities for users to manually enter information about a variety of resources, such as descriptions of books or articles and data on useful contact people. The gathered material is stored as contextualized cases recording information users' consult during their decision-making processes. *CALVIN* provides an interface that proactively and unobtrusively suggests stored content when the user context is similar to the one associated with the stored cases.

The *IntelliZap* system [27] offers context-driven information retrieval. Starting from a text query marked by the user in a document, *IntelliZap* applies semantic keyword extraction and clustering to generate augmented queries automatically. These queries are submitted to general and domain-specific search engines. Finally, results are ranked based on the query context. The system applies natural language processing techniques to the captured context to guide the search process. Only the immediate vicinity of the focus text is considered as search context. A semantic network is used to extract keywords from the context and clustering is applied to identify groups of words that are strongly connected based on a specially defined metric derived from applying the k-means algorithm. The clustering process allows identifying clusters of words that represent different semantic aspects of the query. The most important keywords of each cluster are identified based on their semantic distance to the context text and used to build cluster-specific queries. Re-ranking algorithms are in charge of filtering out irrelevant results.

Stuff I've seen [103] is a system that makes it easy for users to find information they have seen before. The system achieves this by providing a unified index of information (email, Web page, document, media file, calendar appointment, etc.) that a person has seen before. Rich contextual cues such as time, author, thumbnails, and previews can be used as part of the user interface to search for and present information.

Another context-based search tool, the *EXTENDER* system [135, 136, 111, 112], searches the Web based on a knowledge model (in the form of a concept map) that is being constructed by a user. *EXTENDER* applies an incremental technique to build up context descriptions. Its goal is to generate brief descriptions of new topics relevant to the knowledge model under construction. A key feature of the context-based method implemented by *EXTENDER* is the use of topic descriptors (terms that often occur in relevant documents) and topic discriminators (terms that tend to occur only in relevant documents) to describe the topic of the knowledge model under construction. Topic descriptors and discriminators are incrementally identified and used for building queries and filtering results.

Three methods for searching in context are discussed in another early work developed in Yahoo! Inc., [49]. The first one is a *Simple Query Re-Writing* technique that has the purpose of automatically augmenting the user query with terms extracted from the search context. The main limitation of this method is that the more terms are added in a conjunctive way, the more restricted the query is, resulting in low recall. The second approach is a *Rank-Biasing* technique that relies on the use of modified search engines to handle more complex query syntaxes. These queries are made up of two parts, the selection part affecting recall and optional ranking terms that are used to impact the score of the selected documents. The *Rank-Biasing* approach was part of *Y!Q (Yahoo! Contextual Search)* [137], a contextual search system integrated with a search engine with modified data structures, indexing, and query processing code designed to support contextual search in a better way. Finally, a method referred to as *Iterative, Filtering Meta-Search* is introduced, where multiple, simple queries are sent to a standard search engine. This method involves two main steps, to generate subqueries reflecting the context at hand and to aggregate the returned results by filtering and re-ranking them.

Connections [138] is a file system search tool that uses context information collected by monitoring the user activity to expand and reorder traditional content search results. The tool identifies temporal relationships between files by tracing file system calls, which allows exploiting temporal locality.

IACS [139] is an incremental algorithm for context-based information retrieval. The salient feature of this algorithm is its ability to learn richer vocabularies from a partial specification of a topic of interest. *IACS* takes a term-based context description to enter a cycle of mining search engines, performing context-

based filtering of results and refining context descriptions to generate new rounds of queries in an expanded vocabulary.

Context matching [140] is a context-based technique for the retrieval of web documents. The technique offers a context-dependent measure of document term significance during retrieval by relying on the notion that if a term in a given document occurs in that document in the context of the query, then that term is considered significant.

TM4L [141] is an environment for building and using topic-centered learning repositories where an extension of the Topic Maps models is combined with contextual search to enhance retrieval. The proposed environment allows formulating context queries, which enable to specify the boundaries of the contextual space and to filter the result by certain semantic properties.

SenseCam [104] is a wearable device containing a camera and embedded sensors worn around a user's neck, which automatically captures images and other contextual aspects of life events (e.g., ambient light levels, temperature, and movement). The captured context is then used to support people's memory by providing useful links to events in people's past.

In [142], a context-based search platform developed in the frame of the European Integrated Project *VIVACE* is presented. This platform is based on the identification of a user context, the use of a context model to represent the user's context and the subsequent suggestion of applicable knowledge to that particular user.

In [106] a searching paradigm called *Context-Based Search (CBS)* is proposed for searching literature digital collections within digital libraries based on predefined contexts. The approach requires associating publications with pre-specified ontology-based contexts. In the meantime, each publication is assigned a context score, based on its relevance to the given context. Relevant contexts are selected for a given query, and search is performed within the selected contexts. The context scores previously assigned to publications are revised into relevancy scores for the query at hand and the publication context. A family of tools for accessing resources in digital libraries based on thematic contexts is proposed in [107, 108]. One of these tools, called *SciELO Suggester*, provides context-based support for cataloging library resources.

In [24] a semi-supervised incremental method (*IM*) to refine queries based on a given thematic context is proposed. The proposed method strategically selects terms from the given context to define a set of initial queries and incrementally refines these queries based on the notions of topic descriptors and topic discriminators. The retrieved documents that are similar to the original context are used as a source for incrementally refining the query vocabulary.

A family of algorithms for contextualized search based on multi-objective evolutionary algorithms (*MOEA*) is proposed in [37]. These algorithms start from a population of topical queries defined from the context description and evolve these queries guided by multiple objectives with the purpose of retrieving material relevant to the given topic.

In [143], an approach for information re-finding by context (*Context-based re-finding*) inspired by the human memory retrieval mechanism is presented. The proposed method works by building a link between the information available in the current context and its previous accessing context instances, represented as multidimensional vectors. A query is considered a context instance and a ranked list of the matched context instances is an intermediate result that links to previously accessed information. Memory reinforcement is also incorporated by adjusting the degradation speeds of context instances.

The *CONRAD* [39] framework addresses the problem of offering context-dependent notifications based on the user's current activities. The techniques adopted by *CONRAD* operate at different granularity levels by implementing a context filter and a task filter to select the notifications to be delivered. Particular emphasis is placed in unobtrusively informing the user, without interrupting with notifications about irrelevant events or with information that could be conveniently conveyed later on. As a consequence, it applies notification management policies that require recognizing the user's current focus of attention.

DECIDE 2.0 [144] is a framework that combines context-based search with argumentation technologies for intelligent processing of citizens' opinion in Social Media. The framework relies on text mining and opinion mining techniques to filter noise and detect main topics being discussed by citizens in Social Media

The *Context-based recommender for problem-solving in software development (CBRPSSD)* [145] analyzes the software developer's contextual information available through exception traces with the purpose of retrieving question-answering web resources. The objective of the proposed tool is to bridge the gap between the IDE and the Web by automatically searching for relevant information. Another tool for context-based search in software development (*CBSSD*) [146] uses ontologies to represent the source code artifacts that exist in the workspace of the developer. The contextual information of the developer is used to improve the search of source code artifacts in the IDE, helping developers find relevant knowledge more efficiently.

Leibniz [105] is a system developed by Microsoft Research that introduces the notion of "contextual insights" to provide users with information that is contextually relevant to the content that they are consuming

or authoring. Some key aspects of the *Leibniz* system include the prediction of the user's focus of attention, the addition of context terms to the query and the adaptation of the results via post-processing.

CPIR [147] is a context-aware personal information retrieval algorithm, which considers both the participatory and implicit-topical properties of the context to improve the retrieval performance. Users' information needs are implicitly captured in an online conversation context. The participatory context coming from all users participating in a conversation is used for query formulation and expansion while the implicit topical context is used to determine the importance of Web content to the given context. A customized method is developed to extract semantic information from texts that are typically short, while a graph-based algorithm is employed to reveal the implicit relations among the personal Web information of all participating users and extract the user personal Web information that matches the topic of the session.

Grisham [113] is a topic-based search system that allows users to search for and interact with topic-relevant documents in three different ways: (1) keyword-based search and topic-based result re-ranking, (2) topic-based search and ranking based on topic relevancy and (3) visual exploration of topically similar documents. *Grisham* uses LDA [16] to derive topic distributions for articles.

In the *SWELL* project [109], a personal information management system is proposed where the content of documents, keystrokes, mouse clicks, window information and other information objects are used as input variables for context recognition and identification. This information is used to build a network based on the interactive activation and competition model [148], describing associations among topical information and entities in the user context.

Systems for maintaining situational awareness in cyber operations are another example of context-aware proactive aids [149]. Some recent approaches attempt to automatically recognize a user's operational context to aid decision-making proactively. For instance, the *OZONE Widget Framework* [110] is a situation-aware system that captures the objectives, actions, and activities of the analysts and decision maker to provide them with context relevant information proactively. *Sol* [150] is a situation-aware agent-based framework for distributed sensemaking whose goal is to augment human perception and cognition. A cognition-driven decision support system called the *Situation Awareness Support System (SASS)* is proposed in [151]. The goal of *SASS* is to manage abnormal situations in safety-critical environments in which the effect of situational complexity on human decision-makers is a concern. A systematic review of the literature on cyber situational awareness can be found in [152].

A novel strategy for query understanding in Community question answering platforms based on context (*Context-aware cQA*) is proposed in [114]. The proposed methods profit from search sessions for semantically categorizing question-like informational search queries and is implemented to help browse community question answering platforms.

Context also plays a vital role in the area of conversational agents. A conversational agent is a computer system that integrates agent capabilities with computational linguistics [153]. These agents exploit natural language technologies to engage users in information-seeking and task-oriented dialogs for a broad range of applications such as customer service, help desk, website navigation, guided selling and technical support, among others. The role of context is crucial for natural language understanding and to achieve coherence in the conversation [154]. Hence, the exploitation of context while communicating in natural language with conversational agents has been an active area of research for more than two decades. For instance, *TRAINS* [155] is an early conversational agent developed at the University of Rochester that uses natural language to interact with and to assist humans in problem-solving tasks in the route-planning domain. Another conversational agent developed by the same group is *TRIPS* [156, 157], which embodies a more complex model of collaborative problem-solving than *TRAINS*.

HITIQA (High-Quality Interactive Question Answering) [158] is an interactive open-domain question answering technology designed to allow analysts to pose complex exploratory questions in natural language. The system explores the scope and context of the desired answer space using a clarification dialogue based on data-driven semantics. A related system *SmartKom* [159], is a conversational agent that takes advantage of a context model to respond to contextually underspecified questions. A salient feature of *SmartKom* is its symmetric multimodality for dialogue systems in which all input modes (e.g., speech, gesture, facial expression) are also available for output.

*Moir*a (Mobile Intelligence Reporting Agent) [160] is a conversational agent whose main purpose is to mediate interactions with human users by speech or typing. In a subsequent work [161], *Moir*a is integrated with *CISpaces* (Collaborative Intelligence Spaces) [162] to support collaborative sense making among analysts via argumentation-based evidential reasoning.

A significant number of the context-sensitive conversational agents have been developed in the area of intelligent tutoring system, where agents interact with students in natural language while being sensitive to their cognitive and emotional states. Examples of context-sensitive intelligent tutoring systems include *AutoTutor* [163, 164], *RMT* [165] and *Easy with Eve* [166].

Google developed a series of personalized search tools that exploit contextual information at different levels to improve search effectiveness. Initial tools include *Google's Alerts*, *Google's Personalized Search* and *My Search History*. These tools build user profiles by allowing users to express the information in which they are interested or by monitoring user queries and the pages visited via search results.

More recently, major companies such as Google, Apple, Amazon, Microsoft, and Facebook have applied deep learning techniques [167] to many different areas across their products. Context-aware systems with natural language interfaces using voice and text have particularly benefited from these techniques. *Google Now* is a context-aware app for Android and iOS that assists the user on demand, by answering user-initiated queries formulated in natural language, or pro-actively, by delivering information based on the user search habits. *Google Assistant* is an upgrade of *Google Now* designed to keep track of a two-way conversation, determine context and respond with relevant information. *Google Assistant* is part of some Google products such as the voice-activated speaker *Google Home*, the messaging app *Google Allo*, and the *Android Wear* watches. Other context-aware virtual assistants and chatbots with conversational capabilities include *Siri* from Apple, *Alexa* from Amazon, *Cortana* and *XiaoIce* from Microsoft, and *Facebook Messenger bots*.

Figure 3 displays a timeline of systems described in this section. The timeline is organized according to the date that each system was first published or released.

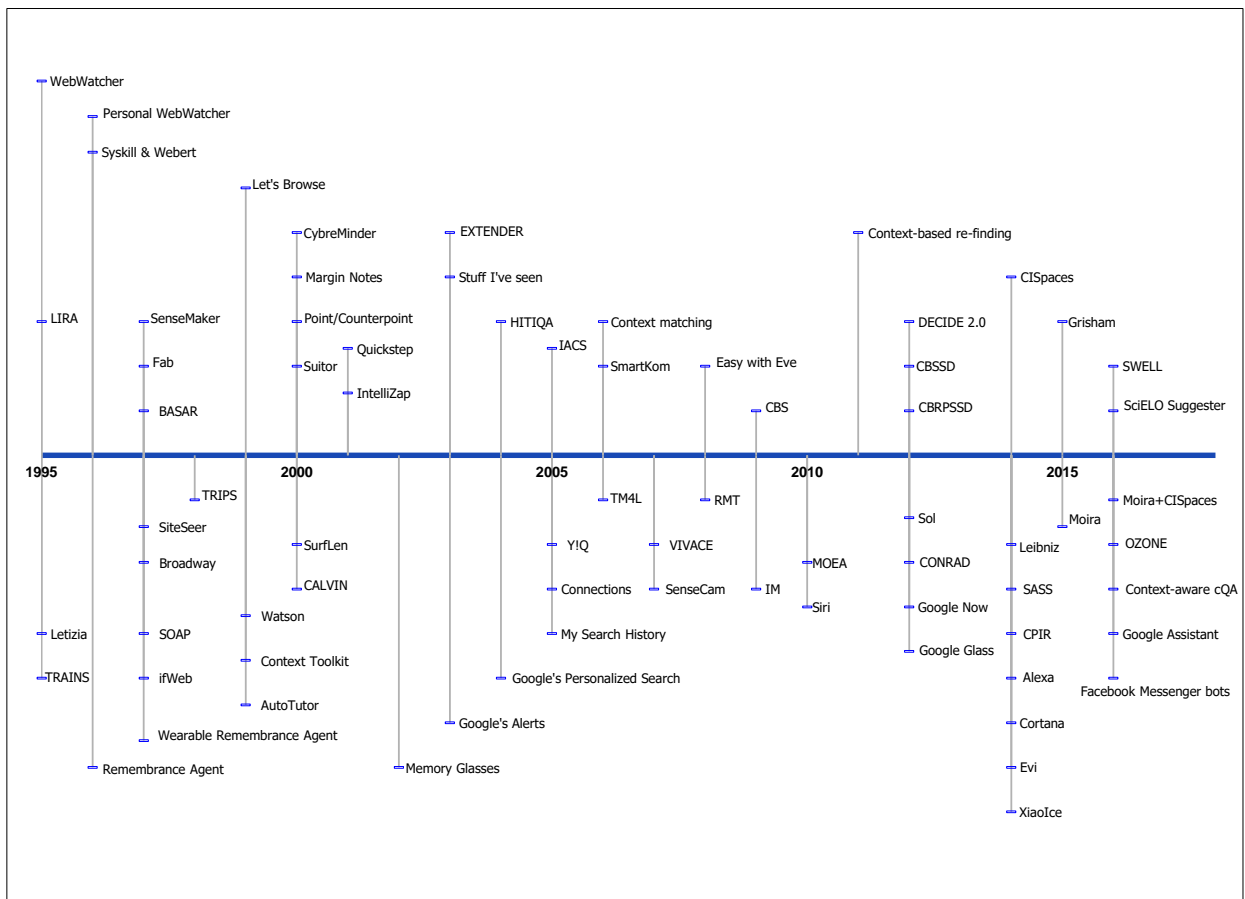


Figure 3: Timeline of context-based search systems.

5 Conclusion and future directions

This article presented a survey of methods and tool for searching in the context of a task. Since search can take place in diverse and rich environments, the problem of incorporating contextual information in the information retrieval process has attracted significant research interest. The methods reviewed in this survey offer solutions to the problem of reflecting a context in the search process through different approaches. These methods take advantage of contextual information to generate, disambiguate or refine queries automatically. Some context-based methods use topical contexts for filtering and ranking results, or for selecting domain-specific search engines. Many context-aware tools illustrate the potential and opportunities that context-based search creates. A major limitation encountered in this research area is the lack of a well-established

evaluation methodology, which is partly due to the context-independent nature of available test collections.

A rapidly evolving research area addresses the problem of topical search in Social Media. Some studies have analyzed the differences between searching for information in Social Media, such as Twitter and searching for information in the Web [168]. A challenging ingredient of the problem of searching in Social Media is its highly dynamic nature. Therefore, tracking and collecting content associated with a particular topic requires learning new vocabularies and adapting the initial representation of the given topic as new relevant subtopics are identified. Another difficulty encountered in Social Media is the problem of topic discovery from short text [169, 170, 171, 172]. A common approach to tackle the sparsity problem that results from text shortness is to aggregate short texts into lengthy pseudo-documents [173, 174, 175]. An interesting problem associated with context-based search in Social Media is the identification of the user social community as an important context dimension [176]. Some approaches have addressed the problem of detecting the user role in Social Media (e.g., [177, 178, 179]), which may offer additional information to identify user intent. Another problem is to understand user information requests in the context of current events around the world, such as social or economic crises, elections, sport events or natural disasters. Being able to integrate event context as part of an information request requires a much ample view of context and higher capabilities to detect these events and their relation to the user information needs. The use of news and Social Media sources, such as Twitter, as sensors of real-world events [180] facilitates the recognition of this important type of context.

The average user is looking for simplicity. Therefore, the tendency of some of the major context-aware search tools is to deliver information proactively, without the explicit formulation of a query by the user. The ability to infer the user's intent is crucial in developing such tools. While several approaches have been proposed to predict the user's intent, research in this area is difficult due to the general lack of realistic data as a result of privacy concerns. Another important issue is how to interact with the user without interrupting the performed task. Detecting information needs, acquiring relevance feedback and unobtrusively delivering information is a major problem with a substantial impact on user satisfaction.

The accelerating pace of innovation is unleashing a wave of technologies that combine symbolic and sub-symbolic approaches to artificial intelligence to learn about the users' needs and to assist them in a great variety of tasks. These technologies include deriving common-sense meaning through massive knowledge bases, such as *Google's Knowledge Base*, and deep learning for natural language understanding. As technology enables the possibility of monitoring user activities and engaging in natural language conversation with the user, artificial intelligence-powered assistants such as *Google Assistant*, *Siri*, *Alexa* and *Cortana* are becoming increasingly usable and are being progressively incorporated as part of devices and home appliances. These innovative tools reveal that there is an ongoing shift from on-demand keyword-based search to high levels of personalization and context-based anticipatory actions.

Intelligent methods with new sensing and searching technologies have already demonstrated that information delivery is becoming ever more common-sense oriented. Conversational interactions are only one way to gather context from the user. Multi-modal interactions in the form of gestures, facial expressions, and other non-verbal behaviors are another source of context that is starting to be usefully exploited. Achieving the ultimate goal of giving people the information they want at the precise moment they need it is to a great degree dependent on being able to recognize and make full use of the many aspects of context. To achieve this goal, context needs to be acquired not only from the user but also from other users, the environment, and the relevant past.

Acknowledgment

This work was supported by CONICET (PIP 11220120100487), MinCyT (PICT 2014-0624) and Universidad Nacional del Sur (PGI-UNS 24/N039).

References

- [1] Q. Liu, H. Ma, E. Chen, and H. Xiong, "A survey of context-aware mobile recommendations," *International Journal of Information Technology & Decision Making*, vol. 12, no. 01, pp. 139–172, 2013.
- [2] P. Makris, D. N. Skoutas, and C. Skianis, "A survey on context-aware mobile and wireless networking: on networking and computing environments' integration," *IEEE Communications Surveys & Tutorials*, vol. 15, no. 1, pp. 362–386, 2013.
- [3] C. Perera, A. Zaslavsky, P. Christen, and D. Georgakopoulos, "Context aware computing for the internet of things: A survey," *IEEE Communications Surveys & Tutorials*, vol. 16, no. 1, pp. 414–454, 2014.

- [4] J. Augusto, A. Aztiria, D. Kramer, and U. Alegre, "A survey on the evolution of the notion of context-awareness," *Applied Artificial Intelligence*, pp. 1–30, 2018.
- [5] A. K. Dey, G. D. Abowd, and D. Salber, "A conceptual framework and a toolkit for supporting the rapid prototyping of context-aware applications," *Human-computer interaction*, vol. 16, no. 2, pp. 97–166, 2001.
- [6] K. A. De Jong, "An analysis of the behavior of a class of genetic adaptive systems." Ph.D. dissertation, University of Michigan, Ann Arbor, MI, USA, 1975, aAI7609381.
- [7] C. Bolchini, C. A. Curino, E. Quintarelli, F. A. Schreiber, and L. Tanca, "A data-oriented survey of context models," *ACM Sigmod Record*, vol. 36, no. 4, pp. 19–26, 2007.
- [8] D. Boughareb and N. Farah, "Context in information retrieval," in *International Conference on Control, Decision and Information Technologies*, ser. CoDIT 2014, 2014, pp. 589–594.
- [9] H. P. Luhn, "A statistical approach to mechanized encoding and searching of literary information," *IBM Journal of research and development*, vol. 1, no. 4, pp. 309–317, 1957.
- [10] S. K. Jones, "A statistical interpretation of term specificity and its application in retrieval," *Journal of Documentation*, vol. 28, pp. 11–21, 1972.
- [11] G. Salton and C. Yang, "On the specification of term values in automatic indexing," *Journal of Documentation*, vol. 29, pp. 351–372, 1973.
- [12] S. C. Deerwester, S. T. Dumais, T. K. Landauer, G. W. Furnas, and R. A. Harshman, "Indexing by latent semantic analysis," *Journal of the American Society of Information Science*, vol. 41, no. 6, pp. 391–407, 1990.
- [13] P. Resnik, "Using information content to evaluate semantic similarity in a taxonomy," in *Proceedings of the 14th International Joint Conference on Artificial Intelligence - Volume 1*, ser. IJCAI'95. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1995, pp. 448–453.
- [14] D. Lin, "An information-theoretic definition of similarity." in *In Proceedings of the Fifteenth International Conference on Machine Learning*, vol. 98, 1998, pp. 296–304.
- [15] P. D. Turney, "Mining the web for synonyms: Pmi-ir versus lsa on toefl," in *Proceedings of the 12th European Conference on Machine Learning*, ser. EMCL '01. London, UK, UK: Springer-Verlag, 2001, pp. 491–502.
- [16] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent dirichlet allocation," *J. Mach. Learn. Res.*, vol. 3, pp. 993–1022, Mar. 2003.
- [17] A. G. Maguitman, F. Menczer, H. Roinestad, and A. Vespignani, "Algorithmic detection of semantic similarity," in *WWW '05: Proceedings of the 14th international conference on World Wide Web*. New York, NY, USA: ACM, 2005, pp. 107–116.
- [18] A. Gruber, M. Rosen-zvi, and Y. Weiss, "Hidden topic markov models," in *In Proceedings of Artificial Intelligence and Statistics*, 2007.
- [19] D. M. Blei, "Probabilistic topic models," *Communications of the ACM*, vol. 55, no. 4, pp. 77–84, Apr. 2012.
- [20] X. Wang, A. McCallum, and X. Wei, "Topical n-grams: Phrase and topic discovery, with an application to information retrieval," in *Seventh IEEE International Conference on Data Mining (ICDM 2007)*. IEEE, 2007, pp. 697–702.
- [21] N. Cristianini, J. Shawe-Taylor, and H. Lodhi, "Latent semantic kernels," in *Proceedings of the Eighteenth International Conference on Machine Learning*, ser. ICML '01. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2001, pp. 66–73.
- [22] N. Liu, B. Zhang, J. Yan, Q. Yang, S. Yan, Z. Chen, F. Bai, and W.-Y. Ma, "Learning similarity measures in non-orthogonal space," in *Proceedings of the thirteenth ACM international conference on Information and knowledge management*. ACM, 2004, pp. 334–341.

- [23] A. Maguitman, D. Leake, T. Reichherzer, and F. Menczer, “Dynamic Extraction of Topic Descriptors and Discriminators: Towards Automatic Context-Based Topic Search,” in *Proceedings of the Thirteenth Conference on Information and Knowledge Management (CIKM)*. New York: ACM Press, November 2004, pp. 463–472.
- [24] C. M. Lorenzetti and A. G. Maguitman, “A semi-supervised incremental algorithm to automatically formulate topical queries,” *Information Sciences*, vol. 179, no. 12, pp. 1881–1892, 2009, including Special Issue on Web Search.
- [25] R. R. Rada, H. Mili, E. Bicknell, and M. Blettner, “Development and application of a metric on Semantic Nets,” *IEEE Transactions on Systems, Man and Cybernetics*, vol. 19, no. 1, pp. 17–30, 1989.
- [26] J. J. Jiang and D. W. Conrath, “Semantic similarity based on corpus statistics and lexical taxonomy,” in *Proceedings of the International Conference on Research in Computational Linguistics (ROCLING X)*, Taiwan, 1998.
- [27] L. Finkelstein, E. Gabrilovich, Y. Matias, E. Rivlin, Z. Solan, G. Wolfman, and E. Ruppín, “Placing search in context: The concept revisited,” in *Proceedings of the 10th international conference on World Wide Web*. ACM, 2001, pp. 406–414.
- [28] R. Kaptein and J. Kamps, “Explicit extraction of topical context,” *Journal of the American Society for Information Science and Technology*, vol. 62, no. 8, pp. 1548–1563, 2011.
- [29] P. A. Chirita, W. Nejdl, R. Paiu, and C. Kohlschütter, “Using odp metadata to personalize search,” in *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, ser. SIGIR ’05. New York, NY, USA: ACM, 2005, pp. 178–185.
- [30] J. Trajkova and S. Gauch, “Improving ontology-based user profiles,” in *Coupling approaches, coupling media and coupling languages for information retrieval*. LE CENTRE DE HAUTES ETUDES INTERNATIONALES D’INFORMATIQUE DOCUMENTAIRE, 2004, pp. 380–390.
- [31] S. Calegari and G. Pasi, “Personal ontologies: Generation of user profiles based on the {YAGO} ontology,” *Information Processing & Management*, vol. 49, no. 3, pp. 640 – 658, 2013, personalization and Recommendation in Information Access.
- [32] N. Hernandez, J. Mothe, C. Chrisment, and D. Egret, “Modeling context through domain ontologies,” *Inf. Retr.*, vol. 10, no. 2, pp. 143–172, Apr. 2007.
- [33] B. J. Rhodes and T. Starner, “The Remembrance Agent: A continuously running automated information retrieval system,” in *Proceedings of The First International Conference on The Practical Application of Intelligent Agents and Multi Agent Technology (PAAM ’96)*. London, UK: AAAI Press, April 1996, pp. 487–495.
- [34] J. Budzik and K. Hammond, “Watson: Anticipating and contextualizing information needs,” in *62nd Annual Meeting of the American Society for Information Science*, Medford, NJ, 1999.
- [35] J.-J. Yang and R. Korfhage, “Query optimization in information retrieval using genetic algorithms,” in *Proceedings of the 5th International Conference on Genetic Algorithms*. Morgan Kaufmann Publishers Inc., 1993, pp. 603–613.
- [36] R. L. Cecchini, C. M. Lorenzetti, A. G. Maguitman, and N. B. Brignole, “Using genetic algorithms to evolve a population of topical queries,” *Information Processing and Management*, vol. 44, no. 6, pp. 1863–1878, 2008.
- [37] —, “Multi-objective Evolutionary Algorithms for Context-based Search,” *Journal of the American Society for Information Science and Technology*, vol. 61, no. 6, pp. 1258–1274, 2010.
- [38] O. Cordón, E. Herrera-Viedma, and M. Luque, “Evolutionary learning of boolean queries by multiobjective genetic programming,” in *International Conference on Parallel Problem Solving from Nature*. Springer, 2002, pp. 710–719.
- [39] L. Ardissono and G. Bosio, “Context-dependent awareness support in open collaboration environments,” *User Modeling and User-Adapted Interaction*, vol. 22, no. 3, pp. 223–254, 2012.
- [40] J. Greenberg, “An examination of the impact of lexical-semantic relationships on retrieval effectiveness during the query expansion process,” Ph.D. dissertation, University of Pittsburgh, 1998.

- [41] M. Chui, “i still haven’t found what i’m looking for: Web searching as query refinement,” Ph.D. dissertation, Indiana University, 2002.
- [42] B. S. Everitt, *Cluster Analysis*. Halsted Press, New York, 1980.
- [43] G. Salton, *Automatic Text Processing*. AddisonWesley, 1989.
- [44] C. Stanfill and D. Waltz, “Toward memory-based reasoning,” *Communications of the ACM*, vol. 29, no. 12, pp. 1213–1228, 1986.
- [45] R. O. Duda and P. E. Hart, *Pattern classification and scene analysis*. New York, Wiley, 1973, b-Swain Q327 .D847.
- [46] E. J. Glover, S. Lawrence, W. P. Birmingham, and C. L. Giles, “Architecture of a metasearch engine that supports user information needs,” in *Proceedings of the eighth international conference on Information and knowledge management*. ACM, 1999, pp. 210–216.
- [47] E. J. Glover, G. W. Flake, S. Lawrence, W. P. Birmingham, A. Kruger, C. L. Giles, and D. M. Pennock, “Improving category specific web search by learning query modifications,” in *Applications and the Internet, 2001. Proceedings. 2001 Symposium on*. IEEE, 2001, pp. 23–32.
- [48] F. M. Coetzee, E. Glover, S. Lawrence, and C. L. Giles, “Feature selection in web applications by roc inflections and powerset pruning,” in *Applications and the Internet, 2001. Proceedings. 2001 Symposium on*. IEEE, 2001, pp. 5–14.
- [49] R. Kraft, F. Maghoul, C. C. Chang, and R. Kumar, “Searching with Context,” in *Proceedings of the 15th International World Wide Web Conference (WWW-15)*, Edinburgh, Scotland, May 2006.
- [50] B. Xiang, D. Jiang, J. Pei, X. Sun, E. Chen, and H. Li, “Context-aware ranking in web search,” in *Proceedings of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, ser. SIGIR ’10. New York, NY, USA: ACM, 2010, pp. 451–458.
- [51] X. Shen, B. Tan, and C. Zhai, “Context-sensitive information retrieval using implicit feedback,” in *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, ser. SIGIR ’05. New York, NY, USA: ACM, 2005, pp. 43–50.
- [52] S. Han, D. He, Z. Yue, and J. Jiang, “Contextual support for collaborative information retrieval,” in *Proceedings of the 2016 ACM on Conference on Human Information Interaction and Retrieval*, ser. CHIIR ’16. New York, NY, USA: ACM, 2016, pp. 33–42.
- [53] C. Lucchese, S. Orlando, R. Perego, F. Silvestri, and G. Tolomei, “Identifying task-based sessions in search engine query logs,” in *Proceedings of the Fourth ACM International Conference on Web Search and Data Mining*, ser. WSDM ’11. New York, NY, USA: ACM, 2011, pp. 277–286.
- [54] W. Hua, Y. Song, H. Wang, and X. Zhou, “Identifying users’ topical tasks in web search,” in *Proceedings of the Sixth ACM International Conference on Web Search and Data Mining*, ser. WSDM ’13. New York, NY, USA: ACM, 2013, pp. 93–102.
- [55] R. Mehrotra and E. Yilmaz, “Towards hierarchies of search tasks & subtasks,” in *Proceedings of the 24th International Conference on World Wide Web*, ser. WWW ’15 Companion. New York, NY, USA: ACM, 2015, pp. 73–74.
- [56] D. Jiang, J. Pei, and H. Li, “Mining search and browse logs for web search: A survey,” *ACM Transactions on Intelligent Systems and Technology (TIST)*, vol. 4, no. 4, p. 57, 2013.
- [57] J. J. Rocchio, “Relevance feedback in information retrieval,” in *The Smart retrieval system - experiments in automatic document processing*, G. Salton, Ed. Englewood Cliffs, NJ: Prentice-Hall, 1971, pp. 313–323.
- [58] B. J. Gao, D. C. Anastasiu, and X. Jiang, “Utilizing User-input Contextual Terms for Query Disambiguation,” in *Proceedings of the 23rd International Conference on Computational Linguistics: Posters*, ser. COLING ’10. Stroudsburg, PA, USA: Association for Computational Linguistics, 2010, pp. 329–337.
- [59] Z. Bar-Yossef and N. Kraus, “Context-sensitive query auto-completion,” in *Proceedings of the 20th International Conference on World Wide Web*, ser. WWW ’11. New York, NY, USA: ACM, 2011, pp. 107–116.

- [60] E. Kharitonov, C. Macdonald, P. Serdyukov, and I. Ounis, "Intent models for contextualising and diversifying query suggestions," in *Proceedings of the 22nd ACM international conference on Conference on information & knowledge management*, ser. CIKM'13. New York, NY, USA: ACM, 2013, pp. 2303–2308.
- [61] F. Cai and M. de Rijke, "Selectively personalizing query auto-completion," in *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval*, ser. SIGIR '16. New York, NY, USA: ACM, 2016, pp. 993–996.
- [62] F. Scholer and H. E. Williams, "Query association for effective retrieval," in *Proceedings of the eleventh international conference on Information and knowledge management*. ACM Press, 2002, pp. 324–331.
- [63] B. Billerbeck, F. Scholer, H. E. Williams, and J. Zobel, "Query expansion using associated queries," in *Proceedings of the twelfth international conference on Information and knowledge management*. ACM Press, 2003, pp. 2–9.
- [64] P. Sondhi, R. Chandrasekar, and R. Rounthwaite, "Using query context models to construct topical search engines," in *Proceedings of the third symposium on Information interaction in context*. ACM, 2010, pp. 75–84.
- [65] C. Carpineto and G. Romano, "A Survey of Automatic Query Expansion in Information Retrieval," *ACM Computing Surveys*, vol. 44, no. 1, pp. 1–50, 2012.
- [66] G. Sukthankar, C. Geib, H. H. Bui, D. Pynadath, and R. P. Goldman, *Plan, activity, and intent recognition: theory and practice*. Newnes, 2014.
- [67] E. Horvitz, J. Breese, D. Heckerman, D. Hovel, and K. Rommelse, "The Lumiere project: Bayesian user modeling for inferring the goals and needs of software users," in *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, Madison, WI, 1998, pp. 256–265. [Online]. Available: citeseer.nj.nec.com/horvitz98lumiegravere.html
- [68] M. Claypool, P. Le, M. Wased, and D. Brown, "Implicit interest indicators," in *Proceedings of the 6th International Conference on Intelligent User Interfaces*, ser. IUI '01. New York, NY, USA: ACM, 2001, pp. 33–40.
- [69] S. Jung, J. L. Herlocker, and J. Webster, "Click data as implicit relevance feedback in web search," *Information Processing & Management*, vol. 43, no. 3, pp. 791–807, 2007.
- [70] Y. Kim, A. Hassan, R. W. White, and I. Zitouni, "Modeling dwell time to predict click-level satisfaction," in *Proceedings of the 7th ACM international conference on Web search and data mining*. ACM, 2014, pp. 193–202.
- [71] Q. Guo and E. Agichtein, "Beyond dwell time: estimating document relevance from cursor movements and other post-click searcher behavior," in *Proceedings of the 21st international conference on World Wide Web*. ACM, 2012, pp. 569–578.
- [72] L. A. Granka, T. Joachims, and G. Gay, "Eye-tracking analysis of user behavior in www search," in *Proceedings of the 27th annual international ACM SIGIR conference on Research and development in information retrieval*. ACM, 2004, pp. 478–479.
- [73] J. Huang, R. White, and G. Buscher, "User see, user point: gaze and cursor alignment in web search," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2012, pp. 1341–1350.
- [74] Q. Guo, H. Jin, D. Lagun, S. Yuan, and E. Agichtein, "Mining touch interaction data on mobile devices to predict web search result relevance," in *Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval*. ACM, 2013, pp. 153–162.
- [75] D. Lagun, C.-H. Hsieh, D. Webster, and V. Navalpakkam, "Towards better measurement of attention and satisfaction in mobile search," in *Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval*. ACM, 2014, pp. 113–122.
- [76] S. Han, I.-H. Hsiao, and D. Parra, "A study of mobile information exploration with multi-touch interactions," in *International Conference on Social Computing, Behavioral-Cultural Modeling, and Prediction*. Springer, 2014, pp. 269–276.

- [77] S. M. Beitzel, E. C. Jensen, A. Chowdhury, D. Grossman, and O. Frieder, “Hourly analysis of a very large topically categorized web query log,” in *Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, ser. SIGIR '04. New York, NY, USA: ACM, 2004, pp. 321–328.
- [78] G. Somlo and A. E. Howe, “QueryTracker: An Agent for Tracking Persistent Information Needs,” in *AAMAS '04: Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems*. Los Alamitos, CA, USA: IEEE Computer Society, 2004, pp. 488–495.
- [79] T. Lau and E. Horvitz, “Patterns of search: analyzing and modeling web query refinement,” in *UM99 User Modeling*. Springer, 1999, pp. 119–128.
- [80] X. Shen, B. Tan, and C. Zhai, “Implicit user modeling for personalized search,” in *Proceedings of the 14th ACM international conference on Information and knowledge management*. ACM, 2005, pp. 824–831.
- [81] T. Štajner, D. Mladenić, and M. Grobelnik, “Exploring contexts and actions in knowledge processes,” in *Proceedings of the 2nd Workshop on Context, Information and Ontologies*, 2010.
- [82] G. Dong and J. Pei, *Sequence data mining*. Springer Science & Business Media, 2007, vol. 33.
- [83] M. Grobelnik, D. Mladenic, and J. Ferlez, “Probabilistic temporal process model for knowledge processes: handling a stream of linked text,” in *Conf. on Data Mining and Data Warehouses (SiKDD 2009)*. Ljubljana, Slovenia, 2009.
- [84] M. Melucci *et al.*, “Contextual search: A computational framework.” *Foundations and Trends in Information Retrieval*, vol. 6, no. 4-5, pp. 257–405, 2012.
- [85] R. Armstrong, D. Freitag, T. Joachims, and T. Mitchell, “WebWatcher: A Learning Apprentice for the World Wide Web,” in *AAAI Spring Symposium on Information Gathering*, 1995, pp. 6–12.
- [86] D. Mladenic, “Personal WebWatcher: Design and implementation,” School of Computer Science, Carnegie-Mellon University, Pittsburgh, USA, Technical Report IJS-DP-7472, 1996.
- [87] H. Lieberman, “Letizia: An agent that assists Web browsing,” in *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence. IJCAI-95*, C. S. Mellish, Ed. Montreal, Quebec, Canada: Morgan Kaufmann publishers Inc.: San Mateo, CA, USA, 1995, pp. 924–929.
- [88] M. J. Pazzani, J. Muramatsu, and D. Billsus, “Syskill & Webert: Identifying interesting Web sites,” in *AAAI/IAAI, Vol. 1*, 1996, pp. 54–61.
- [89] M. Q. W. Baldonado and T. Winograd, “SenseMaker: an information-exploration interface supporting the contextual evolution of a user’s interests,” in *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM Press, 1997, pp. 11–18.
- [90] M. Balabanović and Y. Shoham, “Fab: content-based, collaborative recommendation,” *Communications of the ACM*, vol. 40, no. 3, pp. 66–72, 1997.
- [91] M. Jaczynski and B. Trousse, “BROADWAY: A World Wide Web browsing advisor reusing past navigations from a group of users,” in *Proceedings of the Third UK Case-Based Reasoning Workshop (UKCBBR3)*, Manchester, UK, September 1997.
- [92] J. Rucker and M. J. Polanco, “Siteseer: personalized navigation for the Web,” *Communications of the ACM*, vol. 40, no. 3, pp. 73–76, 1997.
- [93] J. Budzik and K. J. Hammond, “User interactions with everyday applications as context for just-in-time information access,” in *Proceedings of the 2000 International Conference on Intelligent User Interfaces*. New Orleans, Louisiana: ACM Press, 2000.
- [94] J. Budzik, K. J. Hammond, and L. Birnbaum, “Information Access in Context,” *Knowledge-Based Systems*, vol. 14, pp. 37–53, 2001.
- [95] B. Rhodes, “The wearable remembrance agent: A system for augmented memory,” in *Personal Technologies Journal Special Issue on Wearable Computing, Personal Technologies*, vol. 1, 1997, pp. 218–224.

- [96] B. J. Rhodes, N. Minar, and J. Weaver, "Wearable Computing Meets Ubiquitous Computing: Reaping the best of both worlds," in *The Proceedings of The Third International Symposium on Wearable Computers (ISWC '99)*, San Francisco, CA, October 1999.
- [97] R. W. DeVaul and A. Pentland, "Toward the zero attention interface: Wearable subliminal cuing for short term memory support memory," in *Proceedings of ISWC, 2002*, pp. 141–142.
- [98] A. K. Dey and G. D. Abowd, "CybreMinder: A context-aware system for supporting reminders," in *International Symposium on Handheld and Ubiquitous Computing, HUC 2000*, 2000, pp. 172–186.
- [99] P. P. Maglio, R. Barrett, C. S. Campbell, and T. Selker, "SUITOR: an attentive information system," in *Proceedings of the 5th international conference on Intelligent user interfaces*. ACM Press, 2000, pp. 169–176.
- [100] P. P. Maglio and C. S. Campbell, "Attentive agents," *Communications of the ACM*, vol. 46, no. 3, pp. 47–51, 2003.
- [101] D. B. Leake, T. Bauer, A. Maguitman, and D. C. Wilson, "Capture, Storage and Reuse of Lessons about Information Resources: Supporting Task-Based Information Search," in *Proceedings of the AAAI-00 Workshop on Intelligent Lessons Learned Systems. Austin, Texas*. AAAI Press, 2000, pp. 33–37.
- [102] T. Bauer and D. Leake, "A Research Agent Architecture for Real Time Data Collection and Analysis," in *Proceedings of the Workshop on Infrastructure for Agents, MAS, and Scalable MAS*, 2001.
- [103] S. Dumais, E. Cutrell, J. Cadiz, G. Jancke, R. Sarin, and D. C. Robbins, "Stuff I've Seen: A System for Personal Information Retrieval and Re-use," in *Proceedings of the 26th Annual International ACM SIGIR Conference on Research and Development in Informaion Retrieval*, ser. SIGIR '03. New York, NY, USA: ACM, 2003, pp. 72–79.
- [104] A. J. Sellen, A. Fogg, M. Aitken, S. Hodges, C. Rother, and K. Wood, "Do life-logging technologies support memory for the past?: An experimental study using sensecam," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ser. CHI '07. New York, NY, USA: ACM, 2007, pp. 81–90.
- [105] A. Fuxman, P. Pantel, Y. Lv, A. Chandra, P. Chilakamarri, M. Gamon, D. Hamilton, B. Kohlmeier, D. Narayanan, E. Papalexakis, and B. Zhao, "Contextual insights," in *Proceedings of the Companion Publication of the 23rd International Conference on World Wide Web Companion*, ser. WWW Companion '14. Republic and Canton of Geneva, Switzerland: International World Wide Web Conferences Steering Committee, 2014, pp. 265–266.
- [106] N. Ratprasartporn, J. Po, A. Cakmak, S. Bani-Ahmad, and G. Ozsoyoglu, "Context-based literature digital collection search," *The VLDB Journal*, vol. 18, no. 1, pp. 277–301, Jan. 2009.
- [107] P. H. Delgado, A. G. Maguitman, V. M. Ferracutti, and L. A. Herrera, "Using thematic contexts and previous solutions for maintaining and accessing institutional repositories," *Journal of Computer Science and Technology*, vol. 11, no. 2, 2011.
- [108] N. L. Mitzig, M. S. Mitzig, F. A. Martinez, R. A. Piriz, V. M. Ferracutti, M. P. Gonzalez, and A. G. Maguitman, "Scielo suggester: An intelligent support tool for cataloging library resources," *Library & Information Science Research*, vol. 38, no. 1, pp. 39 – 51, 2016.
- [109] M. Sappelli, S. Verberne, and W. Kraaij, "Adapting the interactive activation model for context recognition and identification," *ACM Transactions on Interactive Intelligent Systems (TiiS)*, vol. 6, no. 3, p. 22, 2016.
- [110] B. Newsom, R. Mittu, M. A. Livingston, S. Russell, J. W. Decker, E. Leadbetter, I. S. Moskowitz, A. Gilliam, C. Sibley, J. Coyne *et al.*, "Modeling user behaviors to enable context-aware proactive decision support," in *Context-Enhanced Information Fusion*. Springer, 2016, pp. 231–267.
- [111] D. Leake, A. Maguitman, and T. Reichherzer, "Experience-Based Support for Human-Centered Knowledge Modeling," *Knowledge-Based Systems*, vol. 68, no. 0, pp. 77–87, 2014.
- [112] C. Lorenzetti, A. Maguitman, D. Leake, F. Menczer, and T. Reichherzer, "Mining for topics to suggest knowledge model extensions," *ACM Transactions on Knowledge Discovery from Data*, vol. 11, no. 2, pp. 23:1–23:30, 2016.

- [113] C. E. Grant, C. P. George, V. Kanjilal, S. Nirkiwale, J. N. Wilson, and D. Z. Wang, “A topic-based search, visualization, and exploration system,” in *Proceedings of the Twenty-Eighth International Florida Artificial Intelligence Research Society Conference*, ser. FLAIRS 2015, 2015, pp. 43–48.
- [114] A. Figueroa and G. Neumann, “Context-aware semantic classification of search queries for browsing community question-answering archives,” *Knowledge-Based Systems*, vol. 96, no. C, pp. 1–13, Mar. 2016.
- [115] R. L. Cecchini, C. M. Lorenzetti, A. G. Maguitman, and F. Menczer, “A Semantic Framework for Evaluating Topical Search Methods,” *CLEI Electronic Journal*, vol. 14, no. 1, April 2011.
- [116] J. Budzik, K. J. Hammond, L. Birnbaum, and M. Krema, “Beyond similarity,” in *Proceedings of the 2000 Workshop on Artificial Intelligence and Web Search*. AAAI Press, 2000.
- [117] B. Smyth and P. McClave, “Similarity vs. diversity,” in *Proceedings of the 4th International Conference on Case-Based Reasoning. Vancouver, Canada*, 2001.
- [118] C. Carpineto, M. D’Amico, and G. Romano, “Evaluating subtopic retrieval methods: Clustering versus diversification of search results,” *Information Processing Management*, vol. 48, no. 2, pp. 358–373, Mar. 2012.
- [119] Y. Kim and W. B. Croft, “Diversifying query suggestions based on query documents,” in *Proceedings of the 37th International ACM SIGIR Conference on Research & Development in Information Retrieval*, ser. SIGIR ’14. New York, NY, USA: ACM, 2014, pp. 891–894.
- [120] R. L. T. Santos, C. Macdonald, and I. Ounis, “Search result diversification,” *Foundations and Trends in Information Retrieval*, vol. 9, no. 1, pp. 1–90, Mar. 2015.
- [121] S. Chernov, G. Demartini, E. Herder, M. Kopycki, and W. Nejdl, “Evaluating personal information management using an activity logs enriched desktop dataset,” in *Personal Information Management Workshop at CHI 2008*, April 2008.
- [122] D. Devaurs, A. S. Rath, and S. N. Lindstaedt, “Exploiting the user interaction context for automatic task detection,” *Applied Artificial Intelligence*, vol. 26, no. 1-2, pp. 58–80, 2012.
- [123] L. Tamine-Lechani, M. Boughanem, and M. Daoud, “Evaluation of contextual information retrieval effectiveness: overview of issues and research,” *Knowledge and Information Systems*, vol. 24, no. 1, pp. 1–34, 2010.
- [124] E. Vicente-López, L. M. de Campos, J. M. Fernández-Luna, J. F. Huete, A. Tagua-Jiménez, and C. Tur-Vigil, “An automatic methodology to evaluate personalized information retrieval systems,” *User Modeling and User-Adapted Interaction*, vol. 25, no. 1, pp. 1–37, 2015.
- [125] P. Borlund, “A study of the use of simulated work task situations in interactive information retrieval evaluations: A meta-evaluation,” *Journal of Documentation*, vol. 72, no. 3, pp. 394–413, 2016.
- [126] M. Balabanović, Y. Shoham, and Y. Yun, “An Adaptive Agent for Automated Web Browsing,” Stanford University, Palo Alto, CA, USA, Tech. Rep., 1995.
- [127] C. G. Thomas and G. Fischer, “Using agents to personalize the Web,” in *Proceedings of the 2nd international conference on Intelligent user interfaces*. ACM Press, 1997, pp. 53–60.
- [128] F. Asnicar and C. Tasso, “ifWeb: a Prototype of user models based intelligent agent for document filtering and navigation in the World Wide Web,” in *Sixth International Conference on User Modeling*, Chia Laguna, Sardinia, Italy, June 1997.
- [129] A. Voss and T. Kreifelts, “SOAP: social agents providing people with useful information,” in *Proceedings of the international ACM SIGGROUP conference on Supporting group work: the integration challenge*. ACM Press, 1997, pp. 291–298.
- [130] H. Lieberman, N. W. V. Dyke, and A. S. Vivacqua, “Let’s browse: a collaborative Web browsing agent,” in *Proceedings of the 1999 International Conference on Intelligent User Interfaces (IUI’99)*. Los Angeles, CA, USA: ACM Press, 1999, pp. 65–68.
- [131] X. Fu, J. Budzik, and K. J. Hammond, “Mining navigation history for recommendation,” in *Intelligent User Interfaces*, 2000, pp. 106–112.

- [132] B. J. Rhodes, "Margin Notes: Building a Contextually Aware Associative Memory," in *The Proceedings of the International Conference on Intelligent User Interfaces (IUI '00)*, January 2000.
- [133] S. Middleton, D. DeRoure, and N. Shadbolt, "Capturing knowledge of user preferences: Ontologies in recommender systems," in *Proceedings of the ACM K-CAP'01*. Victoria, Canada: ACM Press, 2001.
- [134] D. Salber, A. K. Dey, and G. D. Abowd, "The context toolkit: Aiding the development of context-enabled applications," in *CHI*, 1999, pp. 434–441.
- [135] D. Leake, A. Maguitman, and T. Reichherzer, "Topic Extraction and Extension to Support Concept Mapping," in *Proceedings of FLAIRS-2003*. AAAI Press, 2003, pp. 325–329.
- [136] D. Leake, A. Maguitman, T. Reichherzer, A. Cañas, M. Carvalho, M. Arguedas, S. Brenes, and T. Eskridge, "Aiding Knowledge Capture by Searching for Extensions of Knowledge Models," in *Proceedings of KCAP-2003*. ACM Press, 2003.
- [137] R. Kraft, F. Maghoul, and C. C. Chang, "Y! q: contextual search at the point of inspiration," in *Proceedings of the 14th ACM international conference on Information and knowledge management*. ACM, 2005, pp. 816–823.
- [138] C. A. N. Soules and G. R. Ganger, "Connections: Using context to enhance file search," in *Proceedings of the Twentieth ACM Symposium on Operating Systems Principles*, ser. SOSP '05. New York, NY, USA: ACM, 2005, pp. 119–132.
- [139] D. Leake, A. Maguitman, and T. Reichherzer, "Exploiting rich context: An incremental approach to context-based web search," *Modeling and Using Context*, pp. 1–10, 2005.
- [140] J. Zakos and B. Verma, "A novel context-based technique for web information retrieval," *World Wide Web*, vol. 9, no. 4, pp. 485–503, 2006.
- [141] C. Dichev and D. Dicheva, "Context-based search in topic-centered digital repositories," *NSDL: Towards Reusable and Shareable Courseware: Topic Maps-Based Digital Libraries*, 2006.
- [142] R. Redon, A. Larsson, R. Leblond, and B. Longueville, "Vivace context based search platform," *Modeling and Using Context*, pp. 397–410, 2007.
- [143] T. Deng, L. Zhao, L. Feng, and W. Xue, "Information re-finding by context: A brain memory inspired approach," in *Proceedings of the 20th ACM International Conference on Information and Knowledge Management*, ser. CIKM '11. New York, NY, USA: ACM, 2011, pp. 1553–1558.
- [144] C. Chesñevar, A. Maguitman, E. Estevez, and R. Brena, "Integrating argumentation technologies and context-based search for intelligent processing of citizens' opinion in social media," in *Proceedings of the 6th International Conference on Theory and Practice of Electronic Governance*, ser. ICEGOV '12. New York, NY, USA: ACM, 2012, pp. 166–170.
- [145] J. Cordeiro, B. Antunes, and P. Gomes, "Context-based recommendation to support problem solving in software development," in *Recommendation Systems for Software Engineering (RSSE), 2012 Third International Workshop on*. IEEE, 2012, pp. 85–89.
- [146] B. Antunes, J. Cordeiro, and P. Gomes, "Context-based search in software development," in *Proceedings of the 20th European Conference on Artificial Intelligence*. IOS Press, 2012, pp. 937–942.
- [147] X. Han, W. Wei, C. Miao, J.-P. Mei, and H. Song, "Context-aware personal information retrieval from multiple social networks," *IEEE Computational Intelligence Magazine*, vol. 9, no. 2, pp. 18–28, 2014.
- [148] J. L. McClelland and D. E. Rumelhart, "An interactive activation model of context effects in letter perception: I. an account of basic findings." *Psychological review*, vol. 88, no. 5, p. 375, 1981.
- [149] P. Barford, M. Dacier, T. G. Dietterich, M. Fredrikson, J. T. Giffin, S. Jajodia, S. Jha, J. H. Li, P. Liu, P. Ning *et al.*, "Cyber sa: Situational awareness for cyber defense," *Cyber Situational Awareness*, vol. 46, no. 1, pp. 3–13, 2010.
- [150] J. M. Bradshaw, M. Carvalho, L. Bunch, T. Eskridge, P. J. Feltovich, M. Johnson, and D. Kidwell, "Sol: An agent-based framework for cyber situation awareness," *KI-Künstliche Intelligenz*, vol. 26, no. 2, pp. 127–140, 2012.

- [151] M. Naderpour, J. Lu, and G. Zhang, “An intelligent situation awareness support system for safety-critical environments,” *Decision Support Systems*, vol. 59, pp. 325–340, 2014.
- [152] U. Franke and J. Brynielsson, “Cyber situational awareness—a systematic review of the literature,” *Computers & Security*, vol. 46, pp. 18–31, 2014.
- [153] J. Lester, K. Branting, and B. Mott, “Conversational agents,” *The Practical Handbook of Internet Computing*, pp. 220–240, 2004.
- [154] R. Porzel and M. Strube, *Towards context-adaptive natural language processing systems*. na, 2002.
- [155] J. F. Allen, L. K. Schubert, G. Ferguson, P. Heeman, C. H. Hwang, T. Kato, M. Light, N. Martin, B. Miller, M. Poesio *et al.*, “The trains project: A case study in building a conversational planning agent,” *Journal of Experimental & Theoretical Artificial Intelligence*, vol. 7, no. 1, pp. 7–48, 1995.
- [156] G. Ferguson, J. F. Allen *et al.*, “Trips: An integrated intelligent problem-solving assistant,” in *AAA/IAAI*, 1998, pp. 567–572.
- [157] J. F. Allen, D. K. Byron, M. Dzikovska, G. Ferguson, L. Galescu, and A. Stent, “Toward conversational human-computer interaction,” *AI magazine*, vol. 22, no. 4, p. 27, 2001.
- [158] S. Small, T. Strzalkowski, T. Janack, T. Liu, S. Ryan, R. Salkin, N. Shimizu, P. Kantor, D. Kelly, R. Rittman *et al.*, “Hitqqa: Scenario based question answering,” in *Proceedings of HLT*, 2004.
- [159] R. Porzel, I. Gurevych, and R. Malaka, “In context: Integrating domain-and situation-specific knowledge,” *SmartKom: Foundations of Multimodal Dialogue Systems*, pp. 269–284, 2006.
- [160] A. Toniolo, T. J. Norman, A. Etuk, F. Cerutti, R. W. Ouyang, M. Srivastava, N. Oren, T. Dropps, J. A. Allen, and P. Sullivan, “Supporting reasoning with different types of evidence in intelligence analysis,” in *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 2015, pp. 781–789.
- [161] A. Toniolo, A. D. Preece, W. Webberley, T. J. Norman, P. Sullivan, and T. Dropps, “Conversational intelligence analysis,” in *Proceedings of the 17th International Conference on Distributed Computing and Networking*, ser. ICDCN '16. New York, NY, USA: ACM, 2016, pp. 42:1–42:6.
- [162] A. Preece, C. Gwilliams, C. Parizas, D. Pizzocaro, J. Z. Bakdash, and D. Braines, “Conversational sensing,” in *SPIE Sensing Technology + Applications*. International Society for Optics and Photonics, 2014, pp. 91 220I–91 220I.
- [163] A. C. Graesser, K. Wiemer-Hastings, P. Wiemer-Hastings, R. Kreuz, T. R. Group *et al.*, “Autotutor: A simulation of a human tutor,” *Cognitive Systems Research*, vol. 1, no. 1, pp. 35–51, 1999.
- [164] A. C. Graesser, H. Li, and C. Forsyth, “Learning by communicating in natural language with conversational agents,” *Current Directions in Psychological Science*, vol. 23, no. 5, pp. 374–380, 2014.
- [165] E. Arnott, P. Hastings, and D. Allbritton, “Research methods tutor: Evaluation of a dialogue-based tutoring system in the classroom,” *Behavior Research Methods*, vol. 40, no. 3, pp. 694–698, 2008.
- [166] A. Sarrafzadeh, S. Alexander, F. Dadgostar, C. Fan, and A. Bigdeli, “How do you know that I don’t understand? A look at the future of intelligent tutoring systems,” *Computers in Human Behavior*, vol. 24, no. 4, pp. 1342–1363, 2008.
- [167] Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [168] J. Teevan, D. Ramage, and M. R. Morris, “#twittersearch: A comparison of microblog search and web search,” in *Proceedings of the Fourth ACM International Conference on Web Search and Data Mining*, ser. WSDM '11. New York, NY, USA: ACM, 2011, pp. 35–44.
- [169] A. McCallum, X. Wang, and A. Corrada-Emmanuel, “Topic and role discovery in social networks with experiments on enron and academic email,” *J. Artif. Int. Res.*, vol. 30, no. 1, pp. 249–272, Oct. 2007.
- [170] D. Ramage, S. Dumais, and D. Liebling, “Characterizing microblogs with topic models,” in *Proc. ICWSM 2010*. American Association for Artificial Intelligence, May 2010.
- [171] X. Yan, J. Guo, Y. Lan, and X. Cheng, “A biterm topic model for short texts,” in *Proceedings of the 22nd international conference on World Wide Web*. ACM, 2013, pp. 1445–1456.

- [172] T. Lin, W. Tian, Q. Mei, and H. Cheng, “The dual-sparse topic model: Mining focused topics and focused terms in short text,” in *Proceedings of the 23rd International Conference on World Wide Web*, ser. WWW '14. New York, NY, USA: ACM, 2014, pp. 539–550.
- [173] J. Weng, E.-P. Lim, J. Jiang, and Q. He, “Titterrank: finding topic-sensitive influential twitterers,” in *Proceedings of the third ACM international conference on Web search and data mining*. ACM, 2010, pp. 261–270.
- [174] L. Hong and B. D. Davison, “Empirical study of topic modeling in twitter,” in *Proceedings of the first workshop on social media analytics*. ACM, 2010, pp. 80–88.
- [175] R. Mehrotra, S. Sanner, W. Buntine, and L. Xie, “Improving lda topic models for microblogs via tweet pooling and automatic labeling,” in *Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval*. ACM, 2013, pp. 889–892.
- [176] A. Zimmermann, A. Lorenz, and R. Oppermann, “An operational definition of context,” in *International and Interdisciplinary Conference on Modeling and Using Context*. Springer, 2007, pp. 558–571.
- [177] A. Bechmann and S. Lomborg, “Mapping actor roles in social media: Different perspectives on value creation in theories of user participation,” *New media & society*, vol. 15, no. 5, pp. 765–781, 2013.
- [178] S. Malinen, “Understanding user participation in online communities: A systematic literature review of empirical studies,” *Computers in Human Behavior*, vol. 46, pp. 228–238, 2015.
- [179] P. N. Terevinto, A. Pont, J. A. Gil, and J. Domenech, “A flexible workload model based on roles of interactive users in social networks,” in *IFIP Networking Conference (IFIP Networking) and Workshops, 2016*. IFIP, 2016, pp. 524–529.
- [180] T. Sakaki, M. Okazaki, and Y. Matsuo, “Earthquake shakes twitter users: Real-time event detection by social sensors,” in *Proceedings of the 19th International Conference on World Wide Web*, ser. WWW '10. New York, NY, USA: ACM, 2010, pp. 851–860.