Personalizing Queries Based on Networks of Composite Preferences

GEORGIA KOUTRIKA Stanford University, USA and YANNIS IOANNIDIS University of Athens, Greece

People's preferences are expressed at varying levels of granularity and detail as a result of partial or imperfect knowledge. One may have some preference for a general class of entities, for example, liking comedies, and another one for a fine-grained, specific class, such as disliking recent thrillers with Al Pacino. In this article, we are interested in capturing such complex, multi-granular preferences for personalizing database queries and in studying their impact on query results. We organize the collection of one's preferences in a *preference network* (a directed acyclic graph), where each node refers to a subclass of the entities that its parent refers to, and whenever they both apply, more specific preferences override more generic ones. We study query personalization based on networks of preferences and provide efficient algorithms for identifying relevant preferences, modifying queries accordingly, and processing personalized queries. Finally, we present results of both synthetic and real-user experiments, which: (a) demonstrate the efficiency of our algorithms, (b) provide insight as to the appropriateness of the proposed preference model, and (c) show the benefits of query personalization based on composite preferences compared to simpler preference representations.

Categories and Subject Descriptors: H.1 [Information Systems]: Models and Principles; H.2.8 [Database Management]: Database Applications

General Terms: Theory, Algorithms, Experimentation

Additional Key Words and Phrases: Preference modeling, preference networks, personalization

ACM Reference Format:

Koutrika, G. and Ioannidis, Y. 2010. Personalizing queries based on networks of composite preferences. ACM Trans. Datab. Syst. 35, 2, Article 13 (April 2010), 50 pages. DOI = 10.1145/1735886.1735892 http://doi.acm.org/10.1145/1735886.1735892

DOI 10.1145/1735886.1735892 http://doi.acm.org/10.1145/1735886.1735892

Authors' addresses: G. Koutrika, Department of Computer Science, Stanford University, Gates Building, Room 425, 353 Serra Mall, Stanford, CA 94305-9040; email: koutrika@stanford.edu; Y. Ioannidis, University of Athens, Greece.

Permission to make digital or hard copies of part or all of this work for personal or classroom use isgranted without fee provided that copies are not made or distributed for profit or commercial-advantage and that copies show this notice on the first page or initial screen of a displayalong with the full citation. Copyrights for components of this work owned by others than ACMmust be honored. Abstracting with credit is permitted. To copy otherwise, to republish,to post on servers, to redistribute to lists, or to use any component of this work in otherworks requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, orpermissions@acm.org. © 2010 ACM 0362-5915/2010/04-ART13 \$10.00

13:2 • G. Koutrika and Y. Ioannidis

1. INTRODUCTION

The World Wide Web has enabled people with varied goals and characteristics to access an ever-growing amount of information. The emergence of hand-held electronic devices, such as palmtops and cellular phones, has increased the possibilities for information access from anywhere, anytime. In this context, finding the "right information" is a difficult problem accentuated by the growing rate at which new information becomes available and the heterogeneity of people searching for information. To cope with this problem, personalization methods, such as recommendations (e.g., [Balabanovic and Shoham 1997; Das et al. 2007; Linden et al. 2003]) and query personalization (e.g., [Koutrika and Ioannidis 2004; Liu et al. 2004; Pitkow et al. 2002]), have been proposed both by industry and academia.

Query personalization is based on the observation that "different users may find different answers relevant when searching" because of different preferences [Pitkow et al. 2002]. A query is personalized by applying related user preferences stored in the user's profile and changing the order and possibly the size of results. Which preferences from a user's profile are related to a query and how they affect the final answer can be determined based on the query, the profile, and the personalization logic applied. Consider as an example a query about comedies. A preference for movies directed by Woody Allen is related since this director has directed comedies. On the other hand, a preference for Andrei Tarkovsky is not related to the same query because this director has never directed any comedies, and, if this preference is conjunctively combined with the query, no results will be returned.

In this article, we are interested in modeling and using user preferences for query personalization in the context of databases. Our interest in databases stems from the fact that databases comprise an important part of the (deep) Web, as witnessed by the proliferation of database-driven Web sites (e.g., e-commerce and social bookmarking sites), where different notions of user information are very important.

Modeling preferences and ranking query results based on preferences hide many challenges. Existing works have studied various aspects of these problems, such as qualitative (e.g., [Chomicki 2003; Kiessling 2002]) versus quantitative preference formulations (e.g., [Agrawal and Wimmers 2000; Koutrika and Ioannidis 2004]), different types of preferences (e.g., likes and dislikes [Kiessling 2002; Koutrika and Ioannidis 2005b]), context and preference combinations for result ranking (e.g., [Holland and Kiessling 2004; Stefanidis et al. 2007]), and so forth. In this article, we study preference modeling at various levels of granularity and result ranking that is aware of and respects the preference importance as well as the natural relationships between multi-granular preferences.

Multi-granular, composite user preferences and ranking. In an ideal world, preferences could exist *for every object* in a domain of interest, for example, for every movie, for every book, etc. Such elaborate preferences would yield a perfectly fine-grained ranking of alternatives and highly personalized content. In another ideal world, one could express preferences *for general but well-defined*,

disjoint sets partitioning the objects of discourse. For example, one could have an explicit preference for movies directed before 1990 and another, explicit, preference for movies released after 1990. Unfortunately, neither of these two extremes is found in practice often. User preferences are typically incomplete and, furthermore, our knowledge of them is imperfect and partial. Different factors may contribute to this phenomenon, as we explain next.

People tend to have a mix of general and specific preferences. This mix may indicate lack of knowledge, lack of elaborate taste, or inability to identify those properties of objects that determine their preferences. For instance, one may have a general liking for adventures but a finer-grained taste only for a subset of them, such as those directed by S. Spielberg, which are characterized as favorites. On the other hand, one may have liked some particular books but may be unable to identify a common characteristic that made them attractive. Implicit collection of user preferences, for instance, by observing the user behavior in the system, suffers from similar problems. For instance, consider a system that builds user profiles from user ratings for movies to provide personalized content. Suppose that a user has given high ratings to all comedies she has seen so far. This knowledge allows the system to conclude that the user likes comedies but does not suffice to differentiate among comedies. Suppose that the user continues to interact with the system and provides some low scores for comedies starring Jack Nicholson. Then, the system may enrich its knowledge of this user's preferences by discriminating against comedies starring Jack Nicholson, which will be ranked lower than other comedies.

Article focus and contribution. Given that general and specific preferences can peacefully coexist, in this article, we study *query personalization based on multi-granular preferences*. We represent preferences as degrees of interest in query conditions following a similar philosophy to previous works [Koutrika and Ioannidis 2004; Stefanidis and Pitoura 2008]. The key difference from earlier works, however, is that we allow a profile to contain any mix of general and specific preferences, where the latter are explicitly stated rather than being implicitly calculated from the former.

For personalizing the results of a query, we adopt the personalization framework presented by Koutrika and Ioannidis [2005b], which is based on the following principles:

- —The decision about when a preference is related to a query is determined with respect to their structure, that is, the relations, attributes, and attribute values they both contain.
- —A personalized answer should satisfy at least 1 of the top k user preferences related to a query, in order of decreasing degree of interest.

Parameters 1 and k provide a quantitative way to describe the desired answer. Parameter k determines which of the related user preferences should be considered for application to a particular query. It also provides a way to control the cost of query personalization, as the fewer the preferences integrated into a query, the more efficient that query may be. Parameter 1 captures the minimum number of user criteria (i.e., preferences) that an answer should meet,

13:4 • G. Koutrika and Y. Ioannidis

thereby offering a degree of flexibility to personalizing a query answer. We discuss possible ways to specify parameters 1 and k in Section 4.1.

At query time, a mix of generic and specific preferences may be related to a query but these may not be freely combined to determine the ranking of the query results. For example, a preference for comedies can be combined with a preference for actor Adam Sandler but they are both overridden by a preference for "comedies with Adam Sandler". We formally define preference overriding relationships based on containment mappings. Given a set of preferences related to a query, we automatically identify such relationships and we organize them in a *preference network* (a directed acyclic graph). Each node in the network refers to a subclass of the entities that its parent refers to. Whenever the more specific preference (e.g., for comedies with Adam Sandler) applies, it overrides the more generic one (e.g., for comedies), whether the former represents a stronger or a weaker preference than the latter. To the best of our knowledge, this is the first systematic study of multi-granular preferences and preference relationships in a quantitative way.

Handling multi-granular preferences and preference relationships can lead to increased complexity and, thus, higher execution times. One approach would be to build an appropriate set of queries that expand the initial query with combinations of preferences that are allowable based on the preference relationships and return ranked results. For example, we could execute a query that returns comedies without Adam Sandler and another query that returns comedies with Adam Sandler. However, such queries may be quite complex and time consuming even for a few preferences with no relationships among them [Koutrika and Ioannidis 2005a]. Organizing the preferences in a network helps preference processing in the following ways. It helps keep track of the preference relationships and the preferences to use for ranking each tuple. Our personalized query answering algorithms can work directly on the network. They exploit the containment mappings implied by the network to decide how to traverse the network, that is, which preferences to process and in what order resulting in reduced preference processing. Exploiting containment mappings for optimizing performance and in general exploiting the characteristics of preferences (not only the characteristics of the queries) to optimize queries with preferences has not been considered in previous works so far.

We evaluate our framework and algorithms both for effectiveness and efficiency. We study the effect that the increased expressive power and freedom offered by our framework has to users through a user evaluation. Naturally, we raise the question of whether increased expressiveness is achieved at the expense of performance and we answer it in the negative through an appropriate experimental performance evaluation.

In summary, our contributions are the following:

—We introduce a framework for multi-granular preference modeling that is based on the concept of networks of preferences, which allow the representation of both generic and specific preferences and their relationships in a concise and flexible way (Section 3).

- We present a system architecture for query personalization based on preference networks (Section 4).
- —We describe efficient algorithms that, given a query and a user profile, derive the top k preferences that are related to the query, identify the relationships among them, and organize them in a network (Section 5).
- -We use the preference network to support algorithms for computing the personalized results of a query. Our algorithms process preferences in an efficient way by exploiting the containment mappings captured in the network (Section 6).
- —We study the overall impact of the increased expressiveness allowed by our framework through both synthetic and real-user experiments. We show that more accurate, finer-grained result rankings can be achieved without losing in performance (Section 7).

2. RELATED WORK

Preference is a fundamental notion in applied mathematics [Fishburn 1999], philosophy [Hansson 2001], and AI [Wellman and Dovle 1991]. In databases, preferences have been used for cooperative query answering, that is, for providing answers with extra or alternative information that may be meaningful to the user [Cuppens and Demolombe 1989; Gaasterland et al. 1992]. Recently, preferences have attracted renewed interest in the database community triggered by the observation that the strict Boolean information access model assumed in databases is some times restrictive. Query criteria are considered as *hard* by default, and a nonempty answer is returned only if it satisfies all criteria. On the contrary, in information retrieval, query criteria (terms) are considered *soft* and an answer contains results ranked according to how well they match the query. Recent years have witnessed the emergence of approaches aiming at the introduction of soft criteria or preferences in database queries, resulting to *preference queries*, with some early efforts going back to the 80's [Lacroix and Lavency 1987]. These approaches are divided into the following categories:

Qualitative approaches aim at a relative formulation of preferences, such as a user prefers comedies over westerns. This formulation is natural for humans and results in partial orders of results. Absolute specification of the significance of a preference is not possible. Preferences between tuples in the answer to a query are specified directly, using binary preference relations [Borzsonyi et al. 2001; Chomicki 2003, 2004; Kiessling 2002; Kiessling and Kostler 2002]. Two frameworks have been independently proposed in which preference relations are defined using logical formulas [Chomicki 2003] or special preference constructors [Kiessling 2002]. Preference relations are embedded into relational query languages through a relational operator that selects from its input the set of the most preferred tuples (winnow [Chomicki 2003], BMO [Kiessling 2002]).

Quantitative approaches aim at an absolute formulation of preferences, such as a user likes comedies very much and westerns to a lesser degree. This formulation leads to a total ordering of results based on user preferences.

13:6 • G. Koutrika and Y. Ioannidis

Preferences in queries are specified indirectly using scoring functions that associate a numeric score with every tuple of the query answer [Agrawal and Wimmers 2000; Hristidis et al. 2001].

There is extensive work on executing and optimizing preference queries, in particular:

- —top-n queries [Bruno et al. 2002; Chang and Hwang 2002; Tao et al. 2007], that is, queries that retrieve the best n objects that minimize a specific function. Core rank-aware query operators [Ilyas et al. 2004] and special structures, such as the P-Cube [Xin and Han 2008], have been proposed. Special cases of top-n queries have been studied, such as spatial top-n queries [Yiu et al. 2007], which rank objects based on the qualities of features in their spatial neighborhood.

Preference applications. Research on preference queries was mainly inspired by the need to adopt an IR-like answer model, where an answer to a query contains results ranked according to how well they match the query, as opposed to the strict Boolean model traditionally assumed in databases. Still, the answer of a query may be the same for all users. The trend for personalizing answers to fit different, ephemeral, or long-term user preferences also emerged in the IR community [Pitkow et al. 2002]. Personalized IR systems use different representations of preferences, such as bags of words [Joachims et al. 1997], vectors of terms [Balabanovic and Shoham 1997], and concept hierarchies [Liu et al. 2004] for information filtering, recommendations, or query personalization.

The idea of using preferences to customize query results has been transferred to databases, and applications include recommendations [Satzger et al. 2006], query personalization [Koutrika and Ioannidis 2004; Koutrika and Ioannidis 2005a, 2005b; Stefanidis et al. 2007], and collaborative query results [Koutrika 2006]. In addition to finding interesting results for a query, a related problem is finding interesting attributes that characterize a set of results [Miah et al. 2008; Wong et al. 2007]. Preference modeling is central to all these approaches. For instance, Satzger et al. [2006] adopt the preference model presented by Kiessling [2002]. Koutrika [2006, 2005a, 2005b, 2004] and Stefanidis et al. [2007] model a preference as a query condition, which can be integrated with a query associated with a degree of interest. Query personalization has been also studied for preferences holding in specific contexts [Agrawal et al. 2006; Holland and Kiessling 2004; Stefanidis and Pitoura 2008; Stefanidis et al. 2007; van Bunningen et al. 2006].

Preference modeling is only one of the challenges in preference applications. Other issues include dynamically identifying the preferences that are related to a query, selecting the most appropriate ones for the specific query and context,

and determining their effect on the final answer. The personalized answer may contain all results of the initial query ranked based on all preferences that are appropriate under a specific context [Stefanidis et al. 2007]. Koutrika and Ioannidis [2005b, 2004] introduced an 1 out of k personalization logic: a personalized answer should satisfy at least 1 out of k preferences from the user profile that are related to a query. A different approach is to view query personalization as an optimization problem using as parameters the query cost, the answer size, and the interest in the answer and dynamically decide the appropriate personalized answer optimizing one parameter given constraints on the others [Koutrika and Ioannidis 2005a]. Acquiring user preferences is another great challenge. In the context of IR and databases, preferences are either entered explicitly through a query interface or acquired at query time using a preference elicitation mechanism [Balke et al. 2007; Lee et al. 2008]. Long-term preferences can be learnt based on user feedback as well. There are studies for learning the order of objects that reflects user preferences [Cohen et al. 1998; Joachims 2002] or for learning preferences in object attributes (such as preferences formulated as described in Kiessling [2002]) [Holland et al. 2003; Jiang et al. 2008].

Comparison to related work. Our work is closer to Koutrika and Ioannidis [2004, 2005b] and Stefanidis et al. [2007] in that we share the idea of representing preferences as query conditions associated with degrees of interest. In addition, we adopt the "1 out of the top k preferences" query personalization logic [Koutrika and Ioannidis 2005b]. However, these efforts focus on problems of preference representation, such as differentiating preferences based on their intensity [Koutrika and Ioannidis 2005b] and studying context dependencies [Stefanidis et al. 2007], and they formulate structurally (and semantically) atomic preferences, that is, over a single attribute. For example, they can capture a preference for Woody Allen and a preference for comedies but would not allow an explicit preference statement for comedies with Woody Allen. Instead, they provide mechanisms for deriving more complex preferences based on simpler ones. However, a preference for comedies with Woody Allen may not be derivable in any natural way by a preference for comedies and a preference for Woody Allen. One may like both comedies and Woody Allen but one may like or dislike comedies with Woody Allen. As a result, the personalized results generated may not always accurately reflect the user preferences.

In this work, we allow complex preferences to be explicitly stated (not be solely derivatives of other preferences) and hence we can generate more accurate results. Furthermore, the central assumption in existing approaches is that preferences hold independently of each other and hence they can be freely combined for result ranking. Allowing preferences of different granularity makes this somewhat unrealistic independence assumption no longer necessary. Our framework captures preference relationships using preference networks, so that interrelated preferences (e.g., one for comedies and one for comedies with Woody Allen) can be treated appropriately (Section 3).

Our work is also different from existing approaches on executing and optimizing queries with preferences, such as skyline and top-*n* preference queries,

13:8 • G. Koutrika and Y. Ioannidis

in several aspects. First, we deal with different types of preferences than skylines. Skylines are formulated in the context of multidimensional Euclidean spaces [Borzsonyi et al. 2001]. The attributes of a relation are partitioned into DIFF, MAX, and MIN attributes. Only tuples with identical values of all DIFF attributes are comparable; among those, MAX, attribute values are maximized and MIN values are minimized. Our approach is also different from top-n queries (e.g., [Bruno et al. 2002; Tao et al. 2007]), where a single scoring function maps tuples to scores. We describe preferences irrespective of the attribute domain by specifying conditions that tuples must satisfy and assigning a degree of interest in them. Then, tuples are ranked depending on which of the conditions they satisfy.

Given these differences in problem formulation, our work addresses several new challenges related to preference integration with queries: identifying preference relationships, constructing a network of related preferences for a query, and using the network to generate personalized answers that respect the semantics of query personalization and the relationships between the preferences. Given any two preferences, we determine their relationship based on conjunctive query containment mappings [Aho et al. 1979; Chekuri and Rajaraman 1997]. Preference mappings are simpler because they are defined on the basis of mappings of the atomic conditions in the preferences considered. When generating personalized results, our algorithms take into account the way preferences are connected in the preference network for result ranking.

By enforcing preference independence, previous works may produce less accurate results, which may be still acceptable given that query personalization using independent preferences is more straightforward and efficient. Clearly, there is a trade-off between efficiency and expressivity in allowing multi-granular preferences. We experimentally compare our algorithms with their simpler counterparts which assume preference independence, both for efficiency and answer accuracy, and we show that personalized answers that capture user preferences more accurately are feasible without losing in efficiency.

3. FRAMEWORK

We consider that for every user there is a profile storing preferences for personalizing queries. We consider SPJ queries over relational databases. In this section, we present our multi-granular preference model (Section 3.1). Then, we define preference relationships and we introduce preference networks that capture a set of preferences and their relationships (Section 3.2). We estimate user preferences for the results of a query taking also into account the way preferences are interrelated (Section 3.3).

3.1 Preference Formulation

A database comprises a set of relations. A relation R (denoted R_i when more than one relation is implied) has a set A of attributes. We will use RA_j to refer to an attribute in A or simply A_j when R is understood.

Personalizing Queries Based on Networks of Composite Preferences •

MOVIE(<u>mid</u> , title, year, duration)	GENRE(<u>mid</u> ,genre)
CAST(<u>mid,aid</u>)	$ACTOR(\underline{aid},name)$
$DIRECTED(\underline{mid},\underline{did})$	$DIRECTOR(\underline{did}, name, nationality)$

Fig. 1. A database.

A preference for a set of tuples of a relation R in the database is expressed as a degree of interest $d \in [0, 1]$ in a condition q that describes this set of tuples, and is denoted

R:(q, d).

Depending on the value of d, we capture weaker or stronger preferences, with d = 0 indicating no interest in the qualifying tuples and d = 1 indicating extreme interest. Depending on the form of condition q, we distinguish the following types of preferences.

- —*Atomic Preference*. If q is a single, atomic, selection, or join condition, then a preference for q is called atomic.
- -*Composite Preference*. If *q* is a conjunction of multiple atomic conditions, then a preference for *q* is called composite.
- -Selection Preference. If q is a conjunction of atomic selections involving a set A of attributes and atomic joins transitively connecting these attributes to R on the database graph, then a preference for q is called a selection preference.
- —Join Preference. If q is a conjunction of atomic join conditions representing the transitive join of relations R and R_j on the database graph, then a preference for q is called a join preference.

A selection preference indicates a user's preference for tuples from R that satisfy the transitive selections in q. A join preference indicates the degree in which preferences for tuples in R are influenced by preferences on related, that is, joining, tuples. Since a preference for a set of tuples is essentially captured as a preference for a condition that describes this set, in what follows, we consider the terms "preference for tuples" and "preference for a condition" equivalent and we use them interchangeably. We will use the symbol p (or p_i , if a set of preferences is discussed) to refer to a preference. We will use doi(p) to refer to the degree of interest for a preference p or simply d if p is understood.

Example. To illustrate preferences, we consider the small database shown in Figure 1. A user, named Julie, has preferences for movies captured in her profile, part of which is depicted in Figure 2. For instance, she prefers movies released after 1990 (p_1) and she likes comedies (p_2) followed by adventures (p_3) . She particularly likes family adventures (p_4) . Preferences p_1 , p_2 and p_3 are atomic selection preferences expressed on different attribute values. We also observe that preference p_1 is stronger than p_2 and p_3 . Preference p_4 is a composite selection preference. Furthermore, Julie considers the director of a movie more important than the genre $(p_5$ and $p_6)$. p_5 is a composite join preference and p_6 is an atomic join preference. Her favorite director is Alfred Hitchcock (p_7) . She also likes sci-fi movies by Steven Spielberg captured by the composite selection preference p_8 .

ACM Transactions on Database Systems, Vol. 35, No. 2, Article 13, Publication date: April 2010.

13:9

13:10 • G. Koutrika and Y. Ioannidis

MOVIE :	MOVIE.year > 1990,	0.9
GENRE :	GENRE.genre = ``comedy",	0.85
GENRE :	GENRE.genre = ``adventure",	0.7
MOVIE :	MOVIE.mid = GENRE1.mid and GENRE1.genre = "family" and MOVIE.mid = GENRE2.mid and GENRE2.genre = "'adventure"	0.8
MOVIE :	$\label{eq:move_model} \begin{split} \text{MOVIE.mid} &= \text{DIRECTED.mid} \ \text{and} \\ \text{DIRECTED.did} &= \text{DIRECTOR.did}, \end{split}$	1.0
MOVIE :	MOVIE.mid = GENRE.mid,	0.7
DIRECTOR :	DIRECTOR.name = "A.Hitchcock",	0.7
MOVIE :	MOVIE.mid = GENRE.mid and GENRE.genre = "scifi" and MOVIE.mid = DIRECTED.mid and DIRECTED.did = DIRECTOR.did and DIRECTOR.name = "S.Spielberg"	0.65
	MOVIE : GENRE : GENRE : MOVIE : MOVIE : DIRECTOR : MOVIE :	MOVIE:MOVIE.year > 1990,GENRE:GENRE.genre = "comedy",GENRE:GENRE.genre = "adventure",MOVIE:MOVIE.mid = GENRE1.mid and GENRE1.genre = "family" and MOVIE.mid = GENRE2.mid and GENRE2.genre = "adventure"MOVIE:MOVIE.mid = DIRECTED.mid and DIRECTED.did = DIRECTOR.did,MOVIE:MOVIE.mid = GENRE.mid,DIRECTOR:DIRECTOR.name = "A.Hitchcock",MOVIE:MOVIE.mid = DIRECTED.mid and DIRECTED.did = DIRECTED.mid and DIRECTOR.name = "S.Fielberg"

Fig. 2. Example preferences.

Preference inference. Given a set of preferences, one can derive implicit preferences by composing preferences. Two preferences, p_i and p_j , are *composeable*, iff:

- $-p_i$ is a join preference, $R_i : (q_i, d_i)$, connecting R_i to a relation R_i , and
- $-p_j$ is a join or selection preference on R_j , that is, of the form $R_j : (q_j, d_j)$.

We now consider a set of preferences: $R_1 : (q_1, d_1), R_2 : (q_2, d_2), \ldots, R_m : (q_m, d_m)$ such that, $\forall 1 \leq i < m$, preferences $R_i : (q_i, d_i), R_{i+1} : (q_{i+1}, d_{i+1})$ are composeable. Then, we can compose the *implicit preference* R : (q, d), as follows:

$$-R \equiv R_1,$$

-q is the conjunction of the conditions q_1, q_2, \ldots, q_m and

-*d* is some function of the degrees of interest of the base preferences, that is, $d = f(d_1, d_2, \ldots, d_m)$.

Inference assumption. The degrees of interest in an implicit preference is a nonincreasing function f of the degrees of interest, d_1, d_2, \ldots, d_m , of its constituent preferences, that is,

$$f(d_1, d_2, \dots, d_m) \le \min(\{d_1, d_2, \dots, d_m\}).$$
 (1)

Example functions for computing the degree of interest in implicit preferences include product and minimum. One advantage of the product is that it captures human intuition and cognitive evidence [Collins and Quillian 1969]: the more the preferences required to derive an implicit preference, the weaker this preference is [Koutrika and Ioannidis 2005b].

Example. From the preferences stored in Julie's profile (Figure 2), we can derive implicit preferences. For instance, since Julie likes director Alfred Hitchcock (p_7) and her movie preferences are affected by who is the director of a movie (p_5) , we can assume that (in lack of any other evidence) Julie would like

to some extent movies directed by Alfred Hitchcock by composing p_5 and p_7 into the following preference.

$(p_{5,7})$ MOVIE :	$MOVIE.mid = DIRECTED.mid \ and \label{eq:movies}$	
	DIRECTED.did = DIRECTOR.did and	
	DIRECTOR.name = ``A.Hitchcock'',	0.7

Similarly, we can derive an implicit preference for movies that are adventures by combining preferences p_3 and p_6 into this preference.

$(p_{6,3})$ MOVIE :	MOVIE.mid = GENRE.mid and	
	GENRE.genre = "adventure".	0.49

For the sake of illustration, in both preferences, we have taken the product of the partial degrees of interest as the degree of interest in the derived preference. Of course, we could have used any other nonincreasing function, such as the minimum.

We observe that both base selection preferences p_3 and p_7 have the same degree of interest, that is, 0.7. However, when combined with their composeable join preferences, the degrees of interest in the resulting implicit preferences, $p_{6,3}$ and $p_{5,7}$, are different. It is join preferences that determine to what extent preferences for tuples in one relation are influenced by preferences on joining tuples. In this case, preference p_5 indicates that directors are more important for movies than genres (preference p_6).

Preference diagram. A preference R : (q, d) can be graphically represented as a rooted directed acyclic graph g(V, E). The preference diagram g(V, E) can be thought of as an extension of the traditional schema graph and has the following characteristics. Nodes in V map to relations, attributes, and values in q, and can be possibly replicated if the corresponding relation, attribute, or value is used more than once in q. Edges in E connect: (a) a relation node to another relation node, representing a join condition in q or (b) an attribute to its container relation or (c) an attribute to a value, representing a selection in q. Join edges are tagged with the corresponding joining attributes, and selection edges are tagged with the operator used in the selection condition. R maps to the graph's root. For selection preferences, the leaves are always values.

A preference diagram can take either of two forms. It can contain one single path mapping to a join or selection preference or it can contain a set of paths connecting its root to a set of values capturing a composite multivalue selection preference. As we will see in Section 5.2, we can determine if there is an overriding relationship between a pair of preferences by comparing their corresponding preference diagrams.

Example. Figure 3 illustrates some example preference diagrams. Figure 3(a) depicts the preference diagram for p_4 . Figure 3(b) shows the preference diagram for p_5 , which comprises a single path in contrast to the diagram for p_4 , which maps to a set of paths from the root to the value nodes. For the sake of presentation, we do not show tags on the edges, since for our small database they are understood. Repeated relations and attributes are mapped to different nodes in the preference diagram. As a notational convention, different

ACM Transactions on Database Systems, Vol. 35, No. 2, Article 13, Publication date: April 2010.

13:12 G. Koutrika and Y. Ioannidis



Fig. 3. Example preference diagrams.

occurrences of a relation (mapping to different nodes in the diagram) are written as a concatenation of the relation name and a sequential number. To distinguish when an attribute is used with a different occurrence of the same relation, the attribute takes the relation's sequential number in their representation.

Preference diagrams can "unfold" complex conditions, such as self-joins, and map them to distinctive paths on the diagram. To illustrate this, consider the following, rather nontypical, preference for movies released in 2008 and directed by French directors who have also directed movies released in 2000. Its preference diagram is shown in Figure 3(c).

MOVIE : MOVIE.year = 2008 and MOVIE.mid = DIRECTED1.mid and DIRECTED1.did = DIRECTOR1.did and DIRECTOR1.nationality = "French" and DIRECTOR1.did = DIRECTED2.did and DIRECTED2.mid = MOVIE2.mid and MOVIE2.year = 2000

3.2 Preference Networks

Depending on the form of the condition, selection preferences can be defined on a relation at varying levels of granularity, ranging from quite generic (e.g., an atomic selection) to very specific that combine multiple (atomic or implicit) selections. We can view a preference R_i : (q_i, d_i) as a possible conjunctive query Q_i that selects tuples from R_i that satisfy q_i . On the basis of this correspondence, we can define preference overriding through conjunctive query subsumption, that is, *query containment*: Given two preferences, $R_i : (q_i, d_i)$ and $R_j : (q_i, d_i)$, with corresponding conjunctive queries Q_i and Q_j , respectively, we say that p_i is overridden by p_j if Q_j is subsumed by Q_i (i.e., if for all database instances,



Fig. 4. Preference overriding.

every answer to Q_j is an answer to Q_i as well, $Q_j \subseteq Q_i$). It is well known [Aho et al. 1979; Chekuri and Rajaraman 1997] that Q_j is subsumed by Q_i iff there is a containment mapping from Q_i to Q_j . A containment mapping is defined on the basis of mapping the predicates involved in the two queries. In what follows, we formally define preference overriding on the basis of how one preference can be mapped to another:

Preference relationships. Given a set P of selection preferences and two preferences $R_i : (q_i, d_i)$ and $R_j : (q_j, d_j)$ (denoted p_i and p_j , respectively) from P, we distinguish the following possible relationships between them.

—*Preference Overriding*. If p_i can be mapped to p_j , such that: (a) $R_i \equiv R_j$ and (b) each atomic condition in q_i is mapped to an atomic condition in q_j with the same relations and attributes, then p_i is overridden by p_j (p_j overrides p_i), denoted $p_i \sqsubset p_j$. In this relationship, p_j is specific, p_i is generic, and the following degree of interest correspondences hold (denoted by \leftarrow):

$$doi(p_i \wedge p_j) \leftarrow doi(p_j)$$
 (2)

$$doi(p_i \wedge \neg p_j) \leftarrow doi(p_i) \tag{3}$$

that is, the generic preference p_i holds for a tuple only if the specific preference p_j does not hold for the same tuple.

Furthermore, if $\nexists p \in \mathbf{P}$ such that $p_i \Box p \Box p_j$, that is, p_j is the most generic specialization of p_i in \mathbf{P} , then p_i is *tightly overridden* by p_j , denoted $p_i \Box p_j$.

—Preference Independence. If $R_i \equiv R_j$ but $p_i \not \sqsubset p_j$ and $p_j \not \sqsubset p_i$, then p_i and p_j are *independent* and the following holds:

$$doi(p_i \wedge p_j) \leftarrow h(doi(p_i), doi(p_j)), \tag{4}$$

where h is some appropriate ranking function, that is, both preferences can hold for the same tuple and the overall degree of interest in this tuple is the combination of their degrees of interest.

In other words, a preference p_i is overridden by a preference p_j if both of them are expressed over the same relation and there is a mapping of the atomic conditions from one preference to the other. Figure 4 shows an example.

Preference network. We can map a set P of selection preferences defined over the same relation R to a preference network, where nodes map to preferences and edges represent preference relationships. A preference network $G_H(V_H, E_H)$ corresponding to a set P of preferences over relation R is a directed acyclic graph. Nodes in V_H map to preferences in P. Given two nodes v_i and v_j in V_H ,

ACM Transactions on Database Systems, Vol. 35, No. 2, Article 13, Publication date: April 2010.

13:14 • G. Koutrika and Y. Ioannidis



Fig. 5. A set of preferences as a network.

mapping to preferences p_i , $p_j \in \mathbf{P}$, respectively, an edge in E_H from v_i to v_j , $e(v_i, v_j)$, denotes that p_i is tightly overridden by p_j . When a preference is not overridden by any other preference, it is called a *leaf*. When it does not override any preferences, it is called a *root*.

Example. Figure 5 illustrates a preference network (on the left side) corresponding to a set of preferences (on the right side). p_a is a root, p_d and p_e are leaves, and p_f is both a root and a leaf. We observe that the degrees of interest at one level of the network are not derived from the degrees of interest at the higher or lower levels of the network. For instance, p_b is a strong preference, whereas the more specific p_e is a weak preference. If there is no directed path in the network that connects two preferences, then these preferences are independent. For instance, p_b , p_d and p_f are independent.

3.3 Combining Preferences

To compute the degree of interest in a particular tuple, we combine the degrees of interest of the preferences expressed for this tuple. For instance, given that Julie likes movies by S. Spielberg and action movies, we can compute a preference for the movie "Minority Report" that satisfies both preferences.

The degree of interest in a tuple $t \in R$ satisfying a set P of selection preferences over relation R is estimated with the help of a ranking function h as follows:

$$doi(t) = doi(\wedge_{p_i \in \mathbf{P}'} p_i) = h(\{doi(p_i) | p_i \in \mathbf{P}'\})$$
(5)

with $\mathbf{P}' \subseteq \mathbf{P}$ such that $(a) \forall p_i, p_j \in \mathbf{P}', p_i, p_j$ are independent and $(b) \forall p_i \in \mathbf{P}', \nexists p_j \in \mathbf{P}$ that is more specific than p_i . In other words, given a set of preferences satisfied by a tuple *t*, only the most specific, independent preferences determine the degree of interest in *t*.

There are many possible ranking functions that could be used depending on the application, the domain, and so forth. For instance, it may suffice to rank movies based on the best preference they meet. On the other hand, when buying a car, many car features matter and a function such as the average of preferences may be more appropriate.

Example. Consider the preferences shown in Figure 5 and a movie *m* that is a 2000 comedy with J. Nicholson. This movie, in principle, satisfies preferences p_a , p_b , p_c , and p_e . Taking into account their relationships, not all preferences

Personalizing Queries Based on Networks of Composite Preferences •

13:15

count and the degree of interest in this movie is computed as follows.

$$doi(m) = doi(p_a \land p_b \land p_c \land p_e) \stackrel{p_a \sqsubseteq p_b}{=} doi((p_a \land p_b) \land p_c \land p_e) \stackrel{Formula (2)}{=} doi(p_b \land p_c \land p_e) \stackrel{p_b \sqsubseteq p_e}{=} doi((p_b \land p_e) \land p_c) \stackrel{Formula (2)}{=} doi(p_e \land p_c) \stackrel{p_c \sqsubseteq p_e}{=} doi(p_e) = 0.4$$

That is, since p_e overrides p_a , p_b , and p_c , it alone determines the degree of interest in the movie.

4. QUERY PERSONALIZATION

4.1 Query Personalization Logic

We consider that a user's preferences are stored in a *user profile*. In order to personalize the results of a query, we need to identify the (explicit or implicit) selection preferences from the profile that are related to the query and determine how they will affect the final answer. Obviously, not all preferences are related to all queries. For instance, if one is looking for movies directed by S. Spielberg, a preference for movies directed by W. Allen is not related. Moreover, user preferences may have different effects on the results of a query. For instance, the system may return all results that match a query ranked according to the user's preferences or only a subset of them that exactly match some preferences.

We define preference relatedness on the basis of the syntactic characteristics of the query and the preferences in the spirit of the work by Koutrika and Ioannidis [2004], extended to our multi-granular preference model.

Related preference. Preference R: (q, d) is related to query Q, if R is referred to in Q.

Conflicting preference. Preference R : (q, d) is conflicting with a query Q, if: (a) it is related to the query; (b) when q is inserted into the original query qualification, replacing any part of the latter that coincides with it, that is, joins or selections on a common attribute, the resulting query qualification and the original one contain at least one common implicit selection; and (c) this selection is specified over a single-value attribute of the query results.

Example. Consider the set of preferences shown in Figure 5, which are all related to movies, and a query about family movies released in 2008. Then, preferences p_a - p_e are clearly conflicting with the query since they are about movies released in 2000. Preference p_f is valid because a movie can have more than one genre.

We adopt the query personalization framework introduced by Koutrika and Ioannidis [2005b] for determining the effect of query personalization. Given a preference R: (q, d), d shows the user interest in considering q for personalizing the results of a query involving R. Hence, preferences that are related to and not conflicting with a query can be ordered based on their degree of interest. We

ACM Transactions on Database Systems, Vol. 35, No. 2, Article 13, Publication date: April 2010.

13:16 • G. Koutrika and Y. Ioannidis

can describe the desired personalized answer with the help of two parameters as follows:

A personalized answer for a query Q and a user profile U is the set of results of Q that satisfy at least 1 of the top k user preferences related to Q and are ranked according to their degree of interest.

Parameters k and 1 together offer a flexible mechanism to capture the desired level of personalization for the results. For example, some special cases are:

- -1 = 0: the personalized answer contains all the results matching the query ranked based on the k preferences.
- -1 = k: the personalized answer contains only the query results that satisfy all k preferences.

In what follows, we discuss several ways for determining the values for parameters 1 and k. A user can explicitly put constraints on the acceptable personalized answer by specifying desired values for 1 and k. The system can also automatically determine the appropriate values.

Parameter k controls the desired *extent* of personalization, that is, how many of the top user preferences should participate in personalizing the query results. For instance, the user may be willing to put no bound to k and see results that satisfy any of her related preferences (i.e., $k \equiv$ 'all'). On the other hand, if the results are too many or the user is very selective, returning results that satisfy a small number k of preferences may be more appropriate. For example, a user interested in dinner after work may want to see only restaurants that satisfy her top 3 preferences.

Parameter 1 controls the desired *forcefulness* of personalization, that is, how many of the k live preferences should the query results be forced to simultaneously satisfy. For instance, when buying a car a user may definitely want offers that meet all her preferences but, when selecting a movie, she may be more flexible and may accept movie recommendations that meet just two of her preferences.

Parameters k and 1 also provide a way for the system to control the size of the personalized answer and the cost of query personalization, hence, the system could automatically determine appropriate values for them. The system may restrict k since considering a smaller set of preferences leads to smaller answers and may be more efficient. For example, if the user is browsing results through a cell phone, then the system may pick few preferences from the user profile in order to provide results on the go. If the user uses a laptop and a high-bandwidth connection, then a more flexible k or no bound for k may be selected. For the case of k = 1, determining automatically appropriate values of k has been studied in the past [Koutrika and Ioannidis 2005a]. In that work, query personalization is formulated as a constrained optimization problem, where constraints are expressed as an upper bound on execution time of the final query and/or a lower or upper bound on its result size.

In a similar vein, the system may automatically determine the right values for 1. For instance, if the initial query is very general and returns many results, then a larger number of preferences (i.e., $1 \rightarrow k$) may be applied on the query

in order to produce an answer of manageable size. Likewise, if the initial query is overly specific and returns very few results, then a smaller number of simultaneous preferences (i.e., $1 \rightarrow 0$) may be chosen to avoid an empty answer. The preference network may also serve as the basis for choosing appropriate values of 1. For instance, if the top preferences are highly correlated, that is, they form a chain, then the system should automatically determine that 1 can only be equal to 1 and return results that satisfy exactly one of the k preferences.

The system may also learn 1 and k for each user by mining user logs. These may contain both explicit choices of 1 and k by users as well as other user actions, that is, paying attention to specific tuples that offer implicit indications for the desirable values of these parameters. For example, if the user always examines results that satisfy up to the top 3 of her preferences and most of the times her final picks satisfy at least 2 of these preferences, then we can provide more targeted results by using k = 3 and 1 = 2 in in future interactions of the user with the system.

Finally, if personalization is an interactive process and the user can specify the values of k and 1, then there are potentially several cases where freelyselected values for k and 1 may not make sense, such as: (a) the top k preferences are highly correlated, that is, there are fewer than k independent preferences, (b) there are not as many as k preferences related to the current query, or (c) there are no results that satisfy as many as 1 preferences. In these cases, the system should guide the user in selecting meaningful values for k and 1.

4.2 Overall Approach

Given a query Q, a user profile U, and values for k and 1, query personalization proceeds as follows: The top k preferences that are related to a query are extracted from the user profile. Given that these may comprise a mix of generic and more specific preferences, we automatically identify the relationships among them and organize the preferences into a network. Next, we use the network to guide preference processing for the computation and ranking of the personalized results. We propose two algorithms for personalized query answering that apply different traversal strategies on the network, exploit the containment mappings captured in it to process the preferences, and rank the results.

The query personalization system architecture is shown in Figure 6. *Preference Construction* extracts the top k preferences from the user profile that are related to a query. *Preference-Network Integration* is responsible for organizing them in a network that captures the relationships among them (if any). This module places each preference found from the first module in the right "position" in the network. It collaborates with the *Relationship Finder* that determines the relationship between a pair of preferences. *Personalized Query Answering* takes into account the network of related preferences and returns results that meet the initial query and at least 1 from the top k preferences. Sections 5 and 6 describe in detail the preceding modules and the algorithms and techniques employed in them.

13:18 • G. Koutrika and Y. Ioannidis



Fig. 6. System architecture.

5. PREFERENCE CONSTRUCTION

Algorithm PFC (preference-construction) returns the top k preferences from the user profile U that are related to and not conflicting with a query Q, organized in a network.

Preferences that are related to a given query include explicit preferences stored in the user profile but also implicit ones that can be derived by composing stored preferences. In order to select the top k preferences, the algorithm starts from the preferences that are stored in the user profile and are related to the query and iteratively considers additional preferences that are composeable with those already known to derive implicit ones that are also related to the query. The set of preferences that are related to the query is kept ordered in decreasing degree of interest. When k selection preferences have been inserted into the network, the *Inference Assumption* guarantees that these are the top k preferences is avoided.

The algorithm is presented in Figure 7. QP is the set of preferences related to the query. Initially, it contains all related preferences explicitly stored in the profile (line 2). At each round, the most significant preference is processed based on its type.

- -A selection preference is added to the preference network that comprises the final output of the algorithm (line 3.2). This process is performed by PNET, which will be described in Section 5.1.
- —A join preference is composed with other preferences (line 3.3). The algorithm considers all stored preferences that are composeable with it to infer new preferences that are related to and not conflicting with the query. These are inserted into QP.

We will use the *Inference Assumption* to prove the correctness of the algorithm.

THEOREM 5.1. The algorithm is correct, that is, it finds the top k preferences related to a query.

Personalizing Queries Based on Networks of Composite Preferences • 13:19

Algorithm PFC

Input: query Q, a user profile U, k **Output:** a preference network $G_H(V_H, E_H)$ Begin 1. $V_H := \emptyset$, $E_H := \emptyset$ 2. $QP := \{(p_i, P_i) \mid p_i \in U \text{ related to } Q\}$ 3. While $QP \ll \emptyset$ and k > 0{ 3.1. remove head (p, P) from QP3.2 If *p* is a selection preference { 3.2.1. add p to V_H ; k := k-1 $G_H := \text{PNET}(G_H, (p, P))$ 3.2.2. } 3.3. Else { 3.3.1. s := the unique path in P3.3.2. **Foreach** $p_i \in U$ composeable with pIf $(p \wedge p_i)$ is not conflicting with Q $P' := \emptyset$ **Foreach** path $s_i \in P_i$ $s' := concatenate(s, s_i)$ add s' in P'add $(p \wedge p_i, P')$ to QP4. output $G_H(V_H, E_H)$ End

Fig. 7. Building a network of top k preferences related to a query.

PROOF. It is sufficient to show that the algorithm produces preferences related to the query Q in decreasing order of their degree of interest. It keeps an ordered list QP of preferences. At each round, the head p of the queue (which has the highest degree of interest) is picked and processed. If p is a selection preference with degree of interest d, we will show that: (a) d is greater than (or equal to) the degree of interest of any other selection preference in QP, and (b) d is greater than (or equal to) the degree of interest of any other selection preference that has not been seen yet (i.e., not in QP).

Since QP is ordered, (a) is self-evident. For (b), we observe that any unseen selection preference can be only derived from the current join preferences in QP. Based on the *Inference Assumption*, for each join preference p_j in QP with degree of interest d_j , any selection preference that contains p_j will have a degree of interest at most equal to d_j . Since it is $d \ge d_j$ (due to (a)), (b) also holds. Consequently, the algorithm produces preferences related to the query Q in decreasing order of their degree of interest. \Box

Example. Let us assume that we want to find the top 3 preferences related to a query about movies given the user profile of Figure 2. (How the preference network is progressively constructed is the focus of the next section.) Figure 8 shows the status of QP and the top preferences found at the end of each round of the algorithm. Newly inserted preferences in both lists are shown instantly in different color. When three selection preferences have been found, the algorithm can safely stop.

13:20 • G. Koutrika and Y. Ioannidis



Fig. 8. Example of deriving top k related preferences.

In addition, for each preference p, algorithm PFC constructs its preference diagram. In particular, it constructs the set P of all root-to-leaf paths on the preference diagram. Each path is encoded as a string for the purposes of finding relationships between pairs of preferences, as we will discuss in Section 5.2. We distinguish two cases.

- —If a preference is stored in a profile, then we have also stored all the root-toleaf paths on its preference diagram. This step is done offline. The algorithm PATHS is sketched in Figure 5 and is described next.
- —If it is an implicit preference composed of other preferences: Algorithm PFC builds a new preference by composing a join preference p with another preference p_j (line 3.3.2). Then, the set P' of paths for the new preference is generated by concatenating each of the paths that map to p_j (the set of paths for p_j is P_j) with the paths in P that maps to p. Note that since p is a join preference, its set P of paths is singleton.

Algorithm PATHS is used for offline processing of the preferences in a profile. For a given preference p, it builds its preference diagram g(V, E) (lines 2–3). The diagram's root is the relation for which the preference is defined and the leaves are the value nodes (in case of selection preferences) or the most distant relation from the root (in case of join preferences). The diagram edges are considered undirected at this stage. Their direction is determined as the algorithm performs a breadth-first traversal of the diagram and generates the set P of root-to-leaf paths (line 4). The paths are stored in the user profile. This pre-processing saves time for the online algorithm PFC. A path is represented as a string (for comparison purposes as we will see in Section 5.2.)

Personalizing Queries Based on Networks of Composite Preferences • 13:21

Algorithm PATHS

Input: a preference p**Output:** its set of root-to-leaf paths P Begin 1. $V := \emptyset, E := \emptyset$ 2. Foreach join condition q_i in p between relations R and R' { If there is no node in V mapping to $R \{ add R to V \}$ 2.1 2.2. If there is no node in V mapping to $R' \{ add R' to V \}$ 2.3. add edge e(R, R') in E 3. Foreach selection condition q_i in p on relation R If there is no node in V mapping to $R \in \{ add R to V \}$ 3.1 3.2. for the attribute A in q_i : add A to V 33 for the value v in q_i : add v to V 3.4. add edge e(R,A) in E 3.5. add edge e(A,v) in E } 4. P := BFS(V, E)5. output P End

Fig. 9. Finding the root-to-leaf paths on a preference diagram.

5.1 Preference-Network Integration

In this section, we focus on the integration of a preference p into a network $G_H(V_H E_H)$ of preferences related to the same query.

Problem statement. Given a preference p and a preference network $G_H(V_H, E_H)$, the result of their integration is a new network $G'_H(V'_H, E'_H)$, such that:

$$\begin{split} &-V'_{H}=V_{H}\cup\{p\},\\ &-\forall p_{i},\,p_{j}\in V'_{H} \text{ with } p_{i} \mathrel{\tilde{\sqsubset}} p_{j}, \exists e(p_{i},\,p_{j})\in E'_{H}. \end{split}$$

In other words, the result of integrating a preference into a preference network is a network, that is, containing only tightly overriding relationships, that captures all such relationships that exist between the new preference and the preferences already in the network.

Algorithm PNET. Integrating a preference into a network comprises two problems: (a) identifying a tightly overriding $(\tilde{-})$ relationship that involves this preference and some preference in the network and (b) traversing the network to establish all such relationships.

In order to find the $\tilde{-}$ -relationships between a new preference p and the preferences in a network $G_H(V_H, E_H)$, algorithm PNET employs a set of patterns, each one identifying one or more $\tilde{-}$ -relationship using $_$ -relationships.

- *—root pattern*: $(p_r \text{ is a root in } G_H \text{ with } p \sqsubset p_r) \Longrightarrow p \overset{\sim}{\sqsubset} p_r$
- *—leaf pattern*: $(p_i \text{ is a node in } G_H \text{ with } p_i \sqsubset p \text{ and } p_i \text{ is a leaf or all its children are independent with } p) \Longrightarrow p_i \subset p$
- *—intermediate node pattern*: $(p_i, p_j \text{ are nodes in } G_H \text{ with } p_i \sqsubset p \sqsubset p_j \text{ and } p_i \tilde{\sqsubset} p_j) \Longrightarrow p_i \tilde{\sqsubset} p \tilde{\sqsubset} p_j$

13:22 • G. Koutrika and Y. Ioannidis

The first pattern shows the case of p taking over the role of root from an earlier root. The second pattern shows the case of p becoming a leaf in the network. The last pattern captures the case of p being added between two other nodes. Based on this pattern, given two preferences p_i and p_j in the network with p_i being tightly overridden by p_j , if the new preference p overrides p_i and is overridden by p_j , then the edge between p_i and p_j is replaced by two edges that essentially place p between the two preferences. Clearly, as algorithm PNET may be called for a sequence of preferences, relationships in the network are added or replaced in light of each new preference considered. If a new preference does not have any relationships with the other preferences, then it is independent from all preferences currently in the network and becomes a new root-leaf in the network.

A network may contain multiple roots, that is, preferences that do not override others (Section 3.2). Hence, preferences that are tightly overridden $(\tilde{\square})$ by p, or vice versa, may be found at different places in the network, under the same or different roots. Algorithm PNET starts from each root and performs a breadth-first traversal of the corresponding subgraph. To avoid an exhaustive traversal of the network, two pruning rules are used:

- -Subgraphs with different roots may overlap. If an edge is already visited then the underlying subgraph is already explored and it can be safely pruned.
- —The subgraph starting from any node in the network is not explored if the preference mapping to this node is independent from or overrides *p*.

Based on the patterns presented, which capture all cases of *p*'s position with respect to other preferences in the network and the traversal strategy, algorithm PNET correctly places a preference in a network by establishing *all* $\tilde{-}$ -relationships with the existing preferences in the network and the result of the integration is a network.

The algorithm is presented in Figure 10. A queue RQ keeps edges to be examined. These are edges that have not been visited before. New edges are always added in RQ's tail. FIND_REL identifies the relationship of a pair of preferences (see Section 5.2).

The algorithm examines the relationship of p with each root p_r in the network (line 3.1). It first tries the root pattern (line 3.2). If p is overridden by p_r , then the property of being root is transferred from p_r to p. If p overrides p_r , then the algorithm tries the leaf pattern (line 3.3.2): if there are no outgoing edges from p_r , then p becomes a leaf under p_r . Otherwise, all edges from p_r to its children are added in the queue along with a dummy edge (p_r , -) and the algorithm will compare p to the preferences below p_r in the network.

The algorithm moves down in the network as long as there is an edge $e(p_s, p_i)$ in RQ (line 3.4). If this is a dummy edge, $e(p_s, -)$, then the leaf pattern is tested: if p has been found independent from p_s 's children (indicated by *inserted* = false), then p becomes p_s 's new child. Hence, a dummy edge is used in combination with the flag *inserted* to show when p should become a new child of an intermediate node in the network. If $e(p_s, p_i)$ is an actual edge in the network, then:

Personalizing Queries Based on Networks of Composite Preferences • 13:23

Algorithm PNET

a preference network $G_H(V_H, E_H)$, preference p Input: **Output:** a preference network $G_H(V_H, E_H)$ Begin 1. $is_root := true$ 2. $RQ := \emptyset$ 3. While \exists root preference $p_r \in V_H$ not examined { *Relationship* := FIND $REL(p, p_r)$ 3.1. 3.2. **If** *Relationship* = '*p* is overridden by p_r ' { 3.2.1. create edge $e(p, p_r)$ in E_H 3.2.2. unmark p_r } If *Relationship* = ' p_r is overridden by p' 3.3. { 3.3.1. *is_root* := false 3.3.2. If \nexists visited edge $e(p_r, *)$ in $E_H = \{$ create edge $e(p_r, p)$ in E_H } Else { **Foreach** not visited edge $e(p_r, p_i)$ in E_H add (p_r, p_i) in RQadd $(p_r, -)$ in RQ*inserted* := false } } } While $RQ <> \emptyset$ 3.4. { 3.4.1. get head element (p_s, p_i) from RQ3.4.2. If $p_i = -$ and *inserted* = false { create edge $e(p_s, p)$ in E_H *inserted* := true } Else { Relationship := FIND_REL (p, p_i) If *Relationship* = '*p* is overridden by p_i ' { remove edge $e(p_s, p_i)$ from E_H create edge $e(p_s, p)$ in E_H create edge $e(p, p_i)$ in E_H *inserted* := true If *Relationship* = ' p_i is overridden by p' { If \nexists visited edge $e(p_i, *)$ in E_H create edge $e(p_i, p)$ in E_H *inserted* := true } Else { **Foreach** not visited edge $e(p_i, p_j)$ in E_H add (p_i, p_j) in RQadd $(p_i, -)$ in RQ} } } } 4. If *is_root* = true 4.1 { mark node *p* as a root } 5. output $G_H(V_H, E_H)$ End

Fig. 10. Integrating a preference into a network.

13:24 • G. Koutrika and Y. Ioannidis



Fig. 11. Examples of preference-network integration.

- —If p is overridden by p_i (intermediate node pattern), then the algorithm "breaks" the edge between p_i and its predecessor, and creates two edges, one connecting the predecessor to p and one from p to p_i . Having found the position of p in the subgraph, *inserted* becomes true and no outgoing edges from p_i are added in RQ, that is, the subgraph under p_i is pruned.
- —If p_i is overridden by p, then this node's outgoing edges are added in RQ along with a dummy edge $(p_i, -)$. If there are no outgoing edges, then p becomes p_i 's child and again *inserted* becomes true.

Finally, if *p* does not override any other preference (*is_root* = true), it becomes a root.

Example. Figure 11 presents the construction of the preference network depicted in Figure 5. Preferences p_a to p_f are presented to the algorithm in

that order. Each block in the figure shows the steps required for inserting one preference in the network. Each step is described by the contents of RQ and the status of the network, with nodes examined already indicated with gray background. For instance, the first step for inserting p_c visits the root p_a , which is overridden by p_c . Hence, its outgoing edges plus a dummy edge are placed in RQ. In the second step, the edge going to p_b is obtained from RQ and as a result p_b is visited. This one is independent from p_c , hence the algorithm will not search below this preference. Finally, pulling out the dummy edge $< p_a, - >$ marks the end of p_a 'children examination and since p_c has been found independent from all of them, it is connected to p_a .

Preference p_e provides an example of how the algorithm moves until it establishes all relationships that p_e participates. The first relationship (i.e., with preference p_b) is found in the second step. The algorithm continues exploring the subgraphs starting from p_b 's siblings, and discovers the second relationship (i.e., with preference p_c). Finally, observe how p_f becomes a root.

5.2 Relationship Finder

A preference is overridden by another preference if they are both defined over the same relation and there is a mapping of the atomic conditions from one preference to the other (Section 3.2). Preferences are represented as graphs. Our approach for identifying the relationship between two selection preferences, pand p', defined over the same relation R, is the following. We consider the sets P and P' of all root-to-leaf paths on their preference diagrams, respectively. |P|and |P'| are the sizes of these sets. For efficient path counting and comparison, we adopt a string representation of a path (string encoding of graphs in general has been proposed in Zaki [2005].) To generate a path representation, we concatenate the names of nodes in the path. For instance, the string representation of p_a in Figure 5, which comprises a single path, is "Myear2000". Two paths $s_i \in P$ and $s'_i \in P'$ match iff their string representations are the same. Then, the relationship of p and p' can be determined by counting the number M of pairs (s_i, s'_i) of matching paths. The following cases are distinguished:

- -If M = |P| it is $p \sqsubset p'$.
- —If M = |P'| then $p' \sqsubset p$.

—If none of these holds, p' and p are independent.

This process is captured in algorithm FIND_REL, shown in Figure 12. In case of selections containing inequalities, the process is slightly different: it matches paths without considering the selection values, and performs an additional check for the atomic selection conditions to determine their relationship. For presentation purposes, in subsequent algorithms, we assume selections with equalities.

The correctness of the algorithm stems directly from the definitions of preference diagram (Section 3.1) and preference overriding (Section 3.2). p is overridden by p' if each atomic condition in p is mapped to an atomic condition in p' with the same relations and attributes. Since selection preferences map to rooted graphs, where the root is a relation and the leaves are always values, the

13:26 • G. Koutrika and Y. Ioannidis

Algorithm FIND REL

preference p, preference p'Input: **Output:** relationship Begin 1. read the sets P and P' of root-to-leaf paths for p, p', resp. {relationship := MATCH(P, P') } 2. If $|P| \leq |P'|$ {relationship := MATCH(P', P) } Else 3. output < p relationship p' >End procedure MATCH Input: set P_1 , set P_2 **Output:** relationship Begin 1. M := 0; found := true 2. While *found* = true and \exists unexamined path $s_i \in P_1$ { $\{M := M+1 \}$ 2.1 If s_i matches some $s'_i \in P_2$ 2.2 Else {found := false } } 3. **If** $M = |P_1|$ 3.1 If $M = |P_2|$ {*relationship* := 'is the same to' } {*relationship* := 'is overridden by' } 3.2 Else 4. Else {*relationship* := 'is independent from' } 4. output relationship End

Fig. 12. Identifying preference relationships.

problem of finding the relationship of two selection preferences is translated to a mapping of the paths that connect the root to the leaves in the preference diagrams.

5.3 Preference Construction Analysis

Preference construction involves the following tasks: extraction of the top k preferences from a user profile that are related to and not conflicting with a given query, identification of any relationships among these preferences, and construction of the preference network.

Algorithm FIND_REL compares two preferences by comparing their respective sets of paths. We store sets of paths in memory as hash tables. The algorithm starts with the "smaller" preference, that is, the one that has the fewest paths (line 2, Figure 12) and it looks in the hash table of the larger set of paths for matching paths. Due to the use of hash tables, the matching cost per path is O(1). We consider that P_o is the maximum number of paths of any preference in the user profile. Hence, the cost for one invocation of FIND_REL is $O(P_o)$. Naturally, we do not expect P_o to take very large values. In the cases we have examined, we found P_o ranging from 1 to 4.

Algorithm PNET compares a preference to a set of preferences, for which their relationships are already known. It is called k times from PFC, once for each of

the top k preferences produced. When PNET examines the j^{th} preference (j = 1 to k), it performs a maximum of j - 1 comparisons. Each comparison is performed with the help of algorithm FIND_REL. Hence, it is $O(j * P_o)$.

Algorithm PFC extracts the top k preferences from a user profile. The algorithm first extracts all related preferences to Q that are found in the profile U (line 2). If relations are not repeated in Q, then it should be O(|Q| + |U|), where |Q| is the size of query Q in terms of relations and |U| is the size of the profile in terms of preferences. Assuming the profile is hashed on the relation for which each preference is defined, then for every query relation the algorithm finds the list of related preferences in O(1) time and then inserts into QP all preferences of that relation. Duplicate insertion can be avoided by maintaining a hash table of all relations examined so far. Hence, there are |Q| probes in the two hash tables and at most |U| preferences inserted into QP (if all preferences are related to the query). If we use $|U_Q|$ to signify the subset of |U| with the related and not conflicting preferences then the complexity is exactly $O(|Q| + |U_Q|)$.

The main loop in PFC (line 3) will be executed as many times as there are preferences added to QP, which are bound by |U| again, or more precisely, by the size of the transitively composeable subset of it $|U_Q^*|$. For k of these iterations, the cost of each one will be that of line 3.2.1, which is O(1) and the cost of a call to PNET, which we have analyzed earlier. For the remaining iterations, which should be $O(|U_Q^*|)$ in number, making the reasonable assumption that $k < |U_Q^*|$, the cost of each one is the sum of:

- —the cost for testing for conflicts, which can be considered equal to $O(P_o)$. Recall that P_o is the maximum number of paths of any preference in the user profile. Since all implicit preferences are generated by composing a join preference (that is always a single path) to another join or selection preference, P_o is the maximum number of paths in any preference, original or generated.
- —the cost for generating the path combinations, which for each iteration is $O(P_o)$, since a join preference from which we compose other preferences has only one path.

Hence, the total cost of PFC is $O(|Q| + |U_Q| + \texttt{k} * cost(\text{PNET}) + |U_Q^*| * P_o)$, which becomes $O(|Q| + |U_Q| + \texttt{k}^2 * P_o + |U_Q^*| * P_o)$ and finally $O(|Q| + (\texttt{k}^2 + |U_Q^*|) * P_o)$.

6. PERSONALIZED QUERY ANSWERING

Given a query Q and a network of k related preferences, the last step of personalization generates *all* query results that: (*a*) satisfy at least some of the preferences, that is, $1 \in [0..k]$ preferences, (*b*) respect the preference relationships, and (*c*) are ranked using a function *h* (Section 3.3). When 1 = 0 then all the results of the initial query are returned whereas when 1 > 0, the personalized answer may be smaller.

Relying directly on SQL to capture these semantics can lead to complex and time-consuming queries. The disadvantages of such approaches to query personalization have been studied in the literature [Koutrika and Ioannidis

ACM Transactions on Database Systems, Vol. 35, No. 2, Article 13, Publication date: April 2010.

13:28 • G. Koutrika and Y. Ioannidis

2005b], and hence are not discussed any further. Here, we describe two new approaches. These approaches work directly on the preference network. They take into account the containment mappings captured in the network in order to generate and rank results with respect to the preference relationships. The first algorithm (EXCLUDE_COMBINE) is based on the observation that an edge from p_i to p_j in the preference network implies a difference operation $Results(p_i) - Results(p_j)$ on the respective result sets in order to find the results that satisfy p_i but not p_j (i.e., with respect to the overriding relationship of p_i and p_j). Hence, instead of complex queries, it executes a set of simpler queries that capture these "local" difference operations and combines the partial results to find those that satisfy at least 1 preferences.

Algorithm EXCLUDE_COMBINE. The algorithm executes a number of simple queries. Each query Q_i is a combination of the initial query Q and a preference $p_i(i = 1...k)$ in the current network and returns the set $Results(p_i)$ of Q's results that match p_i . In this way, we go from a network of preferences to a virtual "network of partial result sets", where nodes and edges have the following meaning:

- —each node $Results(p_i)$ represents the results that satisfy preference p_i
- —for each edge from p_i to p_j in the preference network, there is an edge with the opposite direction in the network of result sets, that is, from node $Results(p_j)$ to node $Results(p_i)$, which implies a difference operation $Results(p_i) - Results(p_j)$ on the respective result sets in order to find the results that satisfy p_i but not p_j (i.e., with respect to the overriding relationship of p_i and p_j).

Then, for each edge from p_i to p_j , the algorithm removes from $Results(p_i)$ any results also found in $Results(p_j)$. These difference operations are sufficient to ensure that the tuples remaining in each result set will be those that satisfy the respective preference (i.e., the preference that initially generated this set) but not any preference more specific than that. The reason is that, by definition, $Results(p_j)$ contains all results corresponding to preferences more specific than p_j . Hence by removing $Results(p_j)$'s tuples from $Results(p_i)$, we satisfy all preference relationships between p_i and any preferences more specific than p_j . In the end, all tuples that occur in at least 1 sets satisfy at least 1 preferences, and they comprise the final answer ranked using a specified ranking function h.

Algorithm EXCLUDE_COMBINE is presented in Figure 13. Given a query Q and a network $G_H(V_H, E_H)$ of related preferences, it proceeds as follows.

- --(*Exclude*) It generates the partial result sets, each one satisfying a single preference (line 2). Each set is ordered on the tuple id *tid*. Then, it traverses the network and excludes from each set all tuples that satisfy other, more specific preferences (line 4).
- --(*Combine*) As long as there are 1 nonempty sets, the algorithm removes the greatest *tid*. If it satisfies at least 1 preferences, then its degree of interest *doi(tid)* is computed from the preferences corresponding to these sets and

Personalizing Queries Based on Networks of Composite Preferences • 13:29

Algorithm EXCLUDE COMBINE

```
Input:
            query Q, a constraint l
            a preference network G_H(V_H, E_H)
Output:
            personalized results Results
Begin
1. Results := \emptyset
2. Foreach p_i in V_H
                                                                 {
      Results(p_i) := execute\_query(Q \land p_i)
2.1.
                                                                 }
3. If l = 0
       Results(Q) := execute\_query(Q)
3.1.
       l := 1
3.2.
3.3.
       l_{copy} := 0
                                              }
4. Foreach not visited e(p_i, p) in E_H
                                                          ł
4.1.
       Results(p_i) := Results(p_i) - Results(p)
4.2.
       mark e visited
5. While at least l non-empty preference result sets exist
                                                                     ł
5.1. remove the head with the highest tid among all result sets
5.2.
      If tid satisfies at least l preferences
                                                   ł
         doi(tid) := h(\{d_i \mid d_i = doi(tid \in Results(p_i)),
5.2.1.
                                    \forall Results(p_i) \text{ having } tid \})
5.2.2. add (tid, doi(tid)) in Results
                                                                    }
6. If l_{copy} = 0 { Results := Results \cup Results(Q) }
7. output Results
End
```

Fig. 13. Generating personalized results - EXCLUDE_COMBINE.

(*tid*, *doi*(*t*)) is added in the *Results*, which is kept ordered on the degree of interest (line 5).

When 1 = 0, all the results of the initial query are returned and are ranked based on the preferences in the network. In that case, the algorithm executes query Q (line 3.1). The tuples returned, Results(Q), have default degree equal to zero (since at this point, we do not know if they satisfy any preferences.) The algorithm also sets 1 equal to 1 (line 3.2). This makes sure that the combine step (line 5) will be executed only if there is at least one preference satisfied, that is, if there is at least one nonempty preference result set $Results(p_i)$. Any results of Q that have not been ranked during the combine step are added to the final results (line 6).

Example. We consider the preference network of Figure 5. The algorithm first generates the partial result sets, each one satisfying a single preference and maps the preference network to a network of results sets shown in Figure 14. For instance, taking into consideration the edge from $Results(p_d)$ to $Results(p_c)$, $Results(p_c) - Results(p_d)$ is the "actual" set of results for which p_c should hold because the more specific p_d does not hold for them. Then, the algorithm proceeds as shown in Figure 15 excluding from each set the nonqualifying tuples, that is, those that satisfy a more specific preference.

13:30 • G. Koutrika and Y. Ioannidis



Fig. 14. Mapping a preference network to a network of result sets.



Fig. 15. Excluding results that satisfy more specific preferences.

Algorithm EXCLUDE_COMBINE executes all possible queries that map to the nodes of the preference network. The second algorithm (REPLICATE_DIFFUSE) that we present shortly tries to minimize the number of such queries based on the following observation: given a preference network, any tuple that satisfies any preference at any level in the network satisfies (at least) one of the root preferences. Hence, we can execute only the queries corresponding to the roots of the network and then perform "look-ups" to find the specific preferences satisfied by each retrieved tuple. In addition, the algorithm works with a hopefully smaller set of tuples than EXCLUDE_COMBINE.

Algorithm REPLICATE_DIFFUSE. This algorithm is based on the idea of visualizing a preference network as a system of pipes with preference nodes acting as safety valves. When a tuple enters the system at some node (as a result of satisfying the corresponding preference), it rolls down the pipes (edges) starting from this node as long as there is a safety valve that can be "opened". A safety valve will remain closed to a tuple, if the latter satisfies the preference

Personalizing Queries Based on Networks of Composite Preferences • 13:31

Algorithm REPLICATE _ DIFFUSE

```
Input:
            query Q, a constraint l
            a preference network G_H(V_H, E_H)
Output:
            personalized results Results
Begin
1. Results := \emptyset
2. If l = 0
                { Results(Q) := execute\_query(Q)
                                                             }
3. Foreach root p_{rt} \in V_H in order of selectivity
                                                                      {
       Results(p_{rt}) := execute\_query(Q \land p_{rt})
3.1.
3.2.
       Foreach tid \in Results(p_{rt}) with tid \notin Results
                                                                     ł
3.2.1.
         P := \{p_{rt}\}
3.2.2.
         execute_query(Q_{rt}(tid))
3.2.3.
         add to P all roots satisfied by tid
3.2.4.
         While \exists p \in P with e(p, p_i) \in E_H
                                                                  {
            overridden := false
            Foreach e(p, p_i) \in E_H with p_i \notin P
                                                          {
               If execute_query(Q_i(tid)) \neq \emptyset  {
                  overridden := true
                  add p_i in P
                                                           }
                                                  }
               If overridden = true {
                  remove p from P \rightarrow \{
                                                             }
3.2.5.
         If there are at least l preferences in P {
            doi(tid) := h(P)
            add (tid, doi(tid)) in Results
3.2.6. remove tid from Results(Q)
                                                                  } }
4. If l = 0
                     { Results := Results \cup Results(Q) }
5. output Results
End
```

Fig. 16. Generating personalized results - REPLICATE_DIFFUSE.

corresponding to the valve, but the valve leads to no other preferences that can be satisfied. Moreover, a tuple satisfying a preference at any level of a network satisfies its ancestors too. This means that any tuple satisfying at least one preference in the network will "enter" the system from the network's roots and roll down following edges from general to more specific preferences, until it is collected by valves mapping to the most specific preferences satisfied by it.

Given a query Q and a preference network $G_H(V_H, E_H)$, algorithm REPLI-CATE_DIFFUSE proceeds in three steps, repeated for each root of the network. It creates a set of queries, each one combining Q with a root preference, in order of increasing selectivity.

- -(*Create*) For each root, the algorithm first executes the respective query and creates the set of results that satisfy this preference (line 3.1).
- -(*Replicate*) For each tuple in this set, it finds which root preferences following the current one in order of selectivity are also satisfied (line 3.2.2).

13:32 G. Koutrika and Y. Ioannidis



Fig. 17. Example of replicate-diffuse steps.

-(*Diffuse*) Then, tuples satisfying the root preferences roll down the network. In the end, all tuples are found only in nodes corresponding to the most specific, independent preferences they satisfy (line 3.2.4).

Example. Figure 17 illustrates these steps for the root p_a . The algorithm retrieves Q's tuples that satisfy p_a , that is, three tuples depicted as different shapes. Next, it examines whether any of these satisfy p_f . Assume that the triangle and square tuples do. To illustrate that, they are shown replicated on this node. Then, all tuples move freely down the network until they are stopped by some valve. For instance, we find that the square tuple satisfies the independent preferences p_d and p_f .

For each root p_{rt} in the network, the algorithm executes a query $Q \wedge p_{rt}$. Each tuple (with id *tid*) returned by this query is processed only once as follows. For *tid*, a set P of preferences satisfied is kept, initially containing p_{rt} . In order to find which other preferences are satisfied by *tid*, the algorithm executes a parameterized query $Q_{rt}(tid)$, which checks whether *tid* satisfies any other root preference following p_{rt} in order of selectivity (line 3.2.2). This query returns zero or more occurrences of *tid*, depending on the number of root preferences that are satisfied by *tid*. All root preferences satisfied by *tid* are placed in P (line 3.2.3).

Then, for each preference p in P, the algorithm checks whether there are more specific preferences that override p. For each edge $e(p, p_i)$, the algorithm executes a parameterized query $Q_i(tid)$ that checks whether tid satisfies p_i . If that happens, then p_i is inserted in P. A preference p stays in P if tid does not satisfy any more specific preference p_i . Note that when considering an edge $e(p, p_i)$, if p_i is already in P, this means that $Q_i(tid)$ has been executed in a previous step (that happens when p_i overrides more than one preference.) In that case, the algorithm does not execute $Q_i(tid)$ again. At the end of this process, P contains the most specific, independent preferences satisfied by tid. If there are at least 1 such preferences, then the degree of interest in *tid* is computed from P using a function h and *tid* is inserted in the final results (line 3.2.5).

When 1 = 0, all the results of the initial query are returned and are ranked based on the preferences in the network. In that case, the algorithm executes query Q (line 2). The tuples returned, Results(Q), have default degree equal to zero as explained for algorithm EXCLUDE_COMBINE. Any tuple id *tid* that is found to satisfy some preference in the network gets removed from Results(Q)

(line 3.2.6). At the end, any remaining results in Results(Q) are added to the final results (line 4).

6.1 Personalized Query Answering Analysis

Let v and r be the number of nodes and edges in a preference network, respectively. We can decompose algorithm $\text{Exclude}_{\text{COMBINE}}$ into a number of operations, some of which involve processing of arbitrary queries (lines 2, 3), whose cost we cannot estimate analytically.

(a) To generate the partial results sets (line 2), the algorithm executes a query $Q \wedge p_i$, for each preference p_i , i = 1...k. In the case of 1 = 0 (line 3), the original query Q is also executed on its own (for ease of reference, we may denote it as $Q \wedge p_0$, where $p_0 = \text{'true'}$). The total cost of these operations is

$$\sum_{i=0}^{ extsf{k}} cost(Q \wedge p_i).$$

- (b) After the results are gathered in main-memory hash tables, the algorithm executes r difference operations (line 4). Assuming that each result set contains at most n tuples, then each such difference can be executed in O(n) time, for a total of O(r*n) time.
- (c) Finally, the algorithm reads the in-memory results to output those satisfying at least 1 preferences. Each iteration of the loop (line 5) requires O(k) time: (line 5.1) examines each one of the k result sets to remove the highest-tid common head, while (line 5.2) uses at most k doi's to compute the overall doi of the tid examined. In the worst case, the loop of (line 5) will consume all tuples in the results, of which there are at most k*n, and will do that by removing the smallest possible number of them each time, which is equal to 1, this way maximizing the number of iterations required. Hence, the worst-case total time required by this operation is $O(k^2/1*n)$. Taking all the preceding into account, the total cost of algorithm EXCLUDE_COMBINE is

$$\sum_{i=0}^{k} cost(Q \wedge p_i) + O(\texttt{r}*\texttt{n} + \texttt{k}^2/\texttt{l}*\texttt{n}).$$

Similarly to the previous analysis, for algorithm REPLICATE_DIFFUSE, we consider the following:

(a) For each root p_{rt} (line 3.1), the algorithm executes a query $Q \wedge p_{rt}$. In the case of nonpersonalizing the output tuples (line 2), the algorithm executes the original query Q on its own as well (again, for ease of reference in this case, let us assume that a preference p = true' is a root as well). The total cost of all these query executions is

$$\sum_{p_{rt} \text{ is root}} cost(Q \wedge p_{rt}).$$

(b) For each tuple returned from the aforesaid queries, the algorithm first executes a parameterized query $Q_{rt}(tid)$ (line 3.2.2) and then parameterized queries $Q_i(tid)$ (within line 3.2.4) to check which other (less selective) root

13:34 • G. Koutrika and Y. Ioannidis

preferences or other preferences down the network, respectively, the tuple satisfies. Note that $Q_{rt}(tid)$ is equivalent to running a $Q_i(tid)$ for each of the roots that follow the one being considered in the main loop of the algorithm. Each *tid* undergoes this process only once and is used as a parameter for each query only once as well. Hence, assuming there are m tuples in all root query results together and that the cost of each parameterized query remains the same independent of the *tid* used as a parameter, in the worst case (when all tuples satisfy all preference queries), the total cost of these operations is

$$\mathtt{m}*\sum_{i=1}^{\mathtt{k}} cost(Q_i(.))$$

(c) Besides query executions, the main-memory operations of the algorithm are inside the loop of (line 3.2). For each one of its m iterations, they are performed either once for every one of the r edges of the network (line 3.2.4 - top) or once for every one of the v preference-nodes of the network (line 3.2.4 - inner foreach-loop). In addition, each iteration uses again at most k degrees of interest to compute the overall degree of interest of the *tid* examined (line 3.2.5). Hence, the worst-case total time required by these operations is O(m * (r + v + k)).

Taking all the preceding into account, the total cost of algorithm REPLI-CATE_DIFFUSE is

$$\sum_{p_{rt} \text{ is root}} cost(Q \land p_{rt}) + \texttt{m} * \sum_{i=1}^{\texttt{k}} cost(Q_i(.)) + O(\texttt{m} * (\texttt{r} + \texttt{v} + \texttt{k})).$$

7. EXPERIMENTS

The novelty of our framework stems from allowing both generic and specific preferences to be explicitly stated in a profile with user-specific, "freely" selected degrees of interest assigned to them. Earlier approaches (e.g., [Koutrika and Ioannidis 2004; Stefanidis et al. 2007]) capture only simple preferences, that is, preferences on single attributes, and use special mechanisms to derive complex preferences. ones Two questions naturally arise.

- —"Simplicity or expressiveness?" One question is whether the proposed model can have a higher impact than an approach to preference modeling that captures simple preferences (i.e., in the form of general rules that hold for a user) and derives more complex preferences using appropriate preference composition and inference mechanisms. One can argue that "simple is beautiful." We test this argument through a user study (Section 7.1).
- ---*"Expressiveness or performance?"* When a profile contains a mix of generic and more specific preferences, preferences cannot be freely combined. We need sophisticated algorithms in order to find which preferences can be combined and how to generate and rank the results of a query taking into account relationships among a set of preferences. Is expressiveness achieved at the

expense of performance and to what extent? Does the trade-off justify the use of this framework instead of one that assumes independent preferences? We investigate these questions through an extensive experimental evaluation of our algorithms (Section 7.2).

7.1 User Study

Existing approaches (e.g., [Koutrika and Ioannidis 2004]) formulate structurally (and semantically) simple preferences and they provide mechanisms for composing complex ones. We believe that these mechanisms cannot always capture the real complex user preferences. We need to enable complex preferences to be explicitly stated (not be solely derivatives of other preferences). The objective of the user study is to provide insight as to the appropriateness of our preference model and its benefits for query personalization compared to simpler preference representations and to highlight the accuracy problem of the latter.

We conducted an empirical evaluation of our approach with 11 human subjects with or working towards a diploma in computer science. We built a database containing information about over 480,000 movie titles that we obtained from IMDB (http://www.imdb.com/interfaces). Our schema contains 18 relations storing information about movies, actors, directors, ratings, writers, and so forth. We built a Web interface that allowed users to manually create their preference profiles. We explained to them that they had two options: either providing a profile (NET_PROFILE) of complex preferences (such as a preference for comedies with Woody Allen) or a profile (FLAT_PROFILE) of simple, independent, preferences (such as a preference for complex preference for woody Allen). In the latter case, we would infer their more complex preferences based on what we knew about a user's generic preferences following the model presented in Koutrika and Ioannidis [2004].

An initial test for the appropriateness of our model took place during the creation of these profiles. 8 out of 11 people created a NET_PROFILE, 2 started with simple preferences and decided that they had more complex preferences to express and only 1 provided a FLAT_PROFILE. Figure 19 shows part of one of the profiles that we have in the system. Once done with their profile (in either form), all users were asked to create a second "view" of their preferences either more elaborate or more generic depending on what was their first profile. Hence, 10 people had to create a FLAT_PROFILE and 1 had to elaborate his preferences in a NET_PROFILE. At the end, for each user there were two profiles in the system, both a NET_PROFILE and a FLAT_PROFILE. The largest group of users complained for "reformulating" their preferences as simpler ones expressing their concerns regarding the system's ability to provide accurate personalization when it relies on simple, partially correct, profiles. These observations seem to indicate that people tend to trust more a system that captures their preferences more accurately.

There is, of course, a psychological dimension in how people perceive preferences and personalization. For example, the user who initially formulated simple preferences explained that he would be probably content with a simple

13:36 • G. Koutrika and Y. Ioannidis

personalization approach separating the chaff from the seed. In addition, the type of networks expected in practice depends on the way user information is (explicitly or implicitly) collected. Different people have different types of preferences, depending on their background, their expectations, etc. In our experiments, "full-fledged" networks that could be derived from the complex profiles had an average depth of 3, and the average number of preferences stored per profile was 21.

We built an interface for searching in the database for movies, actors, or directors based on different criteria, such as the genre and the year of a movie. We discussed with the users about typical searches they would perform over a movie database and we made a pool of possible searches, from which we picked the 4 most popular queries among users. These included a search for short movies (Q_1) , a search for movies based on the movie genre (Q_2) , and a search for favorite directors that have directed movies after 2000 (Q_3) . All subjects were asked to perform these 4 preselected searches plus 2 of their own choice.

In order to evaluate the effectiveness of our approach, we had to take into account several issues. We could not use traditional metrics, such as precision and recall. One reason is that we do not have complete information regarding which of the (possibly too many) results of a search are liked by a user. Even if we showed all results returned for a particular, nonpersonalized, search, users are not willing to browse more than one page of results. In addition, we wanted to get insights into how personalized search results compare to nonpersonalized results. However, there is no single uniform way to compare them all together. For example, in the case of the unranked results, all results are relevant to the query (by definition of the database query model) and there is no way to distinguish between them. On the other hand, when comparing ranked lists we are actually interested in how the results are ranked.

Taking all these issues into account, we proceeded in two steps. First, we asked users to give scores to query answers as a whole, where an answer would contain at most 10 results, in order to have an overall picture of how personalized and nonpersonalized answers compare overall. Then, we focused on investigating how effective the rankings are based on different preference representations.

For the first part of the evaluation, each user submitted the queries three times in arbitrary order. Queries were executed once without personalization (NO_PERS), once using the user's NET_PROFILE and once using the user's FLAT_PROFILE. The system randomly rotated these options so that the user would not be aware of the query processing performed. As default parameters for personalization, we chose k to be half of the preferences in a user's profile and 1 = 1. We ranked results based on the average degree of interest of the preferences satisfied and returned (up to) top-10 results. In the case of NO_PERS, the first 10 results for the query were returned. Users evaluated query answers using two scores measuring [Koutrika and Ioannidis 2004]: (a) the difficulty to find something interesting, if anything was found at all (DEGREE OF DIFFICULTY) and (b) how satisfactory was the answer (ANSWER SCORE) (both scores in the range [0, 10].)



(c) result ranking

Fig. 18. Impact of composite preferences and benefits of query personalization.

Figures 18(a) and 18(b) present the average answer score and degree of difficulty, respectively, per query for each of the three different runs, that is, NO_PERS, NET_PROFILE and FLAT_PROFILE. The use of the complex profiles substantially reduces the difficulty to find interesting tuples within an answer and attracts distinctively higher answer scores. Using the flat profiles improves answers (to a lesser degree than NET_PROFILE). However, we observe that often the improvement is marginal compared to NO_PERS and the degree of difficulty in many cases remains high.

13:38 • G. Koutrika and Y. Ioannidis

MOVIE.year > 1990,	0.8
GENRE.genre = "comedy",	0.7
GENRE.genre = "adventure",	0.85
GENRE.genre = "drama",	0.2
GENRE.genre = "thriller",	0.9
GENRE.genre = "mystery",	0.9
GENRE.genre = "animation",	0.6
DIRECTOR.name = "W. Allen",	0.6
DIRECTOR.name = "S. Spielberg",	0.8
PRODUCER.name = "J. Bruckheimer",	0.6
PRODUCER.name = "S. Spielberg",	0.75
ACTOR.name = "K. Spacey",	0.8
ACTOR.name = "G. Clooney",	0.75
ACTOR.name = "M. Freeman",	0.8
ACTOR.name = "H. Ford",	0.9
ACTOR.name = "A. Sandler",	0.8
ACTOR.name = "B. Stiller",	0.7
MOVIE.mid=GENRE.mid,	0.8
MOVIE.mid = DIRECTED.mid and DIRECTED.did = DIRECTOR.did,	0.9
MOVIE.mid = CAST.mid,	1
MOVIE.mid = CAST.mid and CAST.aid=ACTOR.aid,	1
MOVIE.mid = GENRE.mid and GENRE.genre = "TV series" and MOVIE.mid = CAST.mid and CAST.aid = ACTOR.aid and ACTOR.name = "B. Stiller",	0.1
MOVIE.year > 2000 and MOVIE.mid = GENRE.mid and GENRE.genre = "animation",	0.95
MOVIE.mid = GENRE1.mid and GENRE1.genre = "animation" and MOVIE.mid = GENRE.mid and GENRE.genre = "thriller",	0.1
MOVIE.year > 1995 and MOVIE.mid = DIRECTED.mid and DIRECTED.did = DIRECTOR.did and DIRECTOR .name = "W. Allen",	0.8
CAST.aid=ACTOR.aid and ACTOR.name = "O. Wilson" and CAST.aid=ACTOR1.aid and ACTOR1.name = "B. Stiller",	0.4
CAST.aid=ACTOR.aid and ACTOR.name = "G. Clooney" and CAST.aid=ACTOR1.aid and ACTOR 1.name = "B. Pitt",	1
CAST.aid=ACTOR.aid and ACTOR.name = "G. Clooney" and CAST.aid=ACTOR1.aid and ACTOR 1.name = "E. McGregor" and CAST.aid=ACTOR2.aid and ACTOR 2.name = "K. Spacey",	1
MOVIE.mid=GENRE.mid and GENRE.genre = "comedy" and MOVIE.mid = DIRECTED.mid and DIRECTED.did = DIRECTOR.did and DIRECTOR.name = "G. Clooney",	0.6
MOVIE.mid=GENRE.mid and GENRE.genre = "drama" and MOVIE.mid = DIRECTED.mid and DIRECTED.did = DIRECTOR.did and DIRECTOR.name = "G. Clooney",	0.85

Fig. 19. Part of a real user profile with movie preferences.

During the study, users were also asked to rerank the tuples returned for each query and put them in an order that they thought closer to their preferences. For the second part of the study, we compared the ordering of results using the complex profile or the flat profile to the user's ordering of the same results for each query. There are several standard methods for comparing two lists that are permutations, such as Kendall's tau distance (τ) [Fagin et al. 2003; Kendall and Gibbons 1990]. We used the normalized Kendall tau distance (τ), which lies in the interval [0,1] with 1 indicating maximum disagreement. $\tau_{\rm FU}$ compares the list of results based on the FLAT_PROFILE and the same list reordered by the user. Figure 18(c) plots the average distances for each query. We observe that when the FLAT_PROFILE was used, users often disagreed with the system-based ordering of results because it did not quite capture their actual, elaborate preferences.

A closer look at the user-ordered results revealed that often 1 to 3 tuples appeared out of order in the system answers due to the less accurate profiles

Personalizing Queries Based on Networks of Composite Preferences • 13:39

MOVIE.year > 1990, 0.8	MOVIE.mid=GENRE.mid and GENRE.genre = "animation", 0.48
MOVIE.year > 2000 and MOVIE.mid = GENRE.mid	MOVIE.mid = GENRE1.mid and GENRE1.genre = "animation" and
and GENRE.genre = "animation", 0.95	MOVIE.mid = GENRE.mid and GENRE.genre = "thriller", 0.1

Fig. 20. Part of a network of preferences related to a particular search.

used. On the other hand, finer-grained result rankings could be achieved with the complex profiles that captured more accurately user expectations. The accuracy achieved also depends on the query. For example, for the query Q5, none of the profiles had adequate (or different) information for differentiating the output of this query. The figure brings up another issue: one would expect that since the complex profiles captured more accurately the user preferences, $\tau_{\rm HU}$ would be 0 in all cases. Users still moved tuples in the results for different reasons (e.g., they knew additional information, such as movie reviews, or they evaluated results based on additional properties not captured in their profiles.)

Shortly, we describe some interesting cases of personalized searches that we observed during the study in order to illustrate the impact of our approach. Figure 19 shows part of a participant's profile. This profile was one of the most detailed ones. It contains some general preferences for movie genres, actors and so forth, as well as finer-grained preferences. In one of the searches, the user asked for comedies after 2008 that satisfy at least one of her top 5 preferences. Her top 5 preferences are the following:

MOVIE.mid = CAST.mid and	
CAST.aid = ACTOR.aid and ACTOR.name = "G.Clooney"	
CAST.aid = ACTOR1.aid and ACTOR1.name = "B.Pitt"	1
MOVIE.mid = CAST.mid and	
CAST.aid = ACTOR.aid and ACTOR.name = "G.Clooney	
CAST.aid = ACTOR1.aid and ACTOR1.name = "E.McGregor"	
CAST.aid = ACTOR2.aid and ACTOR2.name = "K.Spacey"	1
MOVIE.year > 2000 and MOVIE.mid $= $ GENRE.mid	
GENRE.genre = "animation"	0.95
MOVIE.mid = CAST.mid and	
CAST.aid = ACTOR.aid and ACTOR.name = "H.Ford"	0.95
MOVIE.year > 2000 and MOVIE.mid $=$ GENRE.mid and	
MOVIE.mid = GENRE.mid and GENRE.genre = "animation"	0.95

It is interesting that the user had preferences for specific actors but she had given very high preferences for groups of actors participating in the same movie as a wishlist. While combining her individual preferences would not lead to these strong preferences, being able to code such fine-grained taste made possible to discover a new movie with title "The Men Who Stare at Goats" that satisfies her wish that if her favorite actors acted together, this would make a not-to-miss movie.

Another search performed by the user was for short movies (duration < 1h 45min) that would satisfy at least 2 preferences. Part of the network built for this query is shown in Figure 20. Short movies satisfying the user preferences were some animation movies. Using this network, we ranked some animation

13:40 • G. Koutrika and Y. Ioannidis

movies, such as the "Ice Age", high but gave low scores to other animation movies, such as the "Corpse Bride". Using the user's simpler profile would not allow us to compute such accurate rankings because the complex preferences for animations are not derivable from the general preferences for animations and thrillers. If we combined these preferences, the "Corpse Bride" would rank high.

7.2 Performance Study

A user may explicitly give only a handful of preferences and implicit sources of preference elicitation, such as log mining and relevance feedback, may indicate many (additional) preferences. In both cases, incomplete information may prevent deriving a small set of generic preferences and a large number of preferences may be recorded that possibly contain many relationships. Handling multi-granular preferences and preference relationships add up to increased complexity and, thus, higher execution times. Therefore, in this section, we evaluate the performance of query personalization using our preference model (specific details are given per experiment.) We considered query loads of 50 random queries representing hypothetic user queries, each one containing one relation and one selection on this relation. Moreover, we divided the query personalization process in two main phases a preference construction phase which builds a network of related preferences for a query, and a preference query answering phase which generates personalized results using this network. We study each phase separately and identify the critical factors that affect their performance before discussing the overall performance of personalization. Two parameters are important.

-the number k of selection preferences manipulated

—the number r of $\tilde{\sqsubset}$ -relationships existing among them

We tested our algorithms and we compared them with simpler algorithms that operate on the assumption of independent preferences. We generated synthetic profiles and networks depending on the requirements of each experiment, as we explain in the following subsections, and we built the indexes required for the processing. Times are shown in ms.

7.2.1 *Preference Construction*. This phase involves two tasks: extraction of the related preferences from a profile and construction of their network. These tasks are interleaved during preference construction (Section 5). We measured:

- —the total time required to extract ${\tt k}$ selection preferences from a user profile ${\tt TIME}_{\tt EXTRACT}$
- —the total time required to find the relationships among k preferences and build their network, that is, the total time required by the *Preference-Network Integration* and *Relationship Finder* modules - TIME_{REL}

The time TIME_{REL} for processing independent preferences, as well as the extra time in $\text{TIME}_{\text{EXTRACT}}$ for preparing the structures required in preference comparisons, such as representing preferences as sets of path strings, are negligible.



(c) extraction & network integration vs. r

Fig. 21. Times for extracting and integrating k preferences with r preference relationships.

This observation allows us to consider that when having only independent preferences, the required execution time for this phase is equal to the time $\text{TIME}_{\text{EXTRACT}}$, and hence, the overhead from the processing of composite preferences is reflected in TIME_{REL} .

 $\text{TIME}_{\text{EXTRACT}}$ depends on the number k of preferences handled, as Figure 21(a) confirms. The task of extracting related preferences does not simply read preferences stored in a profile but it also builds new preferences from the stored ones. Depending on the database schema and the number of selection preferences

13:42 • G. Koutrika and Y. Ioannidis

defined per relation, it may be able to retrieve k stored selection preferences with a few database accesses or it may need to search in relations further away from the initial query relations and compose preferences out of many stored ones. In order to gain insights into the impact of this phenomenon over TIME_{EXTRACT}, we generated the schema of a hypothetic database comprised of 100 relations, each one having 3 attributes, one of which possibly joining this relation to another. Then, we generated two synthetic profile databases of 100 profiles each. Profiles in both databases contained 100 selection preferences each but were generated in a different way: a profile in PROFDB1 was generated with the constraint that each attribute of the database could be used in at most one preference, while a profile in PROFDB2 was generated with the constraint that each attribute of the database should be used in at most three selection preferences. Consequently, profiles in PROFDB1 were sparser than profiles in PROFDB2. We considered a query load of 50 queries over our synthetic data and we measured times for extracting the top k preferences for each query. Figure 21(a) shows TIME_{EXTRACT} as a function of k over all (100) profiles of each profile database for the 50 queries. The difference in execution times can be interpreted as the algorithm's effort to collect k preferences depending on how preferences are distributed over the synthetic data. When preferences are sparsely placed over data (e.g., in PROFDB1), it takes substantially more effort as k increases because it needs to search more in a profile.

Figures 21(b) and 21(c) show execution times for finding k preferences $(TIME_{FXTRACT})$ and constructing a network of k preferences having r relationships ($TIME_{REL}$). For the former, we considered the profiles in PROFDB1 and the query load of 50 queries used in the previous experiment. Measuring the latter, especially with respect to r, is tricky. We wanted to explicitly control the number of relationships r between k preferences in order to measure their impact. For this reason, we ignored the actual relationships among the k preferences found from the extraction step and we built a synthetic network generator, which takes as inputs the number k of preferences and the number r of preference relationships, and generates a network with these characteristics, that is, containing $r \tilde{\perp}$ -relationships between k preferences, generated in a random way but with respect to certain constraints, for example, defining at most one relationship per pair of preferences. The set of preferences and relationships of a synthetic network are fed into the *Relationship Finder*, which can guide Preference-Network Integration to build the same network from scratch. The latter is presented with the set of preferences of the synthetic network in random order and asks the *Relationship Finder* regarding \Box -relationships in order to build the network.

Figure 21(b) shows execution times as a function of k assuming r = 5. Each point in the figure is the average of execution times for 100 networks with the same characteristics. We observe that TIME_{EXTRACT} dominates and that is due to the number of database accesses required for the extraction of the top-k preferences (which increases with k), whereas the construction of the network is an in-memory process. In addition, we observe that although the latter phase depends strongly on k, the overhead of comparing more preferences as k increases is not as large as one might expect. The reason is that although we

increase k, there are always r = 5 preference relationships to be detected. That means that each time a pair of preferences is compared, they are most likely independent and hence the comparison finishes instantly.

Figure 21(c) shows execution times as a function of r assuming k = 50. TIME_{EXTRACT} is constant with r. We observe that as the preference network gets more complicated (i.e., with r increasing), TIME_{REL} deteriorates and exceeds TIME_{EXTRACT}. With more preference relationships, more complicated comparisons between preferences take place.

Consequently, each task's contribution to the total execution time for preference construction is different: one task may dominate the other depending on the parameter settings. Ultimately, the overhead from constructing a network for a query may come not from the number of preferences but from the number of relationships among them.

7.2.2 *Preference Query Answering*. For this phase, we evaluated the performance of the algorithms proposed and the overhead occurred due to the preference networks. Hence, we measured:

- —the execution time required by the EXCLUDE_COMBINE algorithm $TIME_{EXCLUDE}$
- -the execution time required by the REPLICATE_DIFFUSE algorithm TIME_DIFFUSE
- —the execution time required by a SIMPLE algorithm that works with independent preferences $\mathsf{TIME}_{\mathsf{SIMPLE}}$

We built SIMPLE as a simpler version of algorithm EXCLUDE_COMBINE that treats preferences as independent and hence does not have the exclude step of the latter. SIMPLE executes a set of queries that integrate the preferences in the query, and combines their results returning those satisfying at least 1 preferences. For this series of experiments, we used the movie database. We generated a set of 50 random queries representing hypothetic user queries, each one containing one relation and one selection on this relation. Note that experiments with other sets of queries with different features, for instance, with one join and one selection, show similar trends and are not discussed. We generated different sets of preference networks over the movie database, each set containing 50 networks with k of preferences and r edges. Each network was meant to be combined with one query from the preceding set, and was generated as follows: we first composed k independent selection preferences for our movie database related to the query, with the constraint that they contained only one selection (but any joins required.) Then, we ran an iterative procedure that randomly picked r pairs of preferences and combined them to complex preferences. In essence, at each round, in a pair (p_i, p_j) , p_j was replaced by $p_i \wedge p_j$ to form a more specific preference than p_i .

Figure 22(a) shows execution times as a function of k with r = 5. Each point in the figure is the average of execution times for the 50 queries combined with their respective networks for the same k and r values. EXCLUDE_COMBINE requires more time than SIMPLE. Since they are both built on the same philosophy, the overhead observed is due to the additional actions required by preference relationships. Algorithm REPLICATE_DIFFUSE shows the best performance. It executes as few queries as possible and process a smaller number

ACM Transactions on Database Systems, Vol. 35, No. 2, Article 13, Publication date: April 2010.

13:44 • G. Koutrika and Y. Ioannidis



Fig. 22. Times for query answering using k preferences organized in a network with r relationships.

of tuples than EXCLUDE_COMBINE close to the final number of results returned. These queries depend on the number of root preferences in the network. As k increases, more root preferences emerge but they are fewer than k.

Figure 22(b) shows times as a function of r with k = 50. While TIME_{EXCLUDE} deteriorates when more preference relationships need to be resolved, TIME_{DIFFUSE} is actually benefited because the total number of queries executed is lower. That is fewer queries of the type $Q \land p_{rt}$ may be executed (as r increases, fewer root preferences may exist) or fewer parameterized queries Q_i (tuples may be distributed at different nodes in the preference network). Hence, REPLI-CATE_DIFFUSE for generating results based on preference networks exhibits a better behavior in contrast to a naïve approach that works with independent preferences because it exploits preference relationships.

Overall, our results so far indicate that allowing complex preferences in query personalization may make preference construction more expensive but benefit query answering. How do these phenomena shape the total execution time?

7.2.3 *Overall Performance*. To complete the picture regarding the efficiency trade-offs, we compare the contribution of all parts in the total processing time of a query. We consider a query load of 50 queries over the movie database (each one containing one relation and one selection on this relation) and a set



Fig. 23. Overall times for personalizing queries using k preferences organized in a network with r relationships.

of 50 profiles and we plot the time $TIME_{EXTRACT}$ required to extract k selection preferences from a profile related to a query, the time $TIME_{REL}$ required to find the relationships among these preferences and build the network, and the time $TIME_{ANSWER}$ to generate the personalized results. For the latter, we consider the time $TIME_{DIFFUSE}$, since REPLICATE_DIFFUSE is the most efficient.

Figure 23(a) shows the execution times as a function of k for r = 5 and Figure 23(b) shows the execution times as a function of r for k = 50. Each point in the graphs is the average of execution times over the sets of queries and profiles for this experiment. When k increases, the execution time of the preference query answering phase (TIME_{ANSWER}) shapes the overall performance. This time is measured on the right y-axis in Figure 23(a). On the other hand, when r increases (Figure 23(b)), we witness the effect of two opposite forces: decreasing time for generating results tends to "compensate" for the increase in the time required for building the network. Thus, the whole personalization process behaves in a more balanced way.

Finally, Figures 24(a) and 24(b) compare the efficiency of personalization using networks (TIME_{NET}) with personalization using only simple, independent preferences (TIME_{SIM}). Overall, performance is not sacrificed for expressiveness because we can use algorithms that benefit from preference relationships to adapt more smoothly to changes to k or r.

13:46 • G. Koutrika and Y. Ioannidis



(a) impact of expressiveness on performance vs. k



(b) impact of expressiveness on performance vs. r

Fig. 24. Impact of expressiveness.

8. CONCLUSIONS AND FUTURE WORK

In this article, we have studied query personalization based on multi-granular, composite preferences. We have introduced the concept of a preference network to capture preferences and their relationships. We have described algorithms for identifying preference relationships, constructing a preference network for a query, and using the network to generate personalized answers. We have evaluated our framework and algorithms both for efficiency and answer accuracy by comparing with simpler models and algorithms, which assume that preferences are independent, and we have shown that it is feasible to generate personalized answers that more accurately capture preferences without losing in efficiency.

Our approach is applicable to any graph model representing information at the level of entities and relationships. User preferences may be articulated over a higher-level graph model representing the data over the database schema. This is a useful abstraction for using a profile over multiple databases with similar information but possibly different schemas, and for hiding database restructuring and normalization. Preferences expressed at a higher level can be transparently mapped to the underlying schema for executing queries. Schema mappings have been studied in the literature, examining them in the context of query personalization seems an interesting direction.

Processing queries with complex preferences is a challenging problem. Preferences that are related to a query may be redundant. A preference for comedies and a preference for comedies with B. Stiller can never be concurrently met. Hence, the effect of query personalization may vary depending on how we interpret such preferences. For example, a personalized answer could be shaped by the set of top, *mutually independent* preferences related to the query. Obviously, there are different possible definitions of the "desired" set of preferences that need to be investigated and tested through user studies. For instance, one could consider the set of independent preferences that when combined would produce the highest degree of interest.

Furthermore, a number of interesting issues arise regarding the management of complex preferences, for instance, how to exploit preference commonalities and access patterns in the same or different profiles for developing efficient storage and access methods and structures for complex preferences, or for materializing parts of preference networks that are shared among multiple profiles or that are frequently applied for query personalization.

For the long run, a more "holistic" query optimizer that could take both a user query and a structured preference set into account for optimization is an intriguing direction to pursue. There are several possibilities to explore towards this end. For example, one could implement a special operator that ranks an input relation with respect to a preference relation in the spirit of Koutrika et al. [2009]. Alternatively, one could imagine dynamic optimization techniques building on the concept of Eddies [Avnur and Hellerstein 2000]: the preferences act as filter operators and the Eddy routes tuples to them choosing adaptively the way each tuple is routed. A tuple is sent to the output when it has been handled by at least 1 operators.

Designing intuitive GUI's that facilitate defining and editing preferences is an open issue. In the same direction, designing interfaces allow the user to customize the extent of query personalization and get explanations regarding the effect of query personalization is also challenging. Finally, having the freedom to capture multi-granular preferences, we can also tune user profiling methods, such as user log mining, to store richer preferences in profiles instead of being forced to construct simpler profiles.

REFERENCES

- AGRAWAL, R., RANTZAU, R., AND TERZI, E. 2006. Context-Sensitive ranking. In Proceedings of the ACM SIGMOD International Conference on Management of Data. 383–394.
- AGRAWAL, R. AND WIMMERS, E. 2000. A framework for expressing and combining preferences. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*.
- AHO, A., SAGIV, Y., AND ULLMAN, J. D. 1979. Equivalence of relational expressions. *SIAM J. Comput.* 8, 2, 218–246.
- AVNUR, R. AND HELLERSTEIN, J. M. 2000. Eddies: Continuously adaptive query processing. In Proceedings of the ACM SIGMOD International Conference on Management of Data.
- BALABANOVIC, M. AND SHOHAM, Y. 1997. Fab: Content-Based, collaborative recommendation. Comm. ACM 40, 3, 66-72.
- BALKE, W.-T., GUNTZER, U., AND LOFI, C. 2007. User interaction support for incremental refinement of preference-based queries. In *Proceedings of the IEEE International Conference on Research Challenges in Information Science (RCIS).* 511–523.

13:48 • G. Koutrika and Y. Ioannidis

- BORZSONYI, S., KOSSMANN, D., AND STOCKER, K. 2001. The skyline operator. In Proceedings of the International Conference on Data Engineering (ICDE'01). 421–430.
- BRUNO, N., CHAUDHURI, S., AND GRAVANO, L. 2002. Top-k selection queries over relational databases: Mapping strategies and performance evaluation. ACM Trans. Datab. Syst. 27, 2, 153–187.
- CHANG, K. AND HWANG, S. 2002. Minimal probing: Supporting expensive predicates for top-k queries. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*.
- CHEKURI, C. AND RAJARAMAN, A. 1997. Conjunctive query containment revisited. In Proceedings of the International Conference on Database Theory (ICDT'97).
- CHOMICKI, J. 2003. Preference formulas in relational queries. ACM Trans. Datab. Syst. 28, 4, 427–466.
- CHOMICKI, J. 2004. Semantic optimization of preference queries. In Proceedings of the International Symposium on Applications of Constraint Databases. 133–148.
- COHEN, W. W., SCHAPIRE, R. E., AND SINGER, Y. 1998. Efficiently mining frequent trees in a forest. Adv. Neural Inform. Process. Syst. 10.
- Collins, A. and Quillian, M. 1969. Retrieval time from semantic memory. J. Verbal Learn. Verbal Behav. 8, 240–247.
- CUPPENS, D. AND DE MOLOMBE, R. 1989. How to recognize interesting topics to provide cooperative answers. *Inform. Syst.* 14, 2, 163–173.
- DAS, A., DATAR, M., GARG, A., AND RAJARAM, S. 2007. Google news personalization: Scalable online collaborative filtering. In *Proceedings of the International World Wide Web Conference (WWW'07)*. 271–280.
- FAGIN, R., KUMAR, R., AND SIVAKUMAR, D. 2003. Comparing top k lists. SIAM J. Discrete Math. 17, 1, 134–140.
- FISHBURN, P. 1999. Preference structures and their numerical representations. *Theor. Comput. Sci.* 217, 359–383.
- GAASTERLAND, T., GODFREY, P., AND MINKER, J. 1992. An overview of cooperative query answering. J. Intell. Inf. Syst. 1, 2, 123–157.
- HANSSON, S. O. 2001. Preference logic. In Handbook of Philosophical Logic 8, D. Gabbay, Ed.
- HOLLAND, S., ESTER, M., AND KIESSLING, W. 2003. Preference mining: A novel approach on mining user preferences for personalized applications. In *Proceedings of the European Conference on Principles of Data Mining and Knowledge Discovery (PKDD'03)*. 204–216.
- HOLLAND, S. AND KIESSLING, W. 2004. Situated preferences and preference repositories for personalized database applications. In *Proceedings of the ER Workshops*. 511–523.
- HRISTIDIS, V., KOUDAS, N., AND PAPAKONSTANTINOU, Y. 2001. PREFER: A system for the efficient execution of multiparametric ranked queries. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*.
- ILYAS, I., SHAH, R., AREF, W., VITTER, J., AND EL MAGARMID, A. 2004. Rank-Aware query optimization. In Proceedings of the ACM SIGMOD International Conference on Management of Data.
- JIANG, B., PEI, J., LIN, X., CHEUNG, D. W., AND HAN, J. 2008. Mining preferences from superior and inferior examples. In Proceedings of the International SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'08). 390–398.
- JOACHIMS, T. 2002. Optimizing search engines using clickthrough data. In *Proceedings of the International SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'02).*
- JOACHIMS, T., FREITAG, D., AND MITCHELL, T. 1997. Webwatcher: A tour guide for the World Wide Web. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI'97)*.
- $Kendall,\,M.\,and\,Gibbons,\,J.\,D.~\,1990.~\,Rank\,\,Correlation\,\,Methods.\,Edward\,Arnold,\,London.$
- KIESSLING, W. 2002. Foundations of preferences in database systems. In Proceedings of the International Conference on Very Large Databases (VLDB'02). 311–322.
- KIESSLING, W. AND KOSTLER, W. 2002. Preference SQL-Design, implementation, experiences. In Proceedings of the International Conference on Very Large Databases (VLDB'02).
- KOSSMANN, D., RAMSAK, F., AND ROST, S. 2002. Shooting stars in the sky: An online algorithm for skyline queries. In *Proceedings of the International Conference on Very Large Databases* (VLDB'02). 275 286.
- KOUTRIKA, G. 2006. Personalization of structured queries with personal and collaborative preferences. In *Proceedings of the ECAI Workshop about Advances on Preference Handling*.

- KOUTRIKA, G., BERCOVITZ, B., AND GARCIA-MOLINA, H. 2009. FlexRecs: Expressing and combining flexible recommendations. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*.
- KOUTRIKA, G. AND IOANNIDIS, Y. 2004. Personalization of queries in database systems. In Proceedings of the International Conference on Data Engineering (ICDE'04). 597–608.
- KOUTRIKA, G. AND IOANNIDIS, Y. 2005a. Constrained optimalities in query personalization. In Proceedings of the ACM SIGMOD International Conference on Management of Data. 73–84.
- KOUTRIKA, G. AND IOANNIDIS, Y. 2005b. Personalized queries under a generalized preference model. In Proceedings of the International Conference on Data Engineering (ICDE'05).
- LACROIX, M. AND LAVENCY, P. 1987. Preferences: Putting more knowledge into queries. In Proceedings of the International Conference on Very Large Databases (VLDB'87). 217–225.
- LEE, J., WONYOU, G., WONHUANG, S., SELK, J., AND BALKE, W.-T. 2008. Optimal preference elicitation for skyline queries over categorical domains. In *Proceedings of the International Conference on Database and Expert Systems Applications (DEXA'08)*. 610–624.
- LINDEN, G., SMITH, B., AND YORK, J. 2003. Amazon.com recommendations: Item-to-Item collaborative filtering. *IEEE Internet Computing*.
- LIU, F., YU, C., AND MENG, W. 2004. Personalized Web search for improving retrieval effectiveness. *IEEE Trans. Knowl. Data Engin.* 16, 1.
- MIAH, M., DAS, G., HRISTIDIS, V., AND MANNILA, H. 2008. Standing out in a crowd: Selecting attributes for maximum visibility. In Proceedings of the International Conference on Data Engineering (ICDE'08). 356–365.
- PAPADIAS, D., TAO, Y., FU, G., AND SEEGER, B. 2003. An optimal and progressive algorithm for skyline queries. In Proceedings of the ACM SIGMOD International Conference on Management of Data. 467–478.
- PEI, J., JIANG, B., LIN, X., AND YUAN, Y. 2007. Probabilistic skylines on uncertain data. In Proceedings of the International Conference on Very Large Databases (VLDB'07). 15–26.
- PEI, J., JIN, W., ESTER, M., AND TAO, Y. 2005. Catching the best views of skyline: A semantic approach based on decisive subspaces. In Proceedings of the Internaional Conference on Very Large Databases (VLDB'05). 253-264.
- PITKOW, J. E., SCHUTZE, H., CASS, T. A., COOLEY, R., TURNBULL, D., EDMONDS, A., ADAR, E., AND BREUEL, T. M. 2002. Personalized search. Comm. ACM 45, 9, 50–55.
- SARKAS, N., DAS, G., KOUDAS, N., AND TUNG, A. K. H. 2008. Categorical skylines for streaming data. In Proceedings of the ACM SIGMOD International Conference on Management of Data. 239–250.
- SATZGER, B., ENDRES, M., AND KIESSLING, W. 2006. A preference-based recommendation system. In Proceedings of the International Conference on Electronic Commerce and Web Technologies (ECWeb'06).
- STEFANIDIS, K. AND PITOURA, E. 2008. Fast contextual preference scoring of database tuples. In Proceedings of the International Conference on Extending Database Technology (EDBT'08). 344– 355.
- STEFANIDIS, K., PITOURA, E., AND VASSILIADIS, P. 2007. Adding context to preferences. In Proceedings of the International Conference on Data Engineering (ICDE'07).
- TAO, Y., HRISTIDIS, V., PAPADIAS, D., AND PAPAKONSTANTINOU, Y. 2007. Branch-and-Bound processing of ranked queries. *Inform. Syst.* 32, 3, 424–445.
- VAN BUNNINGEN, A., FENG, L., AND APERS, P. M. G. 2006. A context-aware preference model for database querying in an ambient intelligent environment. In *Proceedings of the International Conference on Database and Expert Systems Applications (DEXA'06).* 33–43.
- VLACHOU, A., DOULKERIDIS, C., KOTIDIS, Y., AND VAZIRGIANNIS, M. 2007. SKYPEER: Efficient subspace skyline computation over distributed data. In Proceedings of the International Conference on Data Engineering (ICDE'07). 416–425.
- WELLMAN, M. AND DOYLE, J. 1991. Preferential semantics for goals. In Proceedings of the National Conference on Artificial Intelligence. 698–703.
- WONG, R. C -W., PEI, J., FU, A. W.-C., AND WANG, K. 2007. Mining favorable facets. In Proceedings of the International SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'07). 804–813.

13:50 • G. Koutrika and Y. Ioannidis

- XIN, D. AND HAN, J. 2008. P-Cube: Answering preference queries in multi-dimensional space. In Proceedings of the International Conference on Data Engineering (ICDE'08). 1092–1100.
- YIU, M. L., DAI, X., MAMOULIS, N., AND VAITIS, M. 2007. Top-k spatial preference queries. In Proceedings of the International Conference on Data Engineering (ICDE'07). 1076– 1085.

ZAKI, M. J. 2005. Efficiently mining frequent trees in a forest. Inform. Syst. 17, 8, 1021–1035.

Received December 2008; revised August 2009; accepted January 2010