

Accessing the Deep Web with Keywords: A Foundational Approach

Andrea Cali^{1,2}(✉) and Martín Ugarte³

¹ Department of Computer Science and Information Systems, Birkbeck,
University of London, London, UK

`andrea@dcs.bbk.ac.uk`

² Oxford-Man Institute of Quantitative Finance, University of Oxford, Oxford, UK

³ Laboratory of Web and Information Technologies,

Université Libre de Bruxelles, Brussels, Belgium

`mugartec@ulb.ac.be`

Abstract. The Deep Web is constituted by data that are generated dynamically as the result of interactions with Web pages. The problem of accessing Deep Web data presents many challenges: it has been shown that answering even simple queries on such data requires the execution of recursive query plans. There is a gap between the theoretical understanding of this problem and the practical approaches to it. The main reason behind this is that the problem is to be studied by considering the database as part of the input, but queries can be processed by accessing data according to *limitations*, expressed as so-called access patterns. In this paper we embark on the task of closing the above gap by giving a precise definition that reflects the practical nature of accessing Deep Web data sources. In particular, we define the problem of querying Deep Web sources with keywords. We describe two scenarios: in the first, called *unrestricted*, there query answering algorithm has full access to the data; in the second, called *restricted*, the algorithm can access the data only according to the access patterns. We formalise the associated decision problem associated to that of query answering in the Deep Web, explaining its relevance in both the aforementioned scenarios. We then present some complexity results.

1 Introduction

The *Deep Web* (also called *Hidden Web*) [3, 6, 9] refers to the data content that is created dynamically as the result of interactions with the Web. For example, when we search for a person in a White Pages website, the generated output consists of one or more pages containing the result of a query posed on an underlying database; these pages cannot be indexed by search engines and the underlying database cannot be freely queried by users. When we search in whitepages.com through a form, we are forced to fill in certain fields of the form, for instance the **Name** field; the result is then structured as a table. A Deep Web source can be naturally modeled as a relational table (or a set of relational tables) that can be queried only according to so-called *access patterns*, each of which enforces

the selection on some of the attributes (i.e. the filling of input fields on a form), which are called *input* attributes. Relational tables accessible through access patterns then are said to have *access limitations*.

Obtaining data dynamically from Deep Web sources is the key problem in the integration of such sources. Interestingly, when Deep Web sources are modelled as relations with access limitations, answering a simple query on such sources may require the evaluation of a *recursive* Datalog query plan [3, 7]. In such plans, values obtained as output from one search are used as input for other sources.

In this paper we study the problem of *accessing Deep Web sources via keywords*, that is, given a query, a set of keywords with which to access the sources and a database, computing the answers to the query. Interestingly, the keywords in this context are not used to select the answers or the sources, but to *retrieve* data from the sources. This problem is related to conjunctive query (CQ) answering, an extensively studied topic in the literature [1, 3, 7]; however, in the case of sources with access limitations a thorough theoretical study of the central problems and their computational complexity is surprisingly still lacking. In this paper we present the keyword querying problem in a formal way, distinguishing two variants of it. We then provide results on the computational complexity of the boolean case of both variants.

2 Preliminaries

We consider the relational setting extended with *access limitations* and *abstract domains*. We assume the reader is familiar with the well-known notions of relations, attributes, variables, constants and ground atoms (a.k.a. facts); and the relational setting in general. For formal definitions we refer the reader to [1].

Access limitations on a relation are constraints imposing that certain attributes must be *selected* (that is, bound to a constant) for the relation to be accessed. More formally, a schema with access limitations is a pair $\langle \mathcal{R}, \Lambda \rangle$, where \mathcal{R} is a relational schema (a set of relations) and Λ is a set of access limitations that specifies, for every attribute of every relational predicate, whether it is an *input* or an *output* attribute; in order to access a relation, all input attributes must be selected¹. We indicate the access limitations of each relation as a sequence, of ‘*i*’ and ‘*o*’ symbols written as a superscript in the signature of the relation; an ‘*i*’ (resp., ‘*o*’) indicates that the corresponding argument is an input (resp., output) argument. A signature has therefore the form r^{Λ_r} , where Λ_r represents the access limitation on r . In our setting (see also [3]) some general domains, called *abstract domains*, are associated to attributes; these attributes are used to distinguish, for instance, strings representing names from strings representing addresses. To avoid notational clutter, we assume that attribute names are assigned so that attributes having the same abstract domain also

¹ In general, there could be more than one annotation for each predicate, that is, more than one way of accessing the corresponding relation. However, in this paper we assume there is exactly one access limitation (or pattern) per predicate. Our results can be extended to the general case.

have the same name. The problem that we study in this paper consist of two parts: a database is first accessed through access limitations following the known abstract domains, and the *obtainable* part of the database is then queried by using the most common class of queries, namely *conjunctive queries* (CQs) [1].

In the presence of access limitations on the sources, queries cannot be evaluated as in the traditional case. As we don't have direct access to the database, we need a set I of initial keywords to start *scraping* the database. This has been previously noted in [8], where the authors present an algorithm that extracts all *obtainable* tuples in the answer to the query. This algorithm compiles the evaluation strategy into a suitable Datalog program, which encodes both the access limitations on the sources and the query itself, and is evaluated as follows: starting from a set of initial keywords (that must include those appearing as constants in the query), we access all the relations we can according to the access limitations. With the new facts (if any), we obtain new keywords with which we can repeat the process and access the relations again, until we have no way of making new accesses. The program extracts all facts obtainable while respecting the access limitations, but there may be facts in the sources that cannot be retrieved.

Example 1. Consider the relations $r_1^{i\circ\circ}(N, D, C)$ and $r_2^{i\circ\circ}(C, S, N)$ depicted in Fig. 2. The tuples in r_1 contain a *Nation*, a typical *Dish* of that nation, and a famous *Chef* that prepares it; the tuples in r_2 contain a *Chef*, the amount of Michelin *Stars* he has obtained and his *Nationality*. The access limitations only allow for searching typical dishes by nation (from r_1), and searching chefs by surname (from r_2). Assume we want to obtain the dishes prepared by chefs with three Michelin stars; this is expressed by the conjunctive query $q(D) \leftarrow r_1(N_1, D, C), r_2(C, 3, N_2)$. However, because of the access limitations, we cannot directly pose this query to the database. Instead, we need to recursively access the database starting from a set of keywords known in advance. For example, assume we know that this database contains information about Italy; we then have available the keyword set $\{Italy\}$. We can thus search r_1 using *Italy*, obtaining the tuple t_1 . Now we have the last name of chef Heinz Beck, so we can search r_2 with input *Beck*. This returns s_1 , indicating that *Risotto* is a dish prepared by a chef with three Michelin stars (and thus part of our answer). But s_1 also contains the value *Germany*, which we can use as input to query r_1 again; this time we get t_2 , which contains the last name of a new chef, *Ducasse*. We can now query r_2 with input *Ducasse*, obtaining tuple s_2 and discovering that *Magenbrot* is also part of the answer. The tuple s_2 also contains a new country, *France*. However, when we query r_1 with *France* we get an empty result, and therefore there are no more tuples we can obtain from the database. Notice that *Onigiri* would also be part of the answer if we were computing the answers without access limitations. In our case we could retrieve this value if *Japan* was one of our initial keywords or it was extracted at some point. The recursive procedure illustrated above is formalised by evaluating the Datalog program depicted in Fig. 1. In this program, relations \hat{r}_1 and \hat{r}_2

$\rho_1 : q(D) \leftarrow \hat{r}_1(N, D, C), r_2(C, 3, N)$
 $\rho_2 : \hat{r}_1(N, D, C) \leftarrow r_1(N, D, C), \text{dom}_N(N)$
 $\rho_3 : \hat{r}_2(C, S, N) \leftarrow r_2(C, S, N), \text{dom}_C(C)$
 $\rho_4 : \text{dom}_C(C) \leftarrow \hat{r}_1(N, D, C)$
 $\rho_5 : \text{dom}_A(N) \leftarrow \hat{r}_2(C, S, N)$
 $\rho_6 : \text{dom}_N(\text{Italy})$

Fig. 1. Datalog program of Example 1.

$$r_1^{i\infty}$$

	N	D	C
t_1	Italy	Risotto	Beck
t_2	Germany	Magenbrot	Ducasse
t_3	Japan	Onigiri	Robuchon

$$r_2^{i\infty}$$

	N	S	C
s_1	Beck	3	Germany
s_2	Ducasse	3	France
s_3	Robuchon	3	France

Fig. 2. Database of Example 1.

represent the *obtainable* parts of r_1 and r_2 , respectively (assuming we start only from the keyword set $\{\text{Italy}\}$). Rule ρ_1 represents the original query (over the *obtainable* versions of r_1 and r_2), rules ρ_2 to ρ_5 encode the recursive access to the sources, and ρ_6 simply *initialises* the constant *Italy* by adding it to the abstract domain of nations. ■

The previous example shows the typical way in which Deep Web sources are accessed and queried over the web. We now define the notion of answer and the *obtainable* portion of a database under access limitations; such a portion is determined by the initial keywords. Given a CQ q posed over a schema $\mathcal{S} = \langle \mathcal{R}, A \rangle$, a set of initial keywords $I \subseteq \Delta$, and a database D over schema \mathcal{R} , $\rho_{A,I}(D)$ denotes the set of facts of D that can be recursively obtained under A starting from I . The set of answers to q over D with access limitations A and initial set of keywords I is denoted by $\text{ans}(q, A, D, I)$ and is defined as the set of answers, in the classic sense, to q over $\rho_{A,I}(D)$. If q is a Boolean CQ (i.e., with zero-arity head), we write $D \models_{A,I} q$ when q is true on $\rho_{A,I}(D)$ (denoted $\rho_{A,I}(D) \models q$).

3 The Complexity of Querying Under Access Limitations

In this section we study the complexity of answering queries on Deep Web datasets, where data are to be extracted from an initial set of keywords. Surprisingly, the notions of complexity present in the literature do not seem to fully capture the correct difficulty of the problem. To clarify this problem, we present two variants of the Boolean query answering problem (the extension to the non-Boolean case is straightforward).

Definition 1. *Given a database D , a set of initial keywords $I \subseteq \Delta$, a set A of access limitations and a BCQ q , the problem of query answering with initial keywords I and query q , on database D and under A , is to determine whether $D \models_{A,I} q$. This is defined in two variants:*

- (i) *unrestricted case: this is the problem of determining whether $D \models_{A,I} q$, while having arbitrary access to the relations of D .*
- (ii) *restricted case: this is the problem of determining whether $D \models_{A,I} q$, while having access to the relations of D only according to A .*

Notice that the problem in the restricted case is the “classic” case [2,3], where we are computing the answers to a CQ having only limited access to the data, according to A . The CQ answering problem in the unrestricted case is also relevant in real-world scenarios. Assume for example that access limitations are enforced by an organisation in order to limit access to data by external users (e.g., those outside the organisation). The organisation has arbitrary access to the data, and it is interested in determining what external users, who probably know certain initial keywords and whose access to the data is limited by A , can retrieve from the database. In order to determine this, of course, an algorithm will have the advantage of freely accessing the data, regardless of access limitations.

We need to point out that in the restricted case, if we are to tackle the search problem formally, we need to understand the relations with access limitations as *oracles*; each access (that consists in the processing of an atomic query on a single relation) is a call to an oracle corresponding to the same relation. The execution of such a query takes evidently (at most) linear time in the size of the instance of the relation, therefore the oracle does not really serve to determine a complexity class, as is done, for example, when we have a class $\mathcal{C}_1^{\mathcal{C}_2}$ of problems that can be decided by an algorithm of class \mathcal{C}_1 that can call an oracle solving problems in class \mathcal{C}_2 (each costing 1, given the nature of the oracle). The oracles in our case do not add computational power; instead they *limit* the access to the data rather than allowing the solution to a problem instance in constant time. In this case, the algorithm cannot receive the instance D as (fully accessible) input; yet we want to measure its complexity considering the size of D . Interestingly, this is why the classical notion of complexity does not capture the actual difficulty of the problem. The following example shows that there are instances in which the two variants of the problem actually present different complexity.

Example 2. Consider the schema $\mathcal{R} = \{r^{i \cdots i}\}$, constituted by a single relation with k input arguments and no output arguments, the single-tuple database $D = \{r(a_1, \dots, a_k)\}$, the keyword set $I = \{c_1, \dots, c_m\} \supseteq \{a_1, \dots, a_k\}$, and the atomic Boolean CQ q defined as $q() \leftarrow r(X_1, \dots, X_k)$. In the restricted case, to answer q (checking if $r(a_1, \dots, a_k)$ is in D), one needs to try accessing the relation $r(D)$ with all possible k -tuples of constants of I ; in the worst case, this requires m^k accesses to r . On the contrary, in the unrestricted case the query can be answered trivially. ■

The above example, which uses very simple CQs (atomic Boolean CQs), shows a case where, if using deterministic algorithms, query answering is easy in the unrestricted case, but at least exponential in the restricted case.

We now briefly discuss a result, stated in previous works [4,5], on CQ answering in our setting. The result states that CQ answering is NP-complete both

in the restricted and the unrestricted case. In the light of Example 2, this is somewhat counterintuitive, as the restricted case appears to be computationally more difficult than the unrestricted case. Regarding upper bounds, the most interesting technique is the one used to prove that the problem is in NP in the restricted case: a non-deterministic algorithm is exhibited, whose maximum number of steps, in the worst case, is surprisingly bounded by the number of atoms in the database D . The lower bounds are given instead by the obvious NP lower bound of CQ answering in the case *without* access limitations. However, the lower bound in the restricted case does not constitute a fully satisfactory study of the complexity as it does not make use of the restrictions given by the presence of the oracles; indeed it still remains to *define*: (1) what kind of computational model we need to model the oracles; (2) what kind of reduction would imply a complexity lower bound in this setting. In addition, it remains to understand whether simpler classes of CQs (e.g. atomic, acyclic or bounded-treewidth CQs) enjoy lower complexity. This will be the subject of future investigation.

4 Discussion

In this paper we have introduced two variants of the problem of querying Deep Web sources with a set of initial keywords, namely the *restricted* and an *unrestricted* case. We have shown that the two variants can differ by an exponential factor in very simple cases. However, the problem of CQ answering with keywords under access limitations has been shown to be NP-complete in both the restricted and unrestricted case. As future work we plan to carry out a formal definition of the associated decision problem for the restricted case, and the characterization of those classes of queries for which the complexity of the two variants differ. For instance, we will investigate whether the NP lower bound holds for atomic queries in the restricted case (which is the case of Example 2), or for other restricted classes of CQs such as acyclic or bounded-treewidth CQs.

Acknowledgments. This work was supported by the EU COST Action IC1302 KEYSTONE. Andrea Cali acknowledges partial support by the EPSRC project “Logic-based Integration and Querying of Unindexed Data” (EP/E010865/1).

References

1. Abiteboul, S., Hull, R., Vianu, V.: Foundations of Databases. Addison-Wesley, Reading (1995)
2. Cali, A., Martinenghi, D.: Conjunctive query containment under access limitations. In: Li, Q., Spaccapietra, S., Yu, E., Olivé, A. (eds.) ER 2008. LNCS, vol. 5231, pp. 326–340. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-87877-3_24
3. Cali, A., Martinenghi, D.: Querying data under access limitations. In: Proceedings of ICDE (2008)
4. Cali, A., Martinenghi, D., Razgon, I., Ugarte, M.: Querying the deep web: back to the foundations. In: Proceedings of AMW (2017). To appear

5. Cali, A., Razgon, I.: Complexity of conjunctive query answering under access limitations (preliminary report). In: Proceedings of SEBD (2014)
6. Chang, K.C.-C., He, B., Zhang, Z.: Toward large scale integration: building a meta-querier over databases on the web. In: Proceedings of CIDR (2005)
7. Li, C.: Computing complete answers to queries in the presence of limited access patterns. *Very Large Database J.* **12**(3), 211–227 (2003)
8. Li, C., Chang, E.: Query planning with limited source capabilities. In: Proceedings of ICDE (2000)
9. Madhavan, J., Afanasiev, L., Antova, L., Halevy, A.Y.: Harnessing the deep web: present and future. In: Proceedings of CIDR (2009)