Categorize by: Deductive Aggregation of Semantic Web Query Results

Claudia d'Amato¹, Nicola Fanizzi¹, and Agnieszka Ławrynowicz²

¹ Dipartimento di Informatica, Università degli Studi di Bari, Italy
{claudia.damato,fanizzi}@di.uniba.it

² Institute of Computing Science, Poznan University of Technology, Poland
alawrynowicz@cs.put.poznan.pl

Abstract. Query answering on a wide and heterogeneous environment such as the Web can return a large number of results that can be hardly manageable by users/agents. The adoption of grouping criteria of the results could be of great help. Up to date, most of the proposed methods for aggregating results on the (Semantic) Web are mainly grounded on syntactic approaches. However, they could not be of significant help, when the values instantiating a grouping criterion are all equal (thus creating a unique group) or are almost all different (thus creating one group for each answer). We propose a novel approach that is able to overcome such drawbacks: given a query in the form of a conjunctive query, grouping is grounded on the exploitation of the semantics of background ontologies during the aggregation of query results. Specifically, we propose a solution where answers are deductively grouped taking into account the subsumption hierarchy of the underlying knowledge base. In this way, the results can be shown and navigated similarly to a faceted search. An experimental evaluation of the proposed method is also reported.

1 Introduction

The users of the (Semantic) Web often perform interactive, and exploratory data retrieval, where queries often result in an overwhelming number of the returned answers. However typically, only a small part of the result set is relevant to the user thus making necessary the analysis of the retrieved results to identify those relevant ones. This phenomenon, known as *information overload*, constitutes a ceaseless challenge for researchers. It may occur in situations when the user, at the beginning of a search, submits a broad query to avoid exclusion of possibly interesting results, only further possibly reformulating it into a more precise narrower query. It is also inherent to the task of browsing through a collection of resources instead of searching among them. Manually separating the interesting items from the uninteresting ones is a tedious and time consuming job. For this reason, various services have been set up to manage the information overload. They may be categorized as variations of two complementary approaches: (a) predefined category structures; (b) fully automatic search engines.

Predefined category structures are maintained by humans, who organize resources into a hierarchy, grouping appropriate ones together under a meaningful

category description. This is meant to support user browsing activity. The main drawback of such an approach is that the categories are frequently obsolete, since they are not updated so quickly as new documents and new topics appear. In turn, the mechanisms used by automatic search engines, enable them to stay relatively up-to-date, what makes the search engines the premier choice for most Web users. However, when the number of the results is huge, even despite their ranked order, manually investigating them is a big effort without any additional navigation tools. In order to facilitate browsing and managing results of a Web search/query, methods for grouping answers on the ground of user defined criteria would be exploited.

In this paper, we propose to marry the benefits of two complementary approaches for handling information overload by offering a method that, given a query in the form of a conjunctive query [1,2,3,4,5], produces a dynamic categorization over ranked query results. The key feature of the method is in the exploitation of the *semantics* of knowledge bases of reference (in the form of ontologies) for grouping results with respect to a user defined criteria. Specifically, given a certain grouping criterion, expressed as a (complex) concept from a knowledge base of reference, results are grouped in agreement with (part of) the subsumption hierarchy deductively obtained by considering the specified concept and the given ontology.

Currently, besides of the usual syntactic based approach (grounded on the use of keywords) adopted by the main Web search engines such as Google or Yahoo, the other approaches that are generally exploited for grouping structured query answers are based on the semantics drawn from SQL, the standard language for querying relational databases. An example is given by the implementation of the aggregation operators (COUNT, MIN, MAX, AVG, SUM, GROUP BY) for SPARQL in Virtuoso¹. However, besides of the fact that these implementations are not part of the SPARQL standard, these latter approaches (differently from the method that we propose) do not assume to perform any reasoning involving background knowledge during the query results aggregation, and as such they are purely syntactic as well. The same happens for grouping/aggregating features that are to be included in the ongoing SPARQL 2 standard.

The main contributions of the paper are summarized as follows: (a) we introduce a novel way for grouping query answers (*semantic grouping*) where grouping is done on the ground of a knowledge base of reference; (b) we propose a technique for aggregating query results consisting on a dynamic generation of a navigable hierarchy on top of the retrieved results, that is based on their semantics. Such a hierarchy constitutes a multi-valued classification of the results that may be seen as a novel approach for generating a *dynamic faceted classification* over retrieved results, enabling further faceted search/browsing over them; (c) we present the experimental results of the application of our proposed method to ontologies expressed in *Web Ontology Language (OWL)*², whose semantics is based on *description logic (DL)* [6].

¹ <http://virtuoso.openlinksw.com/>

² <http://www.w3.org/TR/owl-features/>

The rest of the paper is organized as follows. In Sect. 2, a motivating example for our work is presented. In Sect. 3, we introduce the basics of the knowledge representation formalism of choice: description logics and the notion of conjunctive queries over DL knowledge bases. In Sect. 4, we present our proposed approach. Sect. 5 reports on the experimental evaluation of our method while Sect. 6 discusses the work related to ours. In Sect. 7, conclusions are drawn.

2 Motivating Example

Let us discuss the need for and advantage of semantic query answer aggregation on the ground of a motivating example.

Example 1 (Motivating scenario). Maria and Sebastian are searching for a weekend break offer. In order to get faster insight into the results retrieved by their query, they would like to have the answers grouped by a destination criterion. After getting first insight into the results, they perform further exploration, submitting another, most specific query for destinations located in mountains, and for budget accommodations offered in these destinations with a request to group the results w.r.t. destinations, and additionally w.r.t. accommodations offered.

Let us suppose that there is a Web service that exploits a SPARQL endpoint that retrieves the results from Semantic Web datasets. Furthermore the Web service uses an ontology as background knowledge on the given domain (an instance is presented in Example 3). In the following, instances of SPARQL queries for the user needs illustrated in Example 1 are shown.

Example 2 (Example queries in SPARQL syntax).

$Q_{1_group_by}$ = SELECT ?x ?y WHERE { ?x rdf:type :WeekendBreakOffer .
?x :hasDestination ?y } GROUP BY ?y

$Q_{2_group_by}$ = SELECT ?x ?y ?z WHERE { ?x rdf:type :WeekendBreakOffer .
?x :hasDestination ?y . ?y :locatedIn ?v . ?v rdf:type :Mountains .
?x :hasAccommodation ?z . ?z rdf:type :BudgetAccommodation } GROUP BY ?y ?z

Let us assume an application of the classical semantics of the GROUP BY clause known from SQL in the evaluation of the query Q_1 . Let us assume further that instances that bind to variable y are town names. In such a case, the results will be partitioned so that one row for each town name is created. Considering that there may be possibly many towns satisfying the query conditions, there will be also too many groups to provide significant added value.

Note also that classical semantics of GROUP BY clause, which is to partition the results by identical values, disregards the presence of any background knowledge, even if there is available some. Consider for example that the towns are annotated by terms from the ontology such as City, EuropeanDestination, ItalianDestination, PolishDestination. Considering such annotations as a grouping condition, it would be possible to aggregate the results into a smaller number of 'semantic' groups. The classical semantics of GROUP BY disregards also semantic relationships between groups like subclass-superclass relation (e.g. ItalianDestination, and EuropeanDestination). Exploiting the background ontology would enable making such relations explicit.

Another problem with the syntactic approach may happen when all the results in the grouping condition, e.g. all destinations in our case, refer to the same town. This gives as grouping result a unique row which synthesizes all instances, thus giving almost null information.

Summarizing, a merely syntactic approach, as the one used by the **GROUP BY** clause in the database context, could not be of great help when many destinations are found, or when the results refer to the same destination. To manage situations like these, a *semantic group by* could be adopted, namely a method that is able to group query results on the ground of a knowledge base of reference.

3 Preliminaries

3.1 Representation and Inference

Description logics (DLs) [6] are a family of knowledge representation languages (endowed with a model-theoretic semantics and reasoning services) that have been adopted as theoretical foundation for OWL language. Basic elements in DLs are: *atomic concepts* (denoted by A), and *atomic roles* (denoted by R , S). *Complex descriptions* (denoted by C and D) are inductively built by using concept and role *constructors*.

Semantics is defined by *interpretations* $\mathcal{I}=(\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$, where non-empty set $\Delta^{\mathcal{I}}$ is the domain of the interpretation and $\cdot^{\mathcal{I}}$ is an interpretation function which assigns to every atomic concept A a set $A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$, and to every atomic role R a binary relation $R^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$. The interpretation function is extended to complex concept descriptions by the inductive definition as presented in Tab. 1. A DL *knowledge base*, KB , is formally defined as: $\mathcal{KB} = (\mathcal{T}, \mathcal{A})$, where $\mathcal{T}$ is called a TBox, and it contains axioms dealing with how concepts and roles are related to each other, such as those, used further in the paper, defining role domain, and role range, written $\text{Domain}(R, C)$ (where $(a, b) \in R^{\mathcal{I}}$ implies $a \in C^{\mathcal{I}}$), $\text{Range}(R, C)$ (where $(a, b) \in R^{\mathcal{I}}$ implies $b \in C^{\mathcal{I}}$), and where $\mathcal{A}$ is called an ABox, and it contains assertions about individuals such as $C(a)$ (that means the individual a is an instance of the concept C) and $R(a, b)$ (a is R -related to b).

As regards the inference services, we recall the definition of *instance-checking* and *subsumption* that are further used in the paper. *Instance-checking* amounts

Table 1. Syntax and semantics of example *DL* constructors

Constructor	Syntax	Semantics
Universal concept	$\top$	$\Delta^{\mathcal{I}}$
Bottom concept	$\perp$	$\emptyset$
Negation of arbitrary concepts	$(\neg C)$	$\Delta^{\mathcal{I}} \setminus C^{\mathcal{I}}$
Intersection	$(C \sqcap D)$	$C^{\mathcal{I}} \cap D^{\mathcal{I}}$
Union	$(C \sqcup D)$	$C^{\mathcal{I}} \cup D^{\mathcal{I}}$
Value restriction	$(\forall R.C)$	$\{a \in \Delta^{\mathcal{I}} \mid \forall b. (a, b) \in R^{\mathcal{I}} \rightarrow b \in C^{\mathcal{I}}\}$
Full existential quantification	$(\exists R.C)$	$\{a \in \Delta^{\mathcal{I}} \mid \exists b. (a, b) \in R^{\mathcal{I}} \wedge b \in C^{\mathcal{I}}\}$

to determine whether an individual, say a , belongs to a concept extension, i.e. whether $C(a)$ holds for a certain concept C . As regards *subsumption*, given two concept descriptions C and D in a TBox $\mathcal{T}$, C subsumes D (denoted by $D \sqsubseteq C$) if and only if, for every interpretation $\mathcal{I}$ of $\mathcal{T}$ it holds that $D^{\mathcal{I}} \subseteq C^{\mathcal{I}}$. C equivalent to D (denoted by $C \equiv D$) amounts to $C \sqsubseteq D$ and $D \sqsubseteq C$.

3.2 Conjunctive Queries

Queries admitted in this work are expressed in the form of *conjunctive queries* [1,2,3,4,5]. Let N_C, N_R, N_I be the sets of *concept names*, *role names* and *individual names* respectively and let N_V be a countably infinite set of variables disjoint from N_C, N_R , and N_I . Let by $\mathbf{x}$ and $\mathbf{y}$ denote the sets of distinguished and nondistinguished variables, respectively, where $\mathbf{x}, \mathbf{y} \subseteq N_V$. A conjunctive query, denoted with $Q(\mathbf{x}, \mathbf{y})$, is the finite conjunction of a non-empty set of *atoms*. An *atom* is an expression of kind $A(t_1)$ (concept atom) or $R(t_1, t_2)$ (role atom), where A is a concept name, R is a role name, and t_1 and t_2 are individuals from N_I or variables from $\mathbf{x}$ or $\mathbf{y}$.

An answer to a query $Q(\mathbf{x}, \mathbf{y})$ w.r.t. KB is an assignment θ of individuals to distinguished variables such that $KB \models \exists \mathbf{y} : Q(\mathbf{x}\theta, \mathbf{y})$.

Example 3 (Conj. Query). Given the following knowledge base:

$\mathcal{T} = \{ \text{City} \sqsubseteq \text{Destination}, \text{EuropeanDestination} \sqsubseteq \text{Destination}, \text{ItalianDestination} \sqsubseteq \text{EuropeanDestination}, \text{PolishDestination} \sqsubseteq \text{EuropeanDestination}, \text{Hotel} \sqsubseteq \text{Accommodation}, \text{BudgetAccommodation} \sqsubseteq \text{Accommodation}, \text{B\&B} \sqsubseteq \text{BudgetAccommodation}, \text{Hostel} \sqsubseteq \text{BudgetAccommodation}, \text{SkiingSite} \sqsubseteq \text{Site}, \text{SightSeeingSite} \sqsubseteq \text{Site}, \top \sqsubseteq \forall \text{hasSite.Site}, \top \sqsubseteq \forall \text{hasDestination.Destination}, \top \sqsubseteq \forall \text{hasAccommodation.Accommodation} \}.$

$\mathcal{A} = \{ \text{locatedIn}(\text{ZAKOPANE}, \text{TATRA}), \text{hasSite}(\text{ZAKOPANE}, \text{SKI_LIFTS_NOSAL}), \text{SkiingSite}(\text{SKI_LIFTS_NOSAL}), \text{locatedIn}(\text{CHOCHOLOW}, \text{TATRA}), \text{hasSite}(\text{CHOCHOLOW}, \text{HIGHLANDERS_WOODEN_HOUSES}), \text{SightSeeingSite}(\text{HIGHLANDERS_WOODEN_HOUSES}), \text{Mountains}(\text{TATRA}), \text{locatedIn}(\text{TRENTO}, \text{ALPS}), \text{Mountains}(\text{ALPS}), \text{WeekendBreakOffer}(\text{O1}), \text{hasAccommodation}(\text{O1}, \text{A1}), \text{B\&B}(\text{A1}), \text{hasDestination}(\text{O1}, \text{ZAKOPANE}), \text{PolishDestination}(\text{ZAKOPANE}), \text{WeekendBreakOffer}(\text{O2}), \text{hasAccommodation}(\text{O2}, \text{A2}), \text{B\&B}(\text{A2}), \text{hasDestination}(\text{O2}, \text{CHOCHOLOW}), \text{PolishDestination}(\text{CHOCHOLOW}), \text{WeekendBreakOffer}(\text{O3}), \text{hasAccommodation}(\text{O3}, \text{A3}), \text{Hostel}(\text{A3}), \text{hasDestination}(\text{O3}, \text{TRENTO}), \text{ItalianDestination}(\text{TRENTO}), \text{City}(\text{TRENTO}) \}.$

we formalize conjunctive queries corresponding to the ones from Example 2 as:

$$Q_1(x, y) = \text{WeekendBreakOffer}(x) \wedge \text{hasDestination}(x, y)$$

$$Q_2(x, y, z) = \text{WeekendBreakOffer}(x) \wedge \text{hasDestination}(x, y) \wedge \text{locatedIn}(y, v) \wedge \text{Mountains}(v) \wedge \text{hasAccommodation}(x, z) \wedge \text{BudgetAccommodation}(z)$$

4 Deductive Aggregation: Semantic Grouping

Let us consider the knowledge base and the example queries from Example 3. Supposing the case in which many offers are available, the user would be interested in grouping the results with respect to the different destinations. However, as already discussed in the previous section, a merely syntactic approach, as the one used by the *group by* clause in the database context, could not be of great help in cases in which several destinations are found or all results refer to the same destination. In order to manage cases like this, a *semantic group by* could be adopted, namely a method that is able to group query results on the ground of a knowledge base of reference. Specifically, looking at the knowledge base in the example above, results could be grouped on the ground of the pertaining country of a destination (*ItalianDestination*, *PolishDestination*).

In this section we present a method that is able to perform the *semantic group by* on the ground of concepts that are more specific (in the knowledge base of reference) than the concept adopted by the user for grouping the results.

4.1 Basics of the Semantic Aggregation

The general idea behind *semantic group by* is to categorize the results with regard to concept hierarchies inferred for each variable appearing in a grouping condition. Thus, in order to formalize our proposition of a *semantic group by*, we introduce a special second order predicate *categorize_by*.

Definition 1 (categorize_by). *A conjunctive query with a semantic aggregate subgoal is of the form*

$$\text{categorize_by}([X_1, X_2, \dots, X_m]) : Q(\mathbf{x}, \mathbf{y})$$

where $[X_1, X_2, \dots, X_m]$ is a grouping list of variables appearing in $\mathbf{x}$.

In Example 4 we present the conjunctive queries with the *categorize_by* clause for the scenario discussed throughout this paper.

Example 4 (Example queries with the categorize_by clause).

$$Q_{1_categorize_by}(x, y) = \text{categorize_by}(y) : \text{WeekendBreakOffer}(x) \wedge \text{hasDestination}(x, y)$$

$$Q_{2_categorize_by}(x, y, z) = \text{categorize_by}(y, z) : \text{WeekendBreakOffer}(x) \wedge \text{hasDestination}(x, y) \wedge \text{locatedIn}(y, v) \wedge \text{Mountains}(v) \wedge \text{hasAccommodation}(y, z) \wedge \text{BudgetAccommodation}(z)$$

Further, we formally define the notion of *semantic category*.

Definition 2 (Semantic category). *Given is query*

$$Q = \text{categorize_by}([X_1, X_2, \dots, X_m]) : Q(\mathbf{x}, \mathbf{y})$$

A semantic category is a tuple of concepts $\langle C_1, C_2, \dots, C_m \rangle$, where each C_i corresponds to X_i in the grouping variables list of Q .

Semantic categories form a multi-valued classification $\mathcal{H}$, induced by a subsumption relation between concepts appearing in the same place in tuples. Then, the operational semantics for `categorize_by` clause is to first create $\mathcal{H}$, a partially ordered set of semantic categories, based on the inferred semantic types of grouping variables, and then to assign the tuples to the most specific semantic categories.

4.2 Inferring the Type of the Variables

Let v_i be a query variable used in the grouping clause of a query. For each such variable a (complex) concept C^i is derived as follows. Let B_C be a set of the concepts explicitly mentioned in the query atoms of the form $C(v_i)$, $C \in B_C$. Let B_D , and B_{Rn} be sets of the concepts inferred by the role atoms $R(v_i, \cdot)$ or $R(\cdot, v_i)$, by taking into account the domain and range of role R , respectively, $\text{Domain}(R, C_D)$, $C_D \in B_D$, $\text{Range}(R, C_R)$, $C_R \in B_{Rn}$. Now, let $\mathcal{B}_i := B_C \cup B_D \cup B_{Rn}$. The concept C^i is an intersection:

$$C^i := \bigcap_{C_p^i \in \mathcal{B}_i} C_p^i$$

A concept C^i corresponding to query variable v_i , included in the **CATEGORIZE BY** clause, is derived. Then it is classified in the subsumption hierarchy of the concepts in the knowledge base. A sub-hierarchy of concepts rooted at these concepts will be used in the next step, which determines the final multivalued classification of semantic categories used for grouping the query answers.

Example 5 (Determining the type of variable). Consider query $Q_{2_categorize_by}(x, y, z)$ from Example 4, and variable z appearing in the grouping variables list. Then the concept inferred for the variable z is:

$$C^z := \text{Accommodation} \sqcap \text{BudgetAccommodation}$$

4.3 Constructing a Tree of Semantic Categories

The primary motivation for this work is to enable better results exploration for the user by reducing information overload by means of presenting him/her the groupings of the results. From this perspective we are rather interested in some compromise between generating the complete multi-valued classification of semantic categories versus meeting the demands of a real time computation of the results, and improving the user experience, not overloading him/her with too many groupings (the initial problem). Therefore, in this work, we restrict our analysis to the case of DL concept subsumption hierarchies, and semantic category hierarchies in the form of trees. The extension of the proposed ideas to graphs is possible, although not yet considered in this paper.

Concept Hierarchy for a Single Variable. A basic approach for generating a sub-hierarchy of concepts to be rooted at the concept C^i , corresponding to a grouping variable v_i , is to reproduce a part of a classified subsumption hierarchy produced by a DL reasoner.

In some cases, however, this may not be enough to obtain a meaningful hierarchy of groups. Consider for example the situation, where the concept C^i is a leaf in the classified subsumption hierarchy, e.g. `PolishDestination` in the knowledge base from Example 3. Another case where this basic approach would not be of help is when all the retrieved results for variable v_i , e.g. all destinations from the discussed example, fall under the same type, e.g. again `PolishDestination`.

In such cases more advanced approach could be adopted, that is based on an iterative application of a concept refinement operator, for example adapted from such concept specialization operators like the ones proposed in [7,8] to the case of a purely deductive approach. The application of the refinement operator would introduce further, unnamed, complex concepts into the sub-hierarchy rooted at C^i . For example, the concept `PolishDestination` could be then specialized into concepts: `PolishDestination  $\sqcap$   $\exists$ hasSite.SkiingSite` and `PolishDestination  $\sqcap$   $\exists$ hasSite.SightSeeingSite` to differentiate two Polish destinations from the knowledge base, `ZAKOPANE`, and `CHOCHOLOW`, based on the sites they offer.

Multiple Variables: Tree Product. Let us define the subtrees trees by the quadruple (N, E, r, ℓ) , where N is the set of nodes, E is the set of edges between nodes (n_a, n_b) , $r \in N$ stands for the root node and ℓ is the labeling function assigning each node with the corresponding concept.

Now, let us consider the case of two variables, say v_i and v_j in the `CATEGORIZE BY` clause and their related subtrees, $T_i = (N_i, E_i, r_i, \ell_i)$ and $T_j = (N_j, E_j, r_j, \ell_j)$, in the subsumption hierarchy. Multiple variables will be processed by iterating the binary product operation presented in the following. A formal sketch of the algorithm is shown in Figure 1. The product $T_i \times T_j := (N_{ij}, E_{ij}, r_{ij}, \ell_{ij})$ of the two subtrees is computed recursively. The root of the product tree is made up by the node $r_{ij} = \langle r_i, r_j \rangle \in N_{ij}$ labeled with $\ell_{ij}(r_{ij}) = \langle \ell(r_i), \ell(r_j) \rangle$. For each couple of children nodes c_i of r_i and c_j of r_j , a new child $r' = \langle c_i, c_j \rangle$ of r_{ij} is obtained. Hence r' is the root of the subtree that constitutes the product T' of the subtrees T'_i and T'_j , rooted respectively, at c_i and c_j . The product subtree T' is finally connected to the product tree under construction T_{ij} .

An important feature of the product is that its complexity is polynomial (note that if the trees are represented with matrices then the proposed product corresponds to a matrix product).

Figure 2 illustrates the general idea of generation of the tree product, and Figure 3 presents the tree product for our motivating scenario.

Alternative Tree Products. In some cases the level-wise computation of product trees would yield poorly structured trees, when one of them is rooted with a primitive concept, for example. In such cases it may be advisable to use alternative definition of the tree product operator (derived from graph product operators), which is able to retain all of the structural information contained in the subsumption hierarchies.


```

 $\times(T_i, T_j) : (N_{ij}, E_{ij}, r_{ij}, \ell_{ij})$ 
begin
 $E_{ij} \leftarrow \emptyset$ 
 $r_{ij} \leftarrow \langle r_i, r_j \rangle$ 
 $N_{ij} \leftarrow \{r_{ij}\}$ 
 $\ell_{ij}(r_{ij}) \leftarrow \langle \ell(r_i), \ell(r_j) \rangle$ 
for each  $(c_i, c_j)$  where  $(r_i, c_i) \in E_i$  and  $(r_j, c_j) \in E_j$  do
     $T'_i \leftarrow \text{SUBTREE}(T_i, c_i)$ 
     $T'_j \leftarrow \text{SUBTREE}(T_j, c_j)$ 
     $(N', E', r', \ell') \leftarrow \times(T'_i, T'_j)$ 
     $N_{ij} \leftarrow N_{ij} \cup N'$ 
    for each node  $\in N'$  do
         $\ell(\text{node}) \leftarrow \ell'(\text{node})$ 
     $E_{ij} \leftarrow E_{ij} \cup E' \cup \{\langle r_{ij}, r' \rangle\}$ 
return  $(N_{ij}, E_{ij}, r_{ij}, \ell_{ij})$ 
end

```

Fig. 1. Tree product operator

Definition 3 (alternative tree products). *Given two trees $T_1 = (V_1, E_1, r_1, \ell_1)$ and $T_2 = (V_2, E_2, r_2, \ell_2)$, their product, denoted $T_1 \times T_2$, is a tree $T_p = (V_p, E_p, r_p, \ell_p)$ defined as follows:*

- *the vertex set is a subset of the Cartesian product $V_p \subseteq V_1 \times V_2$;*
- $r_p = (r_1, r_2) \in V_p$;
- $\ell_p(u, v) = (\ell_1(u), \ell_2(v))$;
- *two product vertices $(u_1, u_2) \in V_p$ and $(v_1, v_2) \in V_p$ are connected in T_p , i.e. $((u_1, u_2), (v_1, v_2)) \in E_p$, iff u_1, u_2, v_1, v_2 satisfy the conditions specified in the following table:*

Rooted product	$\bigcup_{u \in V_1} \{((u, v), (u, v')) \mid (v, v') \in E_2\}$
Co-normal product	$(u_1, v_1) \in E_1 \vee (u_2, v_2) \in E_2$
Tensor product	$(u_1, v_1) \in E_1 \wedge (u_2, v_2) \in E_2$

Note, that for the last two operators it is not guaranteed that a tree is obtained.

4.4 From Trees to Groups

Having produced a tree of semantic categories, the next step is to assign results (tuples) to the semantic categories. We propose to perform it incrementally in line with the ranking of the results provided by the query answering engine. For each tuple the projection of the bindings of variables appearing in the CATEGORIZE BY clause is taken into account. Then starting from the root of a tree and following paths to its leaves the membership check of instances to concepts in the tree nodes are performed in order to find the most specific nodes under which to place the result tuple. Such technique allows to obtain more populated hierarchy in a longer time frame, where completeness increases over time. A sketch of the algorithm is shown in Figure 4.

Fig. 2. Tree product example: the lower tree is the product of the two upper trees

Fig. 3. Tree product illustration for query $Q_{2_categorize_by}$ from Example 4

The time complexity of the proposed population algorithm depends on a complexity of instance membership checks. The population algorithm itself has the complexity $O(nkb^m)$, where n is the number of answers, k the number of grouping variables, b the branching factor, and m is the maximum depth of the tree. Besides this simple algorithm, also other solutions could be adopted, e.g. such in which a non-standard inference of computing *most specific concept* [9] would

```

POPULATE( $T$ ,  $KB$ ,  $\mathcal{H}$ )
input:
 $T = (a_{ij})_{i=1,r}^{j=1,n}$ : query answer table
 $KB$ : knowledge base
 $\mathcal{H}$ : semantic category hierarchy
 $Q$ : query
output:
 $\mathcal{H}$ : populated semantic category hierarchy

begin
   $i \leftarrow 0$ 
  while !RESOURCES_UP() and  $i < r$ 
     $a_i \leftarrow T.GETANSWER(i)$ ;
    ASSIGNToMOSTSPECIFICSEMANTICCATEGORIES( $a_i$ ,  $\mathcal{H}$ ,  $\mathcal{H}.ROOT()$ ,  $KB$ ,  $Q$ );
return  $\mathcal{H}$ 
end

ASSIGNToMOSTSPECIFICSEMANTICCATEGORIES( $a_i$ ,  $\mathcal{H}$ ,  $\mathcal{H}.NODE()$ ,  $KB$ ,  $Q$ )
begin
   $((C_0), \dots, (C_n)) \leftarrow \mathcal{H}.NODE().GETSEMANTICCATEGORY()$ ;
  assigned  $\leftarrow$  false;
  followsUnderNode  $\leftarrow$  true;
   $j \leftarrow 0$ ;
  while followsUnderNode and  $j < n$ 
     $a_{ij} \leftarrow a_i.GETBINDING(j)$ ;
    if  $\neg(KB \models C_j(a_{ij}))$  then
      followsUnderNode  $\leftarrow$  false;
     $j \leftarrow j + 1$ ;
  if followsUnderNode then
    for each  $\mathcal{H}.NODE().CHILD()$ 
      if ASSIGNToMOSTSPECIFICSEMANTICCATEGORIES( $a_i$ ,  $\mathcal{H}$ ,  $\mathcal{H}.NODE().CHILD()$ ,  $KB$ ,  $Q$ ) then
        assigned  $\leftarrow$  true;
      if assigned=false then
         $\mathcal{H}.NODE().ASSIGN(a_i)$ ;
        assigned  $\leftarrow$  true;
return assigned;
end

```

Fig. 4. Population of a semantic category hierarchy

be employed. It should be noted, however, that for some description logics the most specific concept may not exist.

Example 6 (Assignment of the results to semantic categories). The results of query $Q_{2_categorize_by}$ from Example 4 are assigned to the semantic categories shown in Figure 3 as follows: (O1, ZAKOPANE, A1): $\langle \text{EuropeanDestination}, \text{B\&B} \rangle$, (O2, CHOCHOLOW, A2): $\langle \text{EuropeanDestination}, \text{B\&B} \rangle$, (O3, TRENTO, A3): $\langle \text{City}, \text{Hostel} \rangle$, (O3, TRENTO, A3): $\langle \text{EuropeanDestination}, \text{Hostel} \rangle$.

5 Experimental Evaluation

We have developed a proof-of-concept implementation of our approach in Java, on top of Pellet reasoner³. Please note, that state-of-the-art reasoners (including Pellet) do not support conjunctive queries with truly undistinguished variables. Therefore, our prototype implementation supports only *DL-safe queries* [10], i.e. queries, where all variables are assumed to be bound to named individuals.

³ <http://clarkparsia.com/pellet/>

Table 2. Characteristics of test datasets

dataset	<i>DL</i>	#concepts	#obj. roles	#individuals
FINANCIAL	<i>ALCTF</i>	60	16	17941
NTN	<i>SHIF</i>	49	29	728
VICODI	<i>ALHI</i>	194	10	16942
LUBM	<i>SHI(D)</i>	43	25	17174

Table 3. Tested queries

dataset	query
FINANCIAL	$Q_1(x, y) = \text{categorize_by}(x, y) : \text{Client}(x) \wedge \text{hasCreditCard}(x, y) \wedge \text{CreditCard}(y)$
FINANCIAL	$Q_2(x, y, z) = \text{categorize_by}(x, y, z) : \text{Client}(x) \wedge \text{livesIn}(x, y) \wedge \text{hasAgeValue}(x, z)$
NTN	$Q_3(x, y) = \text{categorize_by}(x, y) : \text{Agent}(x) \wedge \text{relativeOf}(x, y)$
NTN	$Q_4(x, y, z) = \text{categorize_by}(x, y, z) : \text{CognitiveAgent}(x) \wedge \text{collaboratesWith}(x, y) \wedge \text{knows}(x, z)$
VICODI	$Q_5(x, y) = \text{categorize_by}(x, y) : \text{Scientist}(x) \wedge \text{hasRole}(y, x)$
VICODI	$Q_6(x, y, z) = \text{categorize_by}(x, y, z) : \text{Leader}(x) \wedge \text{hasRole}(y, x) \wedge \text{related}(x, z) \wedge \text{Individual}(y) \wedge \text{Location}(z)$
LUBM	$Q_7(x, y) = \text{categorize_by}(x, y) : \text{Faculty}(x) \wedge \text{teacherOf}(x, y)$
LUBM	$Q_8(x, y, z) = \text{categorize_by}(x, y, z) : \text{Person}(x) \wedge \text{memberOf}(x, y) \wedge \text{subOrganizationOf}(x, z)$

We have empirically tested our approach using four benchmark datasets: FINANCIAL⁴, NEW TESTAMENT-NAMES (NTN)⁵, VICODI⁶, and LUBM⁷ (see: Table 2). The experiments were conducted on a laptop computer with a 1.67 GHz Intel processor, 1014 MB of RAM, running Windows Vista.

For each of the datasets we tested 2 queries (see: Table 3). Let us take query Q_1 as an example. Variable x in Q_1 bounds to the names of clients, and variable y to the IDs of credit cards, so the classical, *syntactic group by* would create 892 groups. In turn, our proposed *semantic grouping*, creates the six following semantic categories: $\langle \text{Client}, \text{CreditCard} \rangle$, $\langle \text{Man}, \text{Classic} \rangle$, $\langle \text{Man}, \text{Gold} \rangle$, $\langle \text{Man}, \text{Junior} \rangle$, $\langle \text{Woman}, \text{Classic} \rangle$, $\langle \text{Woman}, \text{Gold} \rangle$, $\langle \text{Woman}, \text{Junior} \rangle$.

For each of the queries we tested how our method scales with the growing number of results for each query. Figure 5 presents the total times of computing, and populating a semantic category hierarchy.

Note that the Web users usually browse not more than a few dozens of the retrieved results. For this number of the results, the basic implementation of our proposed approach performs well, in real-time, on the tested cases, and scales linearly with the number of the processed results.

Note further, that not all the nodes of the generated semantic category hierarchy would become ultimately populated. This observation may lead to an optimization of the proposed algorithm for the cases where some precomputation would be possible to identify concepts without an extension. These concepts could be then stored on an index, and the nodes containing them removed from the semantic category hierarchy in a preprocessing step, thus reducing the further complexity of the hierarchy population. We tested this technique experimentally reaching a significant decrease in execution time (for the hardest query, Q_8 , and

⁴ FINANCIAL, <http://www.cs.put.poznan.pl/alawrynowicz/financial.owl>

⁵ NTN, http://protegewiki.stanford.edu/index.php/Protege_Ontology_Library

⁶ VICODI, http://kaon2.semanticweb.org/download/test_ontologies.zip

⁷ LUBM, <http://swat.cse.lehigh.edu/projects/lubm/>

(a) FINANCIAL

(b) NTN

(c) VICODI

(d) LUBM

Fig. 5. Results of the experiments

for 500 results, the execution time was shorter more than 5 times than this reported in the figure for the basic approach).

6 Related Work

To the best of our knowledge, ours is the first proposal for grouping conjunctive query results on the ground of semantics of the underlying knowledge base.

While aggregate queries were extensively studied for relational databases, there are very few results for aggregate queries over ontologies, especially those that target to meet the peculiarities of the KR formalisms of the Semantic Web. In [11] the syntax and semantics for epistemic aggregate queries over ontologies were proposed, and query answering for typical aggregate functions studied for an ontology language $DL-Lite_A$. That work was motivated by the non-adequacy of certain answer semantics for conditional aggregate queries over ontologies due to the open world assumption. In [12] the meaning and implementation of grouping and aggregate queries over RDF graphs were studied, motivated by the drawback of the previous works where the grouping and aggregate operations had not take the graph structure of the base data into account.

However, none of the above approaches exploited the peculiar feature of the Semantic Web datasets, that is possible availability of the background ontologies expressing semantics of the data.

Some other works may also be considered as relevant to ours. In [13] an approach to automatically categorize the results of SQL queries has been proposed that consisted on a dynamic generation of a labeled, hierarchical category structure. Since this was proposed for relational database model, the constructed categories could only be built based on the values in the retrieved tuples, and not on any extra semantic information, due to the lack of background knowledge. The proposal for clustering the Semantic Web query results [14,15] may be considered as related work. However, an essential difference between clustering query results, and our proposed semantic grouping, is that the former one involves *induction*, while the latter one is based simply on *deductive* reasoning.

7 Conclusions and Future Work

The paper proposes a method for the aggregation of conjunctive query results on the ground of an ontology of reference. We have defined a new type of *categorize by* queries for this purpose. To enable such queries, we have designed a method where in a deductive modality, answers are grouped taking into account dynamically generated subsumption hierarchy induced by the underlying knowledge base. This subsumption hierarchy has the form of a hierarchical, multi-valued classification that may support a faceted search of the results. The research presented in the paper may constitute a proposal for the next generation query languages oriented to knowledge bases annotated with ontology languages.

To the best of our knowledge, this work is the first to propose grouping of conjunctive query results on the ground of semantics of the underlying knowledge base. As such it may be an important first step towards the proposed direction of semantic grouping, and inspiration to conduct future work on the proposed idea. In particular, in the future, alternative versions of the components of the proposed method may be investigated such as different tree/graph product operators, or different techniques of assigning answers to semantic groups. An interesting future research may concern an investigation on a possible integration of semantic grouping into query processing, and query execution plans.

References

1. Calvanese, D., De Giacomo, G., Lenzerini, M.: On the decidability of query containment under constraints. In: Proc. of the 17th ACM SIGACT SIGMOD SIGART Symposium on Principles of Database Systems (PODS'98), pp. 149–158 (1998)
2. Horrocks, I., Tessaris, S.: Querying the semantic web: A formal approach. In: Horrocks, I., Hendler, J. (eds.) ISWC 2002. LNCS, vol. 2342, pp. 177–191. Springer, Heidelberg (2002)

3. Calvanese, D., De Giacomo, G., Lembo, D., Lenzerini, M., Rosati, R.: Data complexity of query answering in description logics. In: Doherty, P., Mylopoulos, J., Welty, C. (eds.) *Proceedings of the Tenth International Conference on Principles of Knowledge Representation and Reasoning (KR'06)*. AAAI Press, Menlo Park (2006)
4. Ortiz, M., Calvanese, D., Eiter, T.: Data complexity of answering unions of conjunctive queries in *SHIQ*. In: Parsia, B., Sattler, U., Toman, D. (eds.) *Proceedings of the International Workshop on Description Logics (DL'06)*. CEUR-WS, vol. 189. CEUR (2006)
5. Glimm, B., Horrocks, I., Lutz, C., Sattler, U.: Conjunctive query answering for the description logic *SHIQ*. *J. Artif. Int. Res.* 31(1), 157–204 (2008)
6. Baader, F., Calvanese, D., McGuinness, D., Nardi, D., Patel-Schneider, P. (eds.): *The Description Logic Handbook*. Cambridge University Press, Cambridge (2003)
7. Iannone, L., Palmisano, I., Fanizzi, N.: An algorithm based on counterfactuals for concept learning in the Semantic Web. *Appl. Intell.* 26(2), 139–159 (2007)
8. Lehmann, J., Hitzler, P.: A refinement operator based learning algorithm for the *ALC* description logic. In: Blockeel, H., Ramon, J., Shavlik, J., Tadepalli, P., et al. (eds.) *ILP 2007. LNCS (LNAI)*, vol. 4894, pp. 147–160. Springer, Heidelberg (2008)
9. Baader, F., Küsters, R.: Computing the least common subsumer and the most specific concept in the presence of cyclic *ALN*-concept descriptions. In: Herzog, O., Günter, A. (eds.) *KI 1998. LNCS*, vol. 1504, pp. 129–140. Springer, Heidelberg (1998)
10. Motik, B., Sattler, U., Studer, R.: Query answering for OWL-DL with rules. *Journal of Web Semantics: Science, Services and Agents on the World Wide Web* 3(1), 41–60 (2005)
11. Calvanese, D., Kharlamov, E., Nutt, W., Thorne, C.: Aggregate queries over ontologies. In: *ONISW '08: Proc. of the 2nd international workshop on Ontologies and Information systems for the Semantic Web*, pp. 97–104. ACM, New York (2008)
12. Seid, D., Mehrotra, S.: Grouping and aggregate queries over Semantic Web databases. In: *International Conference on Semantic Computing*, pp. 775–782. IEEE Computer Society, Los Alamitos (2007)
13. Chakrabarti, K., Chaudhuri, S., Hwang, S.w.: Automatic categorization of query results. In: *SIGMOD '04: Proc. of the 2004 ACM SIGMOD international conference on Management of data*, pp. 755–766. ACM, New York (2004)
14. Lawrynowicz, A.: Grouping results of queries to ontological knowledge bases by conceptual clustering. In: Nguyen, N.T., Kowalczyk, R., Chen, S.-M. (eds.) *ICCCI 2009. LNCS*, vol. 5796, pp. 504–515. Springer, Heidelberg (2009)
15. Lawrynowicz, A.: Query results clustering by extending SPARQL with CLUSTER BY. In: Meersman, R., Herrero, P., Dillon, T.S. (eds.) *OTM 2009 Workshops. LNCS*, vol. 5872, pp. 826–835. Springer, Heidelberg (2009)