

# Preference Queries over Taxonomic Domains

Paolo Ciaccia  
University of Bologna  
Bologna, Italy  
paolo.ciaccia@unibo.it

Davide Martinenghi  
Politecnico di Milano  
Milano, Italy  
davide.martinenghi@polimi.it

Riccardo Torlone  
Roma Tre University  
Roma, Italy  
riccardo.torlone@uniroma3.it

## ABSTRACT

When composing multiple preferences characterizing the most suitable results for a user, several issues may arise. Indeed, preferences can be partially contradictory, suffer from a mismatch with the level of detail of the actual data, and even lack natural properties such as transitivity. In this paper we formally investigate the problem of retrieving the best results complying with multiple preferences expressed in a logic-based language. Data are stored in relational tables with taxonomic domains, which allow the specification of preferences also over values that are more generic than those in the database. In this framework, we introduce two operators that rewrite preferences for enforcing the important properties of transitivity, which guarantees soundness of the result, and specificity, which solves all conflicts among preferences. Although, as we show, these two properties cannot be fully achieved together, we use our operators to identify the only two alternatives that ensure transitivity and minimize the residual conflicts. Building on this finding, we devise a technique, based on an original heuristics, for selecting the best results according to the two possible alternatives. We finally show, with a number of experiments over both synthetic and real-world datasets, the effectiveness and practical feasibility of the overall approach.

### PVLDB Reference Format:

Paolo Ciaccia, Davide Martinenghi, and Riccardo Torlone. Preference Queries over Taxonomic Domains. PVLDB, 14(10): 1859 - 1871, 2021. doi:10.14778/3467861.3467874

## 1 INTRODUCTION

Preferences strongly influence decision making and, for this reason, their collection and exploitation are considered building blocks of content-based filtering techniques [6, 24, 25]. A key issue in this context is the mismatch that usually lies between preferences and data, which often makes it hard to recommend items to customers [19]. Indeed, whether they are collected by tracing the actions of the users or directly elicited from them, preferences are typically expressed in generic terms (e.g., I prefer pasta to beef), whereas available data is more specific (the menu might contain lasagne and hamburger). The problem of automatically suggesting the best solutions becomes even more involved when several preferences at different levels of granularity and possibly conflicting with each other are specified, as shown in the following example that will be used throughout the rest of the paper.

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing [info@vldb.org](mailto:info@vldb.org). Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 14, No. 10 ISSN 2150-8097. doi:10.14778/3467861.3467874

*Example 1.1.* We would like to select some bottles of wine from the list in Figure 1 available in an e-commerce store. We prefer white wines to red ones, yet we prefer Amarone (a famous red wine) to white wine. For the producer, we prefer Tuscan wineries located in the province of Siena to those in the Piedmont province of Asti. Moreover, if the winery lies in the Langhe wine region (which spans different provinces, partially including, among others, Asti and Cuneo) we prefer an aged wine (i.e., produced before 2017) to a more recent one. Finally, we would like to have suggestions only for the “best” possible alternatives. □

Wines			
Wine	Winery	Year	
Arneis	Correggia	2019	<i>a</i>
Amarone	Masi	2014	<i>b</i>
Amarone	Bertani	2013	<i>c</i>
Canaiolo	Montenidoli	2015	<i>d</i>
Barolo	Laficaia	2014	<i>e</i>
Arneis	Ceretto	2019	<i>f</i>

Figure 1: A list of wines

We first observe that further information is needed in this example to identify the solutions that better fit all the mentioned preferences. For instance, we need to know the province and the wine region in which all the wineries are located. In addition, the example shows that there are two important issues that need to be addressed in such scenarios. First, conflicts can occur when preferences are defined at different levels of detail. Indeed, the preference for Amarone, which is a red wine, is in contrast with the more generic preference for white wines. Second, further preferences can be naturally derived from those that are stated explicitly. For instance, from the preference for wines from Siena to those from Asti and the preference for aged wines when they are from the Langhe region, we can also derive, by transitivity, a preference for wines from Siena to young wines from Langhe.

In this paper we address the problem of finding the best data stored in a repository in a very general scenario in which, as in the above example: (i) preferences may not match the level of detail of the available data, (ii) there may be conflicts between different preferences, and (iii) known preferences can imply others. Specifically, unlike previous approaches that have only tackled the problem of mapping preferences to data (see, e.g., [20]), we formally investigate the two main principles that need to be taken into account in this context: *specificity* and *transitivity*. Specificity is a fundamental tool for resolving conflicts between preferences by giving precedence to the most specific ones, as it is natural in practical applications. For instance, in our example, the specific preference for Amarone over

white wines counts more than the generic preference for white wines over red ones. The specificity principle is indeed a pillar of non-monotonic reasoning, where a conclusion derived from a more specific antecedent overrides a conflicting inference based on a less specific antecedent [13]. On the other hand, transitivity, besides being a natural property, is important also from a practical point of view, since non-transitive preferences might induce cycles, a fact that could make it impossible to identify the best solutions [6].

To tackle the problem of dealing with non-monotonic preferences, we rely on a natural extension of the relational model in which we just assume that *taxonomies*, represented by partial orders on values, are defined on some attribute domains [22]. Thus, for instance, in a geographical domain we can establish that the value *Italy* is more generic than the value *Rome*, since the former precedes the latter in the partial order. We then call *t-relations* (i.e., relations over taxonomies) standard relations involving attributes over these taxonomic domains.

We express preferences in this model in a declarative way, by means of first-order *preference formulas* specifying the conditions under which, in a *t-relation*, a tuple  $t_1$  is preferable to a tuple  $t_2$ . By taking advantage of the taxonomies defined over the domains, in a preference formula we can refer to values that are more generic than those occurring explicitly in a *t-relation* (e.g., the fact that we prefer white to red wines, as in Example 1.1). When evaluated over a *t-relation*  $r$ , a preference formula returns a *preference relation* that includes all the pairs of tuples  $(t_1, t_2)$  in  $r$  such that  $t_1$  is preferable to  $t_2$ . Since the input preference formula may not induce a preference relation enjoying both transitivity and specificity, such a formula then needs to be suitably rewritten. Eventually, the rewritten formula is used to select the best tuples in  $r$  by means of the *Best* operator, which filters out all the tuples that are strictly worse than some other tuple [9]. How this rewriting has to be performed is thus the main focus of this paper.

**Problem.** *To study, from both a theoretical and a practical point of view, to which extent the properties of transitivity and specificity can be obtained by suitable rewritings of the initial preference formula.*

We tackle the problem by introducing and formally investigating two operators that rewrite a preference formula: *T* to enforce transitivity and *S* to remove all conflicts between more generic and more specific preferences, thus attaining specificity. In order to try to guarantee both properties, one thus needs to use both operators. The first natural question that arises is whether the order in which they are applied is immaterial. Unfortunately, it turns out that these two operators do not commute. More so, even their repeated application can produce different results, inducing incomparable preference relations. This motivates us to explore the (infinite) space of possible *sequences* of such operators. Based on this analysis, we prove that it is indeed *impossible* to always guarantee at the same time transitivity of the obtained preference relation and a complete absence of conflicts therein, no matter the order in which *T* and *S* are considered and how many times they are applied. Intuitively, the removal of conflicts may compromise transitivity, whereas enforcement of transitivity may (re-)introduce conflicts. We also show that this impossibility result would persist even if one considered a more fine-grained *S* operator that removes conflicts one by one (instead of all at a time). In spite of this intrinsic limitation, we formally show that: (i) the set of all possible

sequences of operators can be reduced to a finite (and small) set, and (ii) there are only two sequences, which we call *minimal-transitive*, that guarantee transitivity and, at the same time, *minimize* residual conflicts between preferences. We also show that the application of the *Best* operator using the rewritten formulas obtained through the two minimal-transitive sequences can lead to very different results. However, in common practical cases, experimental evidence shows that one of the two sequences typically resolves more conflicts, thus returning a more refined set of best tuples.

In order to observe and assess the actual behavior of sequences of operators, we developed an engine for implementing our approach, which rewrites an input preference formula and evaluates it over *t-relations*. We conducted a number of experiments over both synthetic and real-world data and taxonomies in scenarios of different complexities, showing that: (i) the overhead incurred by the rewriting process is low for the considered sequences; (ii) the computation of the best results largely benefits from the minimization of conflicts between preferences, both in terms of execution time and cardinality of results; (iii) the adoption of an original heuristic sorting criterion based on taxonomic knowledge greatly reduces execution times.

In sum, the contributions of this paper are the following:

- a general framework that is able to express, in a logic-based language, preferences over relations with taxonomic domains, as illustrated in Section 2;
- two operators, presented in Section 3, that rewrite, within this framework, the input preferences so as to enforce the important properties of transitivity, which is required for the correctness of the result, and specificity, which solves possible conflicts among preferences;
- the formal investigation, illustrated in Section 4, of the combined and repeated application of these operators to an initial set of preferences;
- a technique based on an original heuristics, presented in Section 5, for selecting the best results associated with given sequences of operators, and the characterization of their differences;
- the experimentation of the overall approach over both synthetic and real-world data, showing its effectiveness and practical feasibility, as illustrated in Section 6.

Related works are reported in Section 7 whereas some conclusions are sketched in Section 8. Formal proofs of our results are omitted in the interest of space.

## 2 PRELIMINARIES

In this section, we introduce our data model, originating from [22], and a logic-based preference model, inspired by [6].

We remind that a *partial order*  $\leq$  on a domain  $V$  is a subset of  $V \times V$ , whose elements are denoted by  $v_1 \leq v_2$ , that is: 1) reflexive ( $v \leq v$  for all  $v \in V$ ), 2) antisymmetric (if  $v_1 \leq v_2$  and  $v_2 \leq v_1$  then  $v_1 = v_2$ ), and 3) transitive (if  $v_1 \leq v_2$  and  $v_2 \leq v_3$  then  $v_1 \leq v_3$ ). A set with a partial order is called a *poset*.

### 2.1 Data Model

We consider a simple extension of the relational model in which the values of an attribute can be arranged in a hierarchical *taxonomy*.

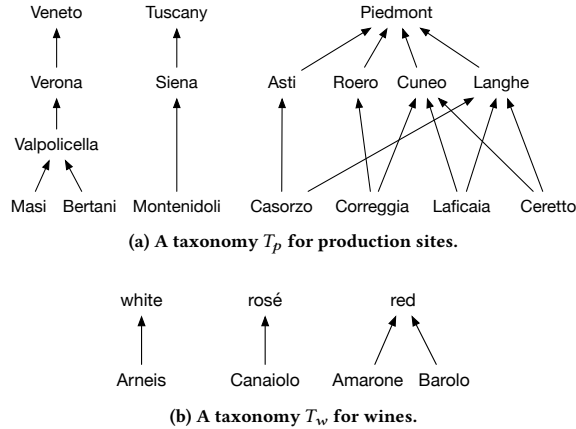


Figure 2: Taxonomies for the running example.

**Definition 2.1 (Taxonomy).** A taxonomy is a poset  $T = (V, \leq_V)$ , where  $V$  is a set of values and  $\leq_V$  is a partial order on  $V$ .

**Example 2.2.** A taxonomy relevant to our working example represents production sites at different levels of granularity. Considering Example 1.1, this taxonomy,  $T_p$ , shown in Figure 2a, includes values representing wineries (as minimal elements of the poset) as well as values representing provinces, wine regions, and regions of Italy. For instance, we can have values like Laficaia (a winery), Cuneo (a province), Langhe (a wine region) and Piedmont (a region of Italy), with Laficaia  $\leq_V$  Cuneo, Laficaia  $\leq_V$  Langhe, Laficaia  $\leq_V$  Piedmont, Cuneo  $\leq_V$  Piedmont, and Langhe  $\leq_V$  Piedmont. Additionally, Figure 2b shows a simple taxonomy  $T_w$  for wines, which associates each wine with a corresponding color. Finally, we assume a taxonomy  $T_y$  mapping production years before 2017 to aged and the other years to young.  $\square$

A *t-relation* is a standard relation of the relational model defined over a collection of taxonomies.

**Definition 2.3 (t-relation, t-schema, t-tuple).** A *t-schema* is a set  $S = \{A_1 : T_1, \dots, A_d : T_d\}$ , where each  $A_i$  is a distinct attribute name and each  $T_i = (V_i, \leq_{V_i})$  is a taxonomy. A *t-relation* over  $S$  is a set of tuples over  $S$  (*t-tuples*) mapping each  $A_i$  to a value in  $V_i$ . We denote by  $t[A_i]$  the restriction of a t-tuple  $t$  to the attribute  $A_i$ .

For the sake of simplicity, in the following we will not make any distinction between the name of an attribute of a t-relation and that of the corresponding taxonomy, when no ambiguities can arise. We observe that our model also accommodates “standard” attributes, in which the domain  $V$  is a set of flat values (i.e.,  $\leq_V$  is empty).

**Example 2.4.** A catalog of Italian wines can be represented by the t-schema  $S = \{\text{Wine} : T_w, \text{Winery} : T_p, \text{Year} : T_y\}$ . A possible t-relation over  $S$  is shown in Figure 1. Then we have  $b[\text{Year}] = 2014$  and  $e[\text{Wine}] = \text{Barolo}$ .  $\square$

## 2.2 Preference Model

Given a set of attribute-taxonomy pairs  $A_1 : T_1, \dots, A_d : T_d$ , in which  $A_1, \dots, A_d$  are all distinct, let  $\mathcal{T}$  denote the set of all possible t-tuples over any t-schema that can be defined using such pairs.

**Definition 2.5 (Preference relation).** A preference relation over the t-tuples in  $\mathcal{T}$  is a relation  $\geq$  on  $\mathcal{T} \times \mathcal{T}$ . Given two t-tuples  $t_1$  and  $t_2$  in  $\mathcal{T}$ , if  $t_1 \geq t_2$  then  $t_1$  is (weakly) preferable to  $t_2$ , also written as  $(t_1, t_2) \in \geq$ . If  $t_1 \geq t_2$  but  $t_2 \not\geq t_1$ , then  $t_1$  is strictly preferable to  $t_2$ , denoted by  $t_1 > t_2$ .

**Definition 2.6 (Incomparability and Indifference).** Given a preference relation on  $\mathcal{T}$  and a pair of t-tuples  $t_1$  and  $t_2$  in  $\mathcal{T}$ , if neither  $t_1 \geq t_2$  nor  $t_2 \geq t_1$ , then  $t_1$  and  $t_2$  are incomparable. When both  $t_1 \geq t_2$  and  $t_2 \geq t_1$  hold,  $t_1$  and  $t_2$  are indifferent, denoted by  $t_1 \approx t_2$ .

Notice that if  $\geq$  is transitive, then  $\approx$  is an equivalence relation (up to reflexivity) and  $>$  is a strict partial order (i.e., transitive and irreflexive). These properties do not hold, in the general case, when  $\geq$  is not transitive.

The transitivity of  $\geq$  implies that all the t-tuples involved in a cycle are indifferent to each other, thus the cycle vanishes when strict preferences are considered.

**Example 2.7.** Let us consider the t-relation in Figure 1 and assume that we have the cycle of preferences:  $a \geq b, b \geq c$ , and  $c \geq a$ . If  $\geq$  is transitive then we also have  $a \geq c$  (from  $a \geq b$  and  $b \geq c$ ),  $b \geq a$  (from  $b \geq c$  and  $c \geq a$ ) and  $c \geq b$  (from  $c \geq a$  and  $a \geq b$ ). Then, since  $a \approx b, b \approx c$ , and  $a \approx c$ , no cycle is present in  $>$ .  $\square$

Given a set of t-tuples  $r \subseteq \mathcal{T}$ , the “best” t-tuples in  $r$  according to the preference relation  $\geq$  can be selected by means of the *Best* operator  $\beta$  [9], which returns the t-tuples  $t_1$  of  $r$  such that there is no other t-tuple  $t_2$  in  $r$  that is strictly preferable to  $t_1$ .

**Definition 2.8 (Best operator).** Given a t-relation  $r$  and a preference relation  $\geq$  on the t-tuples in  $r$ , the best operator  $\beta$  is defined as follows:  $\beta_{>}(r) = \{t_1 \in r \mid \nexists t_2 \in r, t_2 > t_1\}$ .

When  $>$  is a strict partial order,  $\beta_{>}(r)$  is not empty for any non-empty t-relation  $r$ . We remind that, if  $>_1$  and  $>_2$  are such that  $>_1 \supseteq >_2$  then  $\beta_{>_2}(r) \subseteq \beta_{>_1}(r)$  holds for all  $r$  [6].

**Example 2.9.** Let us consider the t-relation in Figure 1 and assume that:  $b \geq a, a \geq f, b \geq f, b \geq d, c \geq e, e \geq c$ . It follows that:  $b > a, a > f, b > f, b > d$  (since the opposite does not hold for those four preferences), but  $c \approx e$  (since both  $c \geq e$  and  $e \geq c$ ). Then, we have  $\beta_{>}(r) = \{b, c, e\}$ .  $\square$

For expressing preferences we consider a logic-based language, in which  $t_1 \geq t_2$  iff they satisfy the first-order preference formula  $F(t_1, t_2): t_1 \geq t_2 \Leftrightarrow F(t_1, t_2)$ . Thus, when considering strict preferences we have:

$$t_1 > t_2 \Leftrightarrow F(t_1, t_2) \wedge \neg F(t_2, t_1). \quad (1)$$

As in [6], we only consider *intrinsic preference formulas* (ipf’s), i.e., first-order formulas in which only built-in predicates are present and quantifiers are omitted, as in Datalog. Predicates have either the form  $(x[A_i] \leq_{V_i} v)$  or  $(x[A_i] \not\leq_{V_i} v)$ , where  $A_i$  is an attribute defined over taxonomy  $T_i = (V_i, \leq_{V_i})$ ,  $x$  is a t-tuple variable over t-schemas including  $A_i$ , and  $v$  is a value in  $V_i$ . The predicate  $(x[A_i] \leq_{V_i} v)$  (resp.  $(x[A_i] \not\leq_{V_i} v)$ ) holds for a t-tuple  $t$  if  $(t[A_i] \leq_{V_i} v)$  (resp.  $(t[A_i] \not\leq_{V_i} v)$ ) holds. For convenience, we get rid of  $\neg$  as needed by transforming  $\leq_{V_i}$  into  $\not\leq_{V_i}$  and vice versa.

For the sake of generality, we consider that formula  $F$  consists of a set of preference statements, where each statement  $P_i$  is in

Disjunctive Normal Form (DNF), each disjunct of  $P_i$  being termed a *preference clause*,  $C_{i,j}$ :

$$P_i(x, y) = \bigvee_{j=1}^{m_i} C_{i,j}(x, y)$$

and where each clause  $C_{i,j}$  is a conjunction of predicates. We assume that each clause  $C_{i,j}$  is non-contradictory, i.e.,  $\exists t_1, t_2 \in \mathcal{T}$  such that  $C_{i,j}(t_1, t_2)$  is true. When a statement consists of a single clause we use the two terms “clause” and “statement” interchangeably.

A formula  $F$  is a disjunction of  $n \geq 1$  preference statements:

$$F(x, y) = \bigvee_{i=1}^n P_i(x, y).$$

*Example 2.10.* The preferences informally stated in Example 1.1 can be expressed by the formula

$$F(x, y) = P_1(x, y) \vee P_2(x, y) \vee P_3(x, y) \vee P_4(x, y)$$

where the 4 preference statements, in which we use  $\leq$  in place of  $\leq_{V_i}$  to improve readability, are:

$$\begin{aligned} P_1(x, y) &= (x[\text{Wine}] \leq \text{white}) \wedge (y[\text{Wine}] \leq \text{red}) \\ P_2(x, y) &= (x[\text{Wine}] \leq \text{Amarone}) \wedge (y[\text{Wine}] \leq \text{white}) \\ P_3(x, y) &= (x[\text{Winery}] \leq \text{Siena}) \wedge (y[\text{Winery}] \leq \text{Asti}) \\ P_4(x, y) &= (x[\text{Winery}] \leq \text{Langhe}) \wedge (x[\text{Year}] \leq \text{aged}) \wedge \\ &\quad (y[\text{Winery}] \leq \text{Langhe}) \wedge (y[\text{Year}] \leq \text{young}) \end{aligned}$$

The above statements, when evaluated over the t-tuples in Figure 1, yield the following preferences, written as pairs of t-tuples in  $\geq$  (for the sake of clarity, for each preference we also show the statement used to derive it):

$$\begin{aligned} P_1 : & (a, b), (a, c), (a, e), (f, b), (f, c), (f, e) \\ P_2 : & (b, a), (b, f), (c, a), (c, f) \\ P_4 : & (e, f) \end{aligned}$$

Notice that  $P_3$  yields no preference, since there is no wine from Asti’s province in the t-relation in Figure 1.  $\square$

In the rest of the paper, with the aim to simplify the notation, preference statements in the examples will be written with a compact syntax, by omitting variables and attributes’ names, and separating with  $\geq$  the “better” part from the “worse” part. For instance, the above statement  $P_4$  will be written as:

$$P_4 = \text{Langhe} \wedge \text{aged} \geq \text{Langhe} \wedge \text{young}.$$

### 3 OPERATIONS ON PREFERENCES

In this section we introduce two operators that can be applied to a preference relation, postponing to the next section the detailed analysis of the possible ways in which they can be combined. The two operators are: Transitive closure (T) and Specificity-based refinement (S). Let  $\geq$  denote the initial preference relation; the resulting relation is indicated  $\geq_{\text{T}}$  for T and  $\geq_{\text{S}}$  for S. Multiple application of operators, e.g., first T and then S, leads to the relation  $(\geq_{\text{T}})_{\text{S}}$ , which we compactly denote as  $\geq_{\text{TS}}$ . In general, for any *sequence*  $X \in \{\text{T}, \text{S}\}^*$ ,  $\geq_X$  is the preference relation obtained from the initial preference relation  $\geq$  by applying the operators in the order in which they appear in  $X$ . Notice that  $\geq_{\varepsilon} = \geq$ , where  $\varepsilon$  denotes the empty sequence.

We describe the behavior of the two operators by means of suitable rewritings of a preference formula. Given a sequence  $X$  of operators, and an initial (input) formula  $F(x, y)$  inducing the preference relation  $\geq$ ,  $F^X(x, y)$  denotes the rewriting of  $F$  that accounts for the application of the  $X$  sequence, thus yielding  $\geq_X$ .

#### 3.1 Transitive Closure

Transitivity of  $\geq$ , and consequently of  $>$ , is a basic requirement of any sound preference-based system. If  $\geq$  is not transitive then  $>$  might contain cycles, a fact that could easily lead either to empty or non-stable results, as the following example shows.

*Example 3.1.* Consider the t-tuples in Figure 3, in which both Sbarbata and Molinara are rosé wines and Vogadori is a winery in the Valpolicella wine region.

Wine	Winery	Year	
Arneis	Correggia	2019	$g$
Barolo	Laficaia	2014	$h$
Sbarbata	Laficaia	2019	$\ell$
Molinara	Vogadori	2014	$m$

Figure 3: A set of wines for Example 3.1.

From the preference statements in Example 2.10, we have  $g \geq h$  (through  $P_1$ ) and  $h \geq \ell$  (through clause  $P_4$ ). However,  $g \not\geq \ell$ . Assume now two additional preference statements

$$\begin{aligned} P_{\alpha} &= \text{rosé} \wedge \text{young} \geq \text{rosé} \wedge \text{aged}, \\ P_{\beta} &= \text{Valpolicella} \geq \text{Roero}, \end{aligned}$$

which, respectively, induce preferences  $\ell \geq m$  and  $m \geq g$ . Overall, since no other preferences hold, we have the non-transitive cycle of strict preferences  $g > h, h > \ell, \ell > m$  and  $m > g$ . So, for a t-relation  $r = \{g, h, \ell, m\}$ , we have  $\beta_{\geq}(r) = \emptyset$ .

Consider now  $r' = \{g, h, \ell\}$ , for which  $\beta_{\geq}(r') = \{g\}$ , and  $r'' = \{g, \ell, m\}$ , for which  $\beta_{\geq}(r'') = \{\ell\}$ . Although both  $r'$  and  $r''$  contain  $g$  and  $\ell$ , the choice of which of these t-tuples is better than the other depends on the presence of other t-tuples (like  $h$  and  $m$ ), thus making the result of the  $\beta$  operator unstable.  $\square$

The transitive closure operator, denoted  $\text{T}$ , given an input preference relation  $\geq_X$  yields the preference relation  $\geq_{X\text{T}}$ . We remind that, as observed in Section 2.2, the transitivity of  $\geq_{X\text{T}}$  entails that of  $>_{X\text{T}}$ . The transitive closure  $F^{X\text{T}}$  of an ipf  $F^X$  with  $n$  statements  $P_1, \dots, P_n$  is still a finite ipf that can be computed via Algorithm 1, along the lines described in [6]. For the sake of conciseness, given a preference clause  $C(x, y)$ , we denote by  $C^b(x)$  (resp.  $C^w(y)$ ) the part of  $C(x, y)$  given by the conjunction of the predicates involving variable  $x$  (resp.  $y$ ). Notice that  $C(x, y) = C^b(x) \wedge C^w(y)$  holds.

In the main loop of the algorithm (lines (2)–(9)) we test the possibility of transitively combining two preference statements at a time (line (4)), by considering each of their clauses (line (6)). Since clauses are assumed to be non-contradictory, the test at line (7), which can also be written as  $C_m^b(t_1) \wedge C_m^w(t_2) \wedge C_q^b(t_2) \wedge C_q^w(t_3)$ , reduces to checking if  $C_m^w(t_2) \wedge C_q^b(t_2)$  is satisfiable in  $\mathcal{T}$ . This can be done by checking whether no contradictory pair of predicates occurs in  $C_m^w(t_2) \wedge C_q^b(t_2)$ . In particular, two predicates of the form

---

**Algorithm 1:** T operator: Transitive closure of  $F^X$ .

---

Input: formula  $F^X = P_1 \vee \dots \vee P_n$ , taxonomies  $T_1, \dots, T_d$ .  
Output:  $F^{XT}$ , the transitive closure of  $F^X$ .

- (1)  $F^{XT} := F^X$
- (2) **repeat**
- (3)    $newPref := \text{false}$
- (4)   **for each** ordered pair  $(P_i, P_j)$ ,  $P_i$  in  $F^{XT}$ ,  $P_j$  in  $F^X$
- (5)      $P := \text{empty}$
- (6)     **for each** ordered pair  $(C_m, C_q)$ ,  $C_m$  in  $P_i$ ,  $C_q$  in  $P_j$
- (7)       **if**  $\exists t_1, t_2, t_3 \in \mathcal{T}$  s.t.  $C_m(t_1, t_2) \wedge C_q(t_2, t_3) = \text{true}$   
          **then**  $P := P \vee (C_m^b(x) \wedge C_q^w(y))$
- (8)     **if**  $P \neq \text{empty}$  **then**  $F^{XT} := F^{XT} \vee P$ ,  $newPref := \text{true}$
- (9) **until**  $newPref = \text{false}$
- (10) **return**  $F^{XT}$

---

$(x[A_i] \leq_{V_i} v_1)$  and  $(x[A_i] \leq_{V_i} v_2)$ , over the same attribute  $A_i$  and using the same variable  $x$ , are contradictory if values  $v_1$  and  $v_2$  are different and have no common descendant in the taxonomy  $V_i$  (Section 6 further discusses how to check the existence of a common descendant). If the predicates are of the form  $(x[A_i] \leq_{V_i} v_1)$  and  $(x[A_i] \not\leq_{V_i} v_2)$ , then they are contradictory in case there is a path from  $v_1$  to  $v_2$  in  $V_i$  (or  $v_1 = v_2$ ).

The fact that the transitive closure is computed with respect to the (possibly infinite) domain  $\mathcal{T}$  of the t-tuples, and *not* with respect to a (finite) t-relation  $r$  of t-tuples, is quite standard for preference relations (see e.g., [6]), and has the advantage of yielding a relation  $\geq_{XT}$  that does not change with  $r$  and avoiding the problems discussed in Example 3.1.

*Example 3.2.* Continuing with Example 2.10, the transitive closure of  $F$  is the formula  $F^T$  that, among others, adds the following statements to  $F$ :

$$\begin{aligned} P_5 &= \text{Amarone} \geq_T \text{red} \\ P_6 &= \text{Siena} \geq_T \text{Langhe} \wedge \text{young} \end{aligned}$$

Statement  $P_5(x, y)$  clearly follows from  $P_2(x, z)$  and  $P_1(z, y)$ . More interesting is statement  $P_6(x, y)$ , obtained from  $P_3(x, z)$  and  $P_4(z, y)$ . Since there exists at least one winery that is both in the Asti province and in the Langhe region (Casorzo is one of them), this allows  $P_3(x, z)$  and  $P_4(z, y)$  to be transitively combined. With reference to the t-tuples in Figure 1, we then have  $d \geq_T f$ .  $\square$

After applying the T operator, we simplify the formula as needed, and, in particular, we remove statements that are subsumed by other statements. Similarly, we also simplify statements by removing contradictory clauses and clauses subsumed within the same statement.

### 3.2 Specificity-based Refinement

The most intriguing of our operators is the *specificity-based refinement*  $S$ . As it is also apparent from Example 2.10, *conflicting preferences*, such as  $(a, b)$  and  $(b, a)$ , may hold. Although these preferences are compatible with the given definition of preference relation, we argue that some of these conflicts need to be resolved in order to derive a preference relation that better represents the stated user preferences. To this end we resort to a *specificity principle*, which we adapt from the one typically used in non-monotonic

reasoning to solve conflicts. According to such a principle, a conclusion derived from a more specific antecedent overrides a conflicting (defeasible) inference based on a less specific antecedent, that is, more specific information overrides more generic information.

*Example 3.3.* In our working example, we have a generic preference for white wines over red wines. With no contradiction with the generic preference, we might have a *more specific* preference stating that a bottle of Amarone (a red wine) is superior to a bottle of Arneis (a white wine). In this case, the more specific preference would entail, among others,  $b \geq a$ ; yet, because of the more generic preference for white wines, we also have  $a \geq b$ , thus  $a$  and  $b$  become indifferent. However, giving the same importance to both preference statements contradicts the intuition, as the more specific preference should take precedence over the more generic one.  $\square$

The specificity principle we adopt for analyzing conflicting preferences is based on the *extension* of preferences statements, i.e., on the set of pairs of t-tuples in  $\mathcal{T}$  for which a statement is true.

*Definition 3.4 (Specificity principle).* Let  $\geq_X$  be a preference relation, and let  $F^X$  be the corresponding formula. Let  $P_i$  and  $P_j$  be two preference statements in  $F^X$ . We say that  $P_i$  is *more specific* than  $P_j$  if, for any pair of t-tuples  $t_1, t_2 \in \mathcal{T}$  such that  $P_i(t_1, t_2)$  is true, then  $P_j(t_2, t_1)$  is also true, and the opposite does not hold.

From Definition 3.4 we can immediately determine how a less specific statement has to be rewritten so as to solve conflicts.

**LEMMA 3.5.** *A preference statement  $P_i(x, y)$  is more specific than  $P_j(y, x)$  iff  $P_i(x, y)$  implies  $P_j(y, x)$  (written  $P_i(x, y) \rightarrow P_j(y, x)$ ) and the opposite does not hold.<sup>1</sup> If  $P_j(y, x)$  is replaced by  $P'_j(y, x) = P_j(y, x) \wedge \neg P_i(x, y)$ , then  $P_i$  and  $P'_j$  do not induce any conflicting preferences.*

Checking whether  $P_i(x, y)$  implies  $P_j(y, x)$  amounts to checking whether  $P_i(x, y) \wedge \neg P_j(y, x)$  is false, i.e., every clause in the resulting formula is contradictory (contradictions can be checked as described for T).

The  $S$  operator, whose behavior is defined by Algorithm 2, removes from the preferences induced by a formula  $F^X$  all those that are conflicting and less specific.

Notice that, after a first analysis of the existing implications among the statements (line (4)) and the rewriting of the implied statements (line (5)), the analysis needs to be repeated, since new implications might arise. For instance, let  $F^X = P_1 \vee P_2 \vee P_3$ , with  $P_1(y, x) \rightarrow P_2(x, y)$  being the only implication. After rewriting  $P_2(x, y)$  into  $P'_2(x, y) = P_2(x, y) \wedge \neg P_1(y, x)$ , it might be the case that  $P'_2(x, y) \rightarrow P_3(y, x)$ , thus  $P_3$  needs to be rewritten.

Although multiple rounds might be needed, Algorithm 2 is guaranteed to terminate. Indeed, if  $P_i(x, y) \rightarrow P_j(y, x)$ , and  $P_j(y, x)$  is consequently replaced by  $P'_j(y, x) = P_j(y, x) \wedge \neg P_i(x, y)$ , the two statements  $P_i$  and  $P'_j$ , as well as their possible further rewritings, have *disjoint* extensions, and therefore will not interact anymore in the rewriting process. Since the number of statements is finite, so is the number of rewritings, which ensures that the algorithm will eventually stop.

<sup>1</sup>The hypothesis that  $P_j(y, x)$  does not imply  $P_i(x, y)$  excludes the case of *opposite preference statements* (e.g., white is better than red, and red is better than white), to which the  $S$  operator clearly does not apply.

---

**Algorithm 2:** S operator: Specificity-based refinement of  $F^X$ .

---

Input: formula  $F^X = P_1 \vee \dots \vee P_n$ , taxonomies  $T_1, \dots, T_d$ .

Output:  $F^{XS}$ , the specificity-based refinement of  $F^X$ .

```
(1) repeat
(2)    $newRound := \text{false}$ 
(3)   for each statement  $P_i$ 
(4)      $Impl(P_i) := \{P_j | P_j(y, x) \rightarrow P_i(x, y) \wedge P_i(x, y) \not\rightarrow P_j(y, x)\}$ 
(5)     if  $Impl(P_i) \neq \emptyset$  then
            $newRound := \text{true}, P'_i := P_i$ 
           for each  $P_j \in Impl(P_i)$ 
              $P'_i(x, y) := P'_i(x, y) \wedge \neg P_j(y, x)$ 
(6)   if  $newRound$  then  $P_i := P'_i, i = 1, \dots, n$ 
(7) until  $newRound = \text{false}$ 
(8) return  $F^{XS} = P_i \vee \dots \vee P_n$ 
```

---

Here too, we simplify the formula resulting from the rewritings according to the same principles used for the T operator.

*Example 3.6.* Continuing with Example 3.2, the application of the S operator amounts to rewriting formula  $F^T$  by replacing the clause  $P_1(x, y)$  with  $P_1(x, y) \wedge \neg P_2(y, x)$ , since  $P_2(y, x) \rightarrow P_1(x, y)$ . This, after distributing  $\neg$  over the two predicates in  $P_2$  and simplifying, leads to the new clause:

$$P_7 = \text{white} \geq_{TS} \text{red} \wedge \neg \text{Amarone}.$$

The preferences that were derived from  $P_1$  can be seen in Example 2.10; we repeat them for the sake of clarity:

$$P_1 : (a, b), (a, c), (a, e), (f, b), (f, c), (f, e).$$

Among them,  $(a, b)$ ,  $(a, c)$ ,  $(f, b)$ , and  $(f, c)$  do not satisfy  $P_7(x, y)$ , since both  $b$  and  $c$  refer to Amarone. Thus,  $P_7 : (a, e), (f, e)$ .  $\square$

It is relevant to observe that the application of the S operator always leads to smaller (i.e., cleaner) results. For instance, considering t-relation  $r$  in Figure 1 and input preference statements  $P_1$  and  $P_2$  from Example 2.10, we have  $\beta_{>}(r) = \{a, b, c, d, f\}$ , whereas  $\beta_{>S}(r) = \{b, c, d\}$ .

**LEMMA 3.7.** For any t-relation  $r$  and any preference relation  $\geq_X$  we have  $\beta_{>XS}(r) \subseteq \beta_{>X}(r)$ .

## 4 MINIMAL-TRANSITIVE SEQUENCES

In this section we analyze the effect of performing the operations described in the previous section, and prove some fundamental properties of the obtained preference relations. After introducing the basic properties and main desiderata in Section 4.1, we explore the space of possible sequences in Section 4.2 and, as a major result, we show that, out of infinitely many candidates, only a finite number of sequences needs to be considered. Finally, in Section 4.3 we identify the only two sequences that meet all our requirements.

### 4.1 Basic properties

In order to clarify the relationships between the results of the different operations, we introduce the notions of equivalence and containment between sequences of operators.

*Definition 4.1 (Equivalence and containment).* Let  $X, Y \in \{T, S\}^*$ ;  $X$  is contained in  $Y$ , denoted  $X \sqsubseteq Y$ , if for every initial preference

relation  $\geq$ ,  $\geq_X \subseteq \geq_Y$ ;  $X$  and  $Y$  are *equivalent*, denoted  $X \equiv Y$ , if both  $X \sqsubseteq Y$  and  $Y \sqsubseteq X$ .

Among the basic properties of our operators, we observe that T and S are idempotent, T is monotone and cannot remove preferences, while S cannot add preferences. In addition, the preference relation obtained after applying T on the initial preference relation  $\geq$  is maximal, in that it includes all other relations obtained from  $\geq$  by applying T and S in any way.

**THEOREM 4.2.** Let  $X, Y \in \{T, S\}^*$ , with  $X \sqsubseteq Y$ . Then:

$$XTT \equiv XT \quad XSS \equiv XS \quad \text{idempotence} \quad (2)$$

$$XT \sqsubseteq YT \quad \text{monotonicity} \quad (3)$$

$$X \sqsubseteq XT \quad XS \sqsubseteq X \quad \text{inflation / deflation} \quad (4)$$

$$X \sqsubseteq T \quad \text{maximality} \quad (5)$$

We now focus on those sequences, that we call complete, that include both T and S, since their corresponding operations are both part of our requirements. In particular, transitivity of the obtained *strict* preference relation is at the core of the computation of the Best ( $\beta$ ) operator, as shown in Example 3.3. To this end, we characterize as transitive those sequences that entail such a transitivity.

*Definition 4.3 (Complete and transitive sequence).* A sequence  $X \in \{T, S\}^*$  is *complete* if  $X$  contains both T and S;  $X$  is *transitive* if, for every initial preference relation  $\geq$ ,  $\geq_X$  is transitive.

Eventually, our goal is to drop conflicting and less specific preferences while preserving transitivity. To this end, we add minimality with respect to  $\sqsubseteq$  as a desideratum. In particular, we want to determine the so-called *minimal-transitive* sequences, i.e., those that are minimal among the transitive sequences. As it turns out, all such sequences are also complete.

*Definition 4.4 (Minimal-transitive sequence).* Let  $\Sigma$  be a set of sequences;  $X \in \Sigma$  is *minimal* in  $\Sigma$  if there exists no other sequence  $Y \in \Sigma, Y \neq X$  such that  $Y \sqsubseteq X$ . A *minimal-transitive* sequence is a sequence that is *minimal* in the set of *transitive* sequences.

### 4.2 The space of possible sequences

We now chart the space of possible sequences so as to understand the interplay between completeness, transitivity and minimality.

We start by observing that any sequence with consecutive repetitions of the same operator is equivalent, through idempotence, to a shorter sequence with no such repetitions; for instance, TSS is equivalent to TS. Since sequences with repetitions play no significant role in our analysis, we shall henceforth disregard them.

Clearly, every sequence is contained in T, due to its maximality. Other containment relationships follow from inflation of T and deflation of S. Further relationships come from the following result, stating that adding ST (i.e., removing conflicts and then transitively closing the resulting preference formula) to a sequence ending with T cannot introduce any new preference.

**LEMMA 4.5.** Let  $X \in \{T, S\}^*$ . Then  $XTST \sqsubseteq XT$ .

Lemma 4.5 induces two chains of inclusions, namely:

$$\dots \sqsubseteq TSTST \sqsubseteq TST \sqsubseteq T \quad (6)$$

$$\dots \sqsubseteq STSTST \sqsubseteq STST \sqsubseteq ST. \quad (7)$$

In addition to that, the following result seems to suggest that the longer sequences in the above chains are preferable, since they lead to larger sets of strict preferences ( $>$ ), which, as was observed in Section 2.2, correspond to smaller (i.e., cleaner) results for the Best  $\beta$  operator.


**PROPOSITION 4.6.** *Let  $X \in \{T, S\}^*$ . Then, for any initial preference relation  $\geq$ , we have  $>_{XT} \subseteq >_{XTST}$ .*

There are, evidently, infinitely many sequences in the chains (6) and (7) and, more generally, in  $\{T, S\}^*$ . However, for any given initial preference formula, a counting argument on the number of formulas obtainable through the operators allows us to restrict to only a finite amount of sequences. Moreover, it turns out that the repeated application of a TS suffix does not change the semantics of a sequence, so we can apply it just once and disregard all other sequences.

**LEMMA 4.7.** *Let  $X \in \{T, S\}^*$ . Then  $XTS \equiv XTSTS$ .*

An immediate consequence of this result is that, through elimination of consecutively repeated operators via idempotence and of consecutively repeated TS sub-sequences via Lemma 4.7, we can restrict our attention to a set of just eight sequences, because any sequence is equivalent to one of those.

**THEOREM 4.8.** *Let  $X \in \{T, S\}^*$ . Then  $\exists Y \in \{\varepsilon, T, S, TS, ST, TST, STS, STST\}$  such that  $X \equiv Y$ .*



**Figure 4: A transitively reduced graph showing containment between sequences. Dashed border for incomplete sequences; grey background for non-transitive sequences; blue background for minimal-transitive sequences. All containment relationships are strict.**

Figure 4 shows a (transitively reduced) graph whose nodes correspond to the eight sequences mentioned in Theorem 4.8 and whose arcs indicate containment. Thanks to the theorem, we have narrowed the space of possible sequences to analyze from an infinite set  $\{T, S\}^*$  to just these eight sequences.

### 4.3 Minimality and transitivity

Now that we have restricted our scope to a small set of representative sequences, we can discuss minimality and transitivity in detail, so as to eventually detect minimal-transitive sequences. Note that incomplete sequences can be immediately ruled out of our analysis: it is straightforward to show that S is not transitive, T is not minimal (it is indeed maximal) and  $\varepsilon$  is neither.

**Minimality.** Generally, any complete sequence not ending with S is non-minimal, in that it may contain conflicting preferences (possibly introduced by T) that turn out to be in contrast with other, more specific preferences. We exemplify this on ST. In the examples to follow, we shall refer to t-tuples with a single attribute on a single taxonomy about time.

**Example 4.9.** Let  $F$  consist of  $P_1$  and the more specific  $P_2$ :

$$P_1 = \text{autumn} \geq \text{sep}, \quad P_2 = \text{sep10} \geq \text{oct10}.$$

By specificity, in  $F^S$ ,  $P_1$  is replaced by the statement  $P_3$  consisting of two clauses (grouped by curly brackets):

$$P_3 = \left\{ \begin{array}{l} \text{autumn} \geq \text{sep} \wedge \neg \text{sep10} \\ \text{autumn} \wedge \neg \text{oct10} \geq \text{sep} \end{array} \right.$$

In  $F^{ST}$ , the clauses in  $P_3$  transitively combine into  $P_1$  again, since, e.g., the value  $\text{sep30}$  is below  $\text{sep}$  but not  $\text{sep10}$  and below  $\text{autumn}$  but not  $\text{oct10}$ ; therefore  $\text{oct10} \geq_{ST} \text{sep10}$  holds. However, in  $F^{STS}$ ,  $P_1$  is again replaced by  $P_3$ , so that  $\text{oct10} \not\geq_{STS} \text{sep10}$ , which shows that ST is not minimal.  $\square$

All the containments indicated in Figure 4 are strict, as can be shown through constructions similar to that of Example 4.9, so no sequence ending with T is minimal in  $\{T, S\}^*$ .

**LEMMA 4.10.** *Let  $X \in \{T, S\}^*$ . Then XT is not minimal in  $\{T, S\}^*$ .*

**Transitivity.** Transitivity is certainly achieved for any sequence ending with T: any relation  $\geq_{XT}$  is transitive by construction, which entails transitivity of  $>_{XT}$ . However, the following result shows that, in the general case, no sequence ending with S is transitive.

**LEMMA 4.11.** *Let  $X \in \{T, S\}^*$ . Then XS is not transitive.*

**Minimal-transitive sequences.** As a consequence of Lemmas 4.10 and 4.11, we can state a major result, showing that transitivity and minimality in  $\{T, S\}^*$  are mutually exclusive.

**THEOREM 4.12.** *No sequence is both transitive and minimal in  $\{T, S\}^*$ .*

Moreover, we observe that all complete sequences starting with S are incomparable (i.e., containment does not hold in any direction) with those starting with T, as stated below (also refer to Figure 4).

**THEOREM 4.13.** *Let  $X \in \{TS, TST\}$  and  $Y \in \{ST, STS, STST\}$ . Then  $X \not\subseteq Y$  and  $Y \not\subseteq X$ .*

This property is shown for TS and STS in the next example.

**Example 4.14.** Let  $F$  consist of the following statements:

$$P_1 = \text{summer} \geq \text{spring}, \quad P_2 = \text{jul21} \geq \text{jun}, \quad P_3 = \text{may} \geq \text{jul}.$$

Then  $F^T$  includes  $P_1, P_3$  and the following 4 statements:

$$P_4 = \text{summer} \geq \text{jul} \quad (P_1 + P_3), \quad P_5 = \text{may} \geq \text{spring} \quad (P_3 + P_1), \\ P_6 = \text{may} \geq \text{jun} \quad (P_3 + P_2), \quad P_7 = \text{summer} \geq \text{jun} \quad (P_1 + P_6),$$

while  $P_2$  is removed, as it is redundant with respect to  $P_7$ . No statement in  $F^T$  is more specific than  $P_4$ , so  $P_4$  is in  $F^{TS}$  and, e.g.,  $\text{jul21} \geq_{TS} \text{jul10}$  holds. In  $F^S$ , instead,  $P_1$  (less specific than  $P_3$ ) is replaced by

$$P_8 = \left\{ \begin{array}{l} \text{summer} \geq \text{spring} \wedge \neg \text{may} \\ \text{summer} \wedge \neg \text{jul} \geq \text{spring} \end{array} \right.$$

So, now, by combining  $P_8$  (instead of  $P_1$ ) and  $P_3$ , in  $F^{\text{ST}}$  we do not obtain  $P_4$  and then  $\text{jul21} \not\geq_{\text{STS}} \text{jul10}$ . With this,  $\text{TS} \not\subseteq \text{STS}$ .

For the other non-containment, consider that, in  $F^{\text{ST}}$ ,  $P_2$  combines with  $P_8$  into the following statement:

$$P_9 = \text{jul21} \geq \text{spring},$$

so that  $\text{jul21} \geq_{\text{ST}} \text{spring}$  may hold. No statement in  $F^{\text{ST}}$  is more specific than  $P_9$ , so  $\text{jul21} \geq_{\text{STS}} \text{spring}$  may also hold. Instead,  $\text{jul21} \not\geq_{\text{TS}} \text{spring}$ , since  $F^{\text{TS}}$  is as  $F^{\text{T}}$ , but with  $P_8$  instead of  $P_1$ . Therefore  $\text{STS} \not\subseteq \text{TS}$ .  $\square$

The notion of minimal-transitive sequence captures the fact that transitivity cannot be waived, since we are indeed looking for the minimal sequences among those that are both complete and transitive. Only three sequences are both complete and transitive: ST, TST and STST, the first of which contains the last one and is therefore not minimal. The remaining two sequences are transitive, incomparable by Theorem 4.13, and, therefore, minimal in the set of complete and transitive sequences, i.e., TST and STST are minimal-transitive sequences.

**THEOREM 4.15.** *The only minimal-transitive sequences are TST and STST.*

As observed in Theorem 4.13, the sequence STST, which removes less specific conflicting preferences before computing the first transitive closure, does not in general entail a set of preferences included in those induced by TST. We shall further characterize the behavior of these two sequences in Section 5, from a theoretical point of view, and, experimentally, in Section 6.

We also observe that the result of Theorem 4.12 is inherent and that no finer granularity in the interleaving of T and S (e.g., by making S resolve one conflict at a time instead of all together) would remove this limitation: as Example 4.9 shows, the presence of one *single* preference ( $\text{oct10} \geq \text{sep10}$ ) is sufficient to make the relation transitive but not minimal, and its absence to make it minimal but not transitive. The atomicity of this conflict is enough to conclude that it is unavoidable and that no method whatsoever (not just those based on the T and S operators) could solve it.

## 5 COMPUTING THE BEST RESULTS

### 5.1 Worst-case difference between TST and STST

As shown in Theorem 4.13, the two minimal-transitive semantics are incomparable, thus there will be t-relations  $r$  and initial preference relations  $\geq$  for which the best results delivered by the two semantics will differ. A legitimate question is: How much can these results be different? In order to answer this question we consider the *maximum* value of the cardinality of the difference of the results delivered by the two minimal-transitive semantics over all t-relations with  $n$  t-tuples and over all input preference relations  $\geq$ . To this end, let us define, for any two sequences  $X$  and  $Y$ :

$$\text{DIFFBEST}(X, Y, n) = \max_{\geq, |r|=n} \{|\beta_{>_X}(r) - \beta_{>_Y}(r)|\}$$

as the worst-case difference in the results delivered by  $X$  with respect to those due to  $Y$ , for any given cardinality of the target t-relation  $r$ . We can prove the following:

**THEOREM 5.1.** *We have both  $\text{DIFFBEST}(\text{TST}, \text{STST}, n) = \Theta(n)$  and  $\text{DIFFBEST}(\text{STST}, \text{TST}, n) = \Theta(n)$ .*

From a practical point of view, Theorem 5.1 shows that there is no all-seasons minimal-transitive semantics. Furthermore, there can be cases (used in the proof of the theorem) in which the number of best results from any of the two semantics is comparable to  $n$ , whereas the other semantics returns  $O(1)$  t-tuples. In Section 6 we will experimentally investigate the actual difference of results delivered by the two minimal-transitive semantics.

### 5.2 A heuristics for computing the best results

In order to compute the best results according to the formula  $F^X$  we adopt the well-known BNL algorithmic pattern [2]. We remind that BNL-like algorithms have worst-case quadratic complexity, although in practice they behave almost linearly [11]. Remind also that, according to Equation (1), given a preference formula  $F^X(x, y)$  defining *weak* preferences, the corresponding *strict* preferences are those induced by the formula  $F^X_{>}(x, y) = F^X(x, y) \wedge \neg F^X(y, x)$ .

The t-tuples that do not match any side of any clause in the preference formula correspond to those objects that the formula does not talk about and that can, thus, be considered irrelevant. As recognized in the germane literature [10], such objects are of little interest and, in the following, we shall therefore compute  $\beta$  so as to only include *relevant* t-tuples (i.e., those that satisfy either side of at least one clause of  $F^X$ , thus of  $F^X_{>}$  as well).

The algorithm keeps the current best t-tuples in the *Best* set. When a new t-tuple  $t$  is read, and  $t$  is found to be relevant,  $t$  is compared to the tuples in *Best*. Given  $t' \in \text{Best}$ , if  $t' >_X t$  then  $t$  is immediately discarded. Conversely,  $t$  is added to *Best* and all t-tuples  $t' \in \text{Best}$  such that  $t >_X t'$  are removed from the *Best* set. Eventually, we have  $\beta_{>_X}(r) = \text{Best}$ .

An improvement to this basic scheme is to pre-sort the t-relation so that the t-tuples matching the left side of a clause and corresponding to lower-level values in the taxonomies come first. The rationale is that lower-level values are likely associated with a smaller amount of t-tuples, so that a smaller *Best* partial result can be found before scanning large amounts of data. Furthermore, such t-tuples are likely to be preferred to many others, in particular when specificity is a concern. More in detail, we scan  $r$  and, for each relevant t-tuple  $t$  (irrelevant t-tuples are immediately discarded) we compute a *height index*,  $hi(t)$ , as follows: For any clause  $C(x, y) = C^b(x) \wedge C^w(y)$  such that  $C^b(t)$  holds, we consider the “height” of each value  $v$  occurring in the clause, computed as the distance of  $v$  from the leaves of its taxonomy.<sup>2</sup> Then, the minimum height over predicates in  $C^b(t)$  and over all other matching clauses is used as value of  $hi(t)$ , and t-tuples are sorted by increasing height index values; conventionally, when  $t$  matches no clauses, we set  $hi(t) = \infty$ .

*Example 5.2.* Consider a formula  $F = P_1 \vee P_2$ , where  $P_1$  and  $P_2$  are taken from Example 2.10. Then, we have  $F^{\text{STST}} = P_3 \vee P_2 \vee P_4$ , where  $P_3 = \text{white} \geq \text{red} \wedge \neg \text{Amarone}$  and  $P_4 = \text{Amarone} \geq \text{red} \wedge \neg \text{Amarone}$ . Out of the t-tuples in Figure 1,  $d$  is irrelevant, while  $e$  does not match any clause, and thus  $hi(e) = \infty$ . Wines  $a$  and  $f$  match white in  $P_3$ , which has height 1 (see Figure 2b), so  $hi(a) = hi(f) = 1$ , while  $b$  and  $c$  match Amarone in both  $P_2$  and  $P_4$ , with  $hi(b) = hi(c) = 0$ . Thus,  $b$  and  $c$  come before  $a$  and  $f$  in the ordering, and  $e$  is last.  $\square$

<sup>2</sup>In case of non-functional taxonomies, in which a node may have more than one parent, we take the minimum distance.



## 6 EXPERIMENTS

In this section, we consider from a practical point of view the sequences of operators T, TST, and STST, discussed in the previous sections. The main goals of the experimental study are: (i) to understand the impact of the rewriting process on the overall query execution time and how this depends on the specific sequence at hand; (ii) to assess the effect of minimal-transitive sequences on (the cardinality of) the results of the  $\beta$  (Best) operator; (iii) to compare overall execution times incurred by minimal-transitive sequences with respect to baseline strategies in which either no rewriting occurs or only the transitive closure of the input formula is computed; (iv) to measure the effects of the heuristics presented in Section 5. In particular, we study how efficiency and effectiveness are affected by taxonomy’s size and morphology, dataset size, number of attributes, and number and type of preferences. The relevant parameters used in our analysis are summarized in Table 1.

In summary, we show that: the rewriting due to the minimal-transitive sequences TST and STST incurs a low overhead across all tested scenarios; such sequences are effective both in reducing the cardinality of  $\beta$  and in achieving substantial speedup with respect to baseline strategies, and that the speedup is further incremented when adopting our heuristics.

**Table 1: Operating parameters for performance evaluation (defaults, when available, are in bold).**

Full name	Tested value
Taxonomy’s depth $\delta$	2, 3, 4, 5, <b>6</b> , 7, 8, 9, 10
Taxonomy’s fanout $f$	2, 3, 4, 5, <b>6</b> , 7, 8, 9, 10
Synthetic taxonomy’s kind	<b>regular</b> , random, scale-free
# of attributes $d$	<b>1</b> , 2, 3, 4, 5
# of input clauses $c$	2, 4, 6, 8, 10
# of maximal values	2, 4, <b>6</b> , 8, 10
Type of preferences	<b>conflicting</b> , contextual
Dataset size $N$	<b>10K</b> , 50K, 100K, 500K, 1M

### 6.1 Taxonomies, datasets, and preferences

We use two families of taxonomies: synthetic and real taxonomies.

We run our tests on three kinds of synthetic taxonomies: regular, random and scale-free. A regular taxonomy is generated as a forest of  $f$  (“fanout”) rooted trees consisting of  $\delta$  levels and  $f$  children for each internal node. The total number of nodes is therefore  $\sum_{i=1}^{\delta} f^i$ , i.e.,  $\frac{f(f^{\delta}-1)}{f-1}$ . A random taxonomy is generated as in the previous case, but the fanout of each node is Poisson distributed with an average of  $f$ . The default values for  $f$  and  $\delta$  are chosen to match the size of the real taxonomies used in the experiments (15-20K nodes). Finally, a scale-free taxonomy targets the same number of nodes, but following a power-law distribution (which is observed to be a recurrent structure, e.g., in the Semantic Web; see [26, 27]), for the fanout. Scale-free taxonomies generated this way (with reasonable exponents around 2.7) are typically very deep (between 30 and 60 levels). All synthetic taxonomies are functional by construction, i.e., every node has exactly one parent. Synthetic datasets of various sizes are generated by drawing values uniformly at random from a different taxonomy for each attribute.

We adopt two real taxonomies and datasets: flipkart<sup>3</sup> and UsedCars<sup>4</sup>. The former lists product categories of various kinds and consists of 15,236 nodes (of which 12,483 leaf categories) and 15,465 arcs spread throughout 10 levels. This taxonomy is non-functional, in that there exist nodes with more than one parent, i.e., some products belong to more than one category. Product info is available as a t-relation consisting of 19,673 t-tuples that also include original price, discounted price, and user rating, rendered here as attributes associated with a “flat” taxonomy with three values (e.g., “high”, “medium”, “low”). UsedCars features a large collection of used vehicles for sale consisting, after cleaning, of 232,470 t-tuples including, among others, price range (as a flat taxonomy) and model. Models are organized in a functional taxonomy, with 14,588 nodes and 14,540 leaves, over three levels (besides model name and make, we obtained country information via the Car Models List DB<sup>5</sup>).

The study of the best taxonomy representation in the general case is orthogonal with respect to the problems we study in this paper (see, e.g., [14]). However, given the taxonomies we deal with, it is convenient to precompute all paths in order to speed up all taxonomy-based computations, e.g., establishing when a value is more specific than another.

For our experiments, we consider two common types of preferences, discussed below: conflicting preferences and contextual preferences. We omit the results concerning other common types of preferences, as their behavior is not essentially different.

A pair of *conflicting preference* statements has the following form:

$$P_1 = v_1 \geq v_2, \quad P_2 = v'_2 \geq v_1,$$

where  $v_1$  and  $v_2$  are maximal values (i.e., tree roots) of the same taxonomy  $T_i$  and  $v'_2 \leq v_i v_2$ . Clearly,  $P_2$  is more specific than  $P_1$ .

The second kind of preferences, used for experiments on multi-attribute relations, are pairs of conflicting *contextual preferences*, i.e., conflicting preferences applied to one attribute, in which the other attributes are used to establish a sort of “context” of applicability. A pair of contextual preferences is of the following form:

$$\begin{aligned} P_1 &= v_1^{(1)} \wedge v_2^{(2)} \wedge \dots \wedge v^{(d)} \geq v_2^{(1)} \wedge v^{(2)} \wedge \dots \wedge v^{(d)}, \\ P_2 &= v_2'^{(1)} \wedge v^{(2)} \wedge \dots \wedge v^{(d)} \geq v_1^{(1)} \wedge v^{(2)} \wedge \dots \wedge v^{(d)}, \end{aligned}$$

where the  $(i)$  superscript denotes values from taxonomy  $T_i$ ,  $v_1^{(1)}$  and  $v_2^{(1)}$  are maximal in  $T_1$ , and  $v_2'^{(1)} \leq v_i v_2^{(1)}$ . Note that, when there are  $d = 1$  attributes, this is just a pair of conflicting preferences. For real data, flat taxonomies are used for context attributes. An example of contextual preference is given by statement  $P_4$  in Example 2.10.

### 6.2 Results: computation of the output formula

In order to assess feasibility of the computation of the preference formula resulting after applying a sequence of operators, we report the corresponding execution time averaged out over 100 different runs (as measured on a machine sporting a 2,3 GHz 8-Core Intel Core i9 with 32 GB of RAM).

Our first experiments test the impact of the characteristics of the taxonomy in the case of synthetic taxonomies and one pair of conflicting preferences. For regular taxonomies, computing  $F^T$

<sup>3</sup><https://www.flipkart.com>

<sup>4</sup><https://www.kaggle.com/austinreese/craigslist-carstrucks-data>

<sup>5</sup><https://www.teoalida.com/cardatabase/car-models-list>

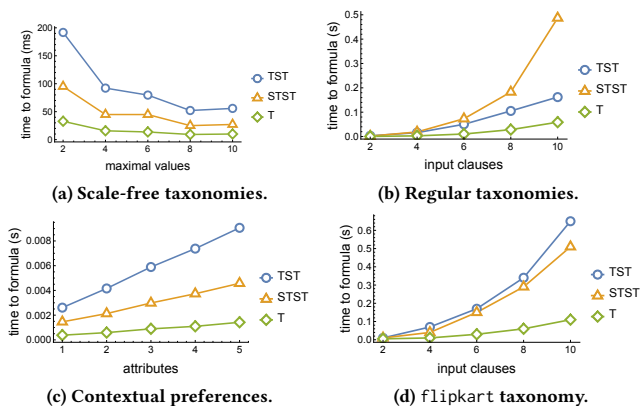


Figure 5: Time for computing the formula: various settings.

(0.5ms on average) is generally faster than computing  $F^{\text{STST}}$  (1.5ms) and  $F^{\text{TST}}$  (2.7ms) and neither  $f$  nor  $\delta$  affect the computation time significantly. Similar times are obtained with random taxonomies. With scale-free taxonomies the same relative costs are kept, but times are slightly higher, due to the much deeper structure, and tend to decrease as the the number of maximal nodes increases, as shown in Figure 5a; still, all times are well under 0.2s and thus negligible with respect to the time required for computation of  $\beta$ , as will be shown in Section 6.3.

Figure 5b shows that the time for computing the formula grows with the number of input clauses, with times always below 0.5s.

For a multi-attribute scenario, Figure 5c shows the behavior with contextual preferences as the number of attributes varies. The resulting formula is always computed in less than 0.01s; times slightly grow as the number of attributes grows, but remain low.

We now turn to the case of real taxonomies. With UsedCars, which is functional, results are very similar to those obtained with synthetic taxonomies and thus not shown here in the interest of space. We then test on flipkart, which is non-functional, the case of conflicting preferences as the number of input clauses  $c$  varies. This has an impact on the overhead for determining redundancies in formulas and for checking clause satisfiability when computing T. Indeed, both require checking whether two values  $v_1$  and  $v_2$  have a common descendant in the taxonomy, which is immediate in the case of functional taxonomies, as it suffices to check whether there is a path from  $v_1$  to  $v_2$  or vice versa. However, for non-functional taxonomies this check may require extracting all descendants of  $v_1$  and  $v_2$ , which may be expensive for large taxonomies, especially when  $v_1$  and  $v_2$  are maximal values. Yet, for taxonomies in which only few nodes have more than one parent (like flipkart, with 170 such nodes), it is convenient to keep track of those nodes at taxonomy load time; with this, we can check the existence of a common descendant between  $v_1$  and  $v_2$  by checking whether there is a path to both from one of those nodes (if they are not the descendant of one another). As Figure 5d shows, the times measured with the flipkart taxonomy are only slightly higher than with synthetic taxonomies (and always sub-second).




Figure 6: Computing  $\beta$  with default parameter values.

### 6.3 Results: computation of $\beta$

As discussed in Section 5, we restrict the  $\beta$  operator to act only on *relevant* t-tuples. In the same vein, we shall only consider preferences inducing a non-empty set of relevant t-tuples.

With conflicting preferences and default parameter values on regular taxonomies, the amount of relevant t-tuples is roughly 40% of the size of a synthetic dataset. Figure 6a shows that both T and  $\epsilon$  retain about half of the relevant t-tuples (which is both the average and the median value we obtained), while TST and STST retain less than 2% in the median case (the average value goes up to 20% due to runs with unfocused input formulas referring to values not in the dataset). This is reflected in the computation times, shown in Figure 6b, which are consistently around 24s for T and 10s for  $\epsilon$ , but nearly two orders of magnitude smaller in the median case for TST and STST. With both scale-free and random taxonomies, the amount of relevant t-tuples varies much more (with an average still around 40%), but times are on average one order of magnitude smaller for TST and STST than for T, with results for the latter covering almost the entire dataset due to the lack of conflict resolution.

We observe that the application of T alone corresponds to the work performed by preference evaluation methods that only aim at guaranteeing transitivity, e.g., [6, 12, 17], which are therefore outperformed by our approach. The inability of T to deal with conflicting preferences, thus generating many indifferent t-tuples, which in turn induce (very) large result sets, indeed applies to all our scenarios. Similar observations apply to  $\epsilon$  (i.e., the empty sequence, corresponding to the input formula), which represents the action of works on preference evaluation using no rewriting whatsoever, such as [5, 20]. Additionally, the results obtained via  $\epsilon$  would be totally unreliable, due to lack of transitivity (see Example 3.1). We thus refrain from considering T and  $\epsilon$  from now on.

We now analyze the cost incurred by the computation of  $\beta$  as we deviate from standard parameter values. In the case of contextual preferences, adding context makes the  $\beta$  set leaner and, thus, easier to compute, so that times are under 1s already with two attributes. As usual, STST is slightly quicker to compute, since it gives rise to a smaller formula (although its strict version coincides with that of TST, and thus their cardinalities coincide).

As already visible in Figure 6, random preference formulas may fail to represent a meaningful specification of preferences, thus leading to very large result sets. For this reason, we disregard such formulas and, in particular, in the next experiments we only retain those “good runs” in which either TST or STST produce less than 2% of the t-tuples in the dataset. Figure 7a shows how the cardinality of  $\beta$  varies, under these hypotheses and default parameter values, as the number of input clauses  $c$  varies, thus confirming that STST typically leads to a smaller result than TST.




Figure 7: Synthetic datasets: conflicting preferences, varying the number of input clauses  $c$  (only good runs).




Figure 8: Synthetic datasets: varying the dataset size  $N$  (only good runs, size in logarithmic scale).




Figure 9: flipkart: conflicting preferences, varying the number of input clauses  $c$  (only good runs).

We now consider the heuristics described in Section 5, which sorts the  $t$ -relation according to increasing height index values. Figure 7b compares times obtained with the heuristic sort strategy (marked with an  $H$  subscript) to those obtained with no heuristics as the number of input clauses  $c$  varies. The sort takes between 3% and 10% of the total time spent for computing  $\beta$ , yet the use of the proposed heuristics largely outperforms standard executions, with times never exceeding 2s; without the heuristics, times diverge to well over 100s, on average, in the more expensive scenarios.

Having ascertained the suitability of the heuristic sort, we demonstrate its scalability with the experiment shown in Figure 8, which shows a linear trend for times as the size  $N$  of the dataset varies, while cardinalities tend to grow logarithmically.

The trends shown with synthetic data are confirmed with real data on flipkart. Figure 9a shows that the cardinality of  $\beta$  typically grows as the number of input clauses grows. Consequently, Figure 9b shows times slightly growing with the number of input clauses, but always under 2.1s. The case of contextual preferences is shown in Figure 10, where times decrease as the number of attributes grows, since the number of relevant  $t$ -tuples decreases with the number of applied contexts. For the same reason, the cardinality




Figure 10: flipkart: contextual preferences, varying the number of attributes  $d$  (only good runs).

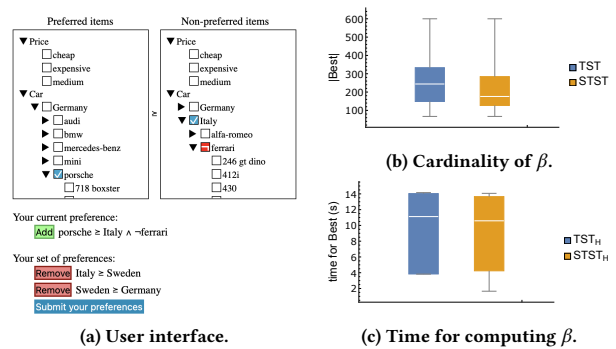


Figure 11: User interface and experiments on UsedCars.

of  $\beta$  is higher with 2 or 3 attributes than with 4; however, with only 1 attribute (and thus no context) the cardinality is the lowest, since the  $t$ -tuples satisfying the most specific preference are not filtered out by contexts.

For experiments on UsedCars, we collected preferences from a set of 107 users by means of a Web interface allowing the specification of statements in the simplified notation presented in Section 2 through an expandable tree-view of the taxonomy (see Figure 11a). After instructing users on how to specify preferences (even conflicting ones), we observed an average of 3.4 statements (from 2 to 9) per query and as many as 78% of cases of conflicts. Figure 11b shows box plots representing the distributions of cardinalities of  $\beta$  obtained with user-defined preferences, which confirms that STST tends to produce slightly smaller results than TST. We observe that such cardinalities, typically corresponding to the number of cars available for a specific model and price range, are very low with respect to the dataset size (and could be further reduced if filters based on other criteria, such as mileage, were applied). Execution times (Figure 11c) are, on average, below 10s for both sequences, and thus overall acceptable and comparable with the measurements obtained with similarly sized synthetic data (Figure 8b).

## 7 RELATED WORKS AND DISCUSSION

In spite of the many works on the use of *qualitative* preferences for querying databases (see, e.g., [25]), only a few address the issues arising when attributes' domains exhibit a hierarchical structure.

Preferences in OLAP systems are considered in [12], where an algebraic language, based on that in [17], is adopted. Preferences on

attributes are only of an *absolute* type, stating which are the most (resp. least) preferred values at a given “level” of a dimensional attribute. Preferences are then propagated along levels, with no concern for the combination of preferences, less so conflicting ones.

Lukasiewicz et al. [20] extend the Datalog+/- ontological language with qualitative preferences, yet they do not address the problems arising from conflicting preferences. In a subsequent work [21], the authors assume that, besides the order generated by the preferences, another linear order exists, originating from probabilistic scores attached to specific objects. Since the two orders may conflict, ad-hoc operators for compromising among the two orders are introduced and evaluated. Although [21] considers conflicts, these are *not* among preferences and their solutions are not applicable to the scenario we consider in this paper.

To the best of our knowledge, no other work addresses the exact same issues we tackle here. Yet, Section 6.3 has shown how existing methods (those that just enforce transitivity as well as works on preference evaluation using no rewritings) would be unsuitable to meet the goals we set in this paper.

The specificity principle on which we have based the definition of our  $S$  operator follows a long-standing tradition in the AI and KR fields, in which conflicts arising from contradictory evidences (antecedents) are solved by means of non-monotonic reasoning. However, in this context, the issue of inheritance of properties, which can be dealt with in different ways according to the adopted reasoning theory (see, e.g., [13]), leads to problems that are quite different from those we have considered in this paper.

The need to address conflicts arising from preferences was also observed in [8]. The framework proposed there allows for a restricted form of taxonomies (with all values organized into distinct, named levels) and hints at an ad hoc procedure with very limited support for conflict resolution; the focus of [8] is, however, on the downward propagation of preferences.

A kind of specificity principle was also considered in [9], albeit on a different preference model (using strict rather than weak preferences) and a different scenario, in which preferences are to be combined across different *contexts*. In that work, given two conflicting preferences, e.g.,  $a > b$ , which is valid in a context  $c$ , and  $b > a$  valid in context  $c'$ , if context  $c$  is more specific than  $c'$  then  $a > b$  wins and  $b > a$  is discarded. Thus, specificity considered in [9] concerns contexts, whereas, in the present paper, specificity has to do with preference statements that involve values at different levels of detail in the taxonomies. Conflicts in [9] are at the level of a single pair of objects (since no language for specifying preferences was considered there), whereas in the present work we deal with conflicts between *preference statements*, which in general involve many pairs of objects - a fact that requires a solution incomparable with those adopted in [9].

A line of research that is only apparently related to ours concerns the problem of propagating preferences across the nodes/terms of an ontology, see, e.g., [3, 4, 23]. Given “interest scores” attached to some terms, these works focus on (numerical) methods to combine and propagate such scores to “similar” terms in the ontology.

A definitely relevant issue, orthogonal to our focus and thus outside the scope of this paper, is that of *preference elicitation*. This problem has been thoroughly studied in various fields, such as Recommender Systems, decision making, marketing, and behavioral

economics, with remarkable recent attention on *relative preferences*, either expressed with pairwise comparisons or inferred from absolute preferences [15, 16].

Common methods to solve conflicts among preferences are based on the use of operators, the most well-known being Pareto and Prioritized composition [6, 9, 17]. Given a conflict between  $a$  and  $b$  originating from two different preference statements, Pareto composition just drops both preferences  $a > b$  and  $b > a$ . Conversely, Prioritized composition *a priori* assumes that one of the two statements is more important than the other, and then solves the conflict by retaining the corresponding preference. We have no such a-priori notion of priority, which might be hard to define in practice; rather, we rely on a definition of specificity that *dynamically* determines if a statement takes precedence over another depending on the available taxonomies.

Many algorithms have been devised to answer preference queries, although most of them work only for numerical attributes [18]. Among the algorithms that can be applied to arbitrary strict partial orders  $>$ , BNL [2] is undoubtedly the most well-known among those that compute the result sets by means of dominance tests. Improvements to the BNL logic, such as those found in the SFS [7] and SaLSa [1] algorithms, require the input relation to be topologically sorted, which in these algorithms is based on the presence of numerical attributes. A different approach, pioneered in [10], avoids (most of the) dominance tests by partitioning the domain of (relevant) tuples into a set of equivalence classes, where each class includes all and only those tuples whose values are the best for a subset of the input preference statements. For instance, a statement like  $(A_i = v) \geq (A_i = v')$  induces two equivalence classes, the first including all tuples with value  $v$  for attribute  $A_i$ , and the second those with value  $v'$ . For each equivalence class a different SQL query is then executed, until it is guaranteed that no further optimal tuples exist. However, since the number of equivalence classes is exponential in the number of input statements, this approach cannot be adopted in our framework, in which the rewritten formula to be evaluated, due to the transitive closure operator, can well contain tens of statements.

## 8 CONCLUSIONS

In this paper we have tackled the problem of finding the best elements from a repository on the basis of preferences referring to values that are more generic than the underlying data and may involve conflicts. To this aim, we have introduced and formally investigated two operators for enforcing, in a given collection of preferences, the properties of specificity, which can solve conflicts, and transitivity, which guarantees the soundness of the final result. We have then characterized the limitations that can arise from their combination and identified the best ways in which they can be used together. We have finally proposed a technique based on an original heuristics for selecting the best results associated with given sequences of operators and shown, with a number of experiments over both synthetic and real-world datasets, the effectiveness and practical feasibility of the overall approach. Future work includes extending our framework to more general scenarios in which domain values are connected by ontological relationships, as is the case in Ontology-Based Data Access [28].

## REFERENCES

- [1] Ilaria Bartolini, Paolo Ciaccia, and Marco Patella. 2008. Efficient sort-based skyline evaluation. *ACM Trans. Database Syst.* 33, 4 (2008), 31:1–31:49. <https://doi.org/10.1145/1412331.1412343>
- [2] Stephan Börzsönyi, Donald Kossmann, and Konrad Stocker. 2001. The Skyline Operator. In *Proceedings of the 17th International Conference on Data Engineering, April 2-6, 2001, Heidelberg, Germany*. 421–430. <https://doi.org/10.1109/ICDE.2001.914855>
- [3] Federica Cena, Silvia Likavec, and Francesco Osborne. 2013. Anisotropic propagation of user interests in ontology-based user models. *Inf. Sci.* 250 (2013), 40–60. <https://doi.org/10.1016/j.ins.2013.07.006>
- [4] Gil Chamiel and Maurice Pagnucco. 2008. Exploiting Ontological Structure for Complex Preference Assembly. In *AI 2008: Advances in Artificial Intelligence, 21st Australasian Joint Conference on Artificial Intelligence, Auckland, New Zealand, December 1-5, 2008. Proceedings (Lecture Notes in Computer Science)*, Wayne Wobcke and Mengjie Zhang (Eds.), Vol. 5360. Springer, 86–92. [https://doi.org/10.1007/978-3-540-89378-3\\_9](https://doi.org/10.1007/978-3-540-89378-3_9)
- [5] Chee Yong Chan, H. V. Jagadish, Kian-Lee Tan, Anthony K. H. Tung, and Zhenjie Zhang. 2006. Finding k-dominant skylines in high dimensional space. In *Proceedings of the ACM SIGMOD International Conference on Management of Data, Chicago, Illinois, USA, June 27-29, 2006*, Surajit Chaudhuri, Vagelis Hristidis, and Neoklis Polyzotis (Eds.). ACM, 503–514. <https://doi.org/10.1145/1142473.1142530>
- [6] Jan Chomicki. 2003. Preference formulas in relational queries. *ACM Trans. Database Syst.* 28, 4 (2003), 427–466. <https://doi.org/10.1145/958942.958946>
- [7] Jan Chomicki, Parke Godfrey, Jarek Gryz, and Dongming Liang. 2003. Skyline with Presorting. In *Proceedings of the 19th International Conference on Data Engineering, March 5-8, 2003, Bangalore, India*, Umeshwar Dayal, Krithi Ramamritham, and T. M. Vijayarajan (Eds.). IEEE Computer Society, 717–719. <https://doi.org/10.1109/ICDE.2003.1260846>
- [8] Paolo Ciaccia, Davide Martinenghi, and Riccardo Torlone. 2019. Finding Preferred Objects with Taxonomies. In *Conceptual Modeling - 38th International Conference, ER 2019, Salvador, Brazil, November 4-7, 2019, Proceedings (Lecture Notes in Computer Science)*, Alberto H. F. Laender, Barbara Pernici, Ee-Peng Lim, and José Palazzo M. de Oliveira (Eds.), Vol. 11788. Springer, 397–411. [https://doi.org/10.1007/978-3-030-33223-5\\_33](https://doi.org/10.1007/978-3-030-33223-5_33)
- [9] Paolo Ciaccia, Davide Martinenghi, and Riccardo Torlone. 2020. Foundations of Context-aware Preference Propagation. *J. ACM* 67, 1 (2020), 4:1–4:43. <https://doi.org/10.1145/3375713>
- [10] Periklis Georgiadis, Ioannis Kapantaidakis, Vassilis Christophides, Elhadji Mamadou Nguer, and Nicolas Spyrtatos. 2008. Efficient Rewriting Algorithms for Preference Queries. In *Proceedings of the 24th International Conference on Data Engineering, ICDE 2008, April 7-12, 2008, Cancún, Mexico*, Gustavo Alonso, José A. Blakeley, and Arbee L. P. Chen (Eds.). IEEE Computer Society, 1101–1110. <https://doi.org/10.1109/ICDE.2008.4497519>
- [11] Parke Godfrey, Ryan Shipley, and Jarek Gryz. 2007. Algorithms and analyses for maximal vector computation. *VLDB J.* 16, 1 (2007), 5–28. <https://doi.org/10.1007/s00778-006-0029-7>
- [12] Matteo Goffarelli, Stefano Rizzi, and Paolo Biondi. 2011. myOLAP: An Approach to Express and Evaluate OLAP Preferences. *IEEE Trans. Knowl. Data Eng.* 23, 7 (2011), 1050–1064. <https://doi.org/10.1109/TKDE.2010.196>
- [13] John F. Horty. 1994. Some Direct Theories of Nonmonotonic Inheritance. In *Handbook of Logic in Artificial Intelligence and Logic Programming (Vol. 3): Nonmonotonic Reasoning and Uncertain Reasoning*. Oxford University Press, Inc., USA, 111–187.
- [14] Ruoming Jin, Yang Xiang, Ning Ruan, and Haixun Wang. 2008. Efficiently answering reachability queries on very large directed graphs. In *Proceedings of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2008, Vancouver, BC, Canada, June 10-12, 2008*, Jason Tsong-Li Wang (Ed.). ACM, 595–608. <https://doi.org/10.1145/1376616.1376677>
- [15] Saikishore Kalloori, Tianyu Li, and Francesco Ricci. 2019. Item Recommendation by Combining Relative and Absolute Feedback Data. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2019, Paris, France, July 21-25, 2019*, Benjamin Piwowarski, Max Chevalier, Éric Gaussier, Yoelle Maarek, Jian-Yun Nie, and Falk Scholer (Eds.). ACM, 933–936. <https://doi.org/10.1145/3331184.3331295>
- [16] Saikishore Kalloori, Francesco Ricci, and Rosella Gennari. 2018. Eliciting pairwise preferences in recommender systems. In *Proceedings of the 12th ACM Conference on Recommender Systems, RecSys 2018, Vancouver, BC, Canada, October 2-7, 2018*, Sole Pera, Michael D. Ekstrand, Xavier Amatriain, and John O'Donovan (Eds.). ACM, 329–337. <https://doi.org/10.1145/3240323.3240364>
- [17] Werner Kießling. 2002. Foundations of Preferences in Database Systems. In *VLDB 2002, Proceedings of 28th International Conference on Very Large Data Bases, August 20-23, 2002, Hong Kong, China*. 311–322. <http://www.vldb.org/conf/2002/S09P04.pdf>
- [18] Georgia Koutrika and Yannis E. Ioannidis. 2004. Personalization of Queries in Database Systems. In *Proceedings of the 20th International Conference on Data Engineering, ICDE 2004, 30 March - 2 April 2004, Boston, MA, USA*. 597–608. <https://doi.org/10.1109/ICDE.2004.1320030>
- [19] G.S. Linoff and M.J.A. Berry. 2001. *Mining the web: Transforming Customer Data into Customer Value*. John Wiley & Sons, New York.
- [20] Thomas Lukasiewicz, Maria Vanina Martinez, and Gerardo Ignacio Simari. 2013. 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*, Francesca Rossi (Ed.). IJCAI/AAAI, 1017–1023. <http://www.aaai.org/ocs/index.php/IJCAI/IJCAI13/paper/view/6505>
- [21] Thomas Lukasiewicz, Maria Vanina Martinez, Gerardo I. Simari, and Oana Tifrea-Marcuska. 2015. Preference-Based Query Answering in Probabilistic Datalog+/- Ontologies. *J. Data Semant.* 4, 2 (2015), 81–101. <https://doi.org/10.1007/s13740-014-0040-x>
- [22] Davide Martinenghi and Riccardo Torlone. 2014. Taxonomy-based relaxation of query answering in relational databases. *VLDB J.* 23, 5 (2014), 747–769. <https://doi.org/10.1007/s00778-013-0350-x>
- [23] Miriam Martinez-Garcia, Aida Valls, and Antonio Moreno. 2019. Inferring preferences in ontology-based recommender systems using WOWA. *J. Intell. Inf. Syst.* 52, 2 (2019), 393–423. <https://doi.org/10.1007/s10844-018-0532-5>
- [24] Francesco Ricci, Lior Rokach, and Bracha Shapira. 2011. Introduction to Recommender Systems Handbook. In *Recommender Systems Handbook*, Francesco Ricci, Lior Rokach, Bracha Shapira, and Paul B. Kantor (Eds.). Springer, 1–35. [https://doi.org/10.1007/978-0-387-85820-3\\_1](https://doi.org/10.1007/978-0-387-85820-3_1)
- [25] Kostas Stefanidis, Georgia Koutrika, and Evaggelia Pitoura. 2011. A survey on representation, composition and application of preferences in database systems. *ACM Trans. Database Syst.* 36, 3 (2011), 19:1–19:45. <https://doi.org/10.1145/2000824.2000829>
- [26] Yannis Theoharis, George Georgakopoulos, and Vassilis Christophides. 2012. PowerGEN: A power-law based generator of RDFs schemas. *Inf. Syst.* 37, 4 (2012), 306–319. <https://doi.org/10.1016/j.is.2011.09.005>
- [27] Yannis Theoharis, Yannis Tzitzikas, Dimitris Kotzinos, and Vassilis Christophides. 2008. On Graph Features of Semantic Web Schemas. *IEEE Trans. Knowl. Data Eng.* 20, 5 (2008), 692–702. <https://doi.org/10.1109/TKDE.2007.190735>
- [28] Guohui Xiao, Diego Calvanese, Roman Kontchakov, Domenico Lembo, Antonella Poggi, Riccardo Rosati, and Michael Zakharyashev. 2018. Ontology-Based Data Access: A Survey. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, Jérôme Lang (Ed.). ijcai.org, 5511–5519. <https://doi.org/10.24963/ijcai.2018/777>