Variable-Strength Conditional Preferences for Matchmaking in Description Logics

Thomas Lukasiewicz*

Dipartimento di Informatica e Sistemistica Università di Roma "La Sapienza" Via Salaria 113, 00198 Rome, Italy lukasiewicz@dis.uniroma1.it

Abstract

We present an approach to variable-strength conditional preferences for matchmaking and ranking objects in description logics. In detail, we introduce conditional preference bases, which consist of a description logic knowledge base and a finite set of variable-strength conditional preferences, and which are associated with a formal semantics based on ranking functions. We then define the notions of consistency and preferential entailment for conditional preference bases, which strictly generalize ε -consistency and entailment in System Z^+ in default reasoning from conditional knowledge bases, respectively. We also describe some semantic properties of preferential entailment. We then show how preferential entailment can be used to define a distance measure between two conditional preference bases. We also define functions for ranking objects relative to a conditional preference base, and we describe an application in the area of literature search. Finally, we provide algorithms for solving the main computational tasks related to conditional preference bases.

Introduction

In their seminal work, Di Noia *et al.* (2003) have explored the problem of matching user profiles in description logics, which is roughly described as follows (with the terminology of Di Noia *et al.*): Given a demand profile P_d and a supplier profile P_s , compute the degree to which P_d (resp., P_s) is matching P_s (resp., P_d). For example, in the domain of dating services, P_d may describe the desired characteristics of a potential partner, while P_s may express one's own characteristics. Such problems are crucial especially in matchmaking in electronic market places and in web service discovery.

Other important works by Poole and Smyth (2004; 2005) deal with the closely related problem of matching instances against models of instances, which are both described at different levels of abstraction and at different levels of detail, using qualitative probability theory. Informally, such problems can be described as follows. Given an instance I and a model of instances M, compute the qualitative probability that the instance I is matching the model M (that is, of I

Jörg Schellhase

Fachgebiet Wirtschaftsinformatik Universität Kassel Nora-Platiel-Str. 4, 34127 Kassel, Germany schellhase@wirtschaft.uni-kassel.de

given M). For example, in a geological exploration domain, we may want to determine whether there might be gold in an area. In this case, an instance I may be given by the description of an area, while a model M may be given by a description of areas where gold can be found, and the qualitative probability that I is matching M describes the likelihood that gold may be found in I.

In this paper, we continue this line of research. A serious drawback of the above works is that they only allow for expressing simple variable-strength preferences of the form "property α is preferred over property $\neg \alpha$ with strength s" in user profiles and models of instances, respectively. In particular, they all do not allow for variable-strength preferences such as "property α is preferred over property β with strength s", and they also do not allow for variable-strength conditional preferences such as "generally, in the context ϕ , property α is preferred over property β with strength s". In this paper, we try to fill this gap. We present a formalism for matchmaking in description logics that allows for expressing both such variable-strength preferences and conditional preferences in user profiles and models of instances.

Like Poole and Smyth's work (2004; 2005), the matching formalism in this paper is also based on qualitative probabilities. Differently from Poole and Smyth's work (2004; 2005), however, it requires a technically much more involved way of computing qualitative probabilities, since our language for encoding models of instances and user profiles is much more expressive. We especially have to suitably handle variable-strength conditional preferences, which are the above statements of the form "generally, in the context ϕ , property α is preferred over property β with strength s". They are strict generalizations of variablestrength conditional desires (Tan & Pearl 1994), which are statements of the form "generally, in the context ϕ , property α is preferred over property $\neg \alpha$ with strength s". The latter bear close similarity to variable-strength defaults of the form "generally, if ϕ then α with strength s" in default reasoning from conditional knowledge bases.

The literature contains several different proposals for default reasoning from conditional knowledge bases and extensive work on its desired properties. The core of these properties are the rationality postulates of System P by Kraus *et al.* (1990), which constitute a sound and complete axiom system for several classical model-theoretic en-

^{*}Alternate address: Institut für Informationssysteme, Technische Universität Wien, Favoritenstraße 9-11, 1040 Vienna, Austria; e-mail: lukasiewicz@kr.tuwien.ac.at.

Copyright © 2006, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

tailment relations under uncertainty measures on worlds. They characterize classical model-theoretic entailment under preferential structures, infinitesimal probabilities, possibility measures (Dubois & Prade 1991), and world rankings (Spohn 1988; Goldszmidt & Pearl 1992). They also characterize an entailment relation based on conditional objects (Dubois & Prade 1994). A survey of all these relationships is given in (Benferhat, Dubois, & Prade 1997; Gabbay & Smets 1998). Mainly to solve problems with irrelevant information, the notion of rational closure as a more adventurous notion of entailment was introduced by Lehmann (1989). It is in particular equivalent to entailment in System Z by Pearl (1990) (which is generalized to variable-strength defaults in System Z^+ by Goldszmidt and Pearl (1991b; 1996)), to the least specific possibility entailment by Benferhat et al. (1992), and to a conditional (modal) logic-based entailment by Lamarre (1992). Another sophisticated quasi-probabilistic formalism for reasoning about variable-strength defaults is Weydert's System JLZ (2003). Recently, also generalizations of many of the above approaches to probabilistic and fuzzy default reasoning have been proposed (see especially (Lukasiewicz 2005) and (de Saint-Cyr & Prade 2006), respectively).

Many previous approaches to conditional desires in the literature have been inspired by default reasoning from conditional knowledge bases. In particular, Boutilier (1994) has presented an approach to conditional desires that is based on the bimodal logic CO, while Tan & Pearl (1994) have proposed an approach to variable-strength conditional desires based on System Z^+ . Recently, work by Lang *et al.* (2002) has presented and explored an approach to conditional desires with variable-strength penalties and rewards.

In this paper, we define a formal semantics for variablestrength conditional preferences (which are strictly more general than variable-strength conditional desires), which is based on a generalization of Goldszmidt and Pearl's System Z^+ (1991b; 1996). Every set of variable-strength conditional preferences is associated with a unique ranking on the set of all objects. In this sense, our approach also differs from CP-nets (Boutilier *et al.* 2004), which are a machinery for dealing with conditional ceteris paribus preferences.

We focus especially on the following three kinds of matching problems: (a) matching objects against descriptions of objects, (b) matching two descriptions of objects, and (c) matching two objects relative to a description of objects. Some examples are (a) ranking all the answers of a web search query, (b) matching two web service descriptions, and (c) ranking all the answers of a web search query, (b) matching two web service descriptions, and (c) ranking all the answers of a web search query relative to a sample answer. Since we are especially interested in the Semantic Web as the main application context of the above matching problems, we assume that objects and description logics SHIF(D) and SHOIN(D), which stand behind the web ontology languages OWL Lite and OWL DL, respectively (Horrocks & Patel-Schneider 2003).

The Semantic Web (Berners-Lee 1999; Fensel *et al.* 2002) aims at an extension of the current World Wide Web by standards and technologies that help machines to understand the information on the Web so that they can support

richer discovery, data integration, navigation, and automation of tasks. The main ideas behind it are to add a machinereadable meaning to Web pages, to use ontologies for a precise definition of shared terms in Web resources, to make use of KR technology for automated reasoning from Web resources, and to apply cooperative agent technology for processing the information of the Web. The Semantic Web consists of several hierarchical layers, where the Ontology layer, in form of the OWL Web Ontology Language (W3C 2004; Horrocks, Patel-Schneider, & van Harmelen 2003) (recommended by the W3C), is currently the highest layer of sufficient maturity. OWL consists of three increasingly expressive sublanguages, namely OWL Lite, OWL DL, and OWL Full. OWL Lite and OWL DL are essentially expressive description logics with an RDF syntax (Horrocks, Patel-Schneider, & van Harmelen 2003). Ontology entailment in OWL Lite (resp., OWL DL) reduces to knowledge base (un)satisfiability in the description logic $SHIF(\mathbf{D})$ (resp., $SHOIN(\mathbf{D})$) (Horrocks & Patel-Schneider 2003).

The main contributions of this paper are as follows:

- We introduce conditional preference bases, which consist of a description logic knowledge base and a finite set of variable-strength conditional preferences, and which are given a formal semantics based on ranking functions. We then introduce the notions of consistency and preferential entailment for conditional preference bases, which are strict generalizations of ε -consistency due to Adams (1975) and Pearl (1989) and of entailment in System Z⁺ by Goldszmidt and Pearl (1991b; 1996) in default reasoning from conditional knowledge bases, respectively.
- We analyze the semantic properties of preferential entailment for conditional preference bases. It turns out that preferential entailment properly generalizes entailment in System Z^+ and thus also has similar properties. In particular, it realizes an inheritance of preference information along subclass relationships, where more specific preference information correctly overrides less specific preference information. Furthermore, we show that the notion of preferential entailment has the desirable properties of Irreflexivity, Asymmetry, Transitivity, and Rational Monotonicity. Finally, we show how the notion of preferential entailment for conditional preference bases and the strength of preferential consequences can be used to define an unsymmetric notion of distance between a demand and a supplier conditional preference base by comparing the entailed strengths of their conditional preferences.
- We then introduce two object rankings that reflect the variable-strength conditional preferences encoded in a conditional preference base, and we describe an application of them in literature search. More precisely, search query languages of current search engines are very restricted in their expressive power. There are scientific search engines on the web, however, that have valuable metadata about research publications, authors, organizations, and scientific events. We show that conditional preference bases allow for a more powerful query language, which can exploit this metadata better than the current approaches do. In particular, we give sample queries

that (i) influence the ranking of the query results, (ii) express quality measures, and (iii) cluster query results.

Finally, we provide algorithms for deciding consistency and computing the ranking z⁺, which generalize previous algorithms for deciding ε-consistency and computing the ranking z⁺ in System Z⁺. We also give an algorithm for rewriting a simple conditional preference base to a non-defeasible equivalent for ranking objects. All the above algorithms are based on a reduction to deciding whether a description logic knowledge base is satisfiable. They require a polynomial number of such satisfiability checks, and thus are all possible in polynomial time when the satisfiability checks are possible in polynomial time.

The rest of this paper is organized as follows. First, we recall the description logics SHIF(D) and SHOIN(D). Then, we introduce conditional preference bases, and we define the notions of consistency and preferential entailment for them, analyze the semantics properties of preferential entailment, and show how it can be used for matchmaking with conditional preference bases. Thereafter, we introduce two object rankings reflecting the conditional preferences of a conditional preference base, and we describe their application in literature search. Finally, we provide algorithms for the main computational tasks related to conditional preference bases, and give an outlook on future research. Note that the proofs of all results are given in the appendix. Further details are given in the extended report (Lukasiewicz & Schellhase 2005).

$\mathcal{SHIF}(D)$ and $\mathcal{SHOIN}(D)$

In this section, we recall the expressive description logics $SHIF(\mathbf{D})$ and $SHOIN(\mathbf{D})$, which stand behind the web ontology languages OWL Lite and OWL DL, respectively (Horrocks & Patel-Schneider 2003). Intuitively, description logics model a domain of interest in terms of concepts and roles, which represent classes of individuals and binary relations between classes of individuals, respectively. Roughly, a description logic knowledge base encodes subset relationships between classes of individuals, the membership of individuals to classes, and the membership of pairs of individuals to binary relations between classes.

Syntax

We first describe the syntax of $\mathcal{SHOIN}(\mathbf{D})$. We assume a set of *elementary datatypes* and a set of *data values*. A *datatype* is an elementary datatype or a set of data values. A *datatype* is an elementary datatype theory $\mathbf{D} = (\Delta^{\mathbf{D}}, \cdot^{\mathbf{D}})$ consists of a *datatype domain* $\Delta^{\mathbf{D}}$ and a mapping $\cdot^{\mathbf{D}}$ that assigns to each elementary datatype a subset of $\Delta^{\mathbf{D}}$ and to each data value an element of $\Delta^{\mathbf{D}}$. We extend $\cdot^{\mathbf{D}}$ to all datatypes by $\{v_1, \ldots\}^{\mathbf{D}} = \{v_1^{\mathbf{D}}, \ldots\}$. Let $\mathbf{A}, \mathbf{R}_A, \mathbf{R}_D$, and \mathbf{I} be pairwise disjoint finite nonempty sets of *atomic concepts, abstract roles, datatype roles*, and *individuals*, respectively. We denote by \mathbf{R}_A^- the set of inverses R^- of all $R \in \mathbf{R}_A$.

A role is any element of $\mathbf{R}_A \cup \mathbf{R}_A^- \cup \mathbf{R}_D$. Concepts are inductively defined as follows. Every $\phi \in \mathbf{A}$ is a concept, and if $o_1, \ldots, o_n \in \mathbf{I}$, then $\{o_1, \ldots, o_n\}$ is a concept (called oneOf). If ϕ , ϕ_1 , and ϕ_2 are concepts and if $R \in \mathbf{R}_A \cup \mathbf{R}_A^-$, then also $\neg \phi$, $(\phi_1 \sqcap \phi_2)$, and $(\phi_1 \sqcup \phi_2)$ are concepts (called *negation*, *conjunction*, and *disjunction*, respectively), as well as $\exists R.\phi, \forall R.\phi, \ge nR$, and $\le nR$ (called *exists*, *value*, *atleast*, and *atmost restriction*, respectively) for an integer $n \ge 0$. If D is a datatype and $U \in \mathbf{R}_D$, then $\exists U.D, \forall U.D, \ge nU$, and $\le nU$ are concepts (called *datatype exists*, *value*, *atleast*, and *atmost restriction*, respectively) for an integer $n \ge 0$. We write \top and \bot to abbreviate $\phi \sqcup \neg \phi$ and $\phi \sqcap \neg \phi$, respectively, and we eliminate parentheses as usual.

An axiom has one of the following forms: (1) $\phi \sqsubseteq \psi$ (called *concept inclusion axiom*), where ϕ and ψ are concepts; (2) $R \sqsubseteq S$ (called *role inclusion axiom*), where either $R, S \in \mathbf{R}_A$ or $R, S \in \mathbf{R}_D$; (3) Trans(R) (called *transitivity axiom*), where $R \in \mathbf{R}_A$; (4) $\phi(a)$ (called *concept membership axiom*), where ϕ is a concept and $a \in \mathbf{I}$; (5) R(a, b) (resp., U(a, v)) (called *role membership axiom*), where $R \in \mathbf{R}_A$ (resp., $U \in \mathbf{R}_D$) and $a, b \in \mathbf{I}$ (resp., $a \in \mathbf{I}$ and v is a data value); and (6) a = b (resp., $a \neq b$) (*equality* (resp., *inequality*) *axiom*), where $a, b \in \mathbf{I}$. A (*description logic*) *knowledge base KB* is a finite set of axioms. For decidability, number restrictions in KB are restricted to simple abstract roles (Horrocks, Sattler, & Tobies 1999).

The syntax of $SHIF(\mathbf{D})$ is as the above syntax of $SHOIN(\mathbf{D})$, but without the oneOf constructor and with the atleast and atmost constructors limited to 0 and 1.

Example 1 An online store (such as *amazon.com*) may use a description logic knowledge base to classify and characterize its products. For example, suppose that (1) textbooks are books, (2) personal computers and laptops are mutually exclusive electronic products, (3) books and electronic products are mutually exclusive products, (4) objects on offer are products, (5) every product has at least one related product, (6) only products are related to each other, (7) *tb_ai* and *tb_lp* are textbooks, (8) which are related to each other, (9) *pc_ibm* and *pc_hp* are personal computers, (10) which are related to each other, and (11) *ibm* and *hp* are providers for *pc_ibm* and *pc_hp*, respectively. These relationships are expressed by the following description logic knowledge base *KB*₁:

- (1) *Textbook* \sqsubseteq *Book*;
- (2) $PC \sqcup Laptop \sqsubseteq Electronics; PC \sqsubseteq \neg Laptop;$
- (3) *Book* \sqcup *Electronics* \sqsubseteq *Product*; *Book* $\sqsubseteq \neg$ *Electronics*;
- (4) *Offer* \sqsubseteq *Product*;
- (5) *Product* $\sqsubseteq \ge 1$ *related*;
- (6) ≥ 1 related $\sqcup \geq 1$ related $\Box = Product;$
- (7) *Textbook*(*tb_ai*); *Textbook*(*tb_lp*);
- (8) related(tb_ai, tb_lp);
- (9) $PC(pc_ibm)$; $PC(pc_hp)$;
- (10) related(pc_ibm, pc_hp);
- (11) *provides*(*ibm*, *pc_ibm*); *provides*(*hp*, *pc_hp*).

Semantics

An *interpretation* $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ with respect to a datatype theory $\mathbf{D} = (\Delta^{\mathbf{D}}, \cdot^{\mathbf{D}})$ consists of a nonempty (*abstract*) *domain* $\Delta^{\mathcal{I}}$ disjoint from $\Delta^{\mathbf{D}}$, and a mapping $\cdot^{\mathcal{I}}$ that assigns

to each atomic concept $\phi \in \mathbf{A}$ a subset of $\Delta^{\mathcal{I}}$, to each individual $o \in \mathbf{I}$ an element of $\Delta^{\mathcal{I}}$, to each abstract role $R \in \mathbf{R}_A$ a subset of $\Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$, and to each datatype role $U \in \mathbf{R}_D$ a subset of $\Delta^{\mathcal{I}} \times \Delta^{\mathbf{D}}$. We extend $\cdot^{\mathcal{I}}$ to all concepts and roles as usual (where #S denotes the cardinality of a set S):

 $\begin{array}{l} \bullet \ \{o_1, \dots, o_n\}^{\mathcal{I}} = \{o_1^{\mathcal{I}}, \dots, o_n^{\mathcal{I}}\}; \ (\neg \phi)^{\mathcal{I}} = \Delta^{\mathcal{I}} - \phi^{\mathcal{I}}; \\ \bullet \ (\phi_1 \sqcap \phi_2)^{\mathcal{I}} = \phi_1^{\mathcal{I}} \cap \phi_2^{\mathcal{I}}; \ (\phi_1 \sqcup \phi_2)^{\mathcal{I}} = \phi_1^{\mathcal{I}} \cup \phi_2^{\mathcal{I}}; \\ \bullet \ (\exists R. \phi)^{\mathcal{I}} = \{x \in \Delta^{\mathcal{I}} \mid \exists y \colon (x, y) \in R^{\mathcal{I}} \land y \in \phi^{\mathcal{I}}\}; \\ \bullet \ (\forall R. \phi)^{\mathcal{I}} = \{x \in \Delta^{\mathcal{I}} \mid \forall y \colon (x, y) \in R^{\mathcal{I}} \rightarrow y \in \phi^{\mathcal{I}}\}; \\ \bullet \ (\geq nR)^{\mathcal{I}} = \{x \in \Delta^{\mathcal{I}} \mid \#(\{y \mid (x, y) \in R^{\mathcal{I}}\}) \geq n\}; \\ \bullet \ (\leq nR)^{\mathcal{I}} = \{x \in \Delta^{\mathcal{I}} \mid \#(\{y \mid (x, y) \in R^{\mathcal{I}}\}) \leq n\}; \\ \bullet \ (\exists U. D)^{\mathcal{I}} = \{x \in \Delta^{\mathcal{I}} \mid \exists y \colon (x, y) \in U^{\mathcal{I}} \land y \in D^{\mathbf{D}}\}; \\ \bullet \ (\forall U. D)^{\mathcal{I}} = \{x \in \Delta^{\mathcal{I}} \mid \forall y \colon (x, y) \in U^{\mathcal{I}} \rightarrow y \in D^{\mathbf{D}}\}; \\ \bullet \ (\geq nU)^{\mathcal{I}} = \{x \in \Delta^{\mathcal{I}} \mid \#(\{y \mid (x, y) \in U^{\mathcal{I}}\}) \geq n\}; \\ \bullet \ (\leq nU)^{\mathcal{I}} = \{x \in \Delta^{\mathcal{I}} \mid \#(\{y \mid (x, y) \in U^{\mathcal{I}}\}) \geq n\}. \\ \text{The satisfaction of an axiom F in an interpretation \mathcal{I} = } \end{array}$

The satisfaction of an axiolit F in an interpretation $\mathcal{I} = (\Delta, \cdot^{\mathcal{I}})$, denoted $\mathcal{I} \models F$, is defined by: (1) $\mathcal{I} \models \phi \sqsubseteq \psi$ iff $\phi^{\mathcal{I}} \subseteq \psi^{\mathcal{I}}$; (2) $\mathcal{I} \models R \sqsubseteq S$ iff $R^{\mathcal{I}} \subseteq S^{\mathcal{I}}$; (3) $\mathcal{I} \models \operatorname{Trans}(R)$ iff $R^{\mathcal{I}}$ is transitive; (4) $\mathcal{I} \models \phi(a)$ iff $a^{\mathcal{I}} \in \phi^{\mathcal{I}}$; (5) $\mathcal{I} \models R(a, b)$ iff $(a^{\mathcal{I}}, b^{\mathcal{I}}) \in R^{\mathcal{I}}$; (6) $\mathcal{I} \models U(a, v)$ iff $(a^{\mathcal{I}}, v^{\mathbf{D}}) \in U^{\mathcal{I}}$; (7) $\mathcal{I} \models a = b$ iff $a^{\mathcal{I}} = b^{\mathcal{I}}$; and (8) $\mathcal{I} \models a \neq b$ iff $a^{\mathcal{I}} \neq b^{\mathcal{I}}$. The interpretation \mathcal{I} satisfies the axiom F, or \mathcal{I} is a model of F, iff $\mathcal{I} \models F$. We say \mathcal{I} satisfies a knowledge base KB, or \mathcal{I} is a model of KB, denoted $\mathcal{I} \models KB$, iff $\mathcal{I} \models F$ for all $F \in KB$. We say KB is satisfiable (resp., unsatisfiable) iff KB has a (resp., no) model. An axiom F is a logical consequence of KB, denoted $KB \models F$, iff each model of KB does not satisfy F.

Conditional Preference Bases

In this section, we define the syntax of conditional preferences, which are statements of the form "generally, if ϕ holds, then α is preferred over β with strength s", and their semantics in terms of object rankings.

Syntax

We assume a finite set of *classification concepts* C (which are the relevant description logic concepts for defining preference relationships). A conditional preference is of the form $(\alpha \succ \beta \mid \phi)[s]$ with concepts $\alpha, \beta, \phi \in C$ and an integer $s \ge 0$. We call ϕ its *body* and *s* its *strength*, also denoted stren($(\alpha \succ \beta \mid \phi)[s]$). Informally, $(\alpha \succ \beta \mid \phi)[s]$ expresses that (i) generally, among the objects satisfying ϕ , the ones satisfying α are preferred over those satisfying β , and (ii) this preference relationship holds with strength s. A conditional preference of the form $(\alpha \succ \neg \alpha \mid \phi)[s]$ is also called a conditional desire and abbreviated as $(\alpha | \phi)[s]$. Conditional preferences and desires of the form $(\alpha \succ \beta \mid \top)[s]$ and $(\alpha \mid \top)[s]$ are also abbreviated as $(\alpha \succ \beta)[s]$ and $(\alpha)[s]$, respectively. A conditional preference base is a triple PB =(T, A, P), where T is a description logic knowledge base, A is a finite set of concepts from C, and P is a finite set of conditional preferences $(\alpha \succ \beta \mid \phi)[s]$ with $s \in \{0, \dots, 100\}$. Informally, T contains terminological knowledge, and A

contains assertional knowledge about an individual o (that is, A actually represents the set of all $\phi(o)$ such that $\phi \in A$), while P contains conditional preferences about the individual o (that is, P actually represents the set of all $(\alpha(o) \succ \beta(o) | \phi(o))[s]$ with $(\alpha \succ \beta | \phi)[s] \in P$). Observe also that the statements in T and A are *strict* (that is, they must always hold), while the statements in P are *defeasible* (that is, they may have exceptions and thus do not always hold), since P may not always be satisfiable as a whole.

Example 2 The preference relationships "generally, PC's are preferred over laptops with strength 20", "generally, laptops on offer are preferred over PC's on offer with strength 70", and "generally, inexpensive objects are preferred over expensive ones with strength 90" can be expressed by the conditional preference base PB = (T, A, P), where T is the knowledge base from Example 1, $A = \emptyset$, and $P = \{(PC \succ Laptop)[20], (Laptop \succ PC | Offer)[70], (Inexpensive)[90]\}$.

Semantics

We now define some basic semantic notions, including objects and object rankings (which are certain functions that map every object to a rank from $\{0, 1, \ldots\} \cup \{\infty\}$), and we then associate with every conditional preference base a set of object rankings as a formal semantics.

An *object* o is a set of concepts from C. We denote by $\mathcal{O}_{\mathcal{C}}$ the set of all objects relative to \mathcal{C} . An object o satis*fies* a description logic knowledge base T, denoted $o \models T$, iff $T \cup \{\phi(i) \mid \phi \in o\}$ is satisfiable and entails (resp., does not entail) every concept membership $\phi(i)$ such that $\phi \in o$ (resp., $\phi \notin o$), where *i* is a new individual. Informally, every object o represents an individual i that is fully specified on C in the sense that *i* belongs (resp., does not belong) to every concept $\phi \in o$ (resp., $\phi \notin o$). An object o satisfies a concept ϕ , denoted $o \models \phi$, iff $\phi \in o$. An object o satisfies a set of concepts A, denoted $o \models A$, iff o satisfies all $\phi \in A$. A concept ϕ is *satisfiable* iff there exists an object $o \in \mathcal{O}_{\mathcal{C}}$ that satisfies ϕ . An object o satisfies a conditional preference $(\alpha \succ \beta \mid \phi)[s]$, denoted $o \models (\alpha \succ \beta \mid \phi)[s]$, iff $o \models \neg \phi \sqcup \neg \beta \sqcup \alpha$. We say *o satisfies* a set of conditional preferences P, denoted $o \models P$, iff o satisfies all $p \in P$. We say o verifies $(\alpha \succ \beta \mid \phi)[s]$ iff $o \models \phi \sqcap \alpha$. We say o falsifies $(\alpha \succ \beta \mid \phi)[s]$, denoted $o \not\models (\alpha \succ \beta \mid \phi)[s]$, iff $o \models \phi \sqcap \beta \sqcap \neg \alpha$. A set of conditional preferences P tolerates a conditional preference p under a description logic knowledge base Tand a set of classification concepts $A \subseteq C$ iff an object o exists that satisfies $T \cup A \cup P$ (that is, o satisfies T, A, and P) and verifies p. We say that P is under T and A in conflict with p iff P does not tolerate p under T and A.

An object ranking κ is a mapping $\kappa \colon \mathcal{O}_{\mathcal{C}} \to \{0, 1, \ldots\} \cup \{\infty\}$ such that $\kappa(o) = 0$ for at least one object $o \in \mathcal{O}_{\mathcal{C}}$. It is extended to all concepts ϕ as follows. If ϕ is satisfiable, then $\kappa(\phi) = \min \{\kappa(o) \mid o \in \mathcal{O}_{\mathcal{C}}, o \models \phi\}$; otherwise, $\kappa(\phi) = \infty$. We say κ is *admissible* with a description logic knowledge base T (resp., a set of concepts A) iff $\kappa(o) = \infty$ for all $o \in \mathcal{O}_{\mathcal{C}}$ such that $o \not\models T$ (resp., $o \not\models A$). We say κ is *admissible* with a conditional preference $(\alpha \succ \beta \mid \phi)[s]$ iff either $\kappa(\phi) = \infty$ or $\kappa(\phi \sqcap \alpha) + s < \kappa(\phi \sqcap \beta)$. We say κ is *ad*-

missible with a conditional preference base PB = (T, A, P) iff κ is admissible with T, A, and all $p \in P$.

A conditional preference base PB = (T, A, P) is *pure* iff $o \models \neg \alpha \sqcup \neg \beta$ for each $(\alpha \succ \beta \mid \phi)[s] \in P$ and each object $o \in \mathcal{O}_{\mathcal{C}}$ that satisfies T and A. Intuitively, each object $o \in \mathcal{O}_{\mathcal{C}}$ that satisfies T and A can belong to at most one concept among α and β . The following theorem shows that every conditional preference base PB has an equivalent (in terms of object rankings) pure conditional preference base PB'. Thereafter, we can thus safely assume that every considered conditional preference base is pure.

Theorem 1 An object ranking κ is admissible with the conditional preference $(\alpha \succ \beta | \phi)[s]$ iff it is admissible with both $(\alpha \succ \neg \alpha \sqcap \beta | \phi)[s]$ and $(\alpha \sqcap \neg \beta \succ \alpha \sqcap \beta | \phi)[s]$.

Reasoning about Conditional Preferences

In this section, we define the notions of consistency and of preferential entailment for conditional preference bases, and we describe semantic properties of preferential entailment.

Consistency

The notion of consistency for conditional preference bases is inspired by the notion of ε -consistency for conditional knowledge bases due to Adams (1975) and Pearl (1989). Formally, a conditional preference base *PB* is *consistent* (resp., *inconsistent*) iff an (resp., no) object ranking κ exists that is admissible with *PB*. We now summarize some results that carry over from conditional knowledge bases.

The following result shows that the existence of an object ranking that is admissible with PB = (T, A, P) is equivalent to the existence of a preference ranking σ on P that is admissible with PB. Here, a *preference ranking* σ on a set of conditional preferences P maps each $p \in P$ to an integer. A preference ranking σ on P is *admissible* with PB = (T, A, P) iff every $P' \subseteq P$ that is under T and A in conflict with some $p \in P$ contains some p' such that $\sigma(p') < \sigma(p)$.

Theorem 2 Let PB = (T, A, P) be a pure conditional preference base. Then, PB is consistent iff there exists a preference ranking σ on P that is admissible with PB.

The next result shows that PB is consistent iff there exists an ordered partition of P with certain properties.

Theorem 3 Let PB = (T, A, P) be a pure conditional preference base. Then, PB is consistent iff an ordered partition (P_0, \ldots, P_k) of P exists such that either (a) every P_i , $0 \le i \le k$, is the set of all $p \in P_i \cup \cdots \cup P_k$ tolerated under T and A by $P_i \cup \cdots \cup P_k$, or (b) for every $i, 0 \le i \le k$, every $p \in P_i$ is tolerated under T and A by $P_i \cup \cdots \cup P_k$.

We call the unique partition in (a) the *z*-partition of *PB*.

Example 3 The conditional preference base *PB* of Example 1 is consistent, and its *z*-partition is given as follows:

$$(P_0, P_1) = (\{(PC \succ Laptop)[20], (Inexpensive)[90]\}, \\ \{(Laptop \succ PC \mid Offer)[70]\}).$$

Preferential Entailment

The notion of preferential entailment for conditional preference bases is based on an object ranking κ^+ and a preference ranking z^+ , which are inspired by Goldszmidt and Pearl's (1991b; 1996) world ranking κ^+ and default ranking z^+ of System Z^+ for default reasoning from conditional knowledge bases, respectively. In the sequel, let PB = (T, A, P)be a consistent pure conditional preference base.

The preference ranking z^+ and the object ranking κ^+ are the unique solution of the following system of equations (1) and (2). For all $p = (\alpha \succ \beta \mid \phi)[s] \in P$ and $o \in \mathcal{O}_{\mathcal{C}}$:

$$z^{+}(p) = s + \kappa^{+}(\phi \wedge \alpha) \tag{1}$$

$$\kappa^{+}(o) = \begin{cases} \infty & \text{if } o \not\models T \cup A \\ 0 & \text{if } o \models T \cup A \cup P \\ 1 + \max_{q \in P: \ o \not\models q} z^{+}(q) & \text{otherwise.} \end{cases}$$
(2)

Observe that, as in default reasoning from conditional knowledge bases, equations (1) and (2) have a unique solution, which satisfies a certain compactness condition, and thus the two rankings z^+ and κ^+ are well-defined.

Example 4 Consider the conditional preference base PB = (T, A, P) given by T and A as in Example 2 and

 $P = \{(PC \succ Laptop)[20], (Laptop \succ PC \mid Offer)[70]\}.$

The object ranking κ^+ of *PB* is shown in Fig. 1, and its preference ranking z^+ is given by $z^+((PC \succ Laptop)[20]) = 20$ and $z^+((Laptop \succ PC \mid Offer)[70]) = 91$.

The following result shows that the preference ranking z^+ and the object ranking κ^+ are both admissible with PB.

Theorem 4 Let PB = (T, A, P) be a consistent pure conditional preference base. Then, z^+ (resp., κ^+) is a preference (resp., an object) ranking that is admissible with PB.

We define the notion of preferential entailment as follows. A conditional preference $p = (\alpha \succ \beta \mid \phi)[s]$ is a *preferential* consequence of PB, denoted PB $\succ p$, iff either $\kappa^+(\phi) = \infty$ or $\kappa^+(\phi \land \alpha) + s < \kappa^+(\phi \land \beta)$. Its strength, denoted stren(PB $\succ p$), is defined as $\kappa^+(\phi \land \beta) - \kappa^+(\phi \land \alpha)$.

This notion of preferential entailment properly generalizes entailment in System Z^+ in default reasoning from conditional knowledge bases, and thus has similar properties. In particular, it realizes some inheritance of conditional preferences along subclass relationships, where conditional preferences of more specific classes override the ones of less specific classes. Similar to System Z^+ , preferential entailment also has the problem of inheritance blocking for *conditional desires* (which are not inherited to more specific subclasses that are exceptional relative to some other conditional desires). However, in the following example, there is no such inheritance blocking for general *conditional preferences*.

Example 5 Let PB = (T, A, P) be defined as in Example 2. It is not difficult to verify that *PB* preferentially entails

 $(PC \succ Laptop)[20], (Laptop \succ PC | Offer)[70], (Inexpensive | Made_By_IBM)[90], (Inexpensive | Offer)[90],$

while the slightly modified conditional preference base $PB = (T, A \cup \{PC \sqcup Laptop\}, P)$ preferentially entails only (*Inexpensive* | Offer)[69], but not (*Inexpensive* | Offer)[90].

Semantic Properties

Preferential entailment has several nice semantic properties, among which there is a direct inference property and an irrelevance property. Moreover, it satisfies the properties of *Irreflexivity*, *Asymmetry*, *Transitivity*, and *Rational Monotonicity*, which is formulated by the following theorem.

Theorem 5 Preferential entailment satisfies the following properties of Irreflexivity, Asymmetry, Transitivity, and Rational Monotonicity, for every conditional preference base PB and all concepts $\alpha, \beta, \gamma, \phi, \phi' \in C$ and integers $s, t \ge 0$:

Irreflexivity. If $PB \sim (\alpha \succ \alpha \mid \phi)[s]$, then ϕ is unsatisfiable under PB.

Asymmetry. If $PB \succ (\alpha \succ \beta \mid \phi)[s]$ and $PB \succ (\beta \succ \alpha \mid \phi)[t]$, then ϕ is unsatisfiable under PB.

Transitivity. If $PB \succ (\alpha \succ \beta \mid \phi)[s]$ and $PB \succ (\beta \succ \gamma \mid \phi)[t]$, then $PB \vdash (\alpha \succ \gamma \mid \phi)[s+t+1]$.

Rational Monotonicity. If $PB \vdash (\alpha \succ \beta \mid \phi)[s]$ and $PB \not\vdash (\neg \phi' \mid \phi \sqcap (\alpha \sqcup \beta))[0]$, then $PB \vdash (\alpha \succ \beta \mid \phi \sqcap \phi')[s]$.

Matching Conditional Preference Bases

Using the notion of preferential entailment for conditional preference bases and the strength of preferential consequences, we can define a measure for the similarity of two conditional preference bases $PB_{\delta} = (T_{\delta}, A_{\delta}, P_{\delta})$ and $PB_{\sigma} = (T_{\sigma}, A_{\sigma}, P_{\sigma})$ (called the *demand* and the *supplier* conditional preference base, respectively) by comparing the entailed strengths of their conditional preferences. Formally, the *distance of* P_{δ} from P_{σ} is defined as

$$\sum_{p \in P_{\sigma}} |stren(PB_{\sigma} \succ p) - stren(PB_{\delta} \succ p)|,$$

while the *distance of* P_{σ} *from* P_{δ} is defined as

$$\sum_{p \in P_{\delta}} |stren(PB_{\sigma} \succ p) - stren(PB_{\delta} \succ p)|.$$

Example 6 Consider a simple matching problem from the domain of dating services, where $PB_{\delta} = (T_{\delta}, A_{\delta}, P_{\delta})$ and $PB_{\sigma} = (T_{\sigma}, A_{\sigma}, P_{\sigma})$ are defined as follows. Suppose that $T_{\delta} = T_{\sigma} = \{poetry \sqsubseteq literature, tennis \sqsubseteq sports\}$ represents the terminological knowledge that somebody interested in poetry resp. tennis is interested in literature resp. sports. Let $A_{\delta} = A_{\sigma} = \emptyset$, and let P_{δ} and P_{σ} represent the variable-strength interests of two persons defined by P_{δ} $= \{(literature)[70], (sports)[20]\}$ and $P_{\sigma} = \{(poetry)[60], (tennis)[20]\}$. Then, the distance of P_{δ} from P_{σ} is given by 82, while the distance of P_{σ} from P_{δ} is given by 10.

Ranking Objects

In this section, we present two object rankings that reflect the variable-strength conditional preferences encoded in a conditional preference base, and we describe an application of these object rankings in literature search.

Ranking Functions

We now define object rankings that reflect the conditional preferences encoded in a consistent pure conditional preference base PB = (T, A, P). Notice first that the object ranking κ^+ of PB encodes only specificity levels of PB for the use in preferential entailment from PB, and thus in general does not properly reflect the conditional preferences in PB.

Example 7 Consider again the conditional preference base PB = (T, A, P) given in Example 4. The object ranking κ^+ of *PB* is shown in Fig. 1. Observe that κ^+ associates with o_3 the same rank as with o_4 . However, o_4 should actually be strictly preferred over o_3 , since o_4 satisfies "its" conditional preference $(Laptop \succ PC \mid Offer)$ [70], while o_3 falsifies "its" conditional preference $(PC \succ Laptop)$ [20].

We first rewrite P from a set of defeasible statements to a set of non-defeasible ones P^* . Roughly, this is done by adding exceptions to the bodies of conditional preferences.

Example 8 Consider the conditional preference base PB = (T, A, P) of Example 4. The rewritten conditional preference base $PB^* = (T, A, P^*)$ is given as follows:

$$P^{\star} = \{(PC \succ Laptop \mid \neg Offer)[20], (Laptop \succ PC \mid Offer)[70]\}.$$

It is obtained from *PB* by adding the exception $\neg Offer$ to the body of the conditional preference $(PC \succ Laptop)[20]$.

A pure conditional preference base PB = (T, A, P) is flat iff its z-partition is given by (P) and thus consists only of one component. Given a pure conditional preference base PB = (T, A, P), a non-defeasible equivalent $PB^* = (T, A, P^*)$ to PB satisfies the properties that (i) PB^* is flat, (ii) $PB^0 \succ p$ for all $p \in P^*$, where PB^0 is obtained from PB by replacing every strength s by the strength 0, and (iii) $P^* = \{(\alpha \succ \beta \mid \phi \sqcap \psi_p)[s] \mid p = (\alpha \succ \beta \mid \phi)[s] \in P\}$, where ψ_p is a conjunction of negated bodies that occur in P. Algorithm **flatten**, which is presented below, transforms a consistent simple conditional preference base PB into a nondefeasible equivalent (see below for details).

We are now ready to define the object rankings κ^{sum} and κ^{lex} . Informally, κ^{sum} associates with every object (as a penalty) the sum of the strengths of all conditional preferences in P^* that are falsified by o. Hence, roughly, objects with smaller values under κ^{sum} are those that satisfy more conditional preferences with larger strengths. More formally, κ^{sum} is defined as follows for all objects $o \in \mathcal{O}_{\mathcal{C}}$:

$$\kappa^{sum}(o) = \begin{cases} \infty & \text{if } o \not\models T \cup A \\ \sum_{p \in P^{\star}: \ o \not\models p} stren(p) + 1 & \text{otherwise.} \end{cases}$$

The object ranking κ^{lex} , in contrast, is based on a lexicographic order. Objects with smaller values under κ^{sum} are those that satisfy more conditional preferences with larger strengths, where satisfying one conditional preference of strength s is strictly preferred to satisfying any set of conditional preferences of strength at most s - 1. Formally, κ^{lex}

	PC Laptop Offer	κ^+	κ^{sum}	κ^{lex}]		PC Laptop Offer	κ^+	κ^{sum}	κ^{lex}
01	false false false	0	0	0]	05	true false false	0	0	0
02	false false true	0	0	0		06	true false true	92	71	2
03	false true false	21	21	1		07	true true false	∞	∞	∞
o_4	false true true	21	0	0		08	true true true	∞	∞	∞

Figure 1: The object rankings κ^+ , κ^{sum} , and κ^{lex} .

is defined as follows for all objects $o \in \mathcal{O}_{\mathcal{C}}$ (where n_j with $j \in \{0, ..., 100\}$ is the number of all $p \in P^*$ of strength j):

$$\kappa^{lex}(o) = \begin{cases} \infty & \text{if } o \not\models T \cup A \\ \sum_{i=0}^{100} |\{p \in P^* \mid o \not\models p\}| \cdot \prod_{j=0}^{stren(p)-1} (n_j+1) & \text{otherwise.} \end{cases}$$

Observe that κ^{lex} corresponds to the following lexicographic order on objects. For each $i \in \{0, ..., 100\}$, let P_i^{\star} denote the set of all $p \in P^{\star}$ such that stren(p) = i. For all objects $o, o' \in \mathcal{O}_{\mathcal{C}}$, we say that o is *lexicographically preferred* to o' iff some $i \in \{0, ..., 100\}$ exists such that $|\{p \in P_i^{\star} \mid o \not\models p\}| < |\{p \in P_i^{\star} \mid o' \not\models p\}|$ and $|\{p \in P_j^{\star} \mid o \not\models p\}| = |\{p \in P_j^{\star} \mid o' \not\models p\}| = |\{p \in P_j^{\star} \mid o' \not\models p\}| = |\{p \in P_j^{\star} \mid o' \not\models p\}|$ for all $j \in \{i + 1, ..., 100\}$.

Example 9 The object rankings κ^{sum} and κ^{lex} for *PB* of Example 4 are shown in Fig. 1. Under both κ^{sum} and κ^{lex} , the object o_4 is strictly preferred over o_3 , as desired.

Summarizing, every object ranking $\kappa \in \{\kappa^{sum}, \kappa^{lex}\}$ of a conditional preference base *PB* represents the variablestrength preference relationships encoded in *PB*. For every (fully specified) object $o \in \mathcal{O}_{\mathcal{C}}$, the rank of o under *PB* is given by $\kappa(o)$. Every object ranking $\kappa \in \{\kappa^{sum}, \kappa^{lex}\}$ can also be used to compare two objects $o, o' \in \mathcal{O}_{\mathcal{C}}$ as follows. The distance between o and o' under *PB* is defined as $|\kappa(o) - \kappa(o')|$. Furthermore, the (credulous) rank of a partially specified object (which is simply a concept) ϕ under *PB* is defined as $\min_{o \in \mathcal{O}_{\mathcal{C}}: o \models \phi} \kappa(o)$. Finally, the (credulous) distance between two partially specified objects ϕ and ϕ' is defined as $\min_{o,o' \in \mathcal{O}_{\mathcal{C}: o} \models \phi, o' \models \phi'} |\kappa(o) - \kappa(o')|$.

Application: Literature Search

Searching for scientific publications is a time-consuming task, for which there exist different instruments, including bibliographic systems, bibliographic web search engines, and general web search engines. The quality of such search engines, however, varies heavily. One high-quality bibliographic search engine is, for example, scholar.google.com, which contains only scientific publications, mainly from conferences and journals. What is generally missing is the possibility to formulate sophisticated search queries. To date, search query languages of most web search engines mainly rely on Boolean operators, and generally have in particular the following deficiencies: (i) little expressive power for formulating semantic queries; (ii) no or little possibilities to influence the ranking of the query results; (iii) no possibilities to formulate ones own quality measures for good/interesting query results; and (iv) no possibilities to influence the clustering of query results.

In this section, we show that our approach to conditional preferences allows for expressing more sophisticated search queries, and avoids the above-mentioned deficiencies. The examples below also show the expressive power of the formalism proposed in this paper. Of course, ordinary users of specialized search engines should be supported by a search query assistant, which helps to formulate such queries without the user having to know the formal syntax.

The strict terminological knowledge is informally described as follows. We assume the concepts *Publication*, *JournalPublication*, *ConferencePublication*, *Person*, and *Keyword*, which are related by the following concept inclusion axioms: *JournalPublication* \sqsubseteq *Publication* and *ConferencePublication* \sqsubseteq *Publication*. We assume the roles *Author* (relating *Publication* and *Person*), *Coauthor* (on *Person*), *Cite* (on *Publication*), and *Keywords* (relating *Publication* and *Keyword*). Moreover, the concept *Publication* has the attributes *year*, *publishedat*, and *type*. Finally, we assume the unary function *in_title* of the type *string* \rightarrow *Publication*.

In the following, some literature search queries are associated with a corresponding conditional preference base PB = (T, A, P), expressed as the conjunction of all the elements in $A \cup P$ (we assume the object ranking κ^{sum}):

- All publications citing papers of ISWC in the year 2000: ∃*Cite.*(*ConferencePublication* □ *publishedat*("ISWC") □=₂₀₀₀(*year*)).
- 2. All publications that cite publications of Ian Horrocks: ∃*Cite*.∃*Author*.{"Ian Horrocks"}.
- 3. All publications of authors who have a joint publication with Tim Berners-Lee:

 \exists Author. \exists Coauthor.{"Tim Berners-Lee"}.

4. All publications that cite journal publications that were cited at least 20 times:

 $\exists Cite.(\geq_{20}Cite^{-} \sqcap JournalPublication).$

5. We are looking for papers with the word "matching" in the title. Among conference papers, we prefer papers of international conferences to papers of national conferences:

 $\begin{array}{l} Publication \sqcap in_title(``matching'') \sqcap \\ (type(``international'') \succ type(``national'') \mid \\ ConferencePublication)[70] \sqcap (ConferencePublication)[80] . \end{array}$

6. All publications of ISWC and KR, preferring those of KR:

 $(publishedat("ISWC") \sqcup publishedat("KR")) \sqcap$ $(publishedat("KR") \succ publishedat("ISWC"))[100].$

 All publications with the keyword "matching". We prefer journal publications that cite at least 4 publications that are cited at least 10 times to non-journal publications that cite at least 5 journal publications cited at least 8 times:

 $\begin{aligned} \exists Keywords. \{``matching''\} \sqcap \\ (\geq_4 Cite. (\geq_{10} Cite^-) \mid Journal Publication) [50] \sqcap \\ (\geq_5 Cite. (\geq_8 Cite^- \sqcap Journal Publication) \mid \\ \neg Journal Publication) [40] \sqcap (Journal Publication) [10]. \end{aligned}$

8. Of all journal publications with the keyword "matching", publications are preferred that cite at least 5 journal publications that are cited at least 8 times to publications that cite at least 4 publications that are cited at least 20 times:

 $\exists Keywords. \{ \{ matching \} \sqcap JournalPublication \sqcap \\ (\geq_5 Cite. (\geq_8 Cite^- \sqcap JournalPublication))[30] \\ (\geq_4 Cite. (\geq_{20} Cite^-))[20].$

9. All publications with the keyword "Semantic Web" that cite Ian Horrocks. We strictly prefer conference publications that are not authored by Ian Horrocks to conference publications by Ian Horrocks:

 $\exists Keywords. \{ \text{``Semantic Web''} \} \sqcap \\ \exists Cite. (\exists Author. \{ \text{``Ian Horrocks''} \}) \sqcap \\ (\neg \exists Author. \{ \text{``Ian Horrocks''} \} \succ \exists Author. \{ \text{``Ian Horrocks''} \}) \mid \\ Conference Publication [70] \sqcap (Conference Publication)[80]. \end{cases}$

10. All publications with the keyword "Semantic Web". The publications should possibly contain the keywords "OWL" and "DAML+OIL". The ranking is influenced by the strength of the keywords "OWL" and "DAML+OIL": ∃*Keywords*.{"Semantic Web"} ⊓ (∃*Keywords*.{"OWL"})[70] ⊓ (∃*Keywords*.{"DAML+OIL"})[20].

The query 5 contains two conditional preferences. An object that fulfills query 5 has to be a publication with the word "matching" in the title and it should possibly satisfy the two conditional preferences. Publications that satisfy the conditional preferences have a lower rank than publications that falsify them. The query 5 therefore divides the publications in the query result into three groups as follows: first international conference publications (lowest rank), second national conference publications (highest rank), and third non-conference publications (highest rank).

The query 7 includes three conditional preferences. An object that fulfills query 7 has to be a publication with the keyword "matching" and it should possibly satisfy the three conditional preferences. A publication cannot be a journal publication and a non-journal publication at the same time, therefore, at worst two of the conditional preferences can be falsified by a publication. The query 7 divides the query result into three groups: first journal publications that cite at least 4 publications that are cited at least 10 times, second non-journal publications that cite at least 5 journal publications that are cited at least 5 journal publications that falsify one of the first two conditional preferences.

Notice that the queries 4, 7, and 8 include a user defined quality measure within the query. The query 10 directly influences the ranking of the query result. Finally, the queries 5, 6, 7, 8, and 9 are clustering the query results.

Algorithms

There are several computational tasks related to conditional preference bases PB = (T, A, P). First, deciding the consis-

Algorithm consistency

Input: pure PB = (T, A, P) with $P \neq \emptyset$. Output: z-partition of PB, if PB is consistent; nil, otherwise. 1. H := P;2. i := -1: 3. repeat 4. i := i + 1;5. $P_i := \{ p \in H \mid p \text{ is tolerated under } T \text{ and } A \text{ by } H \};$ $H := H - P_i$ 6. 7. until $H = \emptyset$ or $P_i = \emptyset$; 8. if $H = \emptyset$ then return (P_0, \ldots, P_i) 9. else return nil.

Figure 2: Algorithm consistency.

Algorithm z^+ -ranking

Input: consistent pure PB = (T, A, P) with $P \neq \emptyset$. **Output**: ranking z^+ of *PB*. 1. for each $p \in P$ do $z^+(p) := 0$; 2. $Z^+ := \{ p \in P \mid p \text{ is tolerated under } T \cup A \text{ by } P \};$ 3. for each $p \in Z^+$ do $z^+(p) := stren(p);$ while $Z^+ \neq P$ do begin 4. $\Delta := \{ p \in P - Z^+ \mid p \text{ is tolerated}$ 5. 6. under $T \cup A$ by $P - Z^+$; 7. update all $z^+(p)$ such that $p \in \Delta$ using equations (1) and (2); 8. 9. $p := \operatorname{argmin}_{p \in \Delta} z^+(p);$ 10. $Z^+ := Z^+ - \{p\}$ end: 11.

12. **return** $(z^+(p))_{p \in P}$.

Figure 3: Algorithm z^+ -ranking.

Algorithm flatten Input: consistent simple PB = (T, A, P) with $P \neq \emptyset$. **Output**: a non-defeasible equivalent PB^* to PB. **Notation**: (P_0, \ldots, P_n) denotes the *z*-partition of *PB*. 1. $D := P_0$; 2. for i := 1 to n do begin 3. $H := \emptyset;$ 4. for each $p = (\alpha \succ \beta \mid \phi)[s] \in D$ do begin 5. $F_p := \{ \psi \mid (\gamma \succ \delta \mid \psi) \mid r \} \in P_i, p \text{ is not tolerated}$ 6. under T and $A \cup \{\psi\}$ by $D \cup P_i \cup \cdots \cup P_n\}$; $H := H \cup \{ (\alpha \succ \beta \mid \phi \sqcap \neg \psi_1 \sqcap \cdots \sqcap \neg \psi_l)[s] \},\$ 7. where $F_p = \{\psi_1, \ldots, \psi_l\}$ 8. 9. end for; 10. $D := H \cup P_i$ end for: 11. 12. return (T, A, D).

Figure 4: Algorithm flatten.

tency of *PB* is done by Algorithm **consistency** (see Fig. 2), which generalizes an algorithm for deciding ε -consistency in default reasoning (Goldszmidt & Pearl 1991a). It takes as input a pure conditional preference base PB = (T, A, P), $P \neq \emptyset$, and it returns as output the z-partition of *PB*, if *PB* is consistent, and *nil*, otherwise. It is essentially based on $O(n^2)$ checks whether a description logic knowledge base is satisfiable, where *n* is the number of elements in *P*.

Computing the preference ranking z^+ of *PB* is done by Algorithm z^+ -ranking (see Fig. 3), which generalizes an algorithm for computing the default ranking z^+ in System Z^+ by Goldszmidt and Pearl (1996). It takes as input a consistent pure conditional preference base PB = (T, A, P), $P \neq \emptyset$, and it returns as output the z^+ -ranking of *PB*. The algorithm is essentially based on $O(n^2 \cdot \log n)$ checks whether a description logic knowledge base is satisfiable.

Another computational task is computing the rank $\kappa^+(\phi)$ of a concept ϕ , which can be done by an algorithm that generalizes the computation of the rank $\kappa^+(\phi)$ of a formula ϕ in System Z^+ by Goldszmidt and Pearl (1996). The extended algorithm is essentially based on $O(\log n)$ checks whether a description logic knowledge base is satisfiable.

Finally, rewriting PB to a non-defeasible equivalent PB^* for the object rankings κ^{sum} and κ^{lex} in the special case in which PB is simple is done by Algorithm **flatten** (see Fig. 4), which is inspired by a rewriting algorithm in fuzzy default reasoning (de Saint-Cyr & Prade 2006). Here, we say PB = (T, A, P) is simple iff $o \models (\neg \alpha \sqcup \beta) \sqcap (\alpha \sqcup \neg \beta)$ for each $(\alpha \succ \beta \mid \phi)[s] \in P$ and each object $o \in \mathcal{O}_{\mathcal{C}}$ that satisfies T and A. Intuitively, PB is simple iff every $p \in P$ is a conditional desire relative to T and A. More precisely, the algorithm takes as input a consistent simple conditional preference base $PB = (T, A, P), P \neq \emptyset$, and it returns as output a non-defeasible equivalent PB^* to PB. The algorithm requires $O(n^2)$ description logic satisfiability checks. The following result shows that Algorithm **flatten** is correct.

Theorem 6 Algorithm **flatten** is correct. That is, given as input a consistent simple conditional preference base PB = (T, A, P) such that $P \neq \emptyset$, Algorithm **flatten** computes a non-defeasible equivalent PB^* to PB.

The above algorithms show that if we restrict the class of description logic expressions in PB in such a way that the above satisfiability checks on description logic knowledge bases can be done in polynomial time (for example, such as in DL-Lite (Calvanese *et al.* 2005)), then all the described computational tasks can also be solved in polynomial time.

Summary and Outlook

We have presented a qualitative probabilistic approach to variable-strength conditional preferences for matchmaking and ranking objects in description logics. More precisely, we have introduced conditional preference bases, which consist of a description logic knowledge base and a finite set of variable-strength conditional preferences, and which are associated with a formal semantics based on ranking functions. We have then defined the notions of consistency and preferential entailment for conditional preference bases, which strictly generalize ε -consistency and entailment in System Z^+ in default reasoning from conditional knowledge bases, respectively. We have also analyzed the semantic properties of preferential entailment, and shown how it can be used for matchmaking with conditional preference bases. Furthermore, we have defined functions for ranking objects relative to a conditional preference base, and we have provided algorithms for solving the main computational tasks related to conditional preference bases.

We have also demonstrated the usefulness of the presented approach to ranking objects under conditional preference bases in the area of literature search. More precisely, search query languages of current search engines are very restricted in their expressive power. There are scientific search engines on the web, however, that have valuable metadata about research publications, authors, organizations, and scientific events. We have shown that conditional preference bases allow for a more powerful query language, which can exploit this metadata better than the current approaches do. In particular, we have given some sample queries that (i) influence the ranking of the query results, (ii) express quality measures, and (iii) cluster query results.

An interesting topic of future research is to explore the concrete application of the presented approach to matchmaking in electronic market places and to matchmaking in web service discovery. Furthermore, it would also be very interesting to explore other important application domains of the presented approach such as, for example, personalization tasks and recommender systems.

Acknowledgments. This work has been partially supported by a Heisenberg Professorship of the German Research Foundation (DFG). We thank the reviewers for their constructive comments, which helped to improve this work.

Appendix: Proofs

Proof of Theorem 1. Recall that κ is admissible with $(\alpha \succ \beta \mid \phi)[s]$ iff $\kappa(\phi \sqcap \alpha) + s < \kappa(\phi \sqcap \beta)$. The latter is equivalent to $\kappa(\phi \sqcap \alpha) + s < \min(\kappa(\phi \sqcap \alpha \sqcap \beta), \kappa(\phi \sqcap \neg \alpha \sqcap \beta))$, which in turn is equivalent to (i) $\kappa(\phi \sqcap \alpha) + s < \kappa(\phi \sqcap \alpha \sqcap \beta)$ and (ii) $\kappa(\phi \sqcap \alpha) + s < \kappa(\phi \sqcap \neg \alpha \sqcap \beta)$. Here, (ii) is equivalent to κ being admissible with $(\alpha \succ \neg \alpha \sqcap \beta \mid \phi)[s]$, while (i) is equivalent to $\min(\kappa(\phi \sqcap \alpha \sqcap \beta), \kappa(\phi \sqcap \alpha \sqcap \neg \beta)) + s < \kappa(\phi \sqcap \alpha \sqcap \beta)$, which in turn is equivalent to $\kappa(\phi \sqcap \alpha \sqcap \neg \beta) + s < \kappa(\phi \sqcap \alpha \sqcap \beta)$. The latter is then equivalent to κ being admissible with $(\alpha \sqcap \neg \beta \succ \alpha \sqcap \beta \mid \phi)[s]$. \Box

Proof of Theorem 2. (\Rightarrow) Suppose first that *PB* is consistent. That is, there exists an object ranking κ that is admissible with *PB*. Let the preference ranking σ on *P* be defined by $\sigma(p) = s + \kappa(\phi \sqcap \alpha)$ for all $p = (\alpha \succ \beta \mid \phi)[s] \in P$. We now show that σ is admissible with *PB*. We first show that for every $P' \subseteq P$, every $p \in P'$ such that $\sigma(p) \leq \sigma(p')$ for all $p' \in P'$ is tolerated under *T* and *A* by *P'*. Towards a contradiction, suppose the contrary. That is, there exists some $p = (\alpha \succ \beta \mid \phi)[s] \in P'$ such that (i) $\sigma(p) \leq \sigma(p')$ for all $p' \in P'$ and (ii) *p* is under *T* and *A* in conflict with *P'*. Consider the object *o* such that $o \models \phi \sqcap \alpha$ and $\kappa(o) = \kappa(\phi \sqcap \alpha)$. Hence, there exists some $p' = (\alpha' \succ \beta' \mid \phi')[s] \in P'$ such that

 $o \models \phi' \sqcap \beta' \sqcap \neg \alpha'$. It thus follows that $s + \kappa(\phi' \sqcap \beta' \sqcap \neg \alpha') \leq s + \kappa(\phi \sqcap \alpha) = \sigma(p) \leq \sigma(p') = s + \kappa(\phi' \sqcap \alpha')$. Since κ is admissible with PB, it follows that $s + \kappa(\phi' \sqcap \alpha') < \kappa(\phi' \sqcap \beta')$. But this is a contradiction, since PB is pure, and thus $\kappa(\phi' \sqcap \beta' \sqcap \neg \alpha') = \kappa(\phi' \sqcap \beta')$. This shows that every $p \in P'$ such that $\sigma(p) \leq \sigma(p')$ for all $p' \in P'$ is tolerated under T and A by P'. Consider now some $P' \subseteq P$ that is under T and A in conflict with some conditional preference $p \in P$. We now show that P' contains some p' such that $\sigma(p') < \sigma(p)$. Towards a contradiction, suppose the contrary. That is, $\sigma(p) \leq \sigma(p')$ for all $p' \in P'$. Then, as argued above, p is tolerated under T and A by P'. But this contradicts P' being under T and A in conflict with p. This shows that P' contains some p' such that $\sigma(p') < \sigma(p)$.

 (\Leftarrow) Suppose next that σ is a preference ranking that is admissible with *PB*. Let the object ranking κ be defined as follows for all objects $o \in \mathcal{O}_{\mathcal{C}}$:

$$\kappa(o) = \begin{cases} \infty & \text{if } o \not\models T \cup A \\ 0 & \text{if } o \models T \cup A \cup P \\ 1 + \max_{p \in P: \ o \not\models p} \sigma(p) \cdot (stren(p) + 1) \text{ otherwise.} \end{cases}$$

We now show that κ is admissible with *PB*. Observe first that $\kappa(o) = \infty$ for all objects $o \in \mathcal{O}_{\mathcal{C}}$ that do not satisfy *T* and *A*. We next prove that $\kappa(\phi \sqcap \alpha) + s < \kappa(\phi \sqcap \beta)$ for all $(\alpha \succ \beta \mid \phi)[s] \in P$. Consider any $p = (\alpha \succ \beta \mid \phi)[s] \in P$. Since σ is admissible with *PB*, it follows that *p* is tolerated under *T* and *A* by the set of all $p' \in P$ such that $\sigma(p') \ge \sigma(p)$. Hence, there exists an object $o \in \mathcal{O}_{\mathcal{C}}$ that satisfies $\phi \sqcap \alpha$ and all $p' \in P$ such that $\sigma(p') \ge \sigma(p)$. Thus, $\kappa(\phi \sqcap \alpha) + s < 1 + \sigma(p) \cdot (s+1)$. Since $o \nvDash p$ for every object *o* that satisfies $\phi \sqcap \beta$, it follows that $1 + \sigma(p) \cdot (s+1) \le \kappa(\phi \sqcap \beta)$. Hence, $\kappa(\phi \sqcap \alpha) + s < \kappa(\phi \sqcap \beta)$. In summary, this shows that κ is admissible with *PB*. \Box

Proof of Theorem 3. Suppose first that *PB* is consistent. By Theorem 2, there exists a preference ranking σ on *P* that is admissible with *PB*. Hence, for every $P' \subseteq P$, the conditional preference $p \in P'$ such that $\sigma(p) \leq \sigma(p')$ for all $p' \in P$ is tolerated under *T* and *A* by *P'*. It thus follows that there exists the unique ordered partition (P_0, \ldots, P_k) of *P* as in (a) (resp., (b)). Conversely, suppose that (P_0, \ldots, P_k) is an ordered partition of *P* as in (a) (resp., (b)). Then, the preference ranking σ defined by $\sigma(p) = i$ for all $p \in P_i$ and $i \in \{0, \ldots, k\}$ is admissible with *PB*. \Box

Proof of Theorem 4. We first prove that κ^+ is admissible with *PB*. Observe that $\kappa(o) = \infty$ for all objects $o \in \mathcal{O}_C$ that do not satisfy *T* and *A*. We next prove that $\kappa^+(\phi \sqcap \alpha) + s < \kappa^+(\phi \sqcap \beta)$ for all $p = (\alpha \succ \beta \mid \phi)[s] \in P$. Consider any $p = (\alpha \succ \beta \mid \phi)[s] \in P$. Since $o \nvDash p$ for every object *o* that satisfies $\phi \sqcap \beta$, it follows that $\kappa^+(\phi \sqcap \beta) \ge 1 + z^+(p) = 1 + s + \kappa^+(\phi \sqcap \alpha) > s + \kappa^+(\phi \sqcap \alpha)$. In summary, κ^+ is admissible with *PB*. Moreover, since $z^+(p) = s + \kappa^+(\phi \sqcap \alpha)$ for all $p = (\alpha \succ \beta \mid \phi)[s] \in P$, by the (\Rightarrow) -part of the proof of Theorem 2, it also follows that z^+ is admissible with *PB*. \Box

Proof of Theorem 5. *Irreflexivity.* Recall that $PB \sim (\alpha \succ \alpha \mid \phi)[s]$ iff $\kappa^+(\phi) = \infty$ or $\kappa^+(\phi \sqcap \alpha) + s < \kappa^+(\phi \sqcap \alpha)$. Since $s \ge 0$, the latter is equivalent to $\kappa^+(\phi) = \infty$, which in turn is equivalent to $o \not\models T \cup A$ for all $o \in \mathcal{O}_{\mathcal{C}}$ that satisfy ϕ . That is, ϕ is unsatisfiable under *PB*.

Asymmetry. Suppose that (a) $PB \triangleright (\alpha \succ \beta \mid \phi)[s]$ and (b) $PB \triangleright (\beta \succ \alpha \mid \phi)[t]$. That is, (a) $\kappa^+(\phi) = \infty$ or $\kappa^+(\phi \sqcap \alpha) + s < \kappa^+(\phi \sqcap \beta)$, and (b) $\kappa^+(\phi) = \infty$ or $\kappa^+(\phi \sqcap \beta) + t < \kappa^+(\phi \sqcap \alpha)$. Towards a contradiction, suppose that $\kappa^+(\phi) \neq \infty$. Hence, $\kappa^+(\phi \sqcap \alpha) + s < \kappa^+(\phi \sqcap \beta)$ and $\kappa^+(\phi \sqcap \beta) + t < \kappa^+(\phi \sqcap \alpha)$, and thus $\kappa^+(\phi \sqcap \alpha) + s + t < \kappa^+(\phi \sqcap \alpha)$, which implies s + t < 0. But this contradicts $s, t \ge 0$. Thus, $\kappa^+(\phi) = \infty$, which is equivalent to $o \nvDash T \cup A$ for all $o \in \mathcal{O}_{\mathcal{C}}$ that satisfy ϕ . That is, ϕ is unsatisfiable under PB.

Transitivity. Suppose that (a) $PB \triangleright (\alpha \succ \beta \mid \phi)[s]$ and (b) $PB \triangleright (\beta \succ \gamma \mid \phi)[t]$. That is, (a) $\kappa^+(\phi) = \infty$ or $\kappa^+(\phi \sqcap \alpha) + s < \kappa^+(\phi \sqcap \beta)$, and (b) $\kappa^+(\phi) = \infty$ or $\kappa^+(\phi \sqcap \beta) + t < \kappa^+(\phi \sqcap \gamma)$. If $\kappa^+(\phi) = \infty$, then $PB \models (\alpha \succ \gamma \mid \phi)[s+t+1]$. Otherwise, $\kappa^+(\phi \sqcap \alpha) + s < \kappa^+(\phi \sqcap \beta)$ and $\kappa^+(\phi \sqcap \beta) + t < \kappa^+(\phi \sqcap \gamma)$ imply that $\kappa^+(\phi \sqcap \alpha) + s + t + 1 < \kappa^+(\phi \sqcap \gamma)$, which in turn implies that $PB \models (\alpha \succ \gamma \mid \phi)[s+t+1]$.

Rational Monotonicity. Suppose that (a) $PB \vdash (\alpha \succ \beta | \phi)[s]$ and (b) $PB \not\vdash (\neg \phi' \mid \phi \sqcap (\alpha \sqcup \beta))[0]$. That is, (a) $\kappa^+(\phi) = \infty$ or $\kappa^+(\phi \sqcap \alpha) + s < \kappa^+(\phi \sqcap \beta)$, and (b) $\kappa^+(\phi \sqcap (\alpha \sqcup \beta)) \neq \infty$ and $\kappa^+(\phi \sqcap (\alpha \sqcup \beta) \sqcap \neg \phi') \ge \kappa^+(\phi \sqcap (\alpha \sqcup \beta) \sqcap \phi')$. By (b), it follows that $\kappa^+(\phi) \neq \infty$. Furthermore, by (a) and (b), it follows that some $o \in \mathcal{O}_{\mathcal{C}}$ exists such that $\kappa^+(o) = \kappa^+(\phi \sqcap \alpha)$ and that o satisfies $T \cup A$ and $\phi \sqcap \alpha \sqcap \phi'$. This shows that $\kappa^+(\phi \sqcap \alpha \sqcap \phi') + s = \kappa^+(\phi \sqcap \alpha) + s < \kappa^+(\phi \sqcap \beta) \le \kappa^+(\phi \sqcap \beta)$

Proof of Theorem 6. We first prove that after the *i*-th iteration step of the for-loop in Algorithm **flatten**, for every $i \in \{0, ..., n\}$, every $p \in D$ is tolerated under $T \cup A$ by $D \cup P_{i+1} \cup \cdots \cup P_n$. In particular, this then shows that after the *n*-th iteration step, every $p \in D$ is tolerated under $T \cup A$ by D, that is, (T, A, D) is flat. We give a proof by induction on the number of iteration steps $i \in \{0, ..., n\}$ as follows (where D_i and $Q_i, i \in \{0, ..., n\}$, denote D and $D \cup P_{i+1} \cup \cdots \cup P_n$ after the *i*-th iteration step, respectively): *Basis:* For i = 0, since $(P_0, ..., P_n)$ is the *z*-partition of *PB*, every $p \in D_0 = P_0$ is tolerated under $T \cup A$ by Q_0 .

Induction: Let i > 0. Suppose that after the i-1-th iteration step, every $p \in D_{i-1}$ is tolerated under $T \cup A$ by Q_{i-1} . We now show that after *i*-th iteration step, every $p \in D_i$ is tolerated under $T \cup A$ by Q_i . Consider first any $p \in D_i - P_i$, which either coincides with some $p' \in D_{i-1}$ or is obtained from some $p' \in D_{i-1}$ by replacing the body ϕ of p' by some $\phi \sqcap \neg \psi_1 \sqcap \cdots \sqcap \neg \psi_l$. By the induction hypothesis, p' is tolerated under $T \cup A$ by Q_{i-1} . That is, there exists an object othat satisfies $T \cup A \cup Q_{i-1}$ and verifies p'. By the construction of the ψ_i 's, it follows that o satisfies every $\neg \psi_i$ and all $q \in D_i - (D_{i-1} \cup P_i)$. That is, o satisfies $T \cup A \cup Q_i$ and verifies p. That is, p is tolerated under $T \cup A$ by Q_i . Consider next any $p \in D_i \cap P_i = P_i$ with body ϕ . Observe first that p is tolerated under $T \cup A$ by $P_i \cup \cdots \cup P_n$, since (P_0, \ldots, P_n) is the z-partition of PB. Thus, if every $q \in D_i - P_i$ has $\neg \phi$ in its body, then p is also tolerated under $T \cup A$ by Q_i . Otherwise, there exists an object o that satisfies $T \cup A \cup Q_{i-1}$ and verifies p. Hence, o satisfies also $T \cup A \cup Q_i$ and verifies p. That is, p is tolerated under $T \cup A$ by Q_i .

In a similar way, by induction on the number of iteration steps $i \in \{0, ..., n\}$ of the for-loop in Algorithm **flatten**, it is not difficult to verify that $PB^0 \succ p$ for every $p \in D_i$. Finally, the constructed (T, A, D) clearly satisfies the condition that $D = \{(\alpha \succ \beta \mid \phi \sqcap \psi_p)[s] \mid p = (\alpha \succ \beta \mid \phi)[s] \in P\}$, where ψ_p is a conjunction of negated bodies that occur in P. \Box

References

Adams, E. W. 1975. *The Logic of Conditionals, Synthese Library* 86. Dordrecht, Netherlands: D. Reidel.

Benferhat, S.; Dubois, D.; and Prade, H. 1992. Representing default rules in possibilistic logic. In *Proceedings KR-1992*, 673–684. Morgan Kaufmann.

Benferhat, S.; Dubois, D.; and Prade, H. 1997. Nonmonotonic reasoning, conditional objects and possibility theory. *Artif. Intell.* 92(1-2):259–276.

Berners-Lee, T. 1999. *Weaving the Web*. San Francisco, CA: Harper.

Boutilier, C.; Brafman, R. I.; Domshlak, C.; Hoos, H. H.; and Poole, D. 2004. CP-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *J. Artif. Intell. Res.* 21:135–191.

Boutilier, C. 1994. Toward a logic for qualitative decision theory. In *Proc. KR-1994*, 75–86. Morgan Kaufmann.

Calvanese, D.; De Giacomo, G.; Lembo, D.; Lenzerini, M.; and Rosati, R. 2005. DL-Lite: Tractable description logics for ontologies. In *Proceedings AAAI-2005*, 602–607.

de Saint-Cyr, F. D., and Prade, H. 2006. Handling uncertainty in a non-monotonic setting, with applications to persistence modeling and fuzzy default reasoning. In *Proceedings KR-2006*. AAAI Press.

Dubois, D., and Prade, H. 1991. Possibilistic logic, preferential models, non-monotonicity and related issues. In *Proceedings IJCAI-1991*, 419–424. Morgan Kaufmann.

Dubois, D., and Prade, H. 1994. Conditional objects as nonmonotonic consequence relationships. *IEEE Trans. Syst. Man Cybern.* 24(12):1724–1740.

Fensel, D.; Wahlster, W.; Lieberman, H.; and Hendler, J., eds. 2002. *Spinning the Semantic Web: Bringing the World Wide Web to Its Full Potential.* MIT Press.

Gabbay, D. M., and Smets, P., eds. 1998. *Handbook on Defeasible Reasoning and Uncertainty Management Systems*. Dordrecht, Netherlands: Kluwer Academic.

Goldszmidt, M., and Pearl, J. 1991a. On the consistency of defeasible databases. *Artif. Intell.* 52(2):121–149.

Goldszmidt, M., and Pearl, J. 1991b. System Z^+ : A formalism for reasoning with variable strength defaults. In *Proceedings AAAI-1991*, 399–404. Morgan Kaufmann.

Goldszmidt, M., and Pearl, J. 1992. Rank-based systems: A simple approach to belief revision, belief update and reasoning about evidence and actions. In *Proceedings KR*-1992, 661–672. Morgan Kaufmann.

Goldszmidt, M., and Pearl, J. 1996. Qualitative probabilities for default reasoning, belief revision, and causal modeling. *Artif. Intell.* 84(1-2):57–112. Horrocks, I., and Patel-Schneider, P. F. 2003. Reducing OWL entailment to description logic satisfiability. In *Proceedings ISWC-2003, LNCS* 2870, 17–29. Springer.

Horrocks, I.; Patel-Schneider, P. F.; and van Harmelen, F. 2003. From SHIQ and RDF to OWL: The making of a web ontology language. *J. Web Semantics* 1(1):7–26.

Horrocks, I.; Sattler, U.; and Tobies, S. 1999. Practical reasoning for expressive description logics. In *Proceedings LPAR-1999*, *LNCS* 1705, 161–180. Springer.

Kraus, S.; Lehmann, D.; and Magidor, M. 1990. Nonmonotonic reasoning, preferential models and cumulative logics. *Artif. Intell.* 14(1):167–207.

Lamarre, P. 1992. A promenade from monotonicity to nonmonotonicity following a theorem prover. In *Proceedings KR-1992*, 572–580. Morgan Kaufmann.

Lang, J.; van der Torre, L.; and Weydert, E. 2002. Utilitarian desires. *Autonomous Agents and Multi-Agent Systems* 5(3):329–363.

Lehmann, D. 1989. What does a conditional knowledge base entail? In *Proceedings KR-1989*, 212–222. Morgan Kaufmann.

Lukasiewicz, T. 2005. Weak nonmonotonic probabilistic logics. *Artif. Intell.* 168(1-2):119–161.

Lukasiewicz, T., and Schellhase, J. 2005. Variable-strength conditional preferences for matchmaking in description logics. Technical Report INFSYS RR-1843-05-11, Institut für Informationssysteme, TU Wien.

Di Noia, T.; Di Sciascio, E.; Donini, F. M.; and Mongiello, M. 2003. Abductive matchmaking using description logics. In *Proceedings IJCAI-2003*, 337–342.

Pearl, J. 1989. Probabilistic semantics for nonmonotonic reasoning: A survey. In *Proceedings KR-1989*, 505–516. Morgan Kaufmann.

Pearl, J. 1990. System Z: A natural ordering of defaults with tractable applications to default reasoning. In *Proceedings TARK-1990*, 121–135. Morgan Kaufmann.

Poole, D., and Smyth, C. 2005. Type uncertainty in ontologically-grounded qualitative probabilistic matching. In *Proc. ECSQARU-2005, LNCS* 3571, 763–774. Springer.

Smyth, C., and Poole, D. 2004. Qualitative probabilistic matching with hierarchical descriptions. In *Proceedings KR*-2004, 479–487. AAAI Press.

Spohn, W. 1988. Ordinal conditional functions: A dynamic theory of epistemic states. In Harper, W., and Skyrms, B., eds., *Causation in Decision, Belief Change, and Statistics*, volume 2. Dordrecht, Netherlands: Reidel. 105–134.

Tan, S.-W., and Pearl, J. 1994. Specification and evaluation of preferences under uncertainty. In *Proceedings KR-1994*, 530–539. Morgan Kaufmann.

W3C. 2004. OWL web ontology language overview. W3C Recommendation (10 February 2004). Available at www. w3.org/TR/2004/REC-owl-features-20040210/.

Weydert, E. 2003. System JLZ — Rational default reasoning by minimal ranking constructions. *J. Applied Logic* 1(3-4):273–308.