

# Handling Qualitative Preferences in SPARQL over Virtual Ontology-Based Data Access

Marlene Goncalves<sup>a,\*</sup>, David Chaves-Fraga<sup>b</sup>, and Oscar Corcho<sup>b</sup>

<sup>a</sup> *Computer Science and Information Technology, Universidad Simón Bolívar, Venezuela*

*E-mail: mgoncalves@usb.ve*

<sup>b</sup> *Ontology Engineering Group, Universidad Politécnica de Madrid, Spain*

*E-mails: dchaves@fi.upm.es, ocorcho@fi.upm.es*

**Abstract.** With the increase of data volume in heterogeneous datasets that are being published following Open Data initiatives, new operators are necessary in order to help users to find the subset of data that best satisfies their preference criteria. Quantitative approaches such as top-k queries cannot be the most appropriate approaches due they require the user to assign weights that may not be known beforehand to a scoring function. Unlike the quantitative approach, under the qualitative approach, which includes the well-known skyline, preference criteria are expressed in a way closer to natural language. In this paper, we address the problem of evaluating SPARQL qualitative preference queries over an Ontology-Based Data Access (OBDA) approach which provides an uniform access over multiple and heterogeneous data sources. Our main contribution is Morph-Pref, a framework for processing SPARQL qualitative preference to directly query relational databases. Our framework implements a technique that translates SPARQL qualitative preference queries directly into queries that can be evaluated by a relational database management system. We evaluate our approach over different scenarios, reporting the effects of data distributions, data size and query complexity on the performance of our proposed technique in comparison with state-of-the-art techniques. Obtained results suggest that the execution time can be reduced by up two orders of magnitude in comparison to current techniques scaling up to larger datasets while identifying precisely the result set.

**Keywords:** Qualitative preference, Skyline, OBDA, Query translation, R2RML

## 1. Introduction

Eliciting and exploiting preferences in query evaluation over relational databases and triple stores has attracted a sustained interest in the last two decades [1–11]. Such interest is motivated by the need of users who are not database experts but are willing to explore large datasets, commonly coming from the integration of multiple and heterogeneous data sources [4–6, 11]. Usually, these users do not know, a priori, what useful information they can extract from this data or they do not have a particular result in mind until it is discovered as an outcome of their data exploration process. During data exploration, they try to identify useful information according to their preferences by means of eliciting queries that best meet their criteria. Typical

examples include travelers looking for the best deals on accommodations to visit a city or users looking for laboratories with the best offers on PCR tests in a range of hours more appropriate for them during the COVID-19 global pandemic. Database engines should be able to filter useful information to support requirements of these non-expert users, i.e., evaluating queries that represent the desirable properties over the final result during the user’s search process. This kind of user expects such queries to be easily pose, correctly interpreted by the engine, and computationally efficient.

Introducing preference criteria into queries has been tackled in the context of relational databases (RDB) [7–10, 12, 13]. They are often used to filter information and thus reduce the volume of data being displayed to the user. At the level of preference criteria, two different approaches can be pursued: qualitative and quantitative. In the quantitative approach [1, 2], preference

---

\*Corresponding author: E-mail: mgoncalves@usb.ve

1 criteria are specified by means of scoring functions  
2 that assign a numerical score with each instance of  
3 the query response. Thus, an instance *A* is preferred to an  
4 instance *B* if the score of *A* is higher than the score  
5 of *B*. The problem of lack of expressiveness of this  
6 approach is well known in utility theory [14] where  
7 preferences are only represented by numerical scor-  
8 ing functions which are not necessarily easy to define  
9 by any user whereas preferences in the qualitative ap-  
10 proach can be expressed closer to natural language. For  
11 example, a typical score function for finding cheap hot-  
12 tels near the sea is weighted average on price and dis-  
13 tance but the user should know how to define weights.  
14 In addition, the user should normalize values of price  
15 and distance because they have different units. In con-  
16 trast, under the qualitative approach, the user may sim-  
17 ply express his preferences as minimizing price and  
18 distance, which is closer to natural language. The qual-  
19 itative approach is strictly more general than the quan-  
20 titative one since preference queries are based on the  
21 observation that expressions, such as "I prefer more *A*  
22 than *B* if *C*", are easily stated by users when asked  
23 about their preferences. For instance, when a customer  
24 wants to book an appointment for a PCR test, it is  
25 easy for him to say that he prefers closer laboratories  
26 with more economical prices and appointments avail-  
27 able after 17:00 on weekdays. There are two reasons  
28 for expressing qualitative preference queries [3]. First,  
29 it is desirable to offer more expressive query languages  
30 that allow the user to express what he is really trying to  
31 specify. Second, a classic query may also return a set  
32 of empty answers, while a qualitative one may produce  
33 at least some answers. Under this scenario, it should be  
34 optimal if a query engine is able to directly optimize  
35 such expressions of preference in a query.

36 Although qualitative preference queries have been  
37 widely studied in relational databases (RDB)[1–3], the  
38 literature on them in the context of Semantic Web is  
39 not as abundant. There are multiple example in the lit-  
40 erature that can be interpreted as preference search [4,  
41 15]. Some extensions of the SPARQL query language  
42 to incorporate user qualitative preferences have been  
43 also studied [4–6]. A particular case of the qualita-  
44 tive preference queries are the skyline ones [16] which  
45 have been widely studied in RDB [17]. Also, a set of  
46 client-server based algorithms have been introduced to  
47 evaluate skyline queries over knowledge graphs using  
48 standard query interfaces, such as SPARQL endpoints  
49 and TPF (Triple Pattern Fragments), with no control  
50 over how the knowledge graph is stored [18]. In our  
51 previous research [19], focused on skyline queries, we

1 have already demonstrated that the use of a physical  
2 operator during the SPARQL skyline query processing  
3 can significantly improve the query performance w.r.t  
4 the evaluation of either queries that were translated to  
5 SPARQL using query rewriting techniques to comply  
6 with the SPARQL standard or the skyline operator on  
7 the top of RDF triplestore. Although this means that  
8 qualitative preference queries can be executed over  
9 SPARQL endpoints, the lack of techniques that exploit  
10 the data storage structures in triplestores to specifically  
11 deal with the complexity of these queries can have a  
12 negative impact over their evaluation. Thus, the per-  
13 formance of these queries over RDF-based knowledge  
14 graphs is still low.

15 In addition, Virtual Knowledge Graphs (VKG) pro-  
16 vide access to data in terms of an existing ontology in  
17 accordance with a set of mapping rules [20], either by  
18 generating materialized views (RDF files) [21, 22] or  
19 by translating SPARQL queries into queries that are  
20 supported by the underlying source, which is known as  
21 virtualization [23, 24]. The latter is specially relevant  
22 when a qualitative preference query has to be evaluated  
23 because: *i*) it ensures up to date results at the moment  
24 of the execution, and *ii*) the preferring clause can be  
25 pushed down to the underlying data management sys-  
26 tems (e.g., an RDBMS) exploiting the benefits of pro-  
27 posed preference physical operators to improve query  
28 performance [12, 13]. *iii*) to the best of our knowl-  
29 edge there are no works in the literature for evaluat-  
30 ing SPARQL qualitative preference queries on VKG.  
31 In this article, we are interested in the evaluation of  
32 SPARQL qualitative preference queries over data that  
33 are not only available in RDF, as in the aforementioned  
34 works, but in relational databases.

35 **Problem Statement and Research Objective:** In this  
36 paper, we focus on the problem of evaluating SPARQL  
37 qualitative preference queries during the process of  
38 constructing a virtual knowledge graph. We are inter-  
39 ested in determining the feasibility of such type of  
40 query evaluation, understanding the correctness and  
41 completeness of our approach, and then, determining  
42 whether its performance is adequate in comparison  
43 with current qualitative preference query approaches  
44 over RDF.

45 **Approach:** We describe Morph-Pref, a virtual knowl-  
46 edge graph approach for qualitative preference queries.  
47 Based on the SPARQL-to-SQL query translation ap-  
48 proach defined in [25], Morph-Pref translates and op-  
49 timizes qualitative preference queries from SPARQL  
50 to SQL by means of a set of R2RML mappings. The  
51 approach includes a novel algorithm QualQT, which

1 pushes down the evaluation of the preferring clause,  
2 delegating its treatment to a physical operator of the  
3 RDBMS.

4 **Evaluation:** We use the Linked Movie Database  
5 (LinkedMDB) and adapt the benchmark for qualita-  
6 tive queries defined in [26], with SPARQL queries.  
7 We compare our work against state-of-the-art tools.  
8 This evaluation allows understanding the impact that  
9 queries of different complexities, and dataset sizes  
10 have on a SPARQL-to-SQL qualitative preference  
11 query evaluation. The results of the experiments sug-  
12 gest that all these variables impact on the total query  
13 execution time and approaches based on VKG over-  
14 come native SPARQL methods up to two orders of  
15 magnitude. Additionally, as an initial study, we run two  
16 of the queries on gas stations from [6] which were de-  
17 fined on an extension of preference operator that cor-  
18 rectly handle any preference relation. To the best of  
19 our knowledge, there is no defined benchmarks nor ex-  
20 perimental studies have been conducted for this type  
21 of queries.

22 **Contributions:** Our contributions can be summarized  
23 as follows: *i)* a formal definition of the problem of  
24 evaluating SPARQL qualitative preference queries on  
25 virtual knowledge graphs; *ii)* Morph-Pref, an approach  
26 based on the SPARQL-to-SQL translation approach  
27 proposed by Chebotko et al. [25] that is able to evalu-  
28 ate SPARQL qualitative preference queries over RDB;  
29 *iii)* the definition and implementation of QualQT, an  
30 algorithm based on SPARQL-to-SQL translations tak-  
31 ing advantage of specific physical operators for these  
32 queries; *iv)* an empirical evaluation of the behavior of  
33 Morph-Pref over two benchmarks with queries of dif-  
34 ferent complexities.

35 The remainder of this article is structured as follows:  
36 Section 2 motivates the problem of evaluating quali-  
37 tative preference queries on virtual knowledge graphs.  
38 Section 3 presents related work in qualitative prefer-  
39 ence queries and OBDA. Sections 4-5 describe our  
40 OBDA-based approach, the background and the pro-  
41 posed solution which is implemented in Morph-Pref.  
42 Section 6 reports and discusses on the results of our  
43 empirical study, and finally, Section 7 concludes and  
44 gives insights for future work.

## 45 2. Motivating Example

46  
47  
48  
49 Suppose a customer will travel from Madrid to Ger-  
50 many and he has decided to get the RT-PCR test in  
51 Madrid on the Friday or Saturday before his trip. The

1 customer is interested in finding the laboratories that  
2 meet the following two conditions in query  $Q_1$ : *i)* are  
3 cheapest, and *ii)* are the closest, and among those lab-  
4 oratories, he prefers the ones that have available ear-  
5 lier appointments *i)* before 10:00 if is Saturday, and  
6 *ii)* before 10:00 or after 17:00 if is Friday, and he also  
7 prefers earlier appointments ( $\leq 10:00$ ) over later ones  
8 ( $\geq 17:00$ ). All of these specifications are equally im-  
9 portant for the customer. Following our customer's cri-  
10 teria, a laboratory will be taken into account if there  
11 is no other laboratory with better price, distance, and  
12 appointments available at the desired time. This set  
13 of laboratories will compose the preferred ones. For-  
14 mally, the set of laboratories preferred by the user is  
15 composed of those that are not dominated by any other  
16 laboratory. Given  $A$  and  $B$  laboratories,  $A$  dominates  $B$ ,  
17 if  $A$  has the same or better values in price, distance, and  
18 appointment schedule as  $B$ , and it has a better value  
19 in at least one of them. To illustrate qualitative prefer-  
20 ences, consider the query  $Q_1$  expressed in SPARQL as  
21 shown in the Fig. 1a and a knowledge graph that con-  
22 tains laboratory properties and their appointment offer-  
23 ings in Fig. 1c.

24 Following these criteria, the appointments app2,  
25 app3, and app4 are the non-dominated ones, i.e., there  
26 is no other appointment with values better than them in  
27 :starts. Also, the appointment app2 dominates app1 be-  
28 cause although its schedule is after 17:00, its schedule  
29 is earlier than the app1. In addition to this, synlab dom-  
30 inates melio because it has better price and distance,  
31 and equal value in appointment.

32 The graph of the Fig. 1c can be seen as an RDF view  
33 defined over a relational database. If we consider most  
34 PCR appointment booking systems are based on RDB  
35 technology and a change in their data model can be  
36 costly and time-consuming, the OBDA approach [20]  
37 can be a solution that ensures the results are always  
38 up to date, converting these data to RDF and to make  
39 them available without renewing the entire software in-  
40 frastructure. Under an OBDA approach, the SPARQL  
41 query from Fig. 1a can be translated according to  
42 R2RML mappings into a SQL query as in Fig. 1b, and  
43 after executing the SQL query on data in Fig. 1d, the  
44 results produced by the relational engine are subse-  
45 quently transformed into a SPARQL resultset which is  
46 described through the concepts of the ontology and de-  
47 fined mappings.


48 Preference-based operators integrated as physical  
49 implementations inside relational database engines  
50 have shown substantial benefits [16, 26, 27]. Figs. 2a-  
51 2b depict query plans for the query  $Q_1$  where the

```

1 SELECT * WHERE {?l :price ?price; :distance ?dist .
2 ?l :offers ?a; :starts ?s; :ends ?e; :day ?day .}
3 PREFERRED LOWEST ?price AND LOWEST ?dist AND
4 (?s <= 1000) PRIOR TO (?s >= 1700) AND
5 LOWEST IF (?day = "Saturday") THEN (?s <= 1000)
6 ELSE IF (?day = "Friday") THEN (?s >= 1700 ||
7 ?s <= 1000) ELSE (false)

```

(a) A SPARQL Qualitative Preference Query



(c) An RDF Graph

```

1 SELECT * FROM laboratory l,offers o,appointment a
2 WHERE l.lab_id = o.lab_id AND o.app_id = a.app_id
3 PREFERRED LOWEST price AND LOWEST distance AND
4 (starts <= 1000) PRIOR TO (starts >= 1700) AND
5 LOWEST IF (day = "Saturday") THEN (starts <= 1000)
6 ELSE IF (day = "Friday") THEN (starts >= 1700 ||
7 starts <= 1000) ELSE (false)

```

(b) A SQL Qualitative Preference Query

laboratory		
lab_id	price	distance
melio	130	2
synlab	120	1

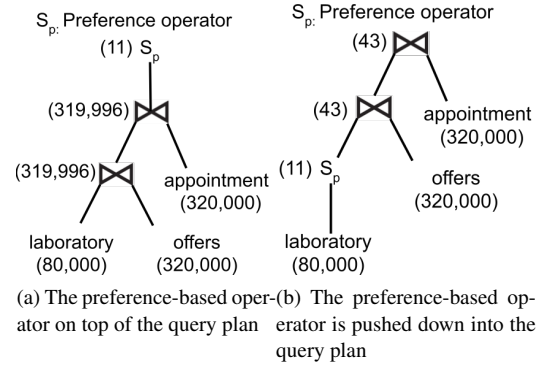
appointment			
app_id	day	starts	ends
app1	Friday	1800	1810
app2	Friday	1700	1710
app3	Saturday	900	910
app4	Saturday	930	940

offers	
lab_id	app_id
melio	app1
synlab	app1
synlab	app2
synlab	app3
synlab	app4

(d) A Relational Database

Fig. 1. Sample qualitative preference queries for booking a PCR appointment. The queries identify cheaper and closer laboratories with earlier appointments before 10:00 or after 17:00 on Friday or Saturday. Under an OBDA approach, the SPARQL query is translated into SQL, and the RDB results are converted to RDF.

preference-based operator is at the plan root or was pushed down, respectively, and the number of instances is placed below the node and the number of intermediate results is placed to the left of the operator. Fig. 2c reports the average execution time in seconds of the query  $Q_1$  only with skyline preferences which was executed 5 times for the state-of-the-art tools [5, 18] over our motivating dataset scaled-up to a 10.000 scale value by using VIG [28]. It is important to highlight that skyTPF was not reported because it returned an empty answer. These state-of-the-art tools evaluate the preference-based operator on top of the database engine in a similar pattern to the plan shown in the Fig. 2a. As well as the projection and selection operators that are pushed down in the query plan, a good heuristic is to apply a preference-based operator early because it reduces the sizes of the intermediate results by being an operator that filters data based on preferences. Pushing down the preference-based operator as in Fig. 2b would result in a lower execution time. Despite state-of-the-art tools have proposed optimizations to improve the performance of a preference query (e.g. [16] avoids either comparing the input instances against themselves in the best of the case), they ignore the engine capabilities underlying data model, which results in higher execution time ( $\geq 32.42$ ). In this article, our objective is to translate the SPARQL-to\_SQL where we take into account the capabilities of



(a) The preference-based operator on top of the query plan (b) The preference-based operator is pushed down into the query plan

SPREFQL	TPF	brTPF	skyTPF
55.45	138.51	32.42	-

(c) Average execution times (sec)

Fig. 2. Two query plans for a preference-based query. When the preference-based operator is integrated into the engine, query performance is improved. Intermediate results are shown over each step of the query plans using brackets.

the underlying data model, exploiting them to improve the preference query performance

### 3. State of the Art

There has been much interest in integrating the notion of preferences into query languages in order to express requirements that accurately reflect user needs [1–11]. The proposed approaches can be classified into two categories according to their qualitative or quantitative nature. In the former, preferences are expressed quantitatively through a scoring function to determine an ordering over query results. Closer to our work, in the latter, preferences are denoted by binary preference relationships yielding a partial order, contrary to the quantitative approach. Chomicki and colleagues [7–9], and Kießling and colleagues [10, 12, 13] defined foundational theories relating preferences to database systems and proposed extensions to SQL that support the specification of qualitative preference queries. These authors defined a logical formalism that allows the use of qualitative preference criteria into SQL where qualitative preferences are represented by strict partial orders through an operator called *winnow*. The *winnow* operator is a generalization of the skyline operator [16]. A skyline operator returns a partially ordered set of instances whose order is induced by criteria composed of conditions on equally important parameters. For example, a typical skyline query is to find hotels cheaper and closer to the beach.

The problem of efficiently computing the skyline has been widely studied in the literature [16, 29–31]. Bentley and colleagues [32] proposed the first skyline algorithm, referred to as the maximum vector problem and based on the divide & conquer principle. Progress continued to be made on how to compute efficiently such queries in a relational system and over large datasets and thus, the skyline approach [16] was defined in the context of databases to distinguish the best tuples that satisfy a given ranking condition. Subsequently, several authors proposed algorithms exploit data ordering properties or index structures in order to improve the skyline computation [29–31, 33, 34]. In the Semantic Web context, [18] presented a set of server-client based algorithms to evaluate skyline queries over knowledge graphs using standard query interfaces for RDF. They are focused on the optimization over client architectures so they assumed no control over how the data is stored (e.g., indexes, internal structures, etc), hence, they cannot exploit them to optimize the performance and quality of the queries at server side. However, in our context, we have control over data stored by the database engine. In this work, preferences more general than skyline are not sup-

ported. In addition, some database engines have been implemented showing significant benefits by including the skyline or *winnow* operators [12, 13, 16, 26, 27].

SPARQL has been extended with qualitative preferences [4–6] based on the *winnow* operator [7]. The authors in [4] propose to add qualitative preference capabilities to SPARQL, focusing on the feasibility of the process. They incorporate the PREFERENCE clause into the SPARQL syntax where preferences are separated by the AND construct. Also, the CASCADE keyword allows prioritizing a preference criterion over another one. They extended the ARQ query engine [35] with BNL as a proof of concept. However, they do not deal with query processing/optimization issues. SPREFQL [5] is another extension of SPARQL for qualitative preferences. Their work comes nearer to Chomicki’s framework [7] than [4] because it allows the expression of extrinsic preferences whose formulas may refer both to built-in predicates on the basis of tuples and to other constructors such as database relations. Unlike [4], they support conditional preferences (if-then-else). At the implementation level, they presented a query rewriting technique that maps from a SPREFQL query to an equivalent SPARQL query by means of the NOT EXISTS operator. They experimentally study the performance of Nested-Loop (NL) algorithm, Block-Nested Loop (BNL) and a query rewriting technique to evaluate preference based queries. BNL iteratively compares each instance with non-dominated instances in a window, and the instance is returned only if it is not dominated by any other while NL is a naive algorithm that compares each input instance against all input instances and whose computational complexity is quadratic. Unfortunately, their solution based on query rewriting does not work correctly due to the fact that it is based on the SPARQL NOT EXISTS, which has many known problems [36]. Thus, [6] identified and fixed the problem in the previous proposals for acyclic and transitive preference relations. Moreover, they proposed an efficient implementation for the preference query evaluation using the union-find algorithm [37]. Datalog +/- was extended in [38] to include preferences and the authors developed algorithms to answer preference queries over Datalog +/- ontologies. In contrast to our work, all these strategies evaluate the preference operator on top of the engine.

OBDA engines are data integration systems usually focused on the transformation of the original sources into a global schema (ontology) [20]. The process can be performed by its corresponding materialization [22]

(i.e., transforming input sources to RDF knowledge graphs) but also on query translation techniques for highly dynamic data sources [23, 24], where mapping rules are used to translate the input SPARQL query to an equivalent one, usually in SQL [25]. To cover the features of different data formats, many different types of OBDA mapping languages have also been proposed in the last few decades, with a wide variety of syntax and formats. Since its W3C recommendation in 2012, R2RML [39] has become the standard mapping specification for accessing relational databases. Many tools support these rules, some of them materialize the data into a knowledge graph (e.g. DB2Triples<sup>1</sup> and R2RMLParser<sup>2</sup>) and others provide virtual RDF views, focusing on formalizing the translation of SPARQL into SQL and optimizing the resulting SQL query (Morph-RDB [24] and Ontop [23]). Despite the efforts in the development of these engines, to the best of our knowledge, current existing OBDA engines do not support SPARQL qualitative queries. The most related work is [40] which considers the problem of evaluating top-k queries in the context of OBDA over relational databases.

#### 4. OBDA-based Qualitative Preference Queries

In this section we describe Morph-Skyline++, an OBDA framework for translating and executing qualitative preference queries from SPARQL-to-SQL. First, we formally define the problem of translating qualitative preference queries from SPARQL-to-SQL. Second, we probe the correctness of translating qualitative preferences from SPARQL-to-SQL. Third, we propose a solution based on the Chebotko et al.'s translation approach [25] which is implemented in Morph-Skyline++.

##### 4.1. Problem Definition

Our formalization is based on OBDA [20], which relies on conceptually representing a domain of interest over data stored in an underlying database system.

**Definition 1** (OBDA Specification [20]). *OBDA is defined as a triple  $\tau = \langle \mathcal{O}, \mathcal{S}, \mathcal{M} \rangle$  where  $\mathcal{O}$  is an ontology that describes the domain,  $\mathcal{S}$  is a schema, and  $\mathcal{M}$  is a mapping between  $\mathcal{O}$  and  $\mathcal{S}$ . In addition, an OBDA*

*instance is defined as a tuple  $\mathcal{PI} = \langle \tau, \mathcal{D} \rangle$  where  $\tau$  is an OBDA specification and  $\mathcal{D}$  is a data instance conforming to  $\mathcal{S}$ .*

A qualitative preference query can be specified on an ontology  $\mathcal{O}$  or on a relational schema  $\mathcal{S}$ . A qualitative preference query returns a set of non-dominated solutions to a SPARQL pattern (resp. a set of non-dominated tuples to a relational schema) where a solution (resp. a tuple) is non-dominated if there is no other solution (resp. other tuple) that is better to it according to a preference relation  $\prec$ .

**Definition 2** (Dominance). *Let  $\Sigma$  be a set of solutions to a SPARQL graph pattern  $gp$  (resp. a set of tuples to  $\mathcal{S}$ ) and let  $\prec$  be a preference relation over  $\Sigma \times \Sigma$ . A solution (resp. a tuple)  $\delta_1 \in \Sigma$  is dominated by a solution (resp. a tuple)  $\delta_2 \in \Sigma$  if  $\delta_2 \prec \delta_1$ .*

A preference relation  $\prec$  can be atomic or combined between them as shown in Definition 3.

**Definition 3** (Atomic and Combined Preference Relations). *A preference formula  $U(\delta_1, \delta_2)$  is a first-order formula defining a preference relation  $U$  in the standard sense, namely,  $\delta_1 \prec_U \delta_2$  if  $U(\delta_1, \delta_2)$  holds. Let  $\prec_U, \prec_V$  be two preference relations over  $\Sigma \times \Sigma$ , we define atomic and combined preference relations as follows:*

- *Boolean (e.g. intersection):  $\delta_1 \prec_{U \wedge V} \delta_2 \equiv (\delta_1 \prec_U \delta_2) \wedge (\delta_1 \prec_V \delta_2)$*
- *Scoring:  $\delta_1 \prec_{LOWEST, <} \delta_2 \equiv \delta_1 < \delta_2$  and  $\delta_1 \prec_{HIGHEST, <} \delta_2 \equiv \delta_2 < \delta_1$*
- *Multidimensional (Skyline):  $\delta_1 \prec_{U \otimes V} \delta_2 \equiv \delta_1 \preceq_U \delta_2 \wedge \delta_1 \preceq_V \delta_2 \wedge (\delta_1 \prec_U \delta_2 \vee \delta_1 \prec_V \delta_2)$*
- *Prioritized:  $\delta_1 \prec_{U \triangleright V} \delta_2 \equiv (\delta_1 \prec_U \delta_2) \vee (\neg(\delta_1 \prec_U \delta_2) \wedge \neg(\delta_2 \prec_U \delta_1) \wedge (\delta_1 \prec_V \delta_2))$ .*
- *Conditional:  $\delta_1 \prec_{\varphi, \phi, \gamma} \delta_2 \equiv (\varphi(\delta_1) \wedge \varphi(\delta_2) \wedge \delta_1 \prec_{\phi} \delta_2) \vee (\neg\varphi(\delta_1) \wedge \neg\varphi(\delta_2) \wedge \delta_1 \prec_{\gamma} \delta_2)$*

*For our motivation example, preferences for the SPARQL query  $Q_1$  can be expressed as the preference formula  $C$  in Fig. 3.*

**Definition 4** (Qualitative Preference Queries). *A qualitative preference query  $Q$  over an ontology  $\mathcal{O}$  (resp. a schema  $\mathcal{S}$ ) is defined as a pair  $Q = (gp, \prec)$  (resp.  $Q' = (\mathcal{S}, \prec)$ ) which produces a subset of  $\Sigma$  such that:  $\{p_1 \in \Sigma \mid \neg \exists p_2 \in \Sigma \wedge p_2 \prec p_1\}$ .*

To compute the desired solutions (resp. tuples) for all preference relations that are acyclic or transitive, a qualitative preference query was extended in [6].

<sup>1</sup><https://github.com/antidot/db2triples>

<sup>2</sup><https://github.com/nkons/r2rml-parser>

$$\begin{aligned}
& (?price, ?distance, ?s, ?day) \prec_C \\
& (?price', ?distance', ?s', ?day) \equiv \\
& (?price, ?distance, ?s, ?day) \prec_{C_1 \otimes C_2 \otimes C_4} (*) \\
& (?price', ?distance', ?s', ?day) \\
& (?price, ?distance) \prec_{C_1} (?price', ?distance') \equiv \\
& ?price \leq ?price' \wedge ?distance \leq distance' \wedge \\
& (?price < ?price' \vee ?distance < distance') \\
& (?s) \prec_{C_2} (?s') \equiv (?s \leq 1000) > (?s' \leq 1000) \\
& \vee (\neg((?s \leq 1000) > (?s' \leq 1000))) \wedge \\
& \neg((?s' \leq 1000) > (?s \leq 1000)) \wedge \\
& (?s \geq 1700) > (?s' \geq 1700)) \\
& (?s, ?day) \prec_{C_3} (?s', ?day') \equiv \\
& (?day = 'Saturday' \wedge (?s \leq 1000)) \wedge \\
& ?day' = 'Saturday' \wedge \neg(?s' \leq 1000)) \vee \\
& (?day = 'Friday' \wedge ((?s \leq 1000) \vee (?s \geq 1700))) \wedge \\
& ?day' = 'Friday' \wedge \neg((?s' \leq 1000) \vee (?s' \geq 1700)))
\end{aligned}$$

(\*)where  $C_4$  is LOWEST on solutions after applying  $\prec_{C_3}$

Fig. 3. Preference formulas for our motivating query

**Definition 5** (Extended Qualitative Preference Queries). If  $\Sigma$  is a finite set of solutions (resp. tuples) and  $\prec$  is a preference relation over  $\Sigma \times \Sigma$  then qualitative preference query  $Q'$  returns:  $\{p_1 \in \Sigma \mid \neg \exists p_2 \in \Sigma \wedge p_2 \prec^* p_1 \wedge \neg(p_1 \prec^* p_2)\}$  where  $\prec^*$  is the transitive closure of  $\prec$  over  $\Sigma$ .

**Problem Definition.** Given an OBDA Specification  $\sigma = \langle \mathcal{O}, \mathcal{S}, \mathcal{M} \rangle$  and an SPARQL qualitative preference query  $Q = (gp, \prec)$  over an ontology  $\mathcal{O}$ , the problem of qualitative preference query translation in an OBDA setting is defined as finding an equivalent qualitative preference query  $Q' = (SQ, \prec)$  over  $\mathcal{S}$  where  $SQ$  is a SQL query with an equivalent relational algebra expression  $RA$  over a relational schema  $\mathcal{S}$ .

#### 4.2. Semantics Preserving translation

In this section, we give a proof of semantics preserving translation based on compositional semantics of the two query languages. A compositional semantics means that an expression is built out of the meanings of its sub-expressions. We define Semantic Algebras and their corresponding evaluation functions for both SPARQL and SQL based on Relational Algebra operators since Relational Algebra forms the basis for query translation and its close relationship with SQL. To simplify the proof, we define its semantics by using tuple relational calculus formulas and the work presented in [41] which probed the equivalence between relational algebra and relational calculus presented.

**Definition 6** (Semantic Algebra for SPARQL). Given:

- RDF terms  $T$  that comprise of pairwise disjoint infinite sets of  $I$ ,  $B$  and  $L$  (IRI's, Blank nodes and Literals, respectively) from  $\mathcal{O}$ .
- A partial function  $\mu \in \mathcal{S} : V \rightarrow T$  where  $V$  is an infinite set of variables that is pairwise disjoint from the sets  $I$ ,  $B$ , and  $L$ .
- Two mappings  $\mu_1, \mu_2$  are compatible when for all  $x \in \text{dom}(\mu_1) \cap \text{dom}(\mu_2)$ , it is the case that  $\mu_1(x) = \mu_2(x)$  where  $\text{dom}(\mu)$  is the subset of  $V$  where  $\mu$  is defined.
- A domain mapping set  $\Omega \equiv \llbracket tp \rrbracket$  where  $\llbracket tp \rrbracket$  is an evaluation function that maps an element of  $gp$  according to Definition 7.
- A well-designed SPARQL graph pattern  $gp$  defined by the following grammar:  $gp ::= tp \mid gp \text{ FILTER } expr \mid gp \text{ UNION } gp \mid gp \text{ AND } gp \mid gp \text{ OPT } gp \mid gp \text{ PREFERRING } crit$ .

The following operators are defined:

1.  $\sigma_{expr} \Omega = \{\mu \mid \mu \in \Omega \wedge \mu \models expr\}$  where  $\mu \models expr$  is evaluated according to [25].
2.  $\Pi_{v_1, \dots, v_n} \Omega = \{\mu_{v_1, \dots, v_n} \mid \mu \in \Omega\}$
3.  $\Omega_1 \cup \Omega_2 = \{\mu \mid \mu \in \Omega_1 \vee \mu \in \Omega_2\}$
4.  $\Omega_1 \bowtie \Omega_2 = \{\mu_1 \cup \mu_2 \mid \mu_1 \in \Omega_1, \mu_2 \in \Omega_2 \text{ are compatible mappings}\}$
5.  $\Omega_1 \setminus \Omega_2 = \{\mu \in \Omega_1 \mid \text{for all } \mu' \in \Omega_2, \mu \text{ and } \mu' \text{ are not compatible}\}$
6.  $\Omega_1 \bowtie \Omega_2 = \{(\Omega_1 \bowtie \Omega_2) \cup (\Omega_1 \setminus \Omega_2)\}$
7. For the preference operator  $\rho$ ,  $\rho_{crit} \Omega = \{\mu \mid \mu \in \Omega \wedge \mu \models crit\}$  where  $crit$  is criteria and  $\mu$  satisfies  $crit$ , denoted by  $\mu \models crit$ , iff: i)  $crit$  is  $expr$  and  $\mu \models expr$  where  $expr$  is a Boolean expression as defined by the FILTER expression; ii)  $crit$  is  $op ?X$  and  $?X \in \text{dom}(\mu)$ , where  $op \rightarrow \text{HIGHEST} \mid \text{LOWEST}$ ; iii)  $crit$  is  $(crit_1 op crit_2)$ ,  $\mu \models crit_1$  and  $\mu \models crit_2$ , where  $op \rightarrow \otimes \mid \triangleright$ ; iv)  $crit$  is if  $crit_1$  then  $crit_2$  else  $crit_3$ ,  $\mu \models crit_1$ ,  $\mu \models crit_2$  and  $\mu \models crit_3$ .

**Definition 7** (Evaluation functions over SPARQL queries).  $\llbracket \cdot \rrbracket : gp \rightarrow \Omega$  denotes an evaluation function that maps an element of  $gp$  to  $\Omega$  and given a triple pattern  $tp$ ,  $\mu(tp)$  denotes the triple obtained by replacing the variables in  $tp$  according to  $\mu$ .

1.  $\llbracket tp \rrbracket = \{\mu \mid \text{dom}(\mu) = \text{var}(tp) \wedge \mu(tp) \in \mathcal{O}\}$
2.  $\llbracket gp \text{ FILTER } expr \rrbracket = \sigma_{expr} \llbracket gp \rrbracket$
3.  $\llbracket gp_1 \cup gp_2 \rrbracket = \llbracket gp_1 \rrbracket \cup \llbracket gp_2 \rrbracket$
4.  $\llbracket gp_1 \text{ AND } gp_2 \rrbracket = \llbracket gp_1 \rrbracket \bowtie \llbracket gp_2 \rrbracket$
5.  $\llbracket gp_1 \text{ OPT } gp_2 \rrbracket = \llbracket gp_1 \rrbracket \bowtie \llbracket gp_2 \rrbracket$

$$6. \llbracket gp \text{ PREFERRED } crit \rrbracket = \rho_{crit} \llbracket gp \rrbracket$$

**Definition 8** (Semantic Algebra for Relational Algebra). *Given:*

- An expression  $E$  in algebra relational defined by the following grammar:  $E ::= \sigma_{Cond} E \mid \Pi_{A_1, \dots, A_n} E \mid E_1 \cup E_2 \mid E_1 \bowtie_{Cond} E_2 \mid E_1 \times E_2 \mid E_1 - E_2 \mid E_1 \cap E_2 \mid \text{PREFERRED}_{crit}$
- A domain relation,  $R = \{t \mid \xi(t) \equiv \xi(R)\}$  with all tuples such that  $t$  and  $R$  are defined over the same schema in  $\mathcal{S}$ .

The following operators are defined:

- $\sigma_{Cond} R = \{t \mid R(t) \wedge Cond(t)\}$ ; where  $Cond(t)$  is a Boolean expression in SQL such that  $Cond$  is satisfied by  $t$ .
- $\Pi_{A_1, \dots, A_n} R = \{t.A_1, \dots, t.A_n \mid R(t)\}$
- $R_1 \cup R_2 = \{t \mid R_1(t) \vee R_2(t)\}$
- $R_1 \bowtie_{Cond} R_2 = \{t \mid \forall t_1, \forall t_2 (R_1(t_1) \wedge R_2(t_2) \rightarrow t.A_1 = t_1.A_1 \wedge \dots \wedge t.A_n = t_1.A_n \wedge t.B_1 = t_2.B_1 \wedge \dots \wedge t.B_n = t_2.B_n \wedge Cond(t))\}$
- $R_1 \times R_2 = \{t \mid \forall t_1, \forall t_2 (R_1(t_1) \wedge R_2(t_2) \rightarrow t.A_1 = t_1.A_1 \wedge \dots \wedge t.A_n = t_1.A_n \wedge t.B_1 = t_2.B_1 \wedge \dots \wedge t.B_n = t_2.B_n)\}$
- $R_1 - R_2 = \{t \mid R_1(t) \wedge \forall t_2 (R_2(t_2) \rightarrow t \neq t_2)\}$
- $R_1 \cap R_2 = \{t \mid R_1(t) \wedge R_2(t)\}$
- $\rho_{crit} R = \{t \mid R(t) \wedge \neg \exists t_1 (R(t_1) \wedge t_1 \prec t)\}$ ; where  $crit$  is a SQL expression that represents the preference relation  $\prec$ .

**Definition 9** (Evaluation functions over Relational Algebra).  $\llbracket \cdot \rrbracket_r : E \rightarrow R$  denotes an evaluation function that maps an element of  $E$  to  $R$ .

1.  $\llbracket \sigma_{Cond} E \rrbracket_r = \sigma_{Cond} \llbracket E \rrbracket_r$
2.  $\llbracket \Pi_{A_1, \dots, A_n} E \rrbracket_r = \Pi_{A_1, \dots, A_n} \llbracket E \rrbracket_r$
3.  $\llbracket E_1 \cup E_2 \rrbracket_r = \llbracket E_1 \rrbracket_r \cup \llbracket E_2 \rrbracket_r$
4.  $\llbracket E_1 \bowtie_{Cond} E_2 \rrbracket_r = \llbracket E_1 \rrbracket_r \bowtie_{Cond} \llbracket E_2 \rrbracket_r$
5.  $\llbracket E_1 \times E_2 \rrbracket_r = \llbracket E_1 \rrbracket_r \times \llbracket E_2 \rrbracket_r$
6.  $\llbracket E_1 - E_2 \rrbracket_r = \llbracket E_1 \rrbracket_r - \llbracket E_2 \rrbracket_r$
7.  $\llbracket E_1 \cap E_2 \rrbracket_r = \llbracket E_1 \rrbracket_r \cap \llbracket E_2 \rrbracket_r$
8.  $\llbracket \rho_{crit} E \rrbracket_r = \rho_{crit} \llbracket E \rrbracket_r$

**Theorem 1.** *Given an OBDA Specification  $\sigma = \langle \mathcal{O}, \mathcal{S}, \mathcal{M} \rangle$  and a SPARQL qualitative preference query  $Q = (gp, \prec)$  over an ontology  $\mathcal{O}$ , the translation of  $Q$  to an equivalent SQL qualitative preference query with an equivalent Relational Algebra expression  $RA \in E$  over a relational schema  $\mathcal{S}$  is a semantics preserving translation.*

*Proof.* Having already been proved in [42] that  $\llbracket RA \rrbracket_r \equiv \llbracket gp \rrbracket$  for all the traditional operators in Relational Algebra where  $RA$  is relational algebra expression of SQL and  $SA$  is relational algebra expression of SPARQL, we prove the operator  $\rho$  in SPARQL and Algebra Relational are equivalent.

$$SA: \rho_{crit} \Omega = \{\mu \mid \mu \in \Omega \wedge \mu \models crit\}$$

$$RA: \rho_{crit} R = \{t \mid R(t) \wedge \neg \exists t_1 (R(t_1) \wedge t_1 \prec t)\}$$

If  $\Omega \equiv \llbracket tp \rrbracket$ ,  $crit \equiv \prec$ ,  $\models \equiv \neg \exists \mu_1 (\mu \in \Omega \wedge \mu_1 \prec \mu)$ , and  $R \in \mathcal{S}$  such that  $tp \stackrel{?}{=} s R ? o$  then  $SA \equiv RA$ .  $\square$

**Proposed Solution** We propose Morph-Pref as an OBDA qualitative preference query translation engine for relational databases. Morph-Pref is based on the Chebotko et al.'s translation approach [25] which formalizes query translation from SPARQL into SQL exploiting the use of a set of mapping rules and their corresponding functions. More specifically, Morph-Pref relies on the formal definition of mappings and functions with R2RML provided by [24] and it exploits the preference operators in databases during query processing. When a SPARQL qualitative preference query is received by Morph-Pref, our solution, QualQT uses R2RML mappings to translate into SQL. Then, the translated query is executed against a qualitative-preference-enabled query engine and the results are returned.

QualQT is a technique based on a physical preference operator. By using QualQT, a SPARQL qualitative preference query  $Q = (gp, \prec)$  is translated into a SQL qualitative preference query  $Q' = (SQ, \prec)$ . First,  $gp$  is translated to  $SQ$  by means of the function  $\mu$  defined as *trans* in [24]. Then, each variable  $v \in \prec$  from  $Q$  is matched to  $v' \in \prec$  from  $Q'$  according to the mapping  $\mathcal{M}$ . Thus, the translated qualitative preference query will be executed into the underlying database system. The SQL qualitative preference query in Fig. 1b is the result of translating  $gp$  to  $SQ$ , and each variable in  $\prec$ .

Algorithm 1 sketches the algorithm QualQT implemented by Morph-Pref to process qualitative preference queries on virtual knowledge graphs. QualQT receives a set of mappings  $\mathcal{M}$  and a SPARQL qualitative preference query  $Q$ . It firstly separates the preferring clause from the query  $Q$  and gets a non-qualitative-preference subquery (lines 1-2). In line 3, QualQT translates the non-qualitative-preference query  $gp$  into SQL by means of the function *trans* defined in [24]. Line 4 builds a hash structure from the graph patterns belonging to  $gp$ ; this hash structure contains the attribute that maps each variable. Note that  $\beta$  returns the column/constant that corresponds to each variable



in a SPARQL query [24]. A mapping  $\beta$  is a function whose input is a set of possible triple patterns  $TP$ , a set of triple pattern positions (subject, predicate or object)  $POS$ , and a set of all possible R2RML mappings  $\mathcal{M}$ , and its output is a set of relational attributes  $ATTR$ . Thus, a mapping  $\beta$  is a many-to-one mapping  $\beta^m : TP \times POS \times \mathcal{M} \rightarrow ATTR$ . Subsequently, for each variable in the preferring clause, QualQT iteratively obtains its corresponding attribute by searching for it in the structure hash  $h$ , and creates the corresponding expression with the mapped attribute (lines 5-7). Each variable in the preferring clause of Fig 1a, LOWEST ?price AND LOWEST ?dist AND (?s <= 1000) PRIOR TO (?s >= 1700) AND LOWEST IF (?day = "Saturday") THEN (?s <= 1000) ELSE IF (?day = "Friday") THEN (?s >= 1700 || ?s <= 1000) ELSE (false) is replaced by its corresponding attribute to LOWEST price AND LOWEST distance AND (starts <= 1000) PRIOR TO (starts >= 1700) AND LOWEST IF (day = "Saturday") THEN (starts <= 1000) ELSE IF (day = "Friday") THEN (starts >= 1700 || starts <= 1000) ELSE (false). Finally, QualQT returns a pair with the translated query and  $ST$ .

**Algorithm 1** Qualitative Preference Query Translation, QualQT:  $\mathcal{M}$  - Mappings,  $Q$  - Qualitative Preference Query

---

```

1: pref  $\leftarrow$  preferring clause from  $Q$ 
2:  $gp \leftarrow Q \setminus \text{pref}$ 
3: query  $\leftarrow \text{trans}(gp, \mathcal{M})$ 
4:  $h \leftarrow \beta(gp, \mathcal{M})$ 
5: for each variable  $v_i \in \text{pref}$  do
6:    $a_i \leftarrow \text{get}(h, v_i)$ 
7:   replace  $v_i$  in pref with  $a_i$ 
8: return (query, pref)

```

---

## 5. Query Language Syntax

In this section, we summarize the grammar supported by the parser of Morph-Skyline++. We rely on the PrefSPARQL grammar [11] to support qualitative preferences in SPARQL queries. The grammar basis is SPARQL 1.1 but the definition of <SolutionModifier> changes. Qualitative preferences are specified as a new solution modifier in the EBNF grammar shown in Grammar 1. This grammar includes rules for skyline, prioritized and conditional preferences. The keyword 'AND' in <Preference> allows to separate inde-

pendent dimensions of skyline queries, and the rules <HighestPref>, <LowestPref> and <DiffPref> correspond to maximizing, minimizing, and grouping criteria in skyline queries. Unlike previous works [4–6, 11], we have included the DIFF directive from Börzsönyi and colleagues's work [16]. Also, Prioritized and Conditional (IF-THEN-ELSE) preferences are included in the grammar as the rules <PrioritizedPref> and <ConditionalPref>, respectively. Finally, all non-terminals that are not defined in this grammar are defined by the standard SPARQL syntax.

```

<SolutionModifier> ::= [ <GroupClause> ]
                       [ <HavingClause> ] [ <PreferClause> ] [ <OrderClause> ]
                       [ <LimitOffsetClauses> ]

<PreferClause> ::= 'PREFERRING' <Preference>

<Preference> ::= <PrioritizedPref> ('AND' <PrioritizedPref>)*

<PrioritizedPref> ::= <ConditionalOrAtomicPref> ('PRIOR TO' <ConditionalOrAtomicPref>)*

<ConditionalOrAtomicPref> ::= <ConditionalPref> | <AtomicPref>

<ConditionalPref> ::= 'IF' <Expression> 'THEN' <ConditionalOrAtomicPref>
                    'ELSE' <ConditionalOrAtomicPref>

<AtomicPref> ::= <Expression> | <HighestPref> | <LowestPref> | <DiffPref>

<HighestPref> ::= 'HIGHEST' <Expression>

<LowestPref> ::= 'LOWEST' <Expression>

<DiffPref> ::= 'DIFF' <Expression>

```

Grammar 1: Grammar for qualitative preferences in SPARQL

In SQL, Preference SQL [26] is a query language to express user wishes. Grammar 2 presents the grammar corresponding to preferring clause for Preference SQL. The preferring clause must be before the group-by clause and after the where clause. The partition-by clause can be seen as the DIFF directive in a skyline query. in contrast to PrefSPARQL, PLUS,

1 CASE-WHEN-THEN-ELSE, HIGH, and LOW are  
2 used instead of AND, IF-THEN-ELSE, HIGHEST,  
3 and LOWEST, respectively.

4  
5  $\langle \text{PreferClause} \rangle ::= \text{'PREFERRING' } \langle \text{Preference} \rangle$   
6  $[\text{'PARTITION' } \langle \text{Expression} \rangle \text{'BY'}$   
7  $\langle \text{Expression} \rangle \text{'(, ' } \langle \text{Expression} \rangle \text{'*)' ]$

8  
9  $\langle \text{Preference} \rangle ::= \langle \text{PrioritizedPref} \rangle \text{'(PLUS'}$   
10  $\langle \text{PrioritizedPref} \rangle \text{'*)'}$

11  
12  $\langle \text{PrioritizedPref} \rangle ::= \langle \text{ConditionalOrAtomicPref} \rangle$   
13  $\text{'(PRIOR TO'}$   
14  $\langle \text{ConditionalOrAtomicPref} \rangle \text{'*)'}$

15  
16  $\langle \text{ConditionalOrAtomicPref} \rangle ::= \langle \text{ConditionalPref} \rangle \text{' |'}$   
17  $\langle \text{AtomicPref} \rangle$

18  
19  $\langle \text{ConditionalPref} \rangle ::= \text{'CASE'}$   
20  $\langle \text{Expression} \rangle \text{'(WHEN'}$   
21  $\langle \text{ConditionalOrAtomicPref} \rangle$   
22  $\text{'THEN' } \langle \text{result} \rangle \text{'*)' 'ELSE'}$   
23  $\langle \text{ConditionalOrAtomicPref} \rangle$

24  
25  $\langle \text{AtomicPref} \rangle ::= \langle \text{Expression} \rangle \text{' |' } \langle \text{HighestPref} \rangle \text{' |'}$   
26  $\langle \text{LowestPref} \rangle \text{' |' } \langle \text{DiffPref} \rangle$

27  
28  $\langle \text{HighestPref} \rangle ::= \text{'HIGH' } \langle \text{Expression} \rangle$

29  
30  $\langle \text{LowestPref} \rangle ::= \text{'LOW' } \langle \text{Expression} \rangle$

31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51

Grammar 2: Grammar for qualitative preferences in PREFERENCE SQL

In both SQL and SPARQL, the keywords AND, PRIOR TO, and IF-THEN-ELSE (resp. CASE-WHEN-THEN-ELSE) are used for the multidimensional, prioritized, and conditional preference relations, respectively.

Lastly, a qualitative preference query can be translated into a nested SQL query while compares each tuple with each other tuple.  $Q \text{ PREFERRING } \prec_{\text{Pref}}$  can be expressed in SQL as  $Q \text{ WHERE NOT EXISTS } (Q' \text{ trans } (Q, Q', \prec_{\text{Pref}}))$  where  $Q'$  is the query  $Q'$  but with the same tables renamed with aliases,  $\prec_{\text{Pref}}$  the preference criteria, and  $\text{trans}$  performs pairwise dominance checks for each pair of tuples. The function  $\text{trans}$  is recursively applied as:

- for  $\text{Pref} = \text{Pref}_1 \text{ AND } \text{Pref}_2, \text{trans } (Q, Q', \prec_{\text{Pref}}) = \text{trans } (Q, Q', \prec_{\text{Pref}_1 \otimes \text{Pref}_2})$
- for  $\text{Pref} = \text{Pref}_1 \text{ PRIOR TO } \text{Pref}_2, \text{trans } (Q, Q', \prec_{\text{Pref}}) = \text{trans } (Q, Q', \prec_{\text{Pref}_1 \triangleright \text{Pref}_2})$

```
1 SELECT * FROM laboratory l, offers o, appointment a
2 WHERE l.lab_id = o.lab_id AND o.app_id = a.app_id AND
3 NOT EXISTS
4 (SELECT * FROM laboratory l1, offers o1, appointment a1
5 WHERE l1.lab_id = o1.lab_id AND o1.app_id = a1.app_id
6 -- Multidimensional preferences
7 AND o1.price <= o.price AND o1.distance <= o.distance
8 AND (o1.price < o.price OR o1.distance <= o.distance)
9 -- Prioritized preferences
10 AND ( (o1.starts <= 1000) > (o.starts <= 1000) OR
11 (NOT ((o1.starts <= 1000) > (o.starts <= 1000)) AND
12 NOT ((o.starts <= 1000) > (o1.starts <= 1000)) AND
13 (o1.starts >= 1700) > (o.starts >= 1700)) AND
14 -- conditional preferences
15 ((o1.day = "Saturday" AND o.day = "Saturday" AND
16 (o1.starts <= 1000) > (o.starts <= 1000)) OR
17 (o1.day = "Friday" AND o.day = "Friday" AND
18 ((o1.starts <= 1000) > (o.starts <= 1000)) OR
19 (o1.starts >= 1000) > (o.starts >= 1000)) ) )
```

Fig. 4. A sample translated qualitative preference query for booking a PCR appointment . The query identifies cheaper and closer laboratories with earlier appointments before 10:00 or after 17:00 on Friday or Saturday by means of a nested SQL query.

– for  $\text{Pref} = \text{IF } E \text{ THEN } \text{Pref}_1 \text{ ELSE } \text{Pref}_2,$   
 $\text{trans } (Q, Q', \prec_{\text{Pref}}) = \text{trans } (Q, Q', \prec_{\implies E, \text{Pref}_1, \text{Pref}_2})$

Our motivating query in Fig. 1b can be translated into a nested SQL query as depicted in the Fig. 4

In our experimental study, as baseline, we compare our work against qualitative preference queries in SQL and their translated into nested SQL query in order to evaluate how costly it is to perform a qualitative preference query on virtual knowledge graphs. We have omitted a detailed description of the SPARQL query rewriting technique [5] to translate a qualitative preference query from SPARQL to SPARQL. The authors have already shown that this technique does not scale. But, intuitively, the triple patterns within an OPTIONAL clause retrieve the dominated solutions, and similar to a left outer join, those solutions not bound to dominated ones are in the answer. For a more detailed description of this technique, please refer to [5].

Lastly, knowing the grammar for preferences in SPARQL and SQL, a qualitative preference query can be translated from SPARQL to SQL.

## 6. Experimental Study

We study the efficiency and effectiveness of our two implementations of Morph-Skyline++. First, we describe the hypotheses that we want to validate as well as the datasets and queries, mappings, metrics and implementation details for our experimental study. A

Table 1  
A Summary of experiments

Section	Objective	Benchmarks	Hypothesis
6.1	A performance comparison between SPARQL preference qualitative query and their equivalent SQL preference qualitative query	TPC-H	$H_1$ , $H_2$ , and $H_3$
6.2	A performance comparison between Morph-Skyline++ and state-of-the-art triplestores. One of the tool only support skyline queries	TPC-H and Linked-MDB	$H_4$ and $H_5$
6.3	A preliminar study over scalability of extended qualitative preference queries	Gas Stations	-

summary of our evaluation setup is presented in Table 1. All the resources used are online<sup>3</sup>.

**Hypotheses.** The impact of the preference criteria and the dataset size on the execution time are studied over different datasets. Our study aims at evaluating the following hypotheses in order to test the correctness of Morph-Skyline++ approach (H1-H4) and its capabilities in comparison with the state of the art:

- $H_1$ . A SPARQL-to-SQL preference-based algorithm executed by an OBDA engine does not spend much more time with respect to a physical operator within a relational database engine.
- $H_2$ . As the preference criteria increase, the answer cardinality augments, hence, the OBDA preference-based algorithms have to perform more dominance checks; therefore, they increase their execution time
- $H_3$ . As the dataset size increases, the answer size also enlarges and the OBDA preferences-based algorithms have to perform more dominance checks; therefore, they increase their execution time.
- $H_4$ . A SPARQL-to-SQL preference-based algorithm executed as a physical operator within a relational database engine has better performance than the operator executed on top of a database engine, even though the engine included techniques to optimize the query.
- $H_5$ . Morph-Skyline++ performs better than the current SPARQL-based approach on materialized RDF for skyline query evaluation.

**Benchmarks and Queries:** *i) TPC-H:* We use the Transaction Processing Council Ad-hoc/decision support benchmark consisting of 15 queries expanded by PREFERRING clauses in [26]. We adapt this benchmark with SPARQL queries. TPC-H datasets are generated in terms of scale factors where a scale factor of 1 corresponds to 1 GB of data. We have generated TPC-

H datasets for scale factors of 1, 5, 10, 20, 30, 40, 50, and 60. A summary of TPC-H datasets and queries is presented in Table 2. The 1GB dataset was transformed into RDF by means of SDM-RDFizer [21] and Table 2 contains the number of triples and size for TPC-H transformed into RDF. To test the methods in [18], we also transform the RDF data to HDT by means of the HDT library<sup>4</sup>. *ii) Linked Movie Database (Linked-MDB):* We run 7 LinkedMDB queries defined by [5] and a summary of this dataset and queries can be observed in Table 3. This dataset was manually transformed to CSV using SQL queries and then it was loaded into an RDB to be evaluated by Morph-Skyline++. *iii) Gas Stations:* In [6], the authors presented a running example for gas stations for 18 instances. We have manually converted this data to CSV and then scaled it with VIG [28]. We have created an equivalent table and we have loaded this CSV file into mysql. We have run VIG with different scales: 5, 10, 50, 100, 500, and 1000. A scale value  $v$  indicates that the table size increases  $v$  times. A summary of the number of triples and size for these datasets can be found in Table 3. This dataset will only be used for the experimental study of the extended qualitative preferences queries.

**Mappings:** *i)* For the TPC-H benchmark, we have created an R2RML mapping document for accessing each relational table with 53 PredicateObjectMaps, 53 Predicates, 53 ObjectsMaps, and 9 JoinConditions. *ii)* For LinkedMDB, an R2RML mapping document was built with 5 TriplesMaps, 10 PredicateObjectMaps, 10 Predicates, 10 ObjectsMaps, and 5 JoinConditions.

**Evaluation Metrics:** We measure performance as query execution time. It is computed as the elapsed time in seconds between the submission of a query to the engine and the delivery of the answers. Each query was executed 5 times in cold mode and timeouts are set up to 1 hour. The quality of preference-based techniques are also measured in terms of precision, recall

<sup>3</sup><https://github.com/marleeng/morph-preferences>

<sup>4</sup><http://www.rdfhdt.org/>

Scale	lineitem	orders	partsupp	part	customer	supplier
1GB	6M	1.5M	0.8M	0.2M	0.15M	0.01M
5GB	30M	7.5M	4M	1M	7.5M	0.05M
10GB	60M	15M	8M	2M	1.5M	0.1M
20GB	120M	30M	16M	4M	3M	0.2M
30GB	180M	45M	24M	6M	4.5M	0.3M
40GB	240M	60M	32M	8M	6M	0.4M
50GB	300M	75M	40M	10M	7.5M	0.5M
60GB	360M	90M	48M	12M	9M	0.6M

Table 2

TPC-H. Number of rows per table according scale (left) and queries summary (right)

Preference	Queries
Skyline	Q0, Q2, Q4, Q6, Q9, Q12, Q14
Partitioning (DIFF)	Q10-Q14
Prioritized composition	Q1, Q3, Q5, Q7, Q10, Q11
Conditional preferences	Q8, Q13

Preference	Queries
Skyline	Q1, Q2 & Q4
Partitioning (DIFF)	Q1
Prioritized composition	Q3 & Q5
Conditional preferences (EXISTS)	Q6 & Q7

Dataset	Triples	Size
TPC-H	111M	20GB
LinkedMDB	6M	851M
1-scaled Gas Stations	18	0.54KB
5-scaled Gas Stations	100	2.7KB
10-scaled Gas Stations	200	5.3KB
50-scaled Gas Stations	1K	27KB
100-scaled Gas Stations	2K	55KB
500-scaled Gas Stations	10K	285KB
1000-scaled Gas Stations	10M	332MB

Table 3

LinkedMDB. Queries summary (left) and summary of RDF datasets (right)

and F-measure. Precision measures the percentage of instances that should be in the preferred set computed in terms of the ground truth; recall measures the percentage of instances produced in terms of the ground truth. Ground truth corresponds to the preferred set directly retrieved from the relational database engine and then materialized in RDF using the R2RML mappings and the SDM-RDFizer [21] engine to ensure their correctness.

**Engines:** Morph-Skyline++ is written in Scala and extends ARQ Jena. The relational database engine used is EXASol because it includes a dedicated preferences-based operator implementing the BNL algorithm [26]. We have compared Morph-Skyline++ against SPREFQL [5], developed in Java within the RDF4J framework<sup>5</sup> and querying data loaded into a locally deployed SPARQL endpoint of Virtuoso Community Edition Version 7.1. We have also evaluated the multi-threaded version of the skyTPF and brTPF-based methods [18], a Java servlet implementation that uses RDF-HDT data sources to process skyline queries over Bindings-restricted Triple Pattern Fragments (brTPF). We considered testing our work against NL [5], and TPF-based method [18], but the authors showed that these algorithms have worse performance with respect to BNL, and the skyTPF and brTPF-based methods, re-

spectively. Experiments are executed on a machine with the following characteristics: 2GHz CPU with 8 cores, 64 GB RAM with Ubuntu 18.04 as its operating system.

### 6.1. Cost of exploiting preferences-based operators over OBDA engines

Based on the state of the art in skyline queries which demonstrate that the skyline operator integrated into a database engine is much more efficient than the corresponding operator but evaluated by translating the skyline query into a SQL query on top of it [16, 26, 27], we wanted to study how expensive a preferences-based operator is in an OBDA engine with respect to the equivalent SQL queries. Thus, our baseline includes a SQL query with preferring clause and a translated SQL query that uses a nested query with NOT EXISTS[16]. We compare the evaluation of the preferences-based queries when the operator is integrated into the RDBMS or when the operator is not supported by the RDBMS and must be translated into an equivalent SQL query in order to be executed. In this section, we analyze how significant the performance degradation can be if a preferences-based query is directly executed by an OBDA engine.

Figure 5 reports the average execution time in seconds on a log scale for the 15 TPC-H queries. These

<sup>5</sup><http://rdf4j.org/>

queries are expressed in SPARQL to be evaluated by QualQT in Morph-Skyline++, and their equivalent queries are expressed in SQL with either the preferring clause (SQL-Pref) or NOT EXISTS (SQL) to be directly executed by EXASol. In Fig. 5a, we can observe that QualQT is slightly worse than SQL for the first query, Q0, which has only 1 criterion, but as the number of criteria increases (Q1-Q9), the QualQT runtime remains between the preferences-based operator runtime (SQL-Pref) and the evaluation time of the SQL query translated with NOT EXISTS (SQL). QualQT is up to two orders of lesser magnitude less than SQL-Pref. These results indicate us that the SPARQL-to-SQL translation and the data conversion from an RDB to triples does not have a high impact on performance of Morph-Skyline++ and even QualQT has a better performance than the evaluation of the SQL query translated with NOT EXISTS (SQL-Pref). For the queries Q10-Q14, QualQT is the worst and even timeout occurs with the last two queries. In particular, these queries are the ones that produce the most results because they identify the preferred solutions per group, i.e., they are grouping queries. To get some idea on the number of results, while the queries Q0-Q9 return less than 73 solutions, the queries Q10-Q12 return between 10K and 20K solutions approximately, and the queries Q13-Q14 return between 58K and 102K solutions approximately. When we analyzed the worst performance of Morph-Skyline++ for these queries, we have observed that the translation time was very small, the SQL query execution time is already included in Fig. 5a, but the data conversion from the RDB was the process that required significantly more time. The generation of results by Morph-Skyline++ is a costly process since it produces the triples in XML format. For most cases, the QualQT time is of the same order of magnitude as the execution time of preferences-based operator (SQL) except for cases where results generated are over 50K solutions. **Therefore, our hypothesis  $H_1$  does not hold: QualQT is not costly w.r.t the physical operator for preferences within an RDB engine.**


Figures. 5b-5d depict the average execution time in seconds on a log scale for the 15 TPC-H queries over dataset sizes from 5GB to 20GB. In order of query complexity, the most complex queries are Q13 and Q14 because they are grouping ones and they have four criteria. After those queries, the queries Q7-Q9 are the next most complex. They have four criteria, and the queries Q8 and Q9 are the same as the queries Q13 and Q14 but without grouping. As the dataset

size and the query complexity augment, it becomes more difficult to answer the query and timeout occurs. The grouping, number of criteria and conditionals increase the query complexity and as a result, the time for the query execution. For dataset sizes greater than 10GB, timeout also occurs in the queries Q13-Q14 for the preferences-based operator and SQL-Pref. Also, Fig. 6 shows the average query time varying the dataset sizes for several configurations of queries. As observed, for the query Q5 characterized by prioritized preferences and skyline, QualQT presents a comparable performance to the preferences-based operator while SQL-Pref significantly worsened by the increase in the dataset size. For the query Q8 which includes skyline and conditional preferences, QualQT begins to degrade its performance from 10GB. For the query Q11 with skyline, grouping and prioritized preferences, timeout occurs from 30GB for QualQT and the equivalent SQL query evaluations. And the worst case for QualQT and the equivalent SQL query evaluations is the query Q13 which involves skyline, grouping and conditional preferences, and timeout occurs from 10GB. Observed results suggest that QualQT and SQL query evaluations are not able to scale up to complex queries for larger datasets. **Thus, our hypothesis  $H_2$  and  $H_3$  holds for QualQT: its execution time increases with the number of criteria and dataset size because the preferences-based operator of EXASol implements BNL and must perform a greater number of dominance checks to determine whether a solution is preferred or not.**


Figures. 5e-5h show the average execution time in seconds on a log scale for the 15 TPC-H queries over dataset sizes from 30GB to 60GB. We can observe that the behavior is similar. For the queries Q0-Q7, the QualQT runtime is between the runtime of preferences-based operator and the SQL-Pref runtime but timeout occurs when the query complexity augments. For dataset sizes greater than 30GB, timeout also occurs for the preferences-based operator and SQL-Pref, and the queries Q10-Q14.

## 6.2. Morph-Skyline++ vs native SPARQL


In this section, We compare Morph-Skyline++ against SPREFQL and the skyTPF and brTPF-based methods in order to analyze their performance with respect to the number of criteria for 1GB TPC-H and Linked-MDB. With this comparison, We want to demonstrate the importance of having specific physical operators integrated into the database engine for preferences-




(a) Performance of Morph-Skyline++ for a 1GB dataset




(b) Performance of Morph-Skyline++ for a 5GB dataset




(c) Performance of Morph-Skyline++ for a 10GB dataset




(d) Performance of Morph-Skyline++ for a 20GB dataset




(e) Performance of Morph-Skyline++ for a 30GB dataset



(f) Performance of Morph-Skyline++ for a 40GB dataset



(g) Performance of Morph-Skyline++ for a 50GB dataset




(h) Performance of Morph-Skyline++ for a 60GB dataset


Fig. 5. **Performance of Morph-Skyline++ for 1-60GB datasets.** We compare QualQT implemented by Morph-Skyline++ against execution of skyline SQL queries and translated skyline SQL queries directly in the database engine for TPC-H in different database sizes.

based queries, and that the evaluation of them on top of the database engine are not a good solution neither to native knowledge graphs (RDF) nor virtual knowledge


graphs. Within the state-of-the-art tools that support preferences-based SPARQL queries, the preferences-based operator is evaluated on top of triplestores [5]




(a) Performance of Morph-Skyline++ for the query Q5



(b) Performance of Morph-Skyline++ for the query Q8



(c) Performance of Morph-Skyline++ for the query Q11



(d) Performance of Morph-Skyline++ for the query Q13

Fig. 6. **Performance of Morph-Skyline++.** We compare QualQT implemented by Morph-Skyline++ against execution of skyline SQL queries and translated skyline SQL queries directly in the database engine for different queries.

or standard query interfaces such as TPF (Triple Pattern Fragments) [18]. It is noteworthy that although we wanted to compare QualQT with the SPARQL

Table 4

**Precision (P), Recall (R) and F-measure (F).** SPREFQL produces incorrect and incomplete answers for 4/15 queries

Query	Precision	Recall	F-Measure
7	8.451	17.143	21.007
9	8.276	16.438	11.009
10	1.250	100.000	2.469
14	1.704	17.518	3.105


preference-based query rewritten technique, Morph-Skyline++ was not able to run the SPARQL translated queries on 1GB TPC-H where the biggest TPC-H table has 6M tuples because Exasol was running out of disk space. Already in our previous work [19], our experimental study showed that from 1M tuples, the translation technique does not scale. The biggest TPC-H table for 1GB has 6M tuples.

Fig. 7a reports the average execution time in seconds for the 15 TPC-H queries on Morph-Skyline++, BNL of SPREFQL, and the rewriting technique on SPREFQL (RW). The first observed result is that RW timed out for the most queries. RW is costly because of expensive not bound, optional and filter clauses. RW is not a good option for such a volume of data. The second observed result is that QualQT was the best for the first 13 queries. For the queries Q13 and Q14, timeout occurs because of the execution time of result generation. It is important to emphasize that SPREFQL was unable to run the queries Q5, Q8, Q11 and Q13. Thus, we have double checked these queries by running them without the preference clause directly in Virtuoso and the result was the successful execution of them; therefore, we think it is because all these queries include functions in their preference clause, in particular, the ABS function.


Additionally, we reported on the SPREFQL quality for the 1GB TPC-H dataset. BNL produces incomplete and incorrect responses for some queries. Table 4 reports precision, recall and F-measure for SPREFQL for 4 queries. For the remaining queries, precision, recall and F-measure of SPREFQL is 100%. We can observe that precision, recall and F-measure are very low when the number of criteria is high (= 4) or the query is a grouping one.

To compare QualQT with TPF, brTPF and skyTPF-methods, Fig. 7b shows the average execution times in seconds on a log scale for the skyline queries in TPC-H: Q0, Q2, Q4, Q6, Q8, and Q9. Although skyTPF is one of the best methods presented in [18], we observed that timeout occurred in 5 queries. This timeout occurs due to skyTPF continues to work despite a Null Pointer Exception. skyTPF only successfully ran for





(a) Performance of Morph-Skyline++ and SPREFQL



(b) Performance of Morph-Skyline++, SPREFQL and skyTPF for skyline queries

Fig. 7. **Performance of Morph-Skyline++.** We compare QualQT implemented by Morph-Skyline++ against SPREFQL, and TPF, brTPF and skyTPF-methods for the 1GB TPC-H dataset.

the first query. And, once more, QualQT is the clear winner. **Thus, our hypothesis  $H_4$  and  $H_5$  hold: the physical operator of QualQT executed within a relational database engine has better performance in view of the fact that the engine optimizes the query.** The execution time for QualQT is at least one order of magnitude less than RDF tools. Given the time required by the RDF tools, it would be an added value to incorporate the preferences-based operator within the RDF engine.

In addition, we reported on the quality of methods presented in [18] for the 1GB TPC-H dataset. These methods produce incomplete and incorrect responses. Table 5 reports precision, recall and F-measure for SPREFQL for 4 queries. For the remaining queries, precision, recall and F-measure is not reported because of timeout. The low recall of the queries Q0 and Q6 is due to the fact that these methods produce different answers and a lower number of instances. With these results, we presume that the methods have problems with duplicated data in TPC-H.

Finally, we also compare Morph-Pref and SPREFQL on LinkedMDB. In this case, we can run the rewriting technique for Morph-Pref because the data is smaller. Fig. 8 reports the average execution time

Table 5

**Precision (P), Recall (R) and F-measure (F).** Methods in [18] produce incorrect and incomplete answers for 4/15queries

Query	Precision	Recall	F-Measure
7	8.451	17.143	21.007
9	8.276	16.438	11.009
10	1.250	100.000	2.469
14	1.704	17.518	3.105

in seconds for the 7 LinkedMDB queries on Morph-Pref and SPREFQL. It should be highlighted that we do not report the times of the methods in [18] because they did not produce results. For this benchmark, the BNL algorithm of SPREFQL has the best performance although QualQT has the same order of magnitude as SPREFQL. The number of results returned by queries is very low and evidently, SPREFQL efficiently process the data through access mechanisms of Virtuoso. Also, the rewriting techniques has the worst performance w.r.t BNL and QualQT. But, the rewriting technique on SPREFQL (RW-S) becomes better than QualQT for the queries Q2, Q3, Q6, and Q7. The queries Q2, Q3, and Q6 are queries with instantiations while the queries Q6 and Q7 only have a simple criterion with the EXISTS condition. In [43], the authors empirically show that RDF engines may be better than RDB ones when the query has instantiations. With reference to the rewriting technique on Morph-Pref (RW-Q), we have analyzed the execution plans of this kind of queries on EXASol. The main complex part of the query plan is a left outer join between the input table  $R$  and the result produced by a theta self-join of the input table  $R$ :  $R \bowtie (R \bowtie R)$ . The engine, after performing the theta self-join operator, builds an on-the-fly index for the left outer join to optimize the execution. This index can be huge, e.g., the index comprises approximately 200 million entries for 100K tuples and a 4-dimensional query. Finally, EXASol executes a left outer join using this index to discard the dominated ones, i.e., all the tuples from  $R$  minus matched tuples from the the theta self-join (dominated ones); the resulting tuples are those ones that have NULL in the case of no matching join predicate. In consequence, RW-Q performance is the worst. In addition, we can observe in Fig. 8 that for the last two queries, times of the rewriting techniques, RW-S and RW-Q, are slightly worse than BNL and QualQT, respectively because the rewriting is simpler and does not require additional self-joins to express the equivalent query, i.e., these queries only check that the criterion is true by means of an exists predicate.







Fig. 8. Performance of Morph-Pref and SPREFQL.

Fig. 9. Performance of  $\omega_{\lambda}^l$ .

### 6.3. Recursion

In this section, we present a preliminary experimental study on the winnow operator  $\omega_{\lambda}^l$ . This winnow operator is evaluated by computing the transitive closure of a preference relation over candidate solutions. We have evaluated the method that computes  $\omega_{\lambda}^l$  by means of recursive queries. Since recursive queries are not supported by Morph-RDB, we have evaluated the equivalent preferences-based SPARQL query directly in Virtuoso. Fig. 9 shows how the query performance is affected scaling the dataset sizes from 1 to 1000. We can observe that the execution time grows exponentially as the dataset size increases. The query runtime is clearly affected by the dataset size.

### 6.4. Discussion

In this section, we have considered the performance issues of implementing a preference-based operator as a physical operator in an OBDA engine. Although the preference-based operator can be expressed in SPARQL [5], a rewriting of qualitative preference queries using SPARQL is expensive [5] while a preference-based physical operator within the OBDA engine that is aware of the properties of preference relations significantly improves the query performance. Firstly, We have defined QualQT, an algorithm that

considers a preference-based physical operator integrated into an OBDA engine, and our results suggest that QualQT achieves a comparable performance to this operator implemented into an RDB engine for most cases.

Secondly, we have empirically shown that current state-of-the-art tools do not scale adequately for large volumes of data and QualQT significantly outperforms the state-of-the-art methods for larger datasets. Very little attention has been paid to the challenges associated with implementing preferences as a physical operator within a triplestore. We have implemented our technique QualQT on EXASol database and we have demonstrated that qualitative preference queries can benefit significantly from such implementation. State of the art tools conceive the preference-based operator as the top-most operator in a query plan, however a preference-based operator can be pushed down, decreasing the cardinality of the intermediate results, causing an important performance benefits and motivating the need for a physical operator.

Finally, a preference-based operator as a physical operator in an OBDA engine also provides the user the ability to express qualitative preference queries where the preference criteria is not the main operator, as in a subquery. State-of-the-art tools would not be able to resolve a query where preferences are expressed in a subquery.

## 7. Conclusions and Future Work

In this article, we have described Morph-Skyline++, an OBDA-based engine that identifies a subset of data in a virtual knowledge graph that best meet criteria of a user request expressed as a SPARQL qualitative preference query. Unlike previous works, we have included the DIFF directive for skyline queries. This directive allows grouping the skyline set by the attribute that comes before the directive. We have also designed and implemented QualQT, an algorithm based on a SPARQL-to-SQL query translation able to compute the preferred set on virtual knowledge graphs. We studied and analyzed the performance of QualQT with respect to the state-of-the-art methods. We reported experimental results with SPREFQL and the TPF-,sky/TPF- and brTPF-based methods, state-of-the-art tools to evaluate skyline queries over RDF data. These results showed that our approach is able to precisely compute the skyline set independently of the sizes of the datasets or the number of dimensions, and

it outperforms state-of-the-art tools for larger datasets. If we consider that the data are already materialized in RDF, SPREFQL and the TPF-, skyTPF- and brTPF-based methods will be better than QualQT when the preferred set is small w.r.t. the dataset size. Nevertheless, SPREFQL is not able to produce complete answers in some cases.

In the future, we plan to consider qualitative preferences into federated queries over SPARQL endpoints extending approaches already proposed such as [43, 44]. We also want to evaluate qualitative SPARQL queries over data sources in other data formats such as CSV or JSON, by extending the idea of enforcing implicit constraints exploiting declarative mapping rules presented in Morph-CSV [45]. Finally, with the results presented in this paper, we devise a research line focused on the development of physical operators in triplestores to enhance the performance of qualitative SPARQL queries over materialized RDF knowledge graphs.

## References

- [1] R. Agrawal and E.L. Wimmers, A Framework for Expressing and Combining Preferences, in: *Proceedings of the 2000 ACM SIGMOD International Conference on Management of Data, May 16-18, 2000, Dallas, Texas, USA*, W. Chen, J.F. Naughton and P.A. Bernstein, eds, ACM, 2000, pp. 297–306. doi:10.1145/342009.335423.
- [2] V. Hristidis, N. Koudas and Y. Papakonstantinou, PREFER: A System for the Efficient Execution of Multi-parametric Ranked Queries, in: *Proceedings of the 2001 ACM SIGMOD international conference on Management of data, Santa Barbara, CA, USA, May 21-24, 2001*, S. Mehrotra and T.K. Sellis, eds, ACM, 2001, pp. 259–270. doi:10.1145/375663.375690.
- [3] D. Dubois, A. Hadjali, H. Prade and F. Touazi, Erratum to: Database preference queries - a possibilistic logic approach with symbolic priorities, *Ann. Math. Artif. Intell.* **73**(3–4) (2015), 359–363. doi:10.1007/s10472-014-9446-2.
- [4] W. Siberski, J.Z. Pan and U. Thaden, Querying the Semantic Web with Preferences, in: *The Semantic Web - ISWC 2006*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2006, pp. 612–624. ISBN 978-3-540-49055-5. doi:10.1007/11926078\_44.
- [5] A. Troumpoukis, S. Konstantopoulos and A. Charalambidis, An Extension of SPARQL for Expressing Qualitative Preferences, in: *The Semantic Web – ISWC 2017*, Springer International Publishing, Cham, 2017, pp. 711–727. ISBN 978-3-319-68288-4. doi:10.1007/978-3-319-68288-4\_42.
- [6] P.F. Patel-Schneider, A. Polleres and D. Martin, Comparative Preferences in SPARQL, in: *Knowledge Engineering and Knowledge Management*, Springer International Publishing, Cham, 2018, pp. 289–305. ISBN 978-3-030-03667-6. doi:10.1007/978-3-030-03667-6/\_19.
- [7] J. Chomicki, Querying with Intrinsic Preferences, in: *Advances in Database Technology — EDBT 2002*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2002, pp. 34–51. ISBN 978-3-540-45876-0. doi:10.1007/3-540-45876-X\_5.
- [8] J. Chomicki, Preference Formulas in Relational Queries, *ACM Trans. Database Syst.* **28**(4) (2003), 427–466. doi:10.1145/958942.958946.
- [9] J. Chomicki, Logical Foundations of Preference Queries, *IEEE Data Eng. Bull.* **34**(2) (2011), 3–10. <http://sites.computer.org/debull/A11june/Chomicki1.pdf>.
- [10] W. Kießling, Foundations of Preferences in Database Systems, in: *Proceedings of the 28th International Conference on Very Large Data Bases, VLDB '02, VLDB Endowment, 2002*, pp. 311–322.
- [11] M. Guerousova, A. Polleres and S.A. McIlraith, SPARQL with Qualitative and Quantitative Preferences, in: *Proceedings of the 2nd International Workshop on Ordering and Reasoning, OrdRing 2013, Co-located with the 12th International Semantic Web Conference (ISWC 2013), Sydney, Australia, October 22nd, 2013*, I. Celino, E.D. Valle, M. Krötzsch and S. Schlobach, eds, CEUR Workshop Proceedings, Vol. 1059, CEUR-WS.org, 2013, pp. 2–8. <http://ceur-ws.org/Vol-1059/ordring2013-paper1.pdf>.
- [12] W. Kießling, M. Endres and F. Wenzel, The Preference SQL System - An Overview, *IEEE Data Eng. Bull.* **34**(2) (2011), 11–18. <http://sites.computer.org/debull/A11june/Kießling.pdf>.
- [13] W. Kießling and G. Köstler, Preference SQL - Design, Implementation, Experiences, in: *Proceedings of 28th International Conference on Very Large Data Bases, VLDB 2002, Hong Kong, August 20-23, 2002*, Morgan Kaufmann, 2002, pp. 990–1001. doi:10.1016/B978-155860869-6/50098-6. <http://www.vldb.org/conf/2002/S30P03.pdf>.
- [14] P.C. Fishburn, Preference Structures and Their Numerical Representations, *Theor. Comput. Sci.* **217**(2) (1999), 359–383. doi:10.1016/S0304-3975(98)00277-1.
- [15] T. Berners-Lee, J. Hendler and O. Lassila, The semantic web, *Scientific american* **284**(5) (2001), 34–43.
- [16] S. Börzsönyi, D. Kossmann and K. Stocker, The Skyline Operator, in: *Proceedings of the 17th International Conference on Data Engineering, April 2-6, 2001, Heidelberg, Germany*, D. Georgakopoulos and A. Buchmann, eds, IEEE Computer Society, 2001, pp. 421–430. doi:10.1109/ICDE.2001.914855.
- [17] C. Kalyvas and T. Tzouramanis, A Survey of Skyline Query Processing, *Computing Research Repository Journal abs/1704.01788* (2017).
- [18] I. Keles and K. Hose, Skyline Queries over Knowledge Graphs, in: *The Semantic Web – ISWC 2019*, Springer International Publishing, Cham, 2019, pp. 293–310. ISBN 978-3-030-30793-6. doi:10.1007/978-3-030-30793-6\_17.
- [19] M. Goncalves, D. Chaves-Fraga and Ó. Corcho, Morph-Skyline: Virtual Ontology-Based Data Access for Skyline Queries, in: *The 2020 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT'20)*, 2020.
- [20] A. Poggi, D. Lembo, D. Calvanese, G.D. Giacomo, M. Lenzerini and R. Rosati, Linking Data to Ontologies, *J. Data Semantics* **10** (2008), 133–173. doi:10.1007/978-3-540-77688-8\_5.
- [21] E. Iglesias, S. Jozashoori, D. Chaves-Fraga, D. Collarana and M.-E. Vidal, SDM-RDFizer: An RML Interpreter for the Efficient Creation of RDF Knowledge Graphs, in: *ACM Intern.*

- 1 Confer. on Information and Knowledge Management, CIKM,  
2 2020.
- 3 [22] A. Dimou, M.V. Sande, P. Colpaert, R. Verborgh, E. Mannens  
4 and R.V. de Walle, RML: A Generic Language for Integrated  
5 RDF Mappings of Heterogeneous Data, in: *LDOW*, CEUR  
6 Workshop Proceedings, Vol. 1184, CEUR-WS.org, 2014.
- 7 [23] D. Calvanese, B. Cogrel, S. Komla-Ebri, R. Kontchakov,  
8 D. Lanti, M. Rezk, M. Rodriguez-Muro and G. Xiao, Ontop:  
9 Answering SPARQL queries over relational databases, *Semantic  
10 Web* **8**(3) (2017), 471–487. doi:10.3233/SW-160217.
- 11 [24] F. Priyatna, Ó. Corcho and J.F. Sequeda, Formalisation and ex-  
12 periences of R2RML-based SPARQL to SQL query translation  
13 using morph, in: *23rd International World Wide Web Confer-  
14 ence, WWW '14, Seoul, Republic of Korea, April 7-11, 2014*,  
15 ACM, 2014, pp. 479–490. doi:10.1145/2566486.2567981.
- 16 [25] A. Chebotko, S. Lu and F. Fotouhi, Semantics preserv-  
17 ing SPARQL-to-SQL translation, *Data Knowl. Eng.* **68**(10)  
18 (2009), 973–1000. doi:10.1016/j.datak.2009.04.001.
- 19 [26] S. Mandl, O. Kozachuk, M. Endres and W. Kießling, Pref-  
20 erence Analytics in EXASolution, in: *Datenbanksysteme für  
21 Business, Technologie und Web (BTW)*, LNI, Vol. P-241, GI,  
22 2015, pp. 613–632. <https://dl.gi.de/20.500.12116/2434>.
- 23 [27] S. Chaudhuri, N. Dalvi and R. Kaushik, Robust Cardinal-  
24 ity and Cost Estimation for Skyline Operator, in: *Data En-  
25 gineering, 2006. ICDE '06. Proceedings of the 22nd In-  
26 ternational Conference on*, IEEE Computer Society, Los  
27 Alamitos, CA, USA, 2006, p. 64. ISBN 0-7695-2570-9.  
28 doi:10.1109/ICDE.2006.131.
- 29 [28] D. Lanti, G. Xiao and D. Calvanese, VIG: Data scaling for  
30 OBDA benchmarks, *Semantic Web* **10**(2) (2019), 413–433.  
31 doi:10.3233/SW-180336.
- 32 [29] J. Chomicki, P. Godfrey, J. Gryz and D. Liang, Skyline with  
33 Presorting., in: *ICDE*, U. Dayal, K. Ramamritham and T.M. Vi-  
34 jayarman, eds, IEEE Computer Society, 2003, pp. 717–  
35 719. ISBN 0-7803-7665-X. [http://dblp.uni-trier.de/db/conf/  
36 icde/icde2003.html#ChomickiGGL03](http://dblp.uni-trier.de/db/conf/icde/icde2003.html#ChomickiGGL03).
- 37 [30] P. Godfrey, R. Shipley and J. Gryz, Maximal Vector Compu-  
38 tation in Large Data Sets., in: *VLDB*, ACM, 2005, pp. 229–  
39 240. ISBN 1-59593-177-5. [http://dblp.uni-trier.de/db/conf/  
40 vldb/vldb2005.html#GodfreySG05](http://dblp.uni-trier.de/db/conf/vldb/vldb2005.html#GodfreySG05).
- 41 [31] I. Bartolini, P. Ciaccia and M. Patella, Efficient sort-based  
42 skyline evaluation, *ACM Trans. Database Syst.* **33**(4) (2008),  
43 31:1–31:49. doi:10.1145/1412331.1412343.
- 44 [32] J.L. Bentley, H.-T. Kung, M. Schkolnick and C.D. Thompson,  
45 On the average number of maxima in a set of vectors and appli-  
46 cations, *Journal of the ACM (JACM)* **25**(4) (1978), 536–543.
- 47 [33] D. Kossmann, F. Ramsak and S. Rost, Shooting Stars in the  
48 Sky: An Online Algorithm for Skyline Queries, in: *Proceed-  
49 ings of 28th International Conference on Very Large Data  
50 Bases, VLDB 2002, Hong Kong, August 20-23, 2002*, Morgan  
51 Kaufmann, 2002, pp. 275–286. doi:10.1016/B978-155860869-6/50032-9. <http://www.vldb.org/conf/2002/S09P01.pdf>.
- [34] M. Endres and E. Glaser, Indexing for Skyline Computation  
- A Comparison Study, in: *Flexible Query Answering Sys-  
tems - 13th International Conference, FQAS 2019, Amantea,  
Italy, July 2-5, 2019, Proceedings*, A. Cuzzocrea, S. Greco,  
H.L. Larsen, D. Saccà, T. Andreassen and H. Christiansen, eds,  
Lecture Notes in Computer Science, Vol. 11529, Springer,  
2019, pp. 31–42. doi:10.1007/978-3-030-27629-4\_6.
- [35] ARQ - A SPARQL Processor for Jena, Accessed: 2019-11-08.
- [36] P.F. Patel-Schneider and D. Martin, EXISTStential Aspects of  
SPARQL, in: *Proceedings of the ISWC 2016 Posters & Demon-  
strations Track co-located with 15th International Semantic  
Web Conference (ISWC 2016), Kobe, Japan, October 19, 2016*,  
T. Kawamura and H. Paulheim, eds, CEUR Workshop Pro-  
ceedings, Vol. 1690, CEUR-WS.org, 2016. [http://ceur-ws.org/  
Vol-1690/paper72.pdf](http://ceur-ws.org/Vol-1690/paper72.pdf).
- [37] R.E. Tarjan, Efficiency of a Good But Not Linear  
Set Union Algorithm, *J. ACM* **22**(2) (1975), 215–225.  
doi:10.1145/321879.321884.
- [38] T. Lukasiewicz, M.V. Martinez and G.I. Simari, Preference-  
Based Query Answering in Datalog+/- Ontologies, in: *IJCAI  
2013, Proceedings of the 23rd International Joint Conference  
on Artificial Intelligence, Beijing, China, August 3-9, 2013*,  
2013, pp. 1017–1023.
- [39] S. Das, S. Sundara and R. Cyganiak, R2RML: RDB to RDF  
Mapping Language. Working group recommendation, *World  
Wide Web Consortium (W3C)(Sep 2012)*, [http://www.w3.  
org/TR/r2rml](http://www.w3.org/TR/r2rml) (2012).
- [40] U. Straccia, On the Top-k Retrieval Problem for Ontology-  
Based Access to Databases, in: *Flexible Approaches in Data,  
Information and Knowledge Management*, Springer Interna-  
tional Publishing, Cham, 2014, pp. 95–114. ISBN 978-3-319-  
00954-4. doi:10.1007/978-3-319-00954-4\_5.
- [41] E.F. Codd, Relational Completeness of Data Base Sublan-  
guages., *Research Report / RJ / IBM / San Jose, California  
RJ987* (1972), republished on "ACM SIGMOD Anthology".
- [42] R. Cyganiak, A relational algebra for SPARQL, *Semantic Web*  
(2005), Technical Report HPL–2005-170, HP Labs.
- [43] K.M. Endris, P.D. Rohde, M.-E. Vidal and S. Auer, Ontario:  
Federated Query Processing Against a Semantic Data Lake, in:  
*Database and Expert Systems Applications*, Springer Interna-  
tional Publishing, Cham, 2019, pp. 379–395. ISBN 978-3-030-  
27615-7. doi:10.1007/978-3-030-27615-7\_29.
- [44] M.N. Mami, D. Graux, S. Scerri, H. Jabeen, S. Auer and  
J. Lehmann, Squerall: Virtual ontology-based access to hetero-  
geneous and large data sources, in: *International Semantic Web  
Conference*, Springer, 2019, pp. 229–245.
- [45] D. Chaves-Fraga, E. Ruckhaus, F. Priyatna, M. Vidal and  
Ó. Corcho, Enhancing Virtual Ontology Based Access over  
Tabular Data with Morph-Csv, *Semantic Web* (2021).
- [46] F. Michel, J. Montagnat and C. Faron Zucker, A survey of  
RDB to RDF translation approaches and tools, Research Re-  
port, I3S, 2014, ISRN I3S/RR 2013-04-FR 24 pages. <https://hal.archives-ouvertes.fr/hal-00903568>.