

A Perspective Look at Keyword-based Search Over Relation Data and its Evaluation (Extended Abstract)*

Sonia Bergamaschi¹, Nicola Ferro², Francesco Guerra¹, and Gianmaria Silvello²

¹ Department of Engineering “Enzo Ferrari”, University of Modena and Reggio Emilia, Italy
{sonia.bergamaschi,francesco.guerra}@unimore.it

² Department of Information Engineering, University of Padua, Italy
{ferro,silvello}@dei.unipd.it

Abstract. This position paper discusses the need for considering keyword search over relational databases in the light of broader systems, where keyword search is just one of the components and which are aimed at better supporting users in their search tasks. These more complex systems call for appropriate evaluation methodologies which go beyond what is typically done today, i.e. measuring performances of components mostly in isolation or not related to the actual user needs, and, instead, able to consider the system as a whole, its constituent components, and their inter-relations with the ultimate goal of supporting actual user search tasks.

1 Introduction

Keyword search is the foremost approach for information searching and in the last decades it has been extensively studied in the field of Information Retrieval (IR). Nevertheless, this model has left out the structured data sources which are typically accessed through structured queries. Structured queries are not suitable for the generic user, since their formulation requires users to know a language and the data source schemas and contents. The large availability of structured data has made of paramount importance the development of keyword search tools for this kind of sources.

In relational databases, keyword search aims at easing and somehow automating the search process by employing two main techniques [8, 9, 31]: schema-based and graph-based. Graph-based techniques model relational databases as graphs, where nodes are tuples, edges foreign-primary key relationships between those tuples and the algorithms are based on the computation of specific structures over the graphs. Whereas, schema-based techniques exploit the schema information to formulate Structured Query Language (SQL) queries determined starting from the user keyword queries. Recently metadata-based approaches have been developed [6, 3, 5]. These approaches are useful when there is not direct access to the database instance or when frequent updates make the process of building and updating indexes too expensive.

On the other hand, it should be considered that search process is part of a wider user task [19, 23] from which the information need arises and, in turn, makes the user resort

* Original publication in [4].

to issuing queries to satisfy it. This makes the whole process quite complex and brings in the accomplishment of the user information need several degrees of uncertainty.

Even if the main focus of keyword search, i.e. getting out the most from the relational data starting from a keyword instead of a structured query, is certainly something that helps users in carrying out their tasks, many factors, beyond algorithmic correctness and completeness, may impair the impact of keyword search and prevent users to fully exploit its potentialities. Therefore, we need to put keyword search into a broader context and envision innovative architectures where keyword search is one of the components, paired with other building blocks to better take into account the variability and uncertainty entailed by the whole search process.

We claim that a conceptual architecture pivoting around keyword search and relational data needs to couple system- and user-oriented components; the former ones aim at augmenting the performances (i.e. efficiency) of the search, whereas the latter ones aim at improving the quality of the search (i.e. effectiveness) from the user perspective. Such system- and user-oriented components already exist and have been demonstrated to be effective for their specific purpose. Nevertheless, their integration into a unique framework for keyword search is still lacking. Evidence of this is the absence of a commercially deployed system.

Moreover, measuring is a key to scientific progress and experimental evaluation is a key means for supporting and fostering the development of systems which are more adherent to the user needs, provide the desired effectiveness and efficiency, guarantee the required robustness and reliability, and operate with the necessary scalability. In light of this, we claim that the current frameworks for the evaluation of keyword search in relational databases [11] need to be re-thought, by moving beyond the evaluation of keyword search components in isolation or not related to the actual user needs, and, instead, by considering the whole system, its constituents, and their inter-relations with the ultimate goal of supporting actual user search tasks.

The paper is organized as follows: Section 2 describes the main components of the proposed approach; Section 3 defines the main characteristics of a framework for evaluating such an innovative keyword search system. Finally, Section 4 draws some final remarks.

2 Conceptual Architecture

In Figure 1 we can see, on the right-hand side, a general architecture of an information access system pivoting on keyword search techniques, and on the left-hand side, the evaluation methodologies that can be employed for evaluating the efficiency and efficacy of the system. Thick solid lines frame modules which are the focus of current keyword search systems, whereas dotted lines frame components which are typically not exploited today and come from neighboring fields [12], such as information retrieval [7], information extraction, data mining, and natural language processing [22].

The main keyword search and relational database layers are surrounded by a Keyword Identification layer, a Results Filtering and Ranking layer, and a Presentation layer. Indeed, an ideal search system has to consider the search task a user desires to conduct, to perform user queries knowing that they may not exactly correspond to

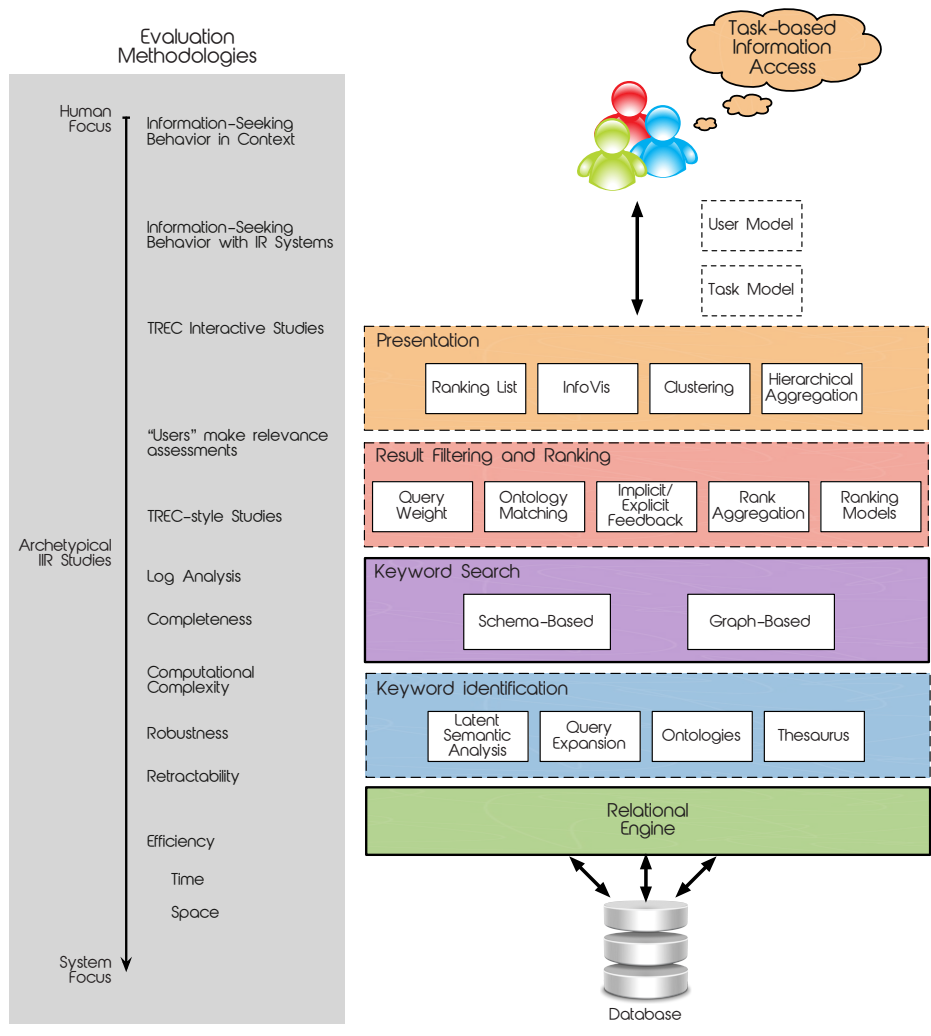


Fig. 1. The Conceptual Architecture of a Keyword-based Search System and Evaluation.

the real user information need, to disambiguate search terms, to rank the results of the search process on the basis of relevance for the user, and to visualize these results in the most proper way for the considered search task.

The *Keyword Identification* layer is aimed at freeing keyword search from an exact match with the keywords present in the relational data and introduce the possibility of matching in multiple ways the keywords expressed by user to the relational data in order to compensate for possible imprecisions or errors in the choice of the keywords.

The *Results Filtering and Ranking* layer accounts for the need of adopting alternative strategies for ranking and selecting the results to be presented to the user by the system. It may concern weighting the results on the basis of the process which generated relational queries from user keywords, or relying on implicit/explicit feedback from the user to filter out some results, or using rank aggregation and data fusion techniques [30] to merge alternative ranking strategies.

The *Presentation* layer regards how the outputs of a system are presented to the user; for instance, we can have traditional ranking lists, results presentation based on advanced Information Visualization techniques [18, 32], or presentations of clusters of results.

3 Evaluation

Innovative proposals for pushing the boundaries of keyword search cannot set aside a proper and shared evaluation methodology which helps in progressing towards the envisioned goals, ensures the soundness and quality of the proposed solutions, and guarantees the repeatability and comparability of the experiments.

As shown in Figure 1, evaluation can be carried out at three levels: at a “task level” for instance by means of user studies [21]; at an “effectiveness level” by means of the test collection methodology [10, 16, 25]; and, at a “system level” by means of benchmarking queries per second, memory and CPU load, correctness and completeness [11], thus moving from a human to a system focus.

Nevertheless, experimental evaluation is hampered by fragmentation – different tasks, different collections, different perspectives from interactive to laboratory evaluation which are usually dealt with in separated ways, without sharing resources and the produced data [29]. This will be even more true for the multidisciplinary approach to the system architecture proposed on the right of Figure 1, whereas the unified and holistic approach to evaluation, proposed on the left, would be needed to assess the different facets of such a complex system and to reconcile the experimental outcomes.

Furthermore, systematic and comparable experimental evaluation is a very demanding activity, in terms of both time and effort needed to perform it. For this reason, in the IR field, it is usually carried out in publicly open and large-scale evaluation campaigns at international level, which allow for sharing the effort, producing large experimental collections, and comparing state-of-the-art systems and algorithms. Relevant and long-lived examples are TREC³ [17] in the United States since 1992, the CLEF initia-

³ <http://trec.nist.gov/>

tive⁴ [13, 15] in Europe since 2000, and NTCIR⁵ in Japan and Asia [20] since 1999, which produced huge amounts of experimental data [1, 14] representing an extremely valuable asset for past, current, and future research. Relying on these experiences would move evaluation of keyword search forward, by sharing resources, providing open fora to compare and discuss approaches, re-use the data and the acquired knowledge, have the possibility of repeating the obtained results, and reduce the overall effort.

Evaluation campaigns not only have a proven scholarly impact [2, 26–28] but also, as reported by [24], for every \$1 that NIST and its partners invested in TREC, at least \$3.35 to \$5.07 in benefits accrued to researchers and industry. During their life-span, large-scale evaluation campaigns have produced huge amounts of scientific data which are extremely valuable for research and development but also from an economic point of view: [24] estimates that the overall investment in TREC of about 30 million dollars in its first 20 years which, as discussed above, produced an estimated return on investment between 90 and 150 million dollars. Therefore, applying experimental evaluation to vision represented in Figure 1 gives the promise not only to advance state-of-the-art techniques, but also to have a concrete economic impact.

4 Use Case

Let us consider a relational database containing tourism information about accommodations. For sake of simplicity, the database is composed of two tables: *Accommodation* and *City*, describing respectively accommodations and the cities where the accommodations are located. Some descriptive attributes are defined for each table and a foreign key - primary key relationship exists between the attribute *City* in table *Accommodation* and the attribute *Name* in table *City*.

By traditional keyword search techniques, even the simple query “Hotel Venice” may not be trivial to answer. Firstly, an index-based approach can find several matches for the keyword “Venice” with instances of the table *Accommodation* (i.e., Venice is the *Name* of an accommodation, or Venice is a value of the attribute *City* in the same table) or with an instance of the table *City* (i.e Venice is the name of the city). This is due to the fact that the user query is ambiguous and it is unclear whether the user is looking for hotels called “Venice” or hotels which are in “Venice”. Moreover, a classical keyword search engine may not find any correspondences between the keyword “Hotel” and an element in the database. In this case, the problem is two-fold: firstly, the chosen keyword may better match metadata instead of a value in a table – indeed, several existing keyword search engines do not consider database metadata as possible targets for user queries; secondly, the chosen keyword may actually refer to a concept represented in the table with a value which is a synonym of the chosen keyword. Since users do not know the information in the data source, this situation frequently occurs and is typically not managed by most of the existing systems.

This problem can be even more complex if some keywords match with foreign key - primary key values. This occurs in our example if the user formulates the query “Venice”, where it is unspecified if the user is interested in accommodations or cities.

⁴ <http://www.clef-initiative.eu/>

⁵ <http://research.nii.ac.jp/ntcir/>

From an algorithmic perspective an answer showing all the hotels in Venice provides an answer as good as the one that reports all the information about the city. Nevertheless, from the user perspective, only one of these two sets of results is relevant to her/his information needs, e.g. the city one, while providing both of them would reduce the effectiveness of the system and hampers performance.

Even in this toy example, it becomes clear that query ambiguity as well as the choice between graph or schema based techniques impacts the system performance. Therefore, the complexity of real scenarios may take even more advantage from the architecture proposed in Figure 1, which complements keyword search with additional components that, in this specific example, analyze the user keywords and disambiguate their meaning. Similarly, the evaluation methodology must be able to detect these different issues in order to properly assess how systems tackle them.

As a further example, it is not clear from the query which is the information that the user would like to receive as a result. Several options are possible: in some cases the user would like to receive all the data about the tuples satisfying the criteria defined by the keyword query (i.e., the “universal relation”), in other cases only the values of some attributes (i.e., a projection of the “universal relation”), or even a boolean value (i.e., the existence of at least an instance). On the other hand, she/he also would like to receive results ordered by the system estimation of their relevance and how much they satisfy her/his information need. The task of understanding the granularity, the ordering and aggregation of the expected results is not typically managed by the existing systems and requires a specific module in the architecture of a complete keyword search system. As above, evaluation must be able to take into account these aspects and assess the systems accordingly.

Acknowledgements

The reported work has been partially supported by the KEYSTONE (<http://www.keystone-cost.eu/keystone/>, IC1302) cost action.

References

1. M. Agosti, G. M. Di Nunzio, and N. Ferro. Scientific Data of an Evaluation Campaign: Do We Properly Deal With Them? In C. Peters, P. Clough, F. C. Gey, J. Karlgren, B. Magnini, D. W. Oard, M. de Rijke, and M. Stempfhuber, editors, *Evaluation of Multilingual and Multimodal Information Retrieval : Seventh Workshop of the Cross-Language Evaluation Forum (CLEF 2006). Revised Selected Papers*, pages 11–20. Lecture Notes in Computer Science (LNCS) 4730, Springer, Heidelberg, Germany, 2007.
2. M. Angelini, N. Ferro, B. Larsen, H. Müller, G. Santucci, G. Silvello, and T. Tsirikra. Measuring and Analyzing the Scholarly Impact of Experimental Evaluation Initiatives. *Procedia Computer Science*, 38:133–137, 2014.
3. S. Bergamaschi, E. Domnori, F. Guerra, R. Trillo-Lado, and Y. Velegrakis. Keyword Search over Relational Databases: A Metadata Approach. In T. Sellis, R. J. Miller, A. Kementsietidis, and Y. Velegrakis, editors, *Proc. ACM SIGMOD International Conference on Management of Data (SIGMOD 2011)*, pages 565–576. ACM Press, New York, USA, 2011.

4. S. Bergamaschi, N. Ferro, F. Guerra, and G. Silvello. Keyword Search and Evaluation over Relational Databases: an Outlook to the Future. In G. Koutrika and M. Theobald, editors, *Proc. 7th International Workshop on Ranking in Databases (DBRank 2013)*, pages 8:1–8:3. ACM Press, New York, USA, 2013.
5. S. Bergamaschi, F. Guerra, S. Rota, and Y. Velegrakis. A Hidden Markov Model Approach to Keyword-Based Search over Relational Databases. In M. A. Jeusfeld, L. M. L. Delcambre, and T. W. Ling, editors, *Proc. of the 30th International Conference on Conceptual Modeling (ER 2011)*, pages 411–420. Lecture Notes in Computer Science (LNCS) 6998, Springer, Heidelberg, Germany, 2011.
6. Lukas Blunski, Claudio Jossen, Donald Kossmann, Magdalini Mori, and Kurt Stockinger. SODA: Generating SQL for Business Users. *PVLDB*, 5(10):932–943, 2012.
7. S. Büttcher, C. L. A. Clarke, and G. V. Cormack. *Information Retrieval: Implementing and Evaluating Search Engines*. The MIT Press, Cambridge (MA), USA, 2010.
8. S. Chaudhuri and G. Das. Keyword Querying and Ranking in Databases. *Proceedings of the VLDB Endowment (PVLDB)*, 2(2):1658–1659, 2009.
9. E. Chu, A. Baid, X. Chai, A. Doan, and J. F. Naughton. Combining Keyword Search and Forms for Ad Hoc Querying of Databases. In U. Çetintemel, S. B. Zdonik, D. Kossmann, and N. Tatbul, editors, *Proc. ACM SIGMOD International Conference on Management of Data (SIGMOD 2009)*, pages 349–360. ACM Press, New York, USA, 2009.
10. C. W. Cleverdon. The Cranfield Tests on Index Languages Devices. In K. Spärck Jones and P. Willett, editors, *Readings in Information Retrieval*, pages 47–60. Morgan Kaufmann Publisher, Inc., San Francisco, CA, USA, 1997.
11. J. Coffman and A. C. Weaver. An Empirical Performance Evaluation of Relational Keyword Search Techniques. *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, 1(26):30–42, 2014.
12. N. Ferro, editor. *Bridging Between Information Retrieval and Databases - PROMISE Winter School 2013, Revised Tutorial Lectures*. Lecture Notes in Computer Science (LNCS) 8173, Springer, Heidelberg, Germany, 2014.
13. N. Ferro. CLEF 15th Birthday: Past, Present, and Future. *SIGIR Forum*, 48(2):31–55, December 2014.
14. N. Ferro, A. Hanbury, H. Müller, and G. Santucci. Harnessing the Scientific Data Produced by the Experimental Evaluation of Search Engines and Information Access Systems. *Proceedia Computer Science*, 4:740–749, 2011.
15. N. Ferro and G. Silvello. CLEF 15th Birthday: What Can We Learn From Ad Hoc Retrieval? In E. Kanoulas, M. Lupu, P. Clough, M. Sanderson, M. Hall, A. Hanbury, and E. Toms, editors, *Information Access Evaluation – Multilinguality, Multimodality, and Interaction. Proceedings of the Fifth International Conference of the CLEF Initiative (CLEF 2014)*, pages 31–43. Lecture Notes in Computer Science (LNCS) 8685, Springer, Heidelberg, Germany, 2014.
16. D. K. Harman. *Information Retrieval Evaluation*. Morgan & Claypool Publishers, USA, 2011.
17. D. K. Harman and E. M. Voorhees, editors. *TREC. Experiment and Evaluation in Information Retrieval*. MIT Press, Cambridge (MA), USA, 2005.
18. M. A. Hearst. *Search User Interfaces*. Cambridge University Press, New York, NY, USA, 1st edition, 2009.
19. P. Ingwersen and K. Järvelin. *The Turn: Integration of Information Seeking and Retrieval in Context*. Springer, Heidelberg, Germany, 2005.
20. N. Kando and K. Kishida, editors. *Proc. 10th NTCIR Conference on Conference on Evaluation of Information Access Technologies*. National Institute of Informatics, Tokyo, Japan, 2013.

21. D. Kelly. Methods for evaluating interactive information retrieval systems with users. *Foundations and Trends in Information Retrieval (FnTIR)*, 3(1–2):1–224, 2009.
22. C. D. Manning and H. Schütze. *Foundations of Statistical Natural Language Processing*. The MIT Press, Cambridge (MA), USA, 4th edition, 2001.
23. G. Marchionini. Exploratory Search: From Finding to Understanding. *Commun. ACM*, 49(4):41–46, 2006.
24. B. R. Rowe, D. W. Wood, A. L. Link, and D. A. Simoni. *Economic Impact Assessment of NIST's Text REtrieval Conference (TREC) Program*. RTI Project Number 0211875, RTI International, USA. <http://trec.nist.gov/pubs/2010.economic.impact.pdf>, July 2010.
25. M. Sanderson. Test Collection Based Evaluation of Information Retrieval Systems. *Foundations and Trends in Information Retrieval (FnTIR)*, 4(4):247–375, 2010.
26. C. V. Thornley, A. C. Johnson, A. F. Smeaton, and H. Lee. The Scholarly Impact of TRECVID (2003–2009). *Journal of the American Society for Information Science and Technology (JASIST)*, 62(4):613–627, April 2011.
27. T. Tsirikika, A. Garcia Seco de Herrera, and H. Müller. Assessing the Scholarly Impact of ImageCLEF. In P. Forner, J. Gonzalo, J. Kekäläinen, M. Lalmas, and M. de Rijke, editors, *Multilingual and Multimodal Information Access Evaluation. Proceedings of the Second International Conference of the Cross-Language Evaluation Forum (CLEF 2011)*, pages 95–106. Lecture Notes in Computer Science (LNCS) 6941, Springer, Heidelberg, Germany, 2011.
28. T. Tsirikika, B. Larsen, H. Müller, S. Endrullis, and E. Rahm. The Scholarly Impact of CLEF (2000–2009). In P. Forner, H. Müller, R. Paredes, P. Rosso, and B. Stein, editors, *Information Access Evaluation meets Multilinguality, Multimodality, and Visualization. Proceedings of the Fourth International Conference of the CLEF Initiative (CLEF 2013)*, pages 1–12. Lecture Notes in Computer Science (LNCS) 8138, Springer, Heidelberg, Germany, 2013.
29. G. Weikum. Where's the Data in the Big Data Wave? ACM SIGMOD Blog, <http://wp.sigmod.org/?p=786>, March 2013.
30. S. Wu. *Data Fusion in Information Retrieval*. Springer-Verlag, Heidelberg, Germany, 2012.
31. J. X. Yu, L. Qin, and L. Chang. Keyword Search in Relational Databases: A Survey. *IEEE Data(base) Engineering Bulletin*, 33(1):67–78, 2010.
32. J. Zhang. *Visualization for Information Retrieval*. Springer-Verlag, Heidelberg, Germany, 2008.