

Many Facets of Reasoning under Uncertainty, Inconsistency, Vagueness, and Preferences— A Brief Survey

Gabriele Kern-Isberner ·
Thomas Lukasiewicz

Received: date / Accepted: date

Abstract In this paper, we give an introduction to reasoning under uncertainty, inconsistency, vagueness, and preferences in artificial intelligence (AI), including some historic notes and a brief survey to previous approaches.

Keywords Uncertainty · inconsistency · vagueness · preferences · nonmonotonic reasoning · probability theory · possibility theory · fuzzy logic · ontology languages · description logics · argumentation

After first successful years and promising applications of artificial intelligence (AI) in the 50s and 60s of the past century, first failures provided a plain, but deep insight: the real world is uncertain and dynamic, and every computer system that is meant to be deployed in real-world scenarios must be able to cope with uncertainty and changes, at least to an extent that must be (reasonably) envisaged for the respective domain. However, not only tackling problems brought about by uncertainty is a challenge, but also to find out and take care of what can *reasonably* be expected. That is, uncertainty is a multi-layered and multi-faceted problem that arises from incomplete and possibly erroneous information in evolving scenarios. Information can also be inherently uncertain, like information about future events, or generic information that can be true or false in different scenarios. Moreover, information can be imprecise, when even the degree of uncertainty is not precisely known, or it can be vague, when clear-cut distinctions would not be appropriate (like in commonsense modelling what a “tall person” is).

However, uncertainty is not only a burden, but also a chance, because representing it explicitly and dealing with it adequately broadens and improves

Gabriele Kern-Isberner
Fakultät für Informatik, TU Dortmund, Germany
E-mail: gabriele.kern-isberner@cs.uni-dortmund.de

Thomas Lukasiewicz
Department of Computer Science, University of Oxford, UK
E-mail: thomas.lukasiewicz@cs.ox.ac.uk

the toolbox for modelling real-world scenarios, making systems more effective and successful (in general, however, less efficient, since handling uncertainty nearly always has a price). From a cognitive point of view, uncertainty leaves room for abstraction and condensing information, making at least the representation of knowledge more efficient and more easily accessible. Humans perform quite well using generic rules with exceptions or estimated probabilities, often ignoring exceptions or impreciseness, until they are forced to take them into account. When it comes to decision making, information about utilities (or preferences) is crucial, which is combined with (precise or imprecise) degrees of uncertainty or vagueness to propose a best option. So, preferences complete the picture of using knowledge for decision making, and since there is also some similarity between uncertainty degrees and priority degrees, it makes sense to look upon both in a joint topical scope.

We would like to emphasize that while uncertainty is often deemed to be numerical, we find it more rewarding to consider both quantitative and qualitative approaches to uncertainty together, since the structures underlying uncertain reasoning in any framework are conceptually similar. Nevertheless, it is still a challenge to establish bridges between both paradigms, because this always demands for understanding and overcoming slight but crucial differences. The same holds for (quantitative) utilities and (qualitative) preferences—although utilities combined with probabilities is the classical approach to decision theory, more general preference relations have become quite popular for modelling human-like decision making [57, 2].

A phenomenon that appears in all approaches to uncertain reasoning is *nonmonotonicity*—previously drawn conclusions are revised when more information arrives: A physician might recommend a therapy A as the most suitable treatment for a specific patient, but coming to know the results of blood tests or genetical tests might completely change his mind. Most of the conclusions that we draw in our everyday lives and on which we base (sometimes important) decisions are defeasible, prone to be given up when further information arrives. This is in clear contrast to classical—e.g., propositional or first-order—logics, which are monotone, i.e., their deductive conclusions are preserved for eternity. It is due to monotonicity that all (correct) proofs in mathematics are still valid, regardless of whatever new theories are being developed, but also that robots based on classical logics fail in uncertain, incompletely specified environments. Nonmonotonicity changes nearly everything that we appreciate in classical logics: Inference rules like modus ponens, transitivity or contraposition do not hold any longer in general. Due to this nonmonotonicity, knowledge is regarded more like subjective beliefs that an agent holds in an uncertain environment.

The first nonmonotonic logics were proposed in the 80s of the past century (cf. [50, 44, 45] for major approaches) focusing on default rules, which are generic rules that allow for exceptions. Also approaches to modern belief revision theory emerged in that time (as a seminal paper, cf. [1]; for a recent overview, see [16]; for a textbook, cf., e.g., [23]). Nonmonotonic logics and belief revision are closely related, since both deal with the rational processing

of further, or new information [18]. Indeed, nonmonotonic reasoning is not a (general) theory of just “jumping to conclusions” (in a maybe uncontrolled way), but rather should be looked upon as logic-based theories to approximate the rationality of human reasoning. Similar remarks apply to belief revision theory: Keeping in mind that changing is easy as long as one does not care about consequences, AGM belief revision (named after the authors of [1]) is a frame for rational belief change under logical closure. In both areas of nonmonotonic logics and belief revision, rationality postulates play a major role. For nonmonotonic logics, system P [32] is currently the most widely used axiom system that is based on well-known classical axioms but replaces monotony by *cautious monotony*: If B can be plausibly inferred from A , then it can be still plausibly inferred from $A \wedge C$, if C is also a plausible consequence of A , i.e., monotony still holds under certain restrictions on the new information. The semantical structures that guide system-P-compliant nonmonotonic reasoning are so-called preference relations on possible worlds that can be interpreted by saying that one world is preferred to another world, if the preferred world is more plausible. So indeed, preferences and plausibilities share some formal common grounds, although they express substantially different notions. Preferential entailment [32, 43] induces nonmonotonic inference relations that are quite intuitive. However, system P is a current standard, but not a gold standard. There is no sufficient proof that human reasoning obeys all axioms of system P, although there is some evidence (cf., e.g., [46]), but one of the most successful and practically relevant approaches to nonmonotonic reasoning, namely extended logic programming under the answer set semantics (cf., e.g., [21, 36, 20]) does not satisfy cautious monotony and so also not system P. The AGM postulates in belief revision theory [1] are closely related to the axioms of system P [18] and can be effectively implemented via ranking preorders such as epistemic entrenchment [19], or Spohn’s ordinal conditional functions [52, 53], which can also be regarded as preference relations.

Much earlier than nonmonotonic logics and the AGM theory, but dealing with similar problems, probability theory has been offering a rich and powerful framework for uncertain reasoning and belief revision: A conditional probability $Prob(B|A)$ can be drastically different from $Prob(B|A \wedge C)$, i.e., conditional probabilities behave nonmonotonically, and Bayes proposed with his definition of conditional probability and his well-known rule immediate and elegant ways to change probabilistic beliefs and perform uncertain reasoning, which has been shown to be in line with the AGM theory [17] and hence with nonmonotonic preferential entailment. Applying the principles of maximum entropy and minimum cross-entropy improves the quality of probabilistic reasoning and probabilistic belief change considerably and induces high-quality inference relations and belief change operators [27, 29, 30, 48]. Paris even argued that an optimal probabilistic commonsense reasoner should reason according to the principle of maximum entropy [49].

There has been an extensive discussion in the literature on the difference between probability/possibility theory and vagueness/fuzziness theory. We recall here some notes from [14] to clarify the role of these theories. A standard

example that explains the difference between degrees of uncertainty and degrees of truth is that of a bottle [14]. In terms of binary truth values, a bottle is either full or empty. However, if one accounts for the quantity of liquid in the bottle, one may, e.g., say that the bottle is half full. Under this way of speaking, “full” becomes a fuzzy predicate [59] and the degree of truth of “the bottle is full” reflects the amount of liquid in the bottle. The situation is very different when we say that the probability of the bottle being full is 0.5, which does not mean that the bottle is half full, but that the bottle is either full or empty, and that we are ignorant about which of the two is true.

More concretely, under (quantitative) *uncertainty theory* fall all those approaches in which statements rather than being either true or false, are true or false to some *probability* or *possibility* (e.g., “it will rain tomorrow”). That is, a statement is true or false in any world, but we are “uncertain” about which world to consider as the right one, and thus we speak about, e.g., a probability distribution or a possibility distribution over the worlds. For example, we cannot exactly establish whether it will rain tomorrow or not, due to our *incomplete* knowledge about our world, but we can estimate to which degree this is probable, possible, or necessary. As for the main differences between probability and possibility theory, the probability of an event is the sum of the probabilities of all worlds that satisfy this event, whereas the possibility of an event is the maximum of the possibilities of all worlds that satisfy the event. Intuitively, the probability of an event aggregates the probabilities of all worlds that satisfy this event, while the possibility of an event is simply the possibility of the “most optimistic” world that satisfies the event. Hence, although both probability and possibility theory allow for quantifying degrees of uncertainty, they are conceptually quite different from each other. That is, probability and possibility theory represent different facets of uncertainty.

On the other hand, under *vagueness/fuzziness theory* fall all those approaches in which statements (e.g., “the tomato is ripe”) are true to some degree, which is taken from a truth space. That is, an interpretation maps a statement to a truth degree, because we are unable to establish whether a statement is completely true or false, due to the involvement of vague concepts, such as “ripe”, which only have an *imprecise* definition. For example, we cannot exactly say whether a tomato is ripe or not, but rather can only say that the tomato is ripe to some degree. Usually, such statements involve so-called *vague/fuzzy predicates* [59]. Note that all vague/fuzzy statements are truth-functional, i.e., the degree of truth of every statement can be calculated from the degrees of truth of its constituents, while uncertain statements cannot be a function of the uncertainties of their constituents [13]. More specifically, in probability theory, only the negation of events is truth-functional, while in possibility theory, only the disjunction (resp., conjunction) of events is truth-functional in possibilities (resp., necessities) of events. Furthermore, fuzzy logics are based on truly many-valued logical operators, while uncertainty logics are defined on top of standard binary logical operators.

While uncertainty and vagueness on the base of propositional logic has been well-explored, the combination of (fragments of) predicate logic with

any form of uncertainty and vagueness is particularly demanding. Subjective uncertainty (of an agent’s world view) collides with objective (or statistical) uncertainty (concerning a population), and coming up with a coherent solution that moreover matches human’s broad understanding of uncertainty is not an easy task. However, the rising popularity and use of ontology languages (such as description logics (DLs)) as (fragments of) predicate logic, and the need to deal with uncertainty and vagueness in their typical application areas (e.g., in the Semantic Web and ontology-based data access (OBDA)) has increasingly attracted the attention of many researchers and practitioners towards ontology languages able to cope with uncertainty and vagueness.

Probabilistic ontology languages [39] can be classified according to the underlying ontology language, the probabilistic semantics, and the supported forms of probabilistic knowledge. Some early approaches [24,26] generalize the central DL \mathcal{ALC} and are based on propositional probabilistic logics, while others [31] generalize tractable DLs and are based on Bayesian networks as probabilistic semantics. The approach in [37] generalizes the expressive DLs behind the sublanguages OWL Lite and OWL DL of the Web Ontology Language (OWL), and is based on probabilistic default logics, while [12] generalizes OWL by probabilistic uncertainty as in multi-entity and standard Bayesian networks. Other recent works explore probabilistic DLs for subjective uncertainty [40] and lift query answering dichotomy results from probabilistic databases to OWL2 QL over probabilistic data [28]. Similarly, there are several extensions of ontology languages using the theory of fuzzy logic [39], which can be classified according to the underlying ontology language, the underlying fuzzy logics, and the allowed fuzzy constructs. Prominent early works extend (fragments of) \mathcal{ALC} based on Zadeh logic [58,56,54]. Recent problems for fuzzy ontology languages are the top- k retrieval problem (which has been adopted from the context of relational databases) [55], as well as the decidability of fuzzy DLs with general concept inclusions [4]. Closely related to degrees of truth in ontology languages, the notion of *similarity* quantifies how close two objects are from a conceptual point of view, and has especially been applied to ontologies [41] and their parts (such as concepts) [9].

In contrast to uncertainty, where one lacks information, the notion of *inconsistency* describes the case where we have “too much” information, i.e., the available information is logically inconsistent, which can especially be the case when we integrate information from many different sources. Due to the “ex falso quodlibet” principle, everything follows from an inconsistent collection of logical sentences, i.e., the classical model-theoretic semantics of logics fails completely in such situations. Early solutions to this problem have been developed in databases, where the areas of *database repairing* and *consistent query answering* (CQA) have gained much attention since the work of [3], which provided the model-theoretic construct of a database *repair*—a consistent database that “minimally” differs from the original (inconsistent) database instance. The most widely accepted semantics for querying a possibly inconsistent database is that of *consistent answers*, which yields the set of tuples (atoms) that appear in the answer to the query over *every* possible

repair. CQA enforces consistency at query time as an alternative to enforcing it at the instance level, as conventional data cleaning techniques do. The work of [11] addresses the basic concepts and results of the area of CQA.

Recently, several works have focused on inconsistency handling for several classes of DLs, adapting and specializing general techniques previously considered for traditional logics. In [22], a comparison between four different approaches to handling inconsistency in DL-based ontologies is presented: *consistent ontology evolution*, *repairing inconsistencies*, *reasoning with inconsistent ontologies*, and *ontology versioning*. In [5], a survey of existing approaches (up to 2007) is presented; the authors analyze different works from the AI literature and their usability on practical problems. In contrast, [34] studies the adaptation of CQA for *DL-Lite* ontologies. The complexity of the problem was studied in [51] for a wider spectrum of DLs. In [35], the authors explore first-order-rewritability of queries over *DL-Lite* ontologies under the IAR and ICAR semantics (intersection of ABox repairs and closed repairs, respectively); both semantics were first introduced in [34] as sound approximations of consistent answers. Other work [38] generalizes the previous work for consistent answers and the IAR semantics to Datalog+/- ontologies, and it proposes the *k*-lazy semantics, which relaxes the notion of minimal repair by resolving conflicts up to a certain budget. In [7], the author studies and formulates general conditions for the existence or non-existence of an FO-rewriting for *DL-Lite_{core}* for consistent instance checking. Following that line of work, [8] provides a fine-grained complexity analysis that aims to understand what makes consistent answers so hard to compute even for simple ontologies consisting of class subsumption and class disjointness. Furthermore, an alternative sound approximation to consistent answers is presented: all answers that are implied by the intersection of the *closure* of the repairs; this is a finer approximation than the IAR semantics (but different from the ICAR semantics in [34]).

Another approach to dealing with inconsistency is based on paraconsistent techniques, which allow contradictory information to be introduced and derived without trivialization. In particular, in [47], the authors define a 3-valued paraconsistent semantics for the DL *SHIQ* [25], while [42] proposes a four-valued semantics for a DL that underlies OWL 2 and explores its impact on several tractable DLs of the OWL 2 standard. The latter approach is reducible to reasoning in DLs under classical semantics, and the transformation required for this is linear in the size of the knowledge base; this is closely related to fuzzy DLs, which can often also be reduced to DLs under classical semantics.

Quite a different methodology that makes use of various aspects of uncertainty, conflict resolution, and preferences is argumentation theory [15, 6]. Argumentation frameworks help to evaluate a claim by elaborating the arguments that are in favor or against the claim. Arguments may form coalitions to support or defend other arguments, and preference relations on arguments induce notions of strength for arguments which are crucial to win an argumentation dialogue. The broad applicability of argumentation frameworks for reasoning and decision making in (multi-)agent scenarios and its closeness to

human dialogues advanced argumentation to one of the most attractive topics in knowledge representation over the past decade.

Given the wide variety of approaches to reasoning under uncertainty, inconsistency, vagueness, and preferences, a major challenge is to find guidelines for how to combine different approaches, and which combinations are particularly useful for which problem classes. Approaches to hybrid reasoning have been developed successfully in recent years (cf., e.g., [33]), and multi-context systems [10] may provide general adequate frameworks for this.

Acknowledgments

This work was partially supported by the UK EPSRC grants EP/J008346/1, EP/L012138/1, and EP/M025268/1, and by The Alan Turing Institute under the EPSRC grant EP/N510129/1.

References

1. Alchourrón, C., Gärdenfors, P., Makinson, D.: On the logic of theory change: Partial meet contraction and revision functions. *J. Symb. Logic* **50**(2), 510–530 (1985)
2. Andreka, H., Ryan, M., Schobbens, P.Y.: Operators and laws for combining preference relations. *J. Logic Comput.* **12**(1), 13–53 (2002)
3. Arenas, M., Bertossi, L.E., Chomicki, J.: Consistent query answers in inconsistent databases. In: *Proc. of PODS*, pp. 68–79 (1999)
4. Baader, F., Borgwardt, S., Peñaloza, R.: On the decidability status of fuzzy \mathcal{ALC} with general concept inclusions. *J. Philos. Logic* **44**(2), 117–146 (2015)
5. Bell, D.A., Qi, G., Liu, W.: Approaches to inconsistency handling in description-logic based ontologies. In: *OTM Workshops (2)*, pp. 1303–1311 (2007)
6. Besnard, P., Hunter, A.: *Elements of Argumentation*. MIT Press (2008)
7. Bienvenu, M.: First-order expressibility results for queries over inconsistent *DL-Lite* knowledge bases. In: *Proc. of DL* (2011)
8. Bienvenu, M.: On the complexity of consistent query answering in the presence of simple ontologies. In: *Proc. of AAAI*, pp. 705–711 (2012)
9. Borgida, A., Walsh, T.J., Hirsh, H.: Towards measuring similarity in description logics. In: *Proc. of DL*, (2005)
10. Brewka, G., Ellmauthaler, S., Gonçalves, R., Knorr, M., Leite, J., Pührer, J.: Inconsistency management in reactive multi-context systems. In: *Proc. of JELIA*, pp. 529–535 (2016)
11. Chomicki, J.: Consistent query answering: Five easy pieces. In: *Proc. of ICDT*, pp. 1–17 (2007)
12. da Costa, P.C.G., Laskey, K.B.: PR-OWL: A framework for probabilistic ontologies. In: *Proc. of FOIS*, pp. 237–249 (2006)
13. Dubois, D., Prade, H.: Can we enforce full compositionality in uncertainty calculi? In: *Proc. of AAAI*, pp. 149–154 (1994)
14. Dubois, D., Prade, H.: Possibility theory, probability theory and multiple-valued logics: A clarification. *Ann. Math. Artif. Intell.* **32**(1–4), 35–66 (2001)
15. Dung, P.: On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artif. Intell.* **77**, 321–357 (1995)
16. Fermé, E., Hansson, S.: AGM 25 years: Twenty-five years of research in belief change. *J. Philos. Logic* **40**(2), 295–331 (2011)
17. Gärdenfors, P.: *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press, Cambridge, MA (1988)

18. Gärdenfors, P.: Belief revision and nonmonotonic logic: Two sides of the same coin? In: Proc. of ECAI, pp. 768–773 (1992)
19. Gärdenfors, P., Makinson, D.: Revisions of knowledge systems using epistemic entrenchment. In: Proc. of TARK, pp. 83–95 (1988)
20. Gebser, M., Kaminski, R., Kaufmann, B., Schaub, T.: Answer Set Solving in Practice. Morgan and Claypool, San Rafael, CA (2013)
21. Gelfond, M., Leone, N.: Logic programming and knowledge representation—The A-Prolog perspective. *Artif. Intell.* **138**, 3–38 (2002)
22. Haase, P., Harmelen, F.V., Huang, Z., Stuckenschmidt, H., Sure, Y.: A framework for handling inconsistency in changing ontologies. In: Proc. of ISWC, pp. 353–367 (2005)
23. Hansson, S.: A Textbook of Belief Dynamics. Kluwer Academic Publishers, Dordrecht, Netherlands (1999)
24. Heinsohn, J.: Probabilistic description logics. In: Proc. of UAI, pp. 311–318 (1994)
25. Horrocks, I., Sattler, U., Tobies, S.: Reasoning with individuals for the description logic *SHIQ*. In: Proc. of CADE, pp. 482–496 (2000)
26. Jaeger, M.: Probabilistic reasoning in terminological logics. In: Proc. of KR, pp. 305–316 (1994)
27. Jaynes, E.: Papers on Probability, Statistics and Statistical Physics. D. Reidel Publishing Company, Dordrecht, Holland (1983)
28. Jung, J.C., Lutz, C.: Ontology-based access to probabilistic data with OWL QL. In: Proc. of ISWC, Part I, pp. 182–197 (2012)
29. Kern-Isberner, G.: Conditionals in Nonmonotonic Reasoning and Belief Revision. *LNAI* 2087, Springer (2001)
30. Kern-Isberner, G.: A thorough axiomatization of a principle of conditional preservation in belief revision. *Ann. Math. Artif. Intell.* **40**(1/2), 127–164 (2004)
31. Koller, D., Levy, A., Pfeffer, A.: P-CLASSIC: A tractable probabilistic description logic. In: Proc. of AAAI, pp. 390–397 (1997)
32. Kraus, S., Lehmann, D., Magidor, M.: Nonmonotonic reasoning, preferential models and cumulative logics. *Artif. Intell.* **44**, 167–207 (1990)
33. Lakemeyer, G., Erdem, E., Kern-Isberner, G., Schaub, T.: Workshop on Hybrid reasoning at IJCAI-2015. <https://www.hybrid-reasoning.org/ijcai15.ws> (2015)
34. Lembo, D., Lenzerini, M., Rosati, R., Ruzzi, M., Savo, D.F.: Inconsistency-tolerant semantics for description logics. In: Proc. of RR, pp. 103–117 (2010)
35. Lembo, D., Lenzerini, M., Rosati, R., Ruzzi, M., Savo, D.F.: Query rewriting for inconsistent *DL-Lite* ontologies. In: Proc. of RR, pp. 155–169 (2011)
36. Leone, N., Eiter, T., Faber, W., Calimeri, F., Dell’Armi, T., Eiter, T., Gottlob, G., Ianni, G., Ielpa, G., Koch, C., Perri, S., Polleres, A.: The DLV system. In: Proc. of JELIA, pp. 537–540 (2002)
37. Lukasiewicz, T.: Expressive probabilistic description logics. *Artif. Intell.* **172**(6/7), 852–883 (2008)
38. Lukasiewicz, T., Martinez, M.V., Simari, G.I.: Inconsistency handling in Datalog+/- ontologies. In: Proc. of ECAI, pp. 558–563 (2012)
39. Lukasiewicz, T., Straccia, U.: Managing uncertainty and vagueness in description logics for the Semantic Web. *J. Web Sem.* **6**(4), 291–308 (2008)
40. Lutz, C., Schröder, L.: Probabilistic description logics for subjective uncertainty. In: Proc. of KR, pp. 393–403 (2010)
41. Maedche, A., Staab, S.: Measuring similarity between ontologies. In: Proc. of EKAW, pp. 251–263 (2002)
42. Maier, F., Ma, Y., Hitzler, P.: Paraconsistent OWL and related logics. *Semantic Web* **4**(4), 395–427 (2013).
43. Makinson, D.: General theory of cumulative inference. In: Proc. of NMR, pp. 1–18 (1989)
44. McDermott, D., Doyle, J.: Non-monotonic logic I. *Artif. Intell.* **13**, 41–72 (1980)
45. Moore, R.: A formal theory of knowledge and action. In: J. Hobbs, R. Moore (eds.) *Formal Theories of the Commonsense World*, pp. 319–358. Ablex Publishing (1985)
46. Neves, R., Bonnefon, J., Raufaste, E.: An empirical test of patterns for nonmonotonic inference. *Ann. Math. Artif. Intell.* **34**(1–3), 107–130 (2002)
47. Nguyen, L.A., Szalas, A.: Three-valued paraconsistent reasoning for Semantic Web agents. In: Proc. of KES-AMSTA, pp. 152–162 (2010)

48. Paris, J.: *The Uncertain Reasoner's Companion—A Mathematical Perspective*. Cambridge University Press (1994)
49. Paris, J.: Common sense and maximum entropy. *Synthese* **117**, 75–93 (1999)
50. Reiter, R.: A logic for default reasoning. *Artif. Intell.* **13**, 81–132 (1980)
51. Rosati, R.: On the complexity of dealing with inconsistency in description logic ontologies. In: *Proc. of IJCAI*, pp. 1057–1062 (2011)
52. Spohn, W.: Ordinal conditional functions: A dynamic theory of epistemic states. In: W. Harper, B. Skyrms (eds.) *Causation in Decision, Belief Change, and Statistics, II*, pp. 105–134. Kluwer Academic Publishers (1988)
53. Spohn, W.: *The Laws of Belief: Ranking Theory and Its Philosophical Applications*. Oxford University Press (2012)
54. Straccia, U.: Reasoning within fuzzy description logics. *J. Artif. Intell. Res.* **14**, 137–166 (2001)
55. Straccia, U.: Top-k retrieval for ontology mediated access to relational databases. *Inf. Sci.* **198**, 1–23 (2012)
56. Tresp, C., Molitor, R.: A description logic for vague knowledge. In: *Proc. of ECAI*, pp. 361–365 (1998)
57. van der Torre, L.W.N.: Defeasible goals. In: *Proc. of ECSQARU*, pp. 274–385 (1999)
58. Yen, J.: Generalizing term subsumption languages to fuzzy logic. In: *Proc. of IJCAI*, pp. 472–477 (1991)
59. Zadeh, L.A.: Fuzzy sets. *Inf. Control* **8**(3), 338–353 (1965)