



Open Archive TOULOUSE Archive Ouverte (OATAO)

OATAO is an open access repository that collects the work of Toulouse researchers and makes it freely available over the web where possible.

This is an author-deposited version published in : <http://oatao.univ-toulouse.fr/>
Eprints ID : 12669

To link to this article : DOI :10.1145/2523817
URL : <http://dx.doi.org/10.1145/2523817>

To cite this version : Kopluku, Arlind and Pinel-Sauvagnat, Karen and Boughanem, Mohand *Aggregated search: a new information retrieval paradigm*. (2014) ACM Computing Surveys, vol. 46 (n° 3). pp. 1-31.
ISSN 0360-0300

Any correspondence concerning this service should be sent to the repository administrator: staff-oatao@listes-diff.inp-toulouse.fr

Aggregated Search: A New Information Retrieval Paradigm

ARLIND KOPLIKU, KAREN PINEL-SAUVAGNAT, and MOHAND BOUGHANEM, IRIT,
University of Paul Sabatier

Traditional search engines return ranked lists of search results. It is up to the user to scroll this list, scan within different documents, and assemble information that fulfill his/her information need. *Aggregated search* represents a new class of approaches where the information is not only retrieved but also assembled. This is the current evolution in Web search, where diverse content (images, videos, etc.) and relational content (similar entities, features) are included in search results.

In this survey, we propose a simple analysis framework for aggregated search and an overview of existing work. We start with related work in related domains such as federated search, natural language generation, and question answering. Then we focus on more recent trends, namely *cross vertical aggregated search* and *relational aggregated search*, which are already present in current Web search.

Categories and Subject Descriptors: H.3.3 [Information Systems]: Information Search and Retrieval

General Terms: Theory, Measurement, Algorithms

Additional Key Words and Phrases: Information retrieval, aggregated search, focused retrieval, result aggregation, vertical search, information extraction, relational search

1. INTRODUCTION

1.1. Context and Definitions

In response to a query, traditional information retrieval systems return a list of ranked documents that match the query. In general, the user looking for relevant material browses and examines the returned documents to find those that are likely to fulfill his need. If lucky, the user will find in this list the document that satisfies completely his/her need. However, when one document alone is not enough (i.e., the relevant information is scattered in different documents), the user has to collect and aggregate information coming from different documents to build the most appropriate response to his/her need. Combining this different information to achieve, at the same time, better focus and better organization within the search results is the scope of *aggregated search*, which is defined to go beyond the uniform ranking of snippets. Its goal is to assemble useful information from one or multiple sources through one interface [Murdock and Lalmas 2008; Kopliku 2011].

If we think of a perfect search engine, we will soon realize that it cannot answer all possible queries in the same way. For some queries, we just need a precise answer

Authors' addresses: IRIT, University Paul Sabatier, 118 Rte de Narbonne, 31062, Toulouse, France; email: {kopliku, sauvagnat, bougha}@irit.fr.

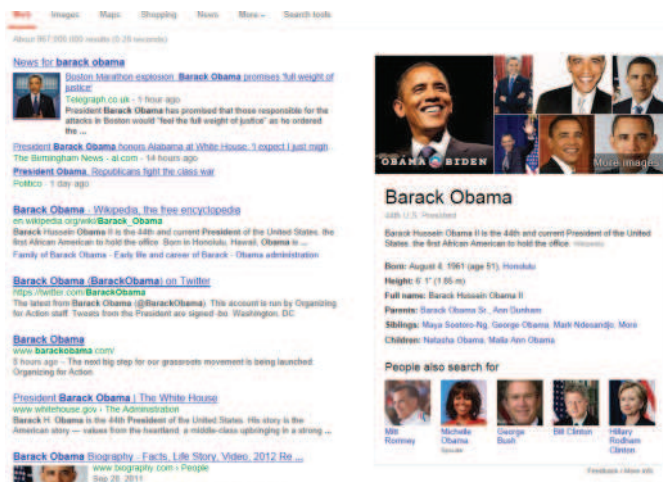


Fig. 1. Google search results on the query “Barack Obama,” accessed on April 2013.

(e.g., weight of Nokia e72). For some others, we need a list of features (e.g., features of Nokia e72). For the query “Nokia e72,” the useful content may include images, videos, a description, features, the Nokia homepage, dedicated pages to the Nokia e72 phone, and so on. Another interesting query can be the name of a person, for example, the US president, Barack Obama. It would be interesting to see in response to the query a short bio of the president, associated news, as well as associated entities with the name of the concerned association (president of the United States, married to Michelle Obama, adversary of Mitt Romney, visiting this month Great Britain, etc.). Answering every query perfectly is impossible because the user will not explain his/her need in an unambiguous and detailed fashion, and often the user does not even know what he/she is exactly looking for. However, for many queries, it is possible to do better than a simple uniform list of document snippets. Major search engines are moving on from this approach toward aggregated search. They pay more attention to search result snippets, they include images, videos, news, and so on, and in recent months, they also group on a separate frame related information. For instance, for the query “Barack Obama,” Google¹ provides a frame that groups a description, images, and a very short bio (see Figure 1). The query is answered with traditional Web search results, vertical search results and relational content provided by the Google Knowledge Graph [Singhal 2012]. Upon query, search engines should be able to integrate diverse content, provide precise answers, put together related content, and so on. In this perspective, aggregated search fits naturally in an effort to group all IR approaches that aim more than ranking. It pushes research effort toward result aggregation and query interpretation, and it helps in understanding current IR evolution.

Aggregated search is not purely a new research area, as it derives from many previous works from many IR subfields such as federated search, metasearch, natural language generation, question answering, semantic search, and entity-oriented search. At first sight, one might find it hard to see that all these domains fit in aggregated search. We will later show how they all can be integrated in one framework that enables having more relevant content, more focus, and better result aggregation. There are, however, some research directions that directly come within the aggregated search *paradigm* and that concern the major recent trends in IR. We will highlight here

¹April 2013.

two research directions that group many existing and current research, namely, *cross-vertical Aggregated Search (cvAS)* and *Relational Aggregated Search (RAS)*.

- Cross-vertical Aggregated Search (cvAS)* is almost omnipresent in Web search nowadays. It consists of including diverse content such as (images, videos, news, etc.) within traditional search results through the use of vertical search engines. Vertical search engines are specialized search engines which work well for a specific kind of content (maps, video, image, news, etc.). Research in this direction includes studies on the interest of the inclusion of vertical content, the proposal of different approaches on reassembling or reranking search results, as well as ideas surrounding evaluation. Federated search is the ancestor of this field and has inspired most of the approaches. However, we find here many unanswered research questions given the specificity of Web search and vertical search engines.
- Relational Aggregated Search (RAS)*. Relational information can often be useful to answer queries right away and/or to structure search results. For instance, the relation triple (“Paris,” *capital of*, “France”) can be used to answer the query “capital of France.” On the other hand, the results for the query “Oscar Wilde” may contain a list of his books shown with a title and description each. This is a simple example where relations such as (“X,” *written by*, “Y”) and (“A,” *description of*, “B”) come into play. Relations can be the implicit output of a vertical search (e.g., “I,” *image of*, “Y”) can be seen as the output of image search), but most of relations are to be extracted explicitly from the documents content (Information Extraction) or to be identified within existing relational sources (knowledge graphs, databases). Forms of relational aggregated search are already met in major search engines, where features, descriptions, and related entities are being shown aside search results (see Figure 1). RAS is thus complementary to cvAS where no explicit result aggregation is done and where search results remain a list of unrelated items from different sources with some grouping (by content type).

In its broadest definition, we say that aggregated search aims at retrieving and assembling *information nuggets*, where an *information nugget* can be of whatever format (text, image, video, etc.) and granularity (document, passage, word, etc.). The key problem that makes the aggregated search complex comes mainly from the fact that the assembly of information is topic dependent (i.e., aggregated results cannot be built a priori). Indeed, one cannot predict and construct all possible combinations that can potentially answer user queries. The composition of information nuggets that meet the query constraints is made upon the query. This leads to several research issues including the identification of candidate information nuggets for aggregation, the ways to assemble content, and a theoretical framework that may support the evaluation of the quality of an aggregated result.

The contribution of this article can be summarized in some main points:

- We group in one class (aggregated retrieval class) multiple approaches that do not fit well in the Boolean or ranked retrieval paradigm.
- We propose a general framework for aggregated search that can be used to analyze and classify the different existing approaches.
- We propose an overview of the broad research domains that are related to aggregated search.
- We propose a detailed survey on cross-vertical aggregated search and relational aggregated search.
- We list different challenging and promising research directions.

The rest of the article is structured as follows. In the next section, we introduce the motivation behind aggregated search. Then, in Section 3, we present a general

framework for aggregated search decomposed in three main components. In Section 4, we provide an overview of existing approaches organized in broad classes such as federated search and natural language generation. From the existing approaches, we select cross-vertical aggregated search and relational aggregated search as most promising and novel. Approaches of these two classes are analyzed in detail in Sections 5 and 6. Then, in Section 7, we discuss how the evaluation of aggregated search is carried out in the literature. The last section is about conclusions and future research directions.

2. MOTIVATION FOR AGGREGATED SEARCH

Most of current Information Retrieval (IR) systems fall into the ranked retrieval paradigm, that is, in response to a query they return a list of ranked documents that match the query. Typically, results are ranked by scoring functions that combine different features generated from the query and the documents. This matching process is driven by different theoretical models such as the vector space model [Salton et al. 1975], the probabilistic model [Robertson and Walker 1994], language models [Ponte and Croft 1998], or fuzzy models [Baziz et al. 2007]. However, in all these models, the relevance score is computed at the document level (i.e., we assume that the entire document is relevant or not).

Different arguments exist with respect to the definition of relevance at document level and its impact on system performance [Boyce 1982; Spärck-Jones et al. 2007; Murdock and Lalmas 2008]. In this section, we will list some of the limitations and challenges of the ranked document retrieval paradigm:

- Data dispersal:** The relevant information can be scattered in different documents [Murdock and Lalmas 2008]. The ranked list for these cases is inadequate, because the user has to scan within different documents to satisfy his information need. This can be a time-consuming and burdensome process.
- Lack of focus:** Ranked retrieval approaches provide a ranked list of uniformly presented results. Typically in Web search, each result is a snippet composed of the result title, a link to the document, and a summary of the linked document. Instead of directing the user to the whole document, for queries when the answer is just a part of document, it might be better to return this part of document right away. The uniform snippets usually returned by search engines might not always be appropriate.
- Ambiguity:** Many queries can be ambiguous in terms of information need. The reference example is *Jaguar*, which can refer to a car, an animal, an operating system, and so on. Diversity can help to disambiguate, especially when multiple sources are used to compose the answer. Ideally, we should return one answer per query interpretation [Spärck-Jones et al. 2007]. This can be multiple ranked lists or linked sets of results.

Data dispersal, lack of focus and ambiguity are just some of the limitations of traditional ranked retrieval. In *aggregated search (aggregated retrieval)*, the final result is not necessarily a ranked list of documents or snippets, but it can be whatever combination of content that can turn useful to the user. Aggregated search can thus propose more focus, more organization, and more diversity in search results.

3. A GENERAL FRAMEWORK FOR AGGREGATED SEARCH

The first definition of aggregated search was given at the SIGIR 2008 workshop [Murdock and Lalmas 2008]:

Definition: *Aggregated search is the task of searching and assembling information from a variety of sources, placing it into a single interface.*

In a broad interpretation, aggregated search concerns at the same time the retrieval and the assembly of search results. Here, we propose a unified framework for




Fig. 2. The general framework of the aggregated search process.

aggregated search that facilitates the analysis of the many and diverse approaches related to aggregated search. The framework, shown in Figure 2, involves three main components, namely, *Query Dispatching (QD)*, *Nuggets Retrieval (NR)*, and *Result Aggregation (RA)*. In other terms, a query is processed and sent to potentially multiple sources of information. Each source returns information nuggets of some granularity and media type, which are to be assembled in the final answer. The answer can be whatever sensed organization of information (though not just the ranked list). This framework has the following multiple purpose: It enables a top-down view of aggregated search, it generalizes well most existing approaches and identifies the main components of aggregated search, it helps to classify and analyze existing approaches; and it makes a clear distinction of the result aggregation process.

Two terms are important to be well-defined at this point for the understanding of this framework: (i) an *information nugget*² is a generalization of content of some granularity and multimedia format, and (ii) a *source* refers to a search engine that relies on at least one collection and one search algorithm. A good definition of the term “source” is important here because we use it to classify aggregated search systems as *monosource* or *multisource*. When there are multiple sources, it is up to the *query dispatching* step to select the sources to be used.

Each of the three components of this framework will be detailed further in the next subsections.

3.1. Query Dispatching

In the query dispatching step, we include the actions that precede query matching, that is, initial interpretation of the query (query analysis) and other actions that depend mainly on the query and knowledge about the collections. We can also see this step as deciding which solutions should be triggered for a given query. We distinguish between the approaches that aim at selecting the right sources to be used, the approaches that try to understand more about the query, and the approaches that try to extend the query to have more chances to find good search results. We will list here briefly well-known problems in this context:

Source selection: Source selection is one of the most well-known problems in aggregated search with multiple sources. Given a set of sources, its goal is to select the sources that are likely to answer the query. Identification of key terms can, for instance,

²The term *information nugget* has been used frequently in research to denote semantic pieces of information [Clarke et al. 2008; Goecks 2002] and, in particular, in question answering [Kelly and Lin 2007; Voorhees 2003], although without any common agreement on its meaning. Here, we give a general definition that suits also the context of aggregated search.

be used to select sources. The presence of special keywords is often useful to understand the type of answer the user is expecting. For example, the presence of words such as “weather” or “definition” is a useful clue that is used by major search engines to trigger a personalized answer. Facebook graph search recognizes terms such as “likes,” “friends,” and “photos.” Named entities have also been shown to be very frequent within search queries [Guo et al. 2009]. Identifying the named entities and their types can assist the IR process. For instance, if the query is “Oscar Wilde,” we can query a knowledge base to provide the user a description and an extract of important features on the writer. As well, we can propose similar authors or authors of the same period. We will discuss this task in more detail when we will discuss multisource aggregated search approaches and, in particular, cross-vertical aggregated search.

Question intent: In question answering, queries are analyzed to find what type of question is being asked (why, where, when, yes/no questions) and to identify syntactic or semantic relations that might exist between the answer and entities or events mentioned in the question [Hirschman and Gaizauskas 2001]. Queries can also be decomposed to ease the retrieval process [Moriceau and Tannier 2010; Lin and Liu 2008; Katz et al. 2005]. Understanding the query intent can also be of interest for some sources: One can, for example, cite some approaches that try to correctly interpret queries with little text on knowledge bases [Pound et al. 2012; Tran et al. 2009].

Query reformulation: Depending on the sources that will be queried, some adaptation/personalization of the query may be needed. For instance, in a data integration perspective, the user only sees the general schema (i.e., the unified view of the different sources of data). A reformulation step is needed: The query on the global schema has to be reformulated in a set of queries on the different sources [Lenzerini 2002].

3.2. Nugget Retrieval

Nugget retrieval, as presented here, is situated between query dispatching and result aggregation, that is, it corresponds basically to the definition of a source which takes as input a query and matches it with a (scored) set of potentially relevant information nuggets. In IR, there are many types of nugget retrieval mostly depending on the type of nugget being returned. We enumerate here some of the most well-known approaches.

As mentioned earlier, it is possible to retrieve entire documents or parts of documents. This corresponds to the distinction between *document retrieval* and *passage retrieval* (also called *focused retrieval* in semistructured collections) [Geva et al. 2009; Kamps et al. 2008; Trotman et al. 2010; Sauvagnat et al. 2006]. The retrieval process also depends on the multimedia type (text, image, video). We distinguish here *textual retrieval* and *multimedia retrieval*. Some search engines such as Web search engines retrieve content of heterogeneous type (Webpages, images, news articles, videos). When the one-size-fits-all solution does not work well, *vertical search* solutions, which are specialized on a media type, query type, task, result type, and so on, are commonly used. We should also mention the approaches that retrieve information from relational databases, knowledge graphs, or extracted information. They deal with relational information, that is, information with some explicit semantic. The results can be tables, table lines, semantic triples, and the like.

In aggregated search systems with multiple sources, each source triggers its own nugget retrieval process. The retrieved results will not necessarily have comparable scores. Assembling search results from multiple sources has been an intensive research area for at least two decades. It starts with *federated search* (Distributed Information Retrieval) [Callan 2000] in the context of distributed data collections (hidden Web, databases, etc.). Then, it evolved into areas such as *metasearch* [Selberg and Etzioni 1995; Srinivas et al. 2011] and more recently into *cross-vertical aggregated search* [Arguello et al. 2009; Lalmas 2011]. Today, the latter represents the most successful

class of approaches and it is applied by almost all major Web search engines. To distinguish between these approaches, it is important to have a clear definition of the term “source.” The terminology in literature is a little messy, as terms such as search engine, resource, and collection are often used as synonyms. We recall that we prefer using the term “source” to refer to a search engine or a component of a search engine that uses at least one collection and one search algorithm. This enables us to classify multisource retrieval approaches such as metasearch, federated search, mashups, and cross-vertical aggregated search. We will describe multisource approaches in detail in Section 4.

To conclude, we can say that the nugget retrieval output depends on the type of sources and the number of sources being used. This affects the focus and the diversity within retrieved information nuggets.

3.3. Result Aggregation

Result aggregation starts at the moment we have a set of potential relevant information nuggets. It involves the different ways of putting content together. In this section, we list the most generic actions that can be applied on candidate search results before showing them to the user. The goal is to show that we can do more than ranking (one possible action). More details on specific result aggregation approaches will be provided later. We could identify five basic aggregation actions: *sorting*, *grouping*, *merging*, *splitting*, and *extracting*. These actions can be met alone or combined (for instance, sorting can be applied on nuggets or group of nuggets). We will detail each of them in the following:

- Sorting:** Given a set of information nuggets $n_1, n_2 \dots n_m$, the *sorting* action produces another list of nuggets $n_{i_1}, n_{i_2} \dots n_{i_m}$ where all elements are ranked with respect to some features. These features can be relevance scores, as well as time, author, location, popularity scores, and so on. Although we mention that the goal of aggregated search is to go beyond the ranked list approach, ranking (sorting) by relevance or by other features remains fundamental for most IR systems.
- Grouping:** Given a set of information nuggets $n_1, n_2 \dots n_m$, the *grouping* action produces groups (sets) of nuggets $G_1, G_2 \dots G_i$ where elements of the same group share some property. A group of results can be composed of results with similar content, results that have happened in the same period in time, or results with a common feature. Among grouping approaches, we can mention clustering [Hearst and Pedersen 1996] and classification [Manning et al. 2008] as special and illustrative cases.
- Merging:** We refer to *merging* as an action that takes a set of information nuggets $n_1, n_2 \dots n_m$ and produces a new cohesive aggregation unit. This is different from grouping because it produces one final grouped unit and not multiple groups. The output of this action can be a new document, a book, a summary, or an object. For instance, multidocument summarization [Dang 2006] takes multiple documents and produces one new aggregate unit, the summary. In object-level search [Nie et al. 2007a, 2007b], information is merged around named entities (e.g., Nokia e72, France, Paris, ...) into so-called *objects*.
- Splitting:** Given some content, we can do the opposite action of merging. We can decompose the content in other smaller nuggets. The result of this decomposition can be one smaller nugget or a set of smaller nuggets. Given some content n , the result of splitting is a set of information nuggets $n_1, n_2 \dots n_m$ with $m \geq 1$. Typically, splitting can produce a total partition over the initial content n such that $n_1 \cup n_2 \cup \dots \cup n_m = n$. Splitting can be as simple as decomposing pure text into passages or sentences, but it quickly becomes a complex task when markup language is present (e.g., HTML pages).
- Extracting:** Extracting is more about identifying one or more semantically sound information nuggets within some content. This can be part-of-speech items

(e.g., nouns, verbs, pronouns), named entities (e.g., Italy, Rome, Roberto Carlos, Hotel Bellagio), images, videos, and so on. The target of this action is not a total partition of the initial content rather than specific type of stand-alone content (i.e., content that can exist alone). The last two actions are not simply useful to decompose some content, but they can precede some sensed re-use of the content.

Whatever the generic actions applied on candidate nuggets, the resulting aggregate should be coherent, with complementary nuggets, and without redundancy. Finding the same nugget several times can, however, help to determine its importance. Probabilistic sampling can, for instance, be used as a way to select appropriate nuggets [Downey et al. 2005].

In the next section, we list different forms of aggregated search listed around research domains.

4. ANALYSIS OF AGGREGATED SEARCH APPROACHES

In this section, we overview work from different IR domains related directly to the definition of aggregated search or may be inspiring to its future development. In natural language generation, we find interesting approaches that focus on result aggregation (generation of documents from scratch). In question answering, there are inspiring case studies with respect to query interpretation and result aggregation. On the other hand, we will see how work from federated search naturally evolves to a specific and promising research direction we call *cross-vertical aggregated search*. In the quest of aggregation, we group together all approaches that retrieve relations in one broad class named relational aggregated search. The latter is claimed complementary to cvAS because relations can indeed bring to real assembly of information. We will not forget to mention here domain-specific case studies. All of the approaches will be analyzed in the light of our framework.

4.1. Natural Language Generation

The goal of Natural Language Generation (NLG) is similar to the broad goal of aggregated search, that is, it aims to avoid the user the burden to browse several documents through the automatic generation of answers in an appropriate linguistic form [Paris et al. 2010]. Let us analyze NLG in the light of our framework. NLG cares less about query dispatching. Indeed, the query is not always present, that is, some NLG approaches start from an information need input as a query and others are designed starting from a known context of use (information need is implicit). In Grice [1975] and Paris et al. [2010], the final goal is the generated document, and this goal is to be decomposed in subgoals. The query can also be manually decomposed in subqueries (beforehand) [Paris et al. 2010]. Nugget retrieval is not the major goal either. In some approaches, the content to be assembled is known beforehand or filtered in a pipeline [Paris and Colineau 2006]. In others, we rely on existing search engines or information coming from databases or XML collections [Paris et al. 2001, 2005]. The assembly of content in a coherent linguistic fashion is indeed the main goal of NLG. The retrieved content is aggregated using prototypical ways of organizing information known also as *discourse strategies* (result aggregation). This prototypical structures can be observed (e.g., from linguists) or learned. Returned facts can be organized based on their relationships. For instance, they can be ordered chronologically, they may have a cause-effect relation, background information is shown first, and so on [Paris 1988]. Discourse strategies can be static (i.e., known beforehand) or dynamic. The dynamic strategies depend not only on the information need but also on the availability of information.

To illustrate this, we provide some examples from literature across different application domains [Paris and Colineau 2006; Paris et al. 2001, 2005]. In Paris et al.

[2005], NLG is applied on a surveillance context. Instead of having monitors poll from different sources, authors propose an NLG approach. Here, data comes from various public-domain information sources (databases, reports, etc.), and it is assembled as a multipage Web site for port surveillance operators. There is no explicit query, but the abundant information that comes from different sources in different interfaces is put together into one aggregate answer. The second application [Paris et al. 2001] concerns traveling. The user can input its travel destination and optionally budget constraints. Using samples of travel guides, a discourse strategy is learned and then applied to generate coherent travel guides from scratch upon query. In Paris and Colineau [2006], the NLG approaches are used to produce brochures that give summaries on given organizations. Another interesting example is met in Sauper and Barzilay [2009]. Authors propose an approach to automatically generate Wikipedia-like documents through machine learning on Wikipedia pages of the same class (e.g., diseases, actors). The learned document structure (e.g., causes, treatment, symptoms) is then used to automatically build new documents for other instances of the class extracting missing passages from the Web.

NLG is an interesting case study in particular for result aggregation. Its limits relate to query dispatching and nugget retrieval, which might be the reason NLG approaches are more successful in domain-specific applications rather than in uses through free natural language querying. We do not focus more on this class of approaches, but we recommend the interested reader a qualitative survey in Paris et al. [2010].

4.2. Question Answering

Question Answering (QA) differs from traditional ranked retrieval, as it does not aim a list of documents, but one or multiple answers. It represents an interesting case study for aggregated search because these answers might not exist—they have to be produced, extracted, assembled [Dang et al. 2007].

In QA, we also find the three main components of our AS framework. Concerning query dispatching, queries in QA are not any free text rather than question-like queries. For questions, there exist different taxonomies, and it is in the interest of the QA system to understand the query type. We can list here some well-known question types such as the “Who,” “What,” “Where,” “When” questions or the “Yes/No” questions. As well, existing approaches look for named entities and other helpful facts within the query, for example, “Where in Africa did the scientists sent by Napoleon find Rosetta Stone?” This question contains different named entities (Napoleon, Africa, Rosetta Stone) and facts (in Africa, scientists). Interesting case studies and observations of the relation between aggregation search and QA can be found more detailed in Moriceau and Tannier [2010].

Nugget retrieval is not the major target of QA, that is, experimenting the different query matching algorithms is not the primary goal. Given a set of documents that match the query or queries derived from the initial question, QA approaches go within the documents to find and assemble answers. This task is not easy. Often, wrong passages are selected and the result is useless.

From the perspective of result aggregation, there are multiple ways to put results together and the result is often dependent on the query. As providing one unique answer may be risky, QA systems often return a list of candidate answers. Answers are often juxtaposed with supporting texts extracted from the documents. A different approach is presented in Wu and Fuller [1997]. Instead of returning a list of documents to answer questions and instead of very focused answers, we propose different intermediary approaches that can assist the answering process. This corresponds to five different result aggregations: (i) an abstract and a document, (ii) sentence fragments with keywords highlighted and a document, (iii) an abstract with one document and other




Fig. 3. Simple schemas for federated search (left), metasearch (center), and cross-vertical aggregated search (right).

related documents, (iv) sentence fragments and multiple documents, and (v) a set of paragraphs.

4.3. Federated Search

Most of the work in IR with multisources is classified as *federated search* [Avrahami et al. 2006; Aditya and Jaya 2008; Callan 2000] also known as *distributed information retrieval*. In federated search, instead of having one central collection indexed by one search engine, there are many distributed text collections each indexed by a search engine [Arguello et al. 2012].³ We may have as many query match algorithms as the number of subcollections (see the left side of Figure 3). At query time, the federated search system has to select the sources that are likely to be useful for the query. To do so, local representations of significantly reduced size of each collection are used. The obtained results from different sources are then assembled with each other. Typically, the final answer is a ranked list.

The FedLemur project [Avrahami et al. 2006] is an illustrate example of federal search engine that is built on top of statistical data issued from 100 US federal agencies. Instead of building a centralized collection, which can quickly get outdated, authors show that it is better to build local search engines in many distributed nodes (one per agency).

Federated search approaches with multiple distributed collections have not been shown to produce very significant improvement with respect to approaches with one unique index. Nevertheless, they are commonly used when the sources are too heterogeneous or when some collection is hidden behind some API. In all approaches with multiple sources, we can easily identify the query dispatching step, nugget retrieval step, and result aggregation step.

Metasearch can be seen as a form of federated search in the context of Web search [Lalmas 2011]. Initial metasearch engines used to wrap different Web search engines with the goal of improving precision and recall of Web search [Selberg and Etzioni 1995]. This was reasonable at the time because the indexes of the existing web search engines covered small fractions of the Web. Though, the chances to have other relevant results in another source were high.

The sources in metasearch are typically black-box search engines that receive queries and return a ranked list of results. The returned results from different sources are then combined into one interface [Selberg and Etzioni 1995; Manoj and Jacob 2008; Srinivas et al. 2011]. Typically results are sorted by source, but they can also be ranked with

³Federated search should not be confused with data fusion or data integration. Data fusion aims at finding optimal combinations of query matching algorithms running on the same index/collection, whereas in data integration, the problem is to combine data from different sources and provide the user a unified view of these data (i.e., a global schema) [Lenzerini 2002].

each other in one list. In contrast with most federated search work, sources can target the same task and collection (see Figure 3, center).

4.4. Mashups

Mashups are an interesting case study from the perspective of aggregated search. A mashup is a Web application that integrates data as well as services (databases, search engines, Web services, etc.) into a single tool. Thus, it allows combining applications and data from different sources. The aim is to provide a new service to the user. For instance, Yahoo! Pipes⁴ is a well-known mashup tool.

Manually coding mashups, we can build documents on demand, where we predefine where to put content coming from each component of a mashup. Although they are an interesting case study, mashups are mostly used for expert tasks rather than for search [Greenspan et al. 2009; Ranganathan et al. 2009].

4.5. Cross-Vertical Aggregated Search

Cross-vertical Aggregated Search (cvAS) [Arguello et al. 2009; Lalmas 2011; Kopliku 2011] is the task of searching and assembling information from vertical search and Web search (see Figure 3, right). This is usually done in a Web search context. A vertical is a specialized subcollection [Diaz 2009a; Zhou et al. 2012], such as images, video, and news. *cvAS* is met in literature as an instance of both federated search [Aditya and Jaya 2008] and metasearch [Gulli and Signorini 2005]. In *cvAS*, however, verticals are a core concept [Arguello et al. 2012]. Since the definition of aggregated search, a lot of recent work [Arguello et al. 2009; Lalmas 2011; Sushmita et al. 2009; Kopliku et al. 2009, 2011c] is classified within this later direction. We will concentrate on cross-vertical aggregated search in Section 5.

4.6. Relational Aggregated Search

We define *Relational Aggregated Search (RAS)* as a broad class of approaches that involve retrieving and aggregating information nuggets using their relations. For example, the query “all about iPhone 4S” would clearly benefit from this approach. We may assemble related attributes (e.g., weight, price), related products, related images, and so on. So defined, RAS can be complementary to *cvAS*. If *cvAS* provides diversity within search results, RAS can find the sense of aggregation and provide more related content within search results.

Let relational information be whatever triple that puts into relation two information nuggets. A relation can be deduced as simply as by performing a vertical search (e.g., image search, video search). However, a huge amount of relations exist in the Web, and it can be extracted (with, for instance, Information Extraction techniques). As well, huge amounts of relational information already exist within knowledge graphs or databases. Being able to extract and use these relations in the context of IR is the goal of RAS.

So defined, RAS has to rely on advances of many well-known fields including information extraction, entity-oriented search, object-level search, semantic search, IR from databases, and so on. The definition of RAS and its situation in research work might seem vague, but we will get back to RAS in a dedicated Section 6.

4.7. Domain-Specific Applications

Instances of aggregated search can be found in domain-specific applications. These approaches can sometimes be too specific, but they remain interesting to study because some of the work can be generalized to larger use.

⁴<http://pipes.yahoo.com/pipes/>, Retrieved September 2013.

Kaptein and Marx [2010] investigate focused retrieval and result aggregation on political data. Their documents are long (50–80 pages), containing transcripts of the meetings of the Dutch Parliament. Instead of returning entire documents, authors choose speeches as best entry points into documents. For a given query, they provide a summary of the results as a graph with three axes: year, political party, and number of search results. Search results can be browsed through three facets: person, political party, and year. Each speech is summarized in both terms of structure and content. The latter corresponds to word clouds and interruption graphs. We may observe at least four different aggregation formats within this work: interruption graph (structure summary), content summarization, facets, and results graph with three axes.

Another application is met in social science [Ou and Khoo 2008]. Here, authors provide search across academic papers. They extract and aggregate research concepts, their relations, research methods, and contextual information. The results can then be browsed by method, relation, or research concept. For each research concept, the user is provided a summary of contextual information.

Strotmann and Zhao [2008] also focus on academic papers. They introduce two graph-based structures to help browsing search results. The first is an aggregated graph on authors. The second is a graph on authors linked through co-citation analysis. Academic research is also considered in sites such as Microsoft Academic Research (<http://academic.research.microsoft.com/>) or Google Scholar citations (<http://scholar.google.com/citations>), which aggregate information about a given researcher: publications, conferences, related keywords, co-authors, global indicators such as number of publications and h-index. *News search aggregators* represent good examples of result aggregation that go beyond the ranked list visualization. News articles with similar topic and publication time are often clustered into news stories [Sahoo et al. 2006; Hennig and Wurst 2006]. This can also help users focus their search within a topic or time interval. In addition, Rohr and Tjondronegoro [2008] propose an interesting timeline visualization showing the evolution of the topic. As well, we can attach to the news stories related multimedia [Rohr and Tjondronegoro 2008]. This is the case for Google News (<http://news.google.com/>) and Ask News (<http://www.ask.com/news>).

Geographic information retrieval is another interesting case study. It corresponds to search on geographic entities mainly through free text queries [Vaid et al. 2005; Jones and Purves 2009]. Usually a list of candidate results is returned unless the query is specific enough to identify exactly one result. The list of results is shown in a map juxtaposed with the corresponding ranked list. This is the case for major search engines such as Google Maps (<http://maps.google.com/>). In addition geographic proximity can be used to group and organize search results [Boudighaghen et al. 2009; McCurley 2001]. If the query is good enough to identify one single geographic entity or the user clicks on one of the proposed results, a geographic search engine can support the user with more information about the entity. For instance, for a hotel, it is possible to provide a phone number, reviews, images, and so on. In fact, geographic entities can be associated with other types of content such as images [Naaman et al. 2006; Kennedy and Naaman 2008], related named entities [Vallet and Zaragoza 2008], news articles, and so on. This is useful to enrich and organize search results.

5. CROSS-VERTICAL AGGREGATED SEARCH

Cross-vertical aggregated search already has a consecrated place within aggregated search to the extent that it is often used as a synonym of aggregated search [Lalmas 2011], [Arguello et al. 2009]. The most well-known examples of cross-vertical aggregated search are met in the major Web search engines. In recent years, these search engines do not return uniform lists of Web pages. They also include results of different




Fig. 4. Google Universal search results on the query “Barcelona,” accessed on April 2012.

type such as images, news, videos, definitions, and so on (see Figure 4). A simple definition of cross-vertical aggregated search is met in Lalmas [2011]:

[Cross-vertical aggregated search] attempts to achieve diversity by presenting search results from different information sources, so-called verticals (image, video, blog, news, etc.), in addition to the standard Web results, on one result page.

Motivations for cross-vertical aggregated search are multiple. Here, we list some of the main highlighted advantages already identified in literature:

- First, Web data are highly heterogeneous, and keeping a fresh index of all real-time data is difficult [Arguello et al. 2012]. Specialized search engines that focus on particular types of media are thus required.
- Second, from an user point of view, it has been shown that **vertical search intent** is often present among Web search queries [Arguello et al. 2009; Liu et al. 2009; Kopliku et al. 2011c]. For example, queries such as “Avatar trailer” or “Eiffel Tower images” can be found in Web search logs, although these queries might have been issued for videos and images. Liu et al. [2009] analyzed 2,153 generic Web queries into verticals, using query logs. They found that 12.3% have an image search intent, 8.5% have a video search intent, and so on. Arguello et al. [2009] classified 25,195 unique queries, randomly sampled from search engine logs, into 18 verticals. Of the queries (mostly navigational), 26% were assigned no vertical, 44% of the queries were assigned one vertical, and the rest of the queries were assigned more than one vertical. The latter were mostly ambiguous. Cross-vertical aggregated search allows the user to search multiple specialized collections with a single-query access.
- At last, it has been shown that *cvAS* increases the **diversity** of relevant results [Sushmita et al. 2009; Kopliku et al. 2011c; Santos et al. 2011]. Results from different relevant sources have been shown to be often **complementary** with each other and they have been shown useful in the presence of **ambiguous queries** [Kopliku et al. 2011c].

Cross-vertical aggregated search can be seen as a “divide and conquer” approach, where customized search engines are used to answer queries on specific collections and where a complete answer is then built to meet the user need.

Cross-vertical aggregated search in the aggregated search framework. In cross-vertical aggregated search it is easy to identify the components of our general

framework for aggregated search. Query dispatching will correspond to the selection of sources. Each source will perform its nugget retrieval process and finally retrieved nuggets (e.g., Web pages, images, videos) will have to be assembled in one interface. Although the tasks are well distributed, the problem is far from being solved. It is not easy to decide which sources should be used and how the retrieved results should be assembled and presented. In particular, we can list some major issues which have the attention of current research:

- Source representation and selection.* Which source should be used? How should they be represented internally in terms of features?
- Result aggregation.* How should search results from different sources be assembled (e.g., ranking by item, ranking by block)?
- Result presentation.* Which are adequate interfaces for cross-vertical aggregated search?

We will describe each of the aforementioned issues in more detail.

5.1. Source Representation and Selection

Some aggregated search systems make use of all sources they can reach for all queries, while others are *source selective* (i.e., they make use only of the sources considered useful for the query). Most of major Web search engines are source selective. This is done to avoid a long answering time: Querying many sources and waiting for all results can be too slow for standard Web search. The goal of source selection is to predict the query intent combining evidence from the query, user profiles, query logs, and sources themselves. This problem is also known as the *vertical selection* problem [Arguello et al. 2011].

To enable efficient source selection, sources have some internal representation in the cvAS systems. This representation can be as simple as a textual description of the source, but in general it contains representative terms and features for the source extracted from sampling or other mining techniques. Some techniques of source selection can be met in federated search [Gravano et al. 1997; Callan 2000; Shokouhi et al. 2007], but we will focus here on the approaches met specifically for cvAS.

Source selection demands the collection of evidence that associates the potential information need behind a query to one source. This evidence is turned into features that can be used with some classification or ranking algorithm. Tables I and II represent an overview of the experimented features in literature. They can be split into preretrieval and postretrieval features, because there exist approaches that select sources after having some search results in return. Within the preretrieval features, we account for simple features such as the presence of terms with strong vertical intent evidence such as “image,” “video” (feature f1), the occurrence of named entities of a given type (location, organization, etc.) (feature f2), the length of the query (feature f3). Evidence is collected from query logs (features f4 and f5) in vertical search engines and click-through analysis in Web search (feature f6). As well, the query can be matched to feature/term sampled representations of the source collection (feature f7) or of given categories/subjects (feature f8). As well, it has been shown useful to apply a classifier that can tell if a query hides a navigational intent (feature f9).

Postretrieval features include different scores computed on search results for a given query. We find relevance scores as computed by each source (feature f10) as well as match scores computed on search results in a uniform fashion across all sources (feature f11). We find also the use of simpler features such as the number of search results (feature f12) and the freshness of search results (feature f13). Features related to the location and context of a user/task (features f14 and f15) can be used to rerank

Table I. Preretrieval Features

	Feature based on	Description
f1	Vertical intent terms [Arguello et al. 2011; Arguello et al. 2009; Ponnuswami et al. 2011]	Some terms indicate vertical intent such as image, video, and photo. This feature combines hard-coded and learned association rules for queries and sources.
f2	Named-entity type [Arguello et al. 2011]	These features indicate the presence of named entities of some type in the query.
f3	Query length [Ponnuswami et al. 2011]	This feature corresponds to the length of the query in terms.
f4	Query logs [Arguello et al. 2011; Arguello et al. 2009]	These features indicate if the query has been met in a source query log.
f5	Recent popularity of the query [Diaz 2009b]	This feature indicates how often the query has been met in a source query log recently.
f6	Clickthrough analysis [Arguello et al. 2011; Ponnuswami et al. 2011; Li et al. 2008]	These features are generated from the documents that have been clicked for the query. The click is considered implicit feedback.
f7	Source representation [Diaz and Arguello 2009]	This feature is computed based on explicit and implicit feedback on the query and its intent.
f8	Category representation [Arguello et al. 2009]	This feature is generated through classification of the query into predefined domains such as sport, arts, and technology.
f9	Evidence for navigational need [Ponnuswami et al. 2011]	This feature indicates the chances of the query to come from navigational needs.

Table II. Postretrieval Features

	Feature based on	Description
f10	Vertical relevance score [Diaz 2009b]	This feature corresponds to the relevance score of the vertical source itself for a result or a block of results.
f11	Uniform Match Score [Arguello et al. 2011; Ponnuswami et al. 2011; Li et al. 2008]	This feature corresponds to match scores on the results computed uniformly across multiple sources.
f12	Number of results [Arguello et al. 2011]	This feature corresponds to the results count of a source.
f13	Freshness of documents [Arguello et al. 2011; Diaz 2009b]	This feature indicates how fresh are search results for the query within a source.
f14	Contextual score of results [Kopliku et al. 2013]	These features indicate the relatedness of search results with respect to some context.
f15	Geographic-context score of results [Kopliku et al. 2013]	These features indicate the relatedness of search results with respect to some geographic context.

or filter search results, that is, the vertical search engine is personalized for a given contextual use and geographical location.

There are different interesting works with respect to source selection. Diaz [2009a] studies the integration of news search results within Web search. He estimates newsworthiness of a query to decide whether to introduce news results on top of the Web search results. They make use of clickthrough feedback to recover from system errors. Li et al. [2008] also rely on clickthrough data to extract implicit feedback for source selection. They represent queries and documents as a bipartite graph, and they propagate implicit feedback in this graph. Their approach is experimented for the integration of product search and job search results. Arguello et al. [2009] list various sources of evidence that can be used to tackle source selection such as query-log features, vertical intent terms, and corpus features. In later work, Diaz and Arguello [2009] show how they can integrate implicit and explicit feedback for vertical selection. From a set of 25,195 labeled queries, they propagate implicit feedback across about

10 million queries. Arguello et al. [2010] show how to adapt source selection methods to new unlabeled verticals.

5.2. Result Aggregation

There are different ways to assemble results in cross-vertical aggregated search. The most simple approach is to rely only on source selection [Li et al. 2008; Diaz 2009b; Arguello et al. 2009], that is, search results from one source are integrated only when the source is considered potentially useful. In this case, results are placed in predefined locations, usually on top of traditional Web results. We have mentioned related approaches in the previous section.

The task becomes more complex when we want to rank results of different sources with each other. Two main approaches exist, namely, block ranking and single-result ranking.

In block-ranking, results are ranked in blocks of results of the same source [Ponnuswami et al. 2011; Arguello et al. 2011] (e.g., three images versus three videos, whereas in single-result ranking, results may be interleaved (e.g., one image versus one Web page for instance) [Liu et al. 2009].

In Ponnuswami et al. [2011] and Arguello et al. [2011], search results are ranked in blocks of vertical content (e.g., three images versus three Web search results). In both cases, ranking functions are trained using pairwise preferences between blocks. Ponnuswami et al. [2011] propose a machine-learning framework that can assist the way search results are ranked in blocks by major search engines. Arguello et al. [2011] compare different algorithms for block-ranking, namely, classification, voting, and learning to rank techniques. The latter (learning to rank) outperforms other algorithms.

Ranking results of different sources one by one (not in blocks) has not been widely adapted and or studied. Although ranking results from very different systems is a difficult problem, interleave results from different vertical search engines may be of interest. Indeed, results in verticals are not equally relevant. A user issuing the query “French military intervention in Mali” may be interested in a result page showing first the last news about Mali intervention and then a Wikipedia article on the Northern Mali conflict, before seeing other less recent news on the subject. Images of the intervention are also more relevant than images of a speech of the French president. Liu et al. [2009] propose a probabilistic framework for this purpose that also lists different useful features for this task. Current evolution of major Web search engines is showing that we can go beyond block ranking. For instance, for some queries, the user is answered with focused information or Wikipedia extracts.

5.3. Result Presentation

Whereas result aggregation deals with the ranking of results, result presentation concerns the presentation interface. One way to aggregate vertical searches is to place content of the same type into predefined layout panels: This type of approach is called *unblended*. This is the approach chosen by Yahoo! Alpha⁵ some years ago and by Kosmix.⁶ Whenever a minimal amount of vertical content is available for a query, the various predefined slots are filled and displayed around the natural search listings. The advantage of such an approach is that the user knows where the vertical content is shown. On the other hand, vertical content is almost always shown even if there is nothing relevant for such a type of content. Other approaches, called *blended*, merge search results from different sources with each other in the same panel (see Figure 4). This was first introduced by major search engines who started integrating vertical

⁵<http://au.alpha.yahoo.com>. See [Lalmas 2011] for a screen shot.

⁶<http://www.kosmix.com>.

content only on the top or bottom of their search results list (Web search results are ranked as usual). When it is likely that the user is looking for other types of content such as images, news, or video, this content is placed in the bottom or the top of the Web results list. One of the advantages of such presentation is that vertical search results are added only when needed and visualization remains simple.

Nowadays, major search engines still use a blended approach, but results are ranked by blocks, that is, groups (blocks) of results from each source are ranked with each other. Within the block, results are usually ranked, that is, they are shown in the order they are returned from the origin source (Figure 4). For instance, we may have a block of three Web search results followed by a block of images, followed by a block of news search results, and then Web search results again. The advantage of this approach is that it is simple and at the same time flexible. We can add as many vertical searches as needed and only show those that are more relevant.

In the Web, we can also find other picturesque applications such as Spezify,⁷ which proposes an original way to exploit visualization space. Spezify aggregates results from different vertical sources with a slight preference for content with immediate visual impact, such as images and videos. Each result fills a rectangle and rectangles are placed side to side to fill the visualization space. The results expand in all directions (up, down, left, right), and it looks as a mosaic filled with content. Google has also launched a new application called “What do you love?”⁸ This application aggregates results from several verticals as well as from other Google applications such as Google Translator, search term popularity measures, and so on.

6. RELATIONAL AGGREGATED SEARCH

Cross-vertical aggregated search does not perform any explicit assembly of search results. Results remain a group of unrelated information nuggets coming from different sources. We define Relational Aggregated Search (RAS) as a complementary set of approaches where relations between information nuggets are taken into account. Relations can be very useful to assemble the search results and can be as simple as (“Paris,” *capital of*, “France”) or (“x.jpg,” *image of*, “Eiffel Tower”). RAS can provide more focus within search results. For example, the query “capital of France” can be answered right away with “Paris.” It can also provide more structure and organization within search results. For example, in the answer to query “Oscar Wilde” we can include a list of attributes (date of birth, date of death, nationality), a list of related images, a list of related books, similar authors, and so on. The user can navigate across the books or the similar authors.

A good example of relational aggregated search is met in Google Squared.⁹ Figure 5 shows its results for the query “arctic explorers.” Each line corresponds to an arctic explorer. For each of the explorers, there are *attributes* such as date of birth, date of death, as well as an image and a description. The user can specify his own attributes to search for, and he can search for a new arctic explorer that is not in the list. To answer this query, it is necessary to rely on relations such as (“Roald Amundsen,” *is a*, “arctic explorer”) or (“Roald Amundsen,” *date of birth*, “1872”). If Google Squared retrieves information from the Web, Wolfram Alpha¹⁰ represents a different approach where relational facts are extracted from an internal prebuilt collection of facts.

⁷<http://www.spezify.com>.

⁸<http://www.wdyl.com/>.

⁹<http://www.google.com/squared/> no longer available online within Google labs. See http://en.wikipedia.org/wiki/Google_Squared.

¹⁰<http://www.wolframalpha.com/>.

The screenshot shows a Google Squared search interface. At the top, there is a search bar with the text 'artic explorers' and buttons for 'Squard', 'Add to His Square', and 'Sign in to save your Squares'. Below the search bar is a table with the following columns: Item Name, Image, Description, Date Of Birth, Date Of Death, and Add to items. The table contains four rows of data:

Item Name	Image	Description	Date Of Birth	Date Of Death	Add to items
Roald Amundsen		Roald Engelbregt Gravning Amundsen (born Roald Engelbregt Gravning Amundsen ; 16 July 1872 – 18 June 1928) was a Norwegian explorer of polar regions. ...	1872	June 1928	<input type="checkbox"/>
Douglas Mawson		'Douglas Mawson' Australian Dictionary of Biography. http://adb.anu.edu.au/biogs/bio/biography/103444/bi.htm . Retrieved on 2007-10-01. ...	5 May 1862	14 October 1956	<input type="checkbox"/>
James Clark Ross		James Clark Ross , born in 1800, entered the Navy at 11 years of age. During his first years of service he was stationed and searched over by his ship. ...	April 15, 1800	April 5, 1862	<input type="checkbox"/>
Henry Hudson		Henry Hudson (d. 1611) was an English sea explorer and navigator in the early 17th century. After several voyages on behalf of ...	1570	1611	<input type="checkbox"/>

Fig. 5. Google Squared result for “artic explorers,” accessed on May 2010.

The Google Knowledge Graph introduced in 2012 by Google [Singhal 2012] is another example of the use of relational facts. The Graph is constructed using public sources such as Freebase or Wikipedia, as well as some real-time sources from the Web. It currently contains 500 million objects and 3.5 billion facts, as well as relationships between them. The Google Knowledge Graph is used by Google to disambiguate results, to present a summary of an entity (the right panel in Figure 1) to help the user to go deeper in his/her query by showing him/her new facts or new connections related to the result (explorative search suggestion).

All these applications show, however, that we can foresee another way of assembling search results based on fine-granularity information and relations.

Relational aggregated search as defined here intersect with many research directions including entity-oriented search, semantic search, object-level search, database information retrieval, ... In the next sections, we will provide main definitions and then we will analyze the related work to show that there is much to rely on in this promising research direction.

6.1. Definitions and Framework

A relation can be seen as a semantic triple of the form (X, R, Y). Such a relation can be well defined and explicit, such as (“Paris,” *capital of*, “France”), or less explicit, such as (“Paris,” *relates to*, “France”). In relational databases, we do also find n-ary relations, but we can observe that most of the useful relations can be mapped into ternary relations which will be the case for this paper. Related items are usually instances¹¹ (also called named entities, class instances). By definition, an instance is an object/concept belonging to some class (e.g., countries, biologists, cities, movies). However, relations can apply to all types of nuggets including longer text, for example, (“The art or science of combining vocal or instrumental sounds (or both) to produce beauty of form, harmony, and expression of emotion,” *definition of*, “music”) and multimedia content (http://en.wikipedia.org/wiki/File:Flag_of_France.svg, *flag of*, “France”).

The RAS paradigm fits in our general aggregated search framework. Its main components are *query dispatching*, *relation retrieval*, and *result aggregation*:

- Query dispatching:** Inspired by the work in Cafarella et al. [2006] and Kopliku et al. [2011b], we distinguish three types of information needs where benefits from relational aggregated search are obvious. We call them relational queries:
 - attribute query** (“GDP of UK,” “address of Hotel Bellagio”)
 - instance query** (“Samsung Galaxy S,” “Scotland,” “Oscar Wilde”)
 - class query** (“Toshiba notebooks,” “British writers”)

¹¹The term instance will be preferentially used instead of named entities.

Each of these queries may demand a different treatment. From this perspective, query dispatching becomes important to detect the query type and to trigger the appropriate solution.

- Relation retrieval:** In relational aggregated search, we are interested in all nugget-nugget relations that can be useful for information retrieval. However, we highlight three types of relations that are particularly interesting and useful, namely:
 - instance-class relation:** for example, (“France,” *is a*, “country”); (“Bill Clinton,” *is a*, “US president”)
 - instance-instance relation:** for example, (“Paris,” *relates to*, “France”); (“Roger Federer,” *plays against*, “Rafael Nadal”)
 - instance-attribute relation:** for example, (“France,” *has attribute*, “(motto:Liberty, Equality, Fraternity)”); (“France,” *has attribute*, “population density: 116/km²”)
- Result aggregation:** Relations can enable new ways to assemble search results. When the query is specifically asking for an attribute, the best choice can be returning its value right away. When the query is an instance, we may show not only a summary of salient attributes but also images, related instances, and so on. When the query is a class of instances, the result can be a comparative table of the class instances with their attributes (as it is done for instance in Google Squared).

We will now detail the novel issues of relational aggregated search focusing on relational queries, relation retrieval, and result aggregation.

6.2. Relational Queries

To build relational aggregated search, we need to identify the queries that benefit the most from this paradigm. Cafarella et al. [2006] present one of the first query taxonomies for relational search. For these queries, named entities and their relations find a crucial role.

- Qualified-list queries* retrieve a list of instances that share multiple properties (e.g., west coast liberal arts college).
- Unnamed-item queries* retrieve single object whose name the user does not know or cannot recall (e.g., the tallest inactive volcano in Africa).
- Relationship queries* retrieve the relation(s) between two objects (e.g., Microsoft ? Bill Gates).
- Tabular queries* retrieve a set of objects annotated by their salient properties (e.g., inventions annotated by their inventor and year of announcement).

In the previous list, the query type and the type of result are bound (i.e., the query definition includes the type of expected result).

In Kopliku et al. [2011b], we proposed a simpler taxonomy of queries inspired by Cafarella et al. [2006]. The type of query is binded to the semantic notion of instance, class and attribute, while it is not binded with the type of expected result:

- attribute query** (“GDP of UK,” “address of Hotel Bellagio”)
- instance query** (“Samsung Galaxy S,” “Scotland,” “Oscar Wilde”)
- class query** (“Toshiba notebooks,” “British writers”)

Most of the aforementioned queries benefit from relations. We will use the last taxonomy because it is simpler and it does not bind with the type of answer.

6.3. Relation Retrieval: How to Acquire Relations?

The major sources of relational content are (i) the Web, (ii) knowledge graphs/ontologies, and (iii) relational databases. We will focus less on relational databases because they

are less frequently used in broad IR applications. The Deep Web is, however, a good provider of databases, and challenges such as data integration are a very important research area. In the following, we detail *Information Extraction (IE)* techniques and *knowledge bases (ontologies, linked data)* as the major sources for relations.

IE techniques correspond to rules that can be applied to sites, documents, or parts of documents to automatically extract classes, instances, attributes, and their relations [Suchanek et al. 2009]. Most of the extraction rules have the form of $LxMyR$, where x and y are meant to be two information extracts and L, M, and R are meant to be patterns that are found, respectively, before, in between, and after the two extracts. For instance, the rule “the x of y is” can be used to identify attributes of instances, for example, “the capital of France is Paris.” Rules that rely only on lexicon (words, phrases) and part-of-speech tags are also referred as *lexico-syntactic rules*, while rules that are based on tag sequences are usually known as *wrappers*. Extraction rules can be hard-coded or learned [Chang and Lin 2001].

Information extraction techniques are quite heterogeneous, and they make use of various evidence such as statistics on terms, tags, decoration markup, part-of-speech tags, and so on. This evidence is then combined to define rules that match classes, instances, attributes and their relations. In the following, we provide just a short taxonomy of features (evidence) that are commonly used for this purpose:

- Word statistics** [Etzioni et al. 2005; Agichtein and Gravano 2000; Crescenzi et al. 2001]: Some terms are more frequent within or around information extracts. Statistics on words (often in conjunction with part-of-speech tags) are helpful to learn common lexico-syntactic patterns for information extraction.
- Part-of-speech tags** [Bellare et al. 2007]: Information extracts are usually nouns or compound noun phrases surrounded by verbs, adjectives, prepositions, and so on. Part-of-speech tags are helpful to learn possible patterns.
- Tags (HTML, XML)** [Crescenzi et al. 2001; Cafarella et al. 2009; Kopliku et al. 2011a, 2011b]: Most of the documents have some structure denoted through tags. The structure of the document is often useful to determine relations. In particular, HTML tables and HTML lists are known to contain relational data.
- Decoration, visual appearance** [Aumann et al. 2006; Meng et al. 2003; Yoshinaga and Torisawa 2007]: Sometimes the structure of a document is easier to learn through its visual aspects, especially when a pattern in terms of tags is difficult to define or learn.
- PMI and search hits** [Church and Hanks 1989; Turney 2001; Popescu and Etzioni 2005]: Pointwise Mutual Information (PMI) is a statistical measure that indicates possible correlation between two expressions. An easy way to estimate PMI is through search hits that indicate the number of search results a search engine has to return on a given query. PMI and similar statistical measures can play a crucial role to determine relations.

Most of the techniques in IE are domain specific, that is, they are designed to work well for some classes or instances or attributes. A noticeable exception concerns Open Information Extraction (OIE) approaches [Banko et al. 2007; Zhu et al. 2009; Etzioni et al. 2011]. Initially introduced by Banko et al. [2007] with their TextRunner system, the OIE paradigm aims at extracting large sets of domain-independent relational tuples (mainly on Web corpora) without requiring any human input. Open IE systems are now a hot subject of interest because they can allow to automatically construct knowledge bases or ontologies [Lin and Etzioni 2010].

Most of the IE techniques are also oriented toward precision. To enable relational aggregated search, we need high recall and reasonable precision. From this perspective, domain-independent methods and high-recall methods become crucial.

A lot of semantic content and their relations are already available in existing *knowledge bases* (e.g., DBPedia, Freebase, YAGO) [Suchanek et al. 2007; Wu et al. 2008; Limaye et al. 2010]. These sources contain semantic (relational) information manually input or automatically extracted from the Web or encyclopaedia-like collections (e.g., Wikipedia). These sources have also been used to learn IE extraction patterns or to reinforce confidence on extraction [Kopliku 2011]. They have attracted the attention of many recent research, and they are already enough to illustrate the potential of RAS.

In the next paragraphs, we list relation-retrieval approaches with respect to the type of relation they target.

6.3.1. Instance-Class Relation. The class instance is often referred in literature as named entity. Initially, there were seven named entity categories defined in Message Understanding Conference [Grishman and Sundheim 1996], while today standardized taxonomies with more than 150 classes can be found such as the extended named entities hierarchy [Sekine et al. 2002]. In reality, we cannot enumerate all possible named entity classes. A class can be as simple as “countries,” but it can also be “members of the UN Security Council” or “my friends.” Sometimes, we might not even be able to name the class in a reasonable way.

The definition of named entities as instances of some class makes the class-instance relation intrinsic for extraction techniques. These techniques are also known as Named Entity Recognition (NER). Hearst [1992, 1998] proposes one of the pioneer domain-independent approaches to extract named entities and their classes. The author identifies six lexico-syntactic patterns which detect this relation. Other approaches can be found in Guo et al. [2009] or Etzioni et al. [2005]. The class-instance relation is also targeted in TREC Entity Tracks [Balog et al. 2009b]. One of the proposed tasks involves returning a list of named entities of a given class.

Existing techniques for this relation are promising, although they are mostly precision oriented.

6.3.2. Instance-Instance Relation. Inspired from Suchanek et al. [2008], we distinguish four main relations that can relate an instance to another, namely, synonymy, sibling relation, meronymy, and nontaxonomic relations. To illustrate, “Big Apple” is a synonym for “New York City.” France and Italy are siblings in that they are instances of the same class “countries.” Meronymy involves part-of-a-whole relations such as (“Italy,” *is a member of*, “NATO”). The nontaxonomic relations are relations between two instances given by textual description such as (“John Travolta,” *plays in*, “Pulp Fiction”) and (“Windows,” *is a product of*, “Microsoft”).

Suchanek et al. [2007] presented YAGO, a large extensible ontology built through careful combination of heuristics on Wordnet and Wikipedia. They extract synonyms, instance-class relations (hyponymy), and nontaxonomic relations. The result is a set of about 1 million instances and about 5 million relations. Instance-instance relations can be extracted in a supervised [Agichtein and Gravano 2000; Etzioni et al. 2008] or unsupervised way [Etzioni et al. 2005]. Instance-instance relations are mostly extracted through lexico-syntactic rules based on term statistics and part-of-speech tags. Some of these techniques are particularly interesting in that they are applicable at a large scale and are domain independent. Nevertheless, some relations are difficult to capture, as their extraction depends on the training data or the technique being used.

6.3.3. Instance-Attribute Relation. Attribute acquisition can be domain independent or domain dependent. From domain-dependent approaches, we may mention approaches that focus on products. In this domain, attributes have been used not only to improve product search and recommendation [Nie et al. 2007a] but also to enable data mining [Wong and Lam 2009]. To acquire attributes from the Web, it is common to use

decoration markup [Yoshinaga and Torisawa 2007; Crescenzi et al. 2001] and text [Ben-Yitzhak et al. 2008; Tokunaga and Torisawa 2005]. A common technique to acquire attributes is through the use of lexico-syntactic rules [Pasca and Durme 2008; Alfonseca et al. 2010; Almuhareb and Poesio 2004; Popescu and Etzioni 2005].

HTML tags (for tables, lists, and emphasis) have also been shown to help for attribute acquisition [Yoshinaga and Torisawa 2007; Wong and Lam 2009]. In particular, HTML tables are known to be a mine for relational data and attributes. An approach to extract attributes from tables using column (or row) similarity is proposed in Chen et al. [2000]. Another common technique is through wrapper induction [Crescenzi et al. 2001; Wong and Lam 2004]. Cafarella et al. [2008, 2009] show that we can identify billions of relational tables in the Web. Kopliku et al. [2011d, 2011b] show how to extract and rank attributes from Web tables for whatever instance query combining table classification with relevance ranking.

At last, knowledge bases (such as DBPedia or Freebase) [Suchanek et al. 2007; Wu et al. 2008; Kopliku 2011] provide millions of relations in the form of linked data, which make it easy to retrieve attributes.

Attribute retrieval from linked data will probably attract more research attention given the fast growth of user generated or automatically extracted linked data [Bizer et al. 2008, 2009].

6.3.4. Other Nugget-Nugget Relations. The nugget-nugget relation is the broader class of relations that can involve whatever type of information nugget (i.e., it includes instance-instance relations, class-instance relations, and attribute-instance relations). We will list here just some other broad examples without aiming to be exhaustive:

- Similarity** (“similar to”): This is a common relation that has already been targeted in clustering, classification, news aggregators, and so on. For instance, news aggregators group similar news articles into stories.
- Diversity** (“different to”): The inverse of similarity represents another useful relation. Studies on novelty and diversity claim that it is sometimes better to promote some diversity among the retrieved results.
- Space-time** (“happens in time/space”): Content can also relate with respect to some features such as time and location. For instance, sometimes it is better to order information chronologically to favor freshness of information [Dong et al. 2010]. In Geographic Information Retrieval, the location feature is fundamental to organize and visualize search results [Jones and Purves 2009].

6.4. Result Aggregation

This section is about ways to aggregate search results in the context of relational aggregated search. We will consider one query type at a time.

Attribute queries (e.g., capital of France) should ideally be answered with the correct attribute value(s). However, for many queries, we might find many candidate values without being certain on their relevance and correctness. With respect to this issue, there is a lot of work in question answering (what is the capital of Congo?). Here, the result is one or more candidate answers, typically associated with supporting text.

Instance queries provide more space for result aggregation. In object-level search, these queries are answered through records and attributes extracted from one or more pages [Nie et al. 2007b]. The answer can also be a set of related instances [Basu Roy et al. 2010; Kato et al. 2009], multimedia content [Taneva et al. 2010], passages [Sauper and Barzilay 2009], or a summary of attributes [Kopliku et al. 2011b, 2011e]. Specific approaches have been adapted for people search [Macdonald 2009], product search [Nie et al. 2007b; Popescu and Etzioni 2005], bibliographic search [Ji et al. 2009], and so forth.

Class queries can be answered with a list of instances. This is the case for most approaches that can be found in Question Answering [Kelly and Lin 2006] and TREC entity tracks [Balog et al. 2009b]. Google Maps¹² answers location-related class queries (e.g., hotels in Las Vegas) with a list of instances in a map sometimes associated with ratings, reviews and images. In Koplaku et al. [2011b] and Krichen et al. [2011], results of class queries are structured in tables with instances in the rows and attributes in the columns.

We can conclude that there are several ways to answer queries in the relational framework. The quality of the result aggregation depends on the quality of the relations and the relevance of the result components.

7. EVALUATION OF AGGREGATED SEARCH

Aggregated search in its broadest definition is difficult to evaluate, that is, an aggregate answer can be organized in many ways and contain whatever information nuggets (images, sentences, Web search results, information extracts, etc.). Aggregated search can provide immediate and organized access to information. The quantity of relevant information, its diversity, coherence, and organization can affect directly user satisfaction.

Although the evaluation of aggregated search is not a solved question, some preliminary works or evaluation frameworks can however be mentioned in the context of relational aggregated search and cross-vertical aggregated search—they are discussed in the following two sections.

7.1. Cross-Vertical Aggregated Search

Although the process of cross-vertical aggregated search is clearly composed of vertical selection, item selection, and results presentation, evaluation does not concern so clearly these different parts. Vertical selection and result presentation are in fact strongly linked [Zhou et al. 2012].

Among approaches aiming at mainly evaluating vertical selection, one can cite Arguello et al. [2009], Li et al. [2008], Diaz [2009a] or Koplaku et al. [2011c]. Evaluation protocols are based either on manual assessments or on user interaction data (clickthrough).

Aggregated search interfaces (blended versus unblended) are evaluated in Sushmita et al. [2010] and Arguello et al. [2012] using user studies. A methodology for evaluating results presentation is also defined in Arguello et al. [2011]; the authors propose a metric-based evaluation of any possible presentation of the query. The idea is to use preference judgments on pairs of blocks to construct a reference presentation for the query. Approaches can then be evaluated by calculating their distance to the reference presentation. Although the proposed approach only considered ranking by block, the methodology can be applied to every approach.

At last, Zhou et al. [2012] propose to evaluate the aggregated search process as a whole, by proposing an evaluation framework using a test collection [Zhou et al. 2011] as well as evaluation metrics that are showed to be more suitable for evaluated cross-vertical aggregated search than traditional IR metrics.

To summarize, today, cross-vertical aggregated search is mainly evaluated by studying the aggregated search behavior, either using laboratory studies [Arguello et al. 2012; Arguello and Capra 2012; Sushmita et al. 2010] or search logs-data [Diaz 2009a]. The community is, however, working on the crucial point of reusable test collection, as shown by the first efforts published in Arguello et al. [2011] and Zhou et al. [2012].

¹²<http://maps.google.com>.

7.2. Relational Aggregated Search

In the context of relational aggregated search, existing frameworks for evaluation mainly concern the evaluation of entity search. Among them, one can cite:

- The INEX Entity Ranking Track (2007–2009)*: The track proposed a framework to evaluate engines that return lists of entities. Two main tasks were considered on structured documents collections: Entity Ranking (ER) and entity List Completion (LC) [Demartini et al. 2010]
- The TREC Entity Track (2009–2011)*: The aim of the track was similar to the INEX track (i.e., to perform entity-oriented search tasks) but on Web data (Web pages or Linked Open Data) [Balog et al. 2009a; Balog et al. 2010, 2011]. In 2011, for example, the ClueWeb 2009 Web corpus and the Sindice-2011 dataset [Campinas et al. 2011] were used as corpora. Two main tasks were evaluated (Related Entity Finding task and Entity List Completion task), whose aim was to find related entities of a given entity regarding a given relation.
- The SemSearch Challenge (2010–2011)*: This challenge aimed at evaluating semantic search systems for entity search on a collection which is a sample of Linked Data (RDF triples) crawled from publicly available sources [Semantic 2010, 2011; Halpin et al. 2010].
- The TAC KBP (Knowledge Base Population) Track*: This track, which first ran in 2009, evaluates the discovering of information about named entities in large corpus and its integration into a knowledge base [TAC 2011]. The number of tasks in the track increase with time. In 2013, for example, seven tasks are proposed, among them one can cite the Entity Linking track, the Slot Filling task, or the Cold Start KBP.
- The TREC KBA Track*: As the TAC KBP track, TREC KBA, which first ran in 2012, aims at helping the construction and maintenance of large Knowledge Bases [Frank et al. 2012]. In 2013, two tasks are proposed: (i) the Cumulative Citation Recommendation (CCR) task, in which systems should filter in a time-ordered corpus documents that are centrally relevant to a given entity; and (ii) the Streaming Slot Filling (SSF) task in which given a slot for a target entity, systems should detect changes to the slot value.

The preceding evaluation campaigns allow us to evaluate approaches targeting the extraction of instance-class or instance-instance relation. The evaluation of attribute extraction do not fit in these campaigns (except maybe for the TAC KBP and TREC KBA SSF tasks, which are interested in attribute values), and existing approaches for attribute retrieval were mainly evaluated on their own test (see Koplaku et al. [2011b]). To conclude, the utility of RAS is not just in the relevance of information, it is also in the focus and organisation. Few studies address the utility of relational information as a whole or as a complement to traditional search results.

8. CONCLUSION AND FUTURE DIRECTIONS

8.1. Conclusion

This survey shows that there is a myriad of approaches that go beyond query matching with an additional effort on result aggregation. We have analyzed many of these approaches from multiple research domains in the light of aggregated search.

Some approaches have strong considerations on the data they process (the sources in the database community should own a schema; RDF triples are queried in semantic search, vertical searches act on very specific data), whereas others are more concerned with textual information (that can or not be structured): This is the case of federated

search, metasearch, information extraction, question answering, or natural language generation.

Three processes are shown to decompose well the broad aggregated search process, namely, *query dispatching*, *nugget retrieval*, and *result aggregation*. They encapsulate well the different existing query matching approaches, with an additional attention at query time and result aggregation time.

Among the listed approaches, there are many that rely on a multiple sources (search engines) architecture. Among them, *cross-vertical aggregated search* is probably the most successful; it has met commercial success in Web search and it is capturing an increasing interest in research. Here, the integration of diverse vertical search engines (e.g., image search, video search) has been shown beneficial, that is, it increases the amount of relevant results and it often provides complementary results.

In addition, we have also analyzed research on relational aggregated search, another promising research direction. The latter makes use of relations to assemble and retrieve search results. This enables retrieving at finer granularity and composing new aggregated objects of related content.

8.2. Future Directions

We can expect aggregated search to attract research direction for the next decades. In this section, we list some challenging (and promising) research directions.

Starting from the broad aggregated search framework, we believe the most challenging issues concern query dispatching and result aggregation. Indeed, nugget retrieval has been the research target for more than 50 years now. Identifying queries that benefit from aggregation as well as new ways to aggregate content remain novel and promising.

Another important issue is evaluation. Evaluation in the context of cvAS and RAS concerns mainly the effectiveness of each sub-task that composes the different approaches. In other words, relevance is mainly assessed on single information nuggets taken separately (even if a noticeable exception can be recently found in the context of cross-vertical aggregated search [Zhou et al. 2012]). The quality of the aggregated results should however be measured by considering the complementarity of nuggets, coherence, completeness, or coverage of the information need, and appropriate metrics are still needed. Moreover, finding a way to measure the quality of any possible aggregate for a given query is a problem that is far from being solved.

At last and as stated in Arguello et al. [2012], a promising perspective concerns the personalization of aggregated results, in order to make users be the “the core content”.

If we keep to specific research directions, we believe that cross-vertical aggregated search and relational aggregated search remain the most promising in the short term.

Despite the success of *cross-vertical aggregated search* in the context of Web search, we believe its potential is far from being completely exploited. There is evidence that we can combine sources such as Wikipedia, DBpedia, images, maps to form aggregate unexisting objects with quite some visualization flexibility. As well, we can target new applications other than Web search. We can foresee enriched map search results, enriches Wikipedia results and so on.

In the context of *relational aggregated search*, there is a lot of work to be done.

We need to improve relation retrieval techniques in both terms of precision and recall. Depending on the targeted application, a good precision/recall is fundamental. We believe that two research directions are more promising at this moment. First, there is an increasing effort towards open-domain and large-scale information extraction techniques. Second, we can observe a fast growth of quality linked data in the Web. We believe that the combination of automatic extraction and user-generated content will produce enough relational and semantic data for multiple commercial uses. Moreover,

a substantial effort is needed to define coherent ways to explore and aggregate results (tables, lists, graphs).

To conclude, we believe that this survey in conjunction with other ongoing research indicate that future IR can integrate more focus, structure, and semantics in search results.

REFERENCES

- Pal Aditya and Kawale Jaya. 2008. Leveraging query association in federated search. In *Proc. of SIGIR 2008 Workshop on Aggregated Search*.
- Eugene Agichtein and Luis Gravano. 2000. Snowball: Extracting relations from large plain-text collections. In *Proc. of the 5th ACM Conference on Digital Libraries*. 85–94. DOI:<http://dx.doi.org/10.1145/336597.336644>
- Enrique Alfonseca, Marius Pasca, and Enrique Robledo-Arnuncio. 2010. Acquisition of instance attributes via labeled and related instances. In *Proc. of SIGIR 2010*. 58–65. DOI:<http://dx.doi.org/10.1145/1835449.1835462>
- Abdulrahman Almuhaireb and Massimo Poesio. 2004. Attribute-based and value-based clustering: An evaluation. In *In EMNLP 2004, ACL*. 158–165.
- Jaime Arguello and Robert Capra. 2012. The effect of aggregated search coherence on search behavior. In *Proc. of CIKM 2012*. 1293–1302.
- Jaime Arguello, Fernando Diaz, and Jamie Callan. 2011. Learning to aggregate vertical results into Web search results. In *Proc. of CIKM 2011*. 201–210.
- Jaime Arguello, Fernando Diaz, Jamie Callan, and Ben Carterette. 2011. A methodology for evaluating aggregated search results. In *Proc. of ECIR 2011*. 141–152.
- Jaime Arguello, Fernando Diaz, Jamie Callan, and Jean-Francois Crespo. 2009. Sources of evidence for vertical selection. In *Proc. of the 32nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 315–322. DOI:<http://dx.doi.org/10.1145/1571941.1571997>
- Jaime Arguello, Fernando Diaz, and Jean-François Paiement. 2010. Vertical selection in the presence of unlabeled verticals. In *Proc. of SIGIR 2010*. 691–698.
- Jaime Arguello, Fernando Diaz, and Milad Shokouhi. 2012. Integrating and ranking aggregated content on the Web. In *Proc. WWW 2012*.
- Jaime Arguello, Wan-Ching Wu, Diane Kelly, and Ashlee Edwards. 2012. Task complexity, vertical display and user interaction in aggregated search. In *Proc. of SIGIR 2012*. 435–444.
- Yonatan Aumann, Ronen Feldman, Yair Liberzon, Benjamin Rosenfeld, and Jonathan Schler. 2006. Visual information extraction. *Knowl. Inf. Syst.* 10, 1 (July 2006), 1–15. DOI:<http://dx.doi.org/10.1007/s10115-006-0014-x>
- Thi Truong Avrahami, Lawrence Yau, Luo Si, and Jamie Callan. 2006. The FedLemur project: Federated search in the real world. *JASIST* 57, 3 (2006), 347–358. DOI:<http://dx.doi.org/10.1002/asi.v57:3>
- K. Balog, A. de Vries, P. Serdyukov, P. Thomas, and T. Westerveld. 2009a. Overview of the TREC 2009 entity track. In *Proc. of TREC 2009*.
- K. Balog, A. P. de Vries, P. Serdyukov, P. Thomas, and T. Westerveld. 2009b. Overview of the TREC 2009 entity track. In *TREC 2009 Working Notes*. NIST.
- K. Balog, P. Serdyukov, and A. de Vries. 2010. Overview of the trec 2010 entity track. In *Proc. of TREC 2010*.
- K. Balog, P. Serdyukov, and A. de Vries. 2011. Overview of the trec 2010 entity track. In *Proc. of TREC 2011*.
- Michele Banko, Michael J. Cafarella, Stephen Soderland, Matt Broadhead, and Oren Etzioni. 2007. Open information extraction from the web. In *Proc. of IJCAI 2007*. 2670–2676.
- Senjuti Basu Roy, Sihem Amer-Yahia, Ashish Chawla, Gautam Das, and Cong Yu. 2010. Constructing and exploring composite items. In *Proc. of SIGMOD 2010*. 843–854.
- Mustapha Baziz, Mohand Boughanem, Yannick Loiseau, and Henri Prade. 2007. Fuzzy logic and ontology-based information retrieval. In *Studies in Fuzziness and Soft Computing*. Vol. 215/2007. 193–218.
- Kedar Bellare, Partha Pratim Talukdar, Giridhar Kumaran, O. Pereira, Mark Liberman, Andrew Mccallum, and Mark Dredze. 2007. Lightly-supervised attribute extraction for Web search. In *Proc. of Machine Learning for Web Search Workshop, NIPS 2007*.
- Ori Ben-Yitzhak, Nadav Golbandi, Nadav Har’El, Ronny Lempel, Andreas Neumann, Shila Ofek-Koifman, Dafna Sheinwald, Eugene Shekita, Benjamin Sznajder, and Sivan Yogev. 2008. Beyond basic faceted search. In *Proc. of WSDM 2008*. 33–44.

- C. Bizer, T. Heath, and T. Berners-Lee. 2009. Linked data—the story so far. *International Journal Semantic Web and Information Systems* 5, 3 (2009), 1–22.
- Christian Bizer, Tom Heath, Kingsley Idehen, and Tim Berners-Lee. 2008. Linked data on the Web (LDOW2008). In *Proc. of WWW 2008*. 1265–1266.
- Ourdia Bouidghaghen, Lynda Tamine, and Mohand Boughanem. 2009. Dynamically Personalizing Search Results for Mobile Users. In *Proc. Flexible Query Answering (FQAS)*. 99–110.
- Bert R. Boyce. 1982. Beyond topicality: A two stage view of relevance and the retrieval process. *Inf. Process. Manage.* 18, 3 (1982), 105–109.
- Michael J. Cafarella, Michele Banko, and Oren Etzioni. 2006. *Relational Web Search*. Technical Report. University of Washington.
- Michael J. Cafarella, Alon Halevy, Daisy Zhe Wang, Eugene Wu, and Yang Zhang. 2008. WebTables: Exploring the power of tables on the Web. *Proc. VLDB Endow.* 1, 1 (2008), 538–549. DOI: <http://dx.doi.org/10.1145/1453856.1453916>
- Michael J. Cafarella, Alon Y. Halevy, and Nodira Khossainova. 2009. Data Integration for the Relational Web. *PVLDB* 2, 1 (2009), 1090–1101.
- Jamie Callan. 2000. Distributed information retrieval. In *Advances in Information Retrieval*, W. Bruce Croft (Ed.). Kluwer Academic Publishers, Dordrecht, 235–266.
- S. Campinas, D. Ceccarelli, T. E. Perry, R. Delbru, K. Balog, and G. Tummarello. 2011. The Sindice-2011 dataset for entity-oriented search in the web of data. In *1st International Workshop on Entity-Oriented Search (EOS)*. 26–32.
- Chih-Chung Chang and Chih-Jen Lin. 2001. *LIBSVM: A Library for Support Vector Machines*. Technical Report.
- Hsin-Hsi Chen, Shih-Chung Tsai, and Jin-He Tsai. 2000. Mining tables from large scale HTML texts. In *Proc. of COLING 2000*. 166–172.
- Kenneth Ward Church and Patrick Hanks. 1989. Word association norms, mutual information, and lexicography. In *Proc. of ACL 1998*. 76–83.
- Charles L. A. Clarke, Maheedhar Kolla, Gordon V. Cormack, Olga Vechtomova, Azin Ashkan, Stefan Büttcher, and Ian MacKinnon. 2008. Novelty and diversity in information retrieval evaluation. In *Proc. of SIGIR 2008*. 659–666.
- Valter Crescenzi, Giansalvatore Mecca, and Paolo Merialdo. 2001. RoadRunner: Towards Automatic Data Extraction from Large Web Sites. In *Proc. of VLDB 2001*. 109–118.
- H. T. Dang, D. Kelly, and J. Lin. 2007. Overview of the TREC 2007 Question Answering Track. In *Proc. TREC 2007*.
- Hoa Tang Dang. 2006. Overview of DUC 2006. In *Proc. of the 2006 Document Understanding Conference*.
- Gianluca Demartini, Tereza Iofciu, and Arjen P. Vries. 2010. Overview of the INEX 2009 Entity Ranking Track. In *Focused Retrieval and Evaluation*. 254–264.
- Fernando Diaz. 2009a. Integration of news content into web results. In *Proc. of WSDM 2009*. 182–191.
- Fernando Diaz. 2009b. Integration of news content into web results. In *Proc. WSDM*. 182–191.
- Fernando Diaz and Jaime Arguello. 2009. Adaptation of offline vertical selection predictions in the presence of user feedback. In *Proc. of SIGIR 2009*. 323–330.
- Anlei Dong, Yi Chang, Zhaohui Zheng, Gilad Mishne, Jing Bai, Ruiqiang Zhang, Karolina Buchner, Ciya Liao, and Fernando Diaz. 2010. Towards recency ranking in Web search. In *Proc. of WSDM 2010*. 11–20.
- Doug Downey, Oren Etzioni, and Stephen Soderland. 2005. A probabilistic model of redundancy in information extraction. In *Proceedings of IJCAI*. 1034–1041.
- Oren Etzioni, Michele Banko, Stephen Soderland, and Daniel S. Weld. 2008. Open information extraction from the web. *Commun. ACM* 51 (December 2008), 68–74. Issue 12. DOI: <http://dx.doi.org/10.1145/1409360.1409378>
- Oren Etzioni, Michael Cafarella, Doug Downey, Ana-Maria Popescu, Tal Shaked, Stephen Soderland, Daniel S. Weld, and Alexander Yates. 2005. Unsupervised named-entity extraction from the web: An experimental study. *Artif. Intell.* 165, 1 (2005), 91–134. DOI: <http://dx.doi.org/10.1016/j.artint.2005.03.001>
- O. Etzioni, A. Fader, J. Christensen, S. Soderland, and Mausam. 2011. Open information extraction: The second generation. In *Proc. of IJCAI 2011*, Barcelona, Spain. 3–10.
- John R. Frank, Max Kleiman-Weiner, Daniel A. Roberts, Feng Niu, Ce Zhang, Christopher Re, and Ian Soboroff. 2012. Building an entity-centric stream filtering test collection for TREC 2012. In *Proc. of TREC 2012*.
- Shlomo Geva, Jaap Kamps, and Andrew Trotman (Eds.). 2009. *Proc. of INEX 2008*.

- Jeremy Goecks. 2002. NuggetMine: Intelligent groupware for opportunistically sharing information nuggets. In *Proc. of IUI 2002*. 87–94.
- Luis Gravano, Chen-Chuan K. Chang, Héctor García-Molina, and Andreas Paepcke. 1997. STARTS: Stanford proposal for Internet meta-searching. *SIGMOD Rec.* 26 (June 1997), 207–218. Issue 2. DOI: <http://dx.doi.org/10.1145/253262.253299>
- Ohad Greenshpan, Tova Milo, and Neoklis Polyzotis. 2009. Autocompletion for mashups. *Proc. VLDB Endow.* 2, 1 (2009), 538–549.
- H. P. Grice. 1975. Logic and conversation. In *Syntax and Semantics: Vol. 3: Speech Acts*, P. Cole and J. L. Morgan (Eds.). Academic Press, San Diego, CA, 41–58.
- Ralph Grishman and Beth Sundheim. 1996. Message Understanding Conference-6: A brief history. In *Proc. of the 16th Conference on Computational Linguistics*. 466–471.
- A. Gulli and A. Signorini. 2005. Building an open source meta-search engine. In *Proc. of WWW '05: Special Interest Tracks and Posters*. 1004–1005.
- Jiafeng Guo, Gu Xu, Xueqi Cheng, and Hang Li. 2009. Named entity recognition in query. In *Proc. of SIGIR 2009*. 267–274.
- H. Halpin, D. M. Herzig, P. Mika, R. Blanco, J. Pound, H. S. Thompson, and D. T. Tran. 2010. Evaluating ad-hoc object retrieval. In *Proc. of the International Workshop on Evaluation of Semantic Technologies (IWEST 2010)*
- Marti A. Hearst. 1992. Automatic acquisition of hyponyms from large text corpora. In *Proc. of the 14th Conference on Computational Linguistics*. 539–545.
- Marti A. Hearst. 1998. Automated discovery of WordNet relations. In *WordNet: An Electronic Lexical Database*, C. Fellbaum (Ed.). MIT Press, 131–153. Retrieved from <http://www.sims.berkeley.edu/hearst/papers/wordnet98.pdf>.
- Marti A. Hearst and Jan O. Pedersen. 1996. Reexamining the cluster hypothesis: Scatter/gather on retrieval results. In *SIGIR*. 76–84.
- Sascha Hennig and Michael Wurst. 2006. Incremental clustering of newsgroup articles. In *IEA/AIE*. 332–341.
- L. Hirschman and R. Gaizauskas. 2001. Natural language question answering: The view from here. *Natural Language Engineering* 7, 4 (11 2001), 275–300. DOI: <http://dx.doi.org/10.1017/S1351324901002807>
- Lei Ji, Jun Yan, Ning Liu, Wen Zhang, Weiguo Fan, and Zheng Chen. 2009. ExSearch: A novel vertical search engine for online barter business. In *Proc. of CIKM 2009*. 1357–1366.
- Christopher B. Jones and Ross S. Purves. 2009. Geographical information retrieval. In *Encyclopedia of Database Systems*. 1227–1231.
- Jaap Kamps, Shlomo Geva, and Andrew Trotman. 2008. Report on the SIGIR 2008 workshop on focused retrieval. *SIGIR Forum* 42, 2 (2008), 59–65. DOI: <http://dx.doi.org/10.1145/1480506.1480517>
- Rianne Kaptein and Maarten Marx. 2010. Focused retrieval and result aggregation with political data. *Inf. Retr.* 13, 5 (October 2010), 412–433. DOI: <http://dx.doi.org/10.1007/s10791-010-9130-z>
- Makoto P. Kato, Hiroaki Ohshima, Satoshi Oyama, and Katsumi Tanaka. 2009. Query by analogical example: Relational search using Web search engine indices. In *Proc. of CIKM 2009*. 27–36.
- B. Katz, G. Borchardt, and S. Felshin. 2005. Syntactic and semantic decomposition strategies for question answering from multiple resources. In *Proc. of the AAAI 2005 Workshop on Inference for Textual Question Answering*. Pittsburgh, Pennsylvania, USA.
- Diane Kelly and Jimmy Lin. 2006. Overview of the TREC 2006 question answering task. In *Proc. of the Text REtrieval Conference 2006*.
- Diane Kelly and Jimmy Lin. 2007. Overview of the TREC 2007 question answering task. In *Proc. of the Text REtrieval Conference 2007*.
- Lyndon S. Kennedy and Mor Naaman. 2008. Generating diverse and representative image search results for landmarks. In *Proc. of WWW 2008*. 297–306.
- Arlind Kopliku. 2011. *Approaches to Implement and Evaluate Aggregated Search*. Thèse de doctorat. Université Paul Sabatier, Toulouse, France.
- Arlind Kopliku, Mohand Boughanem, and Karen Pinel-Sauvagnat. 2011a. Mining the Web for lists of named entities. In *Proc. of CORIA 2011*. 113–120.
- Arlind Kopliku, Mohand Boughanem, and Karen Pinel-Sauvagnat. 2011b. Towards a framework for attribute retrieval. In *Proc. of CIKM 2011*. 515–524.
- Arlind Kopliku, Firas Damak, Karen Pinel-Sauvagnat, and Mohand Boughanem. 2011c. Interest and evaluation of aggregated search. In *Proc. of IEEE/WIC/ACM International Conference on Web Intelligence, Lyon, France*. 154–161.

- Arlind Kopliku, Karen Pinel-Sauvagnat, and Mohand Boughanem. 2009. *Aggregated Search: Potential, Issues and Evaluation*. Technical Report. Institut de Recherche en Informatique de Toulouse, France.
- Arlind Kopliku, Karen Pinel-Sauvagnat, and Mohand Boughanem. 2011d. Attribute Retrieval from Relational Web tables. In *Proc. of SPIRE 2011*. 117–128.
- Arlind Kopliku, Karen Pinel-Sauvagnat, and Mohand Boughanem. 2011e. Retrieving attributes using Web tables. In *Proc. of JDCL 2011*. 397–398.
- Arlind Kopliku, Paul Thomas, Stephen Wan, and Cecile Paris. 2013. Filtering and ranking for social media monitoring. In *Proc. of CORIA 2013*.
- Ines Krichen, Arlind Kopliku, Karen Pinel-Sauvagnat, and Mohand Boughanem. 2011. Une approche de recherche d'attributs pertinents pour l'agrégation d'information. In *Proc. of INFORSID 2009*. 385–400.
- Mounia Lalmas. 2011. Advanced topics on information retrieval. Springer, Chapter Aggregated search.
- Maurizio Lenzerini. 2002. Data integration: A theoretical perspective. In *Proc. of PODS 2002*. 233–246.
- Xiao Li, Ye-Yi Wang, and Alex Acero. 2008. Learning query intent from regularized click graphs. In *Proc. of SIGIR 2008*. 339–346.
- Girija Limaye, Sunita Sarawagi, and Soumen Chakrabarti. 2010. Annotating and searching Web tables using entities, types and relationships. *Proc. VLDB Endow.* 3, 1–2 (September 2010), 1338–1347.
- C.-J. Lin and R.-R. Liu. 2008. An analysis of multi-focus questions. In *Proc. of SIGIR Workshop on Focused Retrieval*.
- Thomas Lin and Oren Etzioni. 2010. Identifying functional relations in web text. In *Proc. of EMNLP 2010*.
- Ning Liu, Jun Yan, and Zheng Chen. 2009. A probabilistic model based approach for blended search. In *Proc. of WWW 2009*. 1075–1076.
- Craig Macdonald. 2009. The voting model for people search. *SIGIR Forum* 43, 1 (June 2009), 73. DOI : <http://dx.doi.org/10.1145/1670598.1670616>
- Christopher Manning, Prabhakar Raghavan, and Heinrich Schütze. 2008. *Introduction to Information Retrieval*. Cambridge University Press.
- M. Manoj and Elisabeth Jacob. 2008. Information retrieval on Internet using meta-search engines: A review. *Journal of Scientific & Industrial Research* 67, 10 (2008), 739–746.
- Kevin S. McCurley. 2001. Geospatial mapping and navigation of the web. In *Proc. of the 10th International Conference on WWW*. ACM, New York, NY, 221–229. DOI : <http://dx.doi.org/10.1145/371920.372056>
- Xiaofeng Meng, Haiyan Wang, Dongdong Hu, and Chen Li. 2003. A supervised visual wrapper generator for Web-data extraction. In *Proc. of the 27th Annual International Conference on Computer Software and Applications (COMPSAC '03)*. IEEE Computer Society, Washington, DC, 657.
- Véronique Moriceau and Xavier Tannier. 2010. FIDJI: Using syntax for validating answers in multiple documents. *Inf. Retr.* 13, 5 (October 2010), 507–533. DOI : <http://dx.doi.org/10.1007/s10791-010-9131-y>
- Vanessa Murdock and Mounia Lalmas. 2008. Workshop on aggregated search. *SIGIR Forum* 42, 2 (2008), 80–83. DOI : <http://dx.doi.org/10.1145/1480506.1480520>
- Mor Naaman, Yee Jiun Song, Andreas Paepcke, and Hector Garcia-Molina. 2006. Assigning textual names to sets of geographic coordinates. *Computers, Environment and Urban Systems* 30, 4 (2006), 418–435.
- Zaiqing Nie, Yunxiao Ma, Shuming Shi, Ji-Rong Wen, and Wei-Ying Ma. 2007a. Web object retrieval. In *Proc. of WWW 2007*. 81–90.
- Zaiqing Nie, Ji-Rong Wen, and Wei-Ying Ma. 2007b. Object-level vertical search. In *Proc. of CIDR*. 235–246.
- Shiyao Ou and Christopher S. G. Khoo. 2008. Aggregating search results for social science by extracting and organizing research concepts and relations. In *SIGIR 2008 Workshop on Aggregated Search*.
- C. Paris and N. Colineau. 2006. *Scify: Tailored Corporate Brochures on Demand*. Technical Report. CSIRO ICT Centre.
- C. Paris, A. Lampert, S. Lu, and M. Wu. 2005. *Enhancing Dynamic Knowledge Management Services—Tailored Documents*. Technical Report 05/034, Commercial-in-Confidence. CSIRO ICT Centre.
- Cécile Paris, Stephen Wan, and Paul Thomas. 2010. Focused and aggregated search: A perspective from natural language generation. *Information Retrieval Journal* 44, 3 (2010).
- Cécile Paris, Stephen Wan, Ross Wilkinson, and Mingfang Wu. 2001. Generating personal travel guides—and who wants them? In *Proc. User Modeling 2001*. 251–253.
- Cécile L. Paris. 1988. Tailoring object descriptions to a user's level of expertise. *Comput. Linguist.* 14, 3 (1988), 64–78.
- Marius Pasca and Benjamin Van Durme. 2008. Weakly-supervised acquisition of open-domain classes and class attributes from Web documents and query logs. In *ACL*. 19–27.

- Ashok Kumar Ponnuswami, Kumaresh Pattabiraman, Qiang Wu, Ran Gilad-Bachrach, and Tapas Kanungo. 2011. On composition of a federated web search result page: Using online users to provide pairwise preference for heterogeneous verticals. In *Proc. of WSDM 2011*. 715–724.
- Jay M. Ponte and W. Bruce Croft. 1998. A language modeling approach to information retrieval. In *Proc. of SIGIR 1998*. 275–281.
- Ana-Maria Popescu and Oren Etzioni. 2005. Extracting product features and opinions from reviews. In *Proc. of HLT 2008*. 339–346.
- Jeffrey Pound, Alexander K. Hudek, Ihab F. Ilyas, and Grant Weddell. 2012. Interpreting keyword queries over web knowledge bases. In *Proc. of CIKM 2012*. 305–314.
- Anand Ranganathan, Anton Riabov, and Octavian Udrea. 2009. Mashup-based information retrieval for domain experts. In *Proc. of CIKM 2009*. 711–720.
- Stephen E. Robertson and Steve Walker. 1994. Some simple effective approximations to the 2-Poisson model for probabilistic weighted retrieval. In *Proc. of SIGIR 1994 (Special Issue of the SIGIR Forum)*, W. Bruce Croft and C. J. van Rijsbergen (Eds.). 232–241.
- Cyril Rohr and Dian Tjondronegoro. 2008. Aggregated cross-media news visualization and personalization. In *Proc. of MIR 2008*. 371–378.
- Nachiketa Sahoo, Jamie Callan, Ramayya Krishnan, George Duncan, and Rema Padman. 2006. Incremental hierarchical clustering of text documents. In *Proc. of CIKM 2006*. 357–366.
- G. Salton, A. Wong, and C. S. Yang. 1975. A vector space model for automatic indexing. *Commun. ACM* 18 (November 1975), 613–620. Issue 11. DOI: <http://dx.doi.org/10.1145/361219.361220>
- Rodrygo L. T. Santos, Craig Macdonald, and Iadh Ounis. 2011. Aggregated search result diversification. In *Proc. of the 3rd International Conference on the Theory of Information Retrieval*. Springer, Bertinoro, Italy.
- Christina Sauper and Regina Barzilay. 2009. Automatically generating Wikipedia articles: A structure-aware approach. In *Proc. of ACL-IJCNLP*. 208–216.
- Karen Sauvagnat, Mohand Boughanem, and Claude Chrisment. 2006. Answering content-and-structure-based queries on XML documents using relevance propagation. *Information Systems, Special Issue SPIRE 2004* 31 (2006), 621–635.
- Satoshi Sekine, Kiyoshi Sudo, and Chikashi Nobata. 2002. Extended named entity hierarchy. In *Proc. of LREC 2002*.
- Erik Selberg and Oren Etzioni. 1995. Multi-service search and comparison using the MetaCrawler. In *Proc. of the 4th International World Wide Web Conference*. 195–208.
- Semantic. 2010. Semantic Search Challenge 2010. Retrieved April 2013 from <http://km.aifb.kit.edu/ws/semsearch10/>.
- Semantic. 2011. Semantic Search Challenge 2011. Retrieved April 2013 from <http://semsearch.yahoo.com/>.
- Milad Shokouhi, Justin Zobel, Saied Tahaghoghi, and Falk Scholer. 2007. Using query logs to establish vocabularies in distributed information retrieval. *Inf. Process. Manage.* 43, 1 (January 2007), 169–180. DOI: <http://dx.doi.org/10.1016/j.ipm.2006.04.003>
- Amit Singhal. 2012. Introducing the Knowledge Graph: Things, Not Strings. Official Blog of Google Retrieved April 2013 from <http://googleblog.blogspot.co.uk/2012/05/introducing-knowledge-graph-things-not.html>.
- Karen Spärck-Jones, Stephen E. Robertson, and Mark Sanderson. 2007. Ambiguous requests: Implications for retrieval tests, systems and theories. *SIGIR Forum* 41, 2 (2007), 8–17. DOI: <http://dx.doi.org/10.1145/1328964.1328965>
- K. Srinivas, P. V. S. Srinivas, and A. Govardhan. 2011. A survey on the performance evaluation of various meta search engines. *Int. Journal of Computer Science Issues* 8, 3 (2011), 359–364.
- Andreas Strotmann and Dangzhi Zhao. 2008. Bibliometric maps for aggregated visual browsing in digital libraries. In *Proc. of the SIGIR 2008 Workshop on Aggregated Search*.
- Fabian M. Suchanek, Gjergji Kasneci, and Gerhard Weikum. 2007. YAGO: A core of semantic knowledge. In *Proc. of WWW 2007*. 697–706.
- Fabian M. Suchanek, Gjergji Kasneci, and Gerhard Weikum. 2008. YAGO: A large ontology from Wikipedia and WordNet. *Web Semant.* 6, 3 (2008), 203–217. DOI: <http://dx.doi.org/10.1016/j.websem.2008.06.001>
- Fabian M. Suchanek, Mauro Sozio, and Gerhard Weikum. 2009. SOFIE: A self-organizing framework for information extraction. In *Proc. of the WWW Conference*. 631–640.
- Shanu Sushmita, Hideo Joho, and Mounia Lalmas. 2009. A task-based evaluation of an aggregated search interface. In *Proc. of SPIRE 2009*. 322–333.
- Shanu Sushmita, Hideo Joho, Mounia Lalmas, and Robert Villa. 2010. Factors affecting click-through behavior in aggregated search interfaces. In *Proc. of CIKM 2010*. 519–528.

- TAC. 2011. *Proc. of the 4th Text Analysis Conference*. National Institute of Standards and Technology, Gaithersburg, MD.
- Bilyana Taneva, Mouna Kacimi, and Gerhard Weikum. 2010. Gathering and ranking photos of named entities with high precision, high recall, and diversity. In *Proc. of WSDM 2010*. 431–440.
- Kosuke Tokunaga and Kentaro Torisawa. 2005. Automatic discovery of attribute words from web documents. In *Proc. of the 2nd International Joint Conference on Natural Language Processing (IJCNLP-05)*. 106–118.
- T. Tran, H. Wang, S. Rudolph, and P. Cimiano. 2009. Top-k exploration of query candidates for efficient keyword search on graph-shaped (rdf) data. In *Proc. of ICDE 2009*. 405–416.
- Andrew Trotman, Shlomo Geva, Jaap Kamps, Mounia Lalmas, and Vanessa Murdock. 2010. Current research in focused retrieval and result aggregation. *Journal of Information Retrieval* (2010).
- Peter D. Turney. 2001. Mining the web for synonyms: PMI-IR versus LSA on TOEFL. In *Proc. of EMCL 2001*. 491–502.
- Subodh Vaid, Christopher B. Jones, Hideo Joho, and Mark Sanderson. 2005. Spatio-textual indexing for geographical search on the web. In *Proc. of SSTD*. 218–235.
- David Vallet and Hugo Zaragoza. 2008. Inferring the most important types of a query: A semantic approach. In *Proc. of SIGIR 2008e, Singapore*. 857–858.
- Ellen M. Voorhees. 2003. Evaluating answers to definition questions. In *Proceedings of NAACL'03*. 109–111.
- Tak-Lam Wong and Wai Lam. 2004. A probabilistic approach for adapting information extraction wrappers and discovering new attributes. In *Proc. of ICDM 2004*. 257–264.
- Tak-Lam Wong and Wai Lam. 2009. An unsupervised method for joint information extraction and feature mining across different Web sites. *Data Knowl. Eng.* 68, 1 (January 2009), 107–125. DOI: <http://dx.doi.org/10.1016/j.datak.2008.08.009>
- Fei Wu, Raphael Hoffmann, and Daniel S. Weld. 2008. Information extraction from Wikipedia: Moving down the long tail. In *Proc. of KDD 2008*. 731–739.
- M. Wu and M. Fuller. 1997. Supporting the Answering Process. In *Proc. of the 2nd Australian Document Computing Symposium*. 65–73.
- Naoki Yoshinaga and Kentaro Torisawa. 2007. Open-domain attribute-value acquisition from semi-structured texts. In *Proc. of the Workshop on Ontolex*. 55–66.
- K. Zhou, R. Cummins, M. Lalmas, and J. M. Jose. 2011. Evaluating large-scale distributed vertical search. In *Proc. of LSDS-IR Workshop in CIKM*.
- Ke Zhou, Ronan Cummins, Mounia Lalmas, and Joemon M. Jose. 2012. Evaluating aggregated search pages. In *Proc. of SIGIR 2012*. 115–124.
- J. Zhu, Z. Nie, X. Liu, B. Zhang, and J.-R. Wen. 2009. Statsnowball: a statistical approach to extracting entity relationships. In *Proc. of WWW 2009, Madrid, Spain*. 101–110.