

Preferences in Artificial Intelligence

Gabriella Pigozzi · Alexis Tsoukiàs ·
Paolo Viappiani

Received: date / Accepted: date

Abstract The paper presents a focused survey about the presence and the use of the concept of “preferences” in Artificial Intelligence. Preferences are a central concept for decision making and have extensively been studied in disciplines such as economy, operational research, decision analysis, psychology and philosophy. However, in the recent years it has also become an important topic both for research and applications in Computer Science and more specifically in Artificial Intelligence, in fields spanning from recommender systems to automatic planning, from non monotonic reasoning to computational social choice and algorithmic decision theory. The survey essentially covers the basics of preference modelling, the use of preference in reasoning and argumentation, the problem of compact representations of preferences, preference learning and the use of non conventional preference models based on extended logical languages. It aims at providing a general reference for all researchers both in Artificial Intelligence and Decision Analysis interested in this exciting interdisciplinary topic.

Gabriella Pigozzi
PSL, University Paris Dauphine
CNRS, UMR 7243, LAMSADE, Paris, France
E-mail: gabriella.pigozzi@dauphine.fr

Alexis Tsoukiàs
CNRS, UMR 7243, LAMSADE, Paris, France
PSL, University Paris Dauphine
E-mail: alexis.tsoukias@dauphine.fr

Paolo Viappiani
CNRS, UMR 7606, LIP6, F-75005, Paris, France
Sorbonne Universités, UPMC Univ Paris 06, UMR 7606, LIP6
4 Place Jussieu, 75005 Paris, France E-mail: paolo.viappiani@lip6.fr

1 Introduction and Notation

Preferences are a central concept of decision making and have extensively been studied in disciplines such as economy, operations research, psychology, and philosophy. As preferences are fundamental for the analysis of human choice behaviour, they are becoming of increasing importance for computational fields such as artificial intelligence, databases, and human-computer interaction. Preference models are needed in decision-support systems such as web-based recommender systems, in automated problem solvers such as configurators, in autonomous systems such as Mars rovers. Nearly all areas of artificial intelligence deal with choice situations and can thus benefit from computational methods for handling preferences. Moreover, recommender systems, personal assistants, and other interactive systems need to elicit and satisfy the user's preferences in order to be able to give truly satisfactory recommendations. Social choice methods are also becoming of importance in computational domains such as multi-agent systems.

The field of “preferences” became an emerging area of scientific investigation for several research groups in computer science and in the recent years we assist to a number of workshops, conferences and editorial initiatives aiming at promoting this area at the edge of fields such as decision analysis, artificial intelligence, social choice and economics (see for instance [244]). We mention:

- the special issue (vol. 20) of the journal *Computational Intelligence* in 2004, (see for instance [95]),
- the Dagstuhl seminar¹ (04271) on “Preferences: Specification, Inference, Applications” in 2004, starting from which a number of seminars and workshops have been organised every year since (notably the MPREF series of Multidisciplinary Workshops on Advances in Preference Handling),
- the special issue of the *AI Magazine* on “Preference Handling in AI” published in Winter 2008 (that included a general tutorial [60], and survey papers covering preferences in interactive systems [223], planning [24], conversational recommender systems and electronic commerce applications [237], constraint satisfaction [249], social choice [79] and multiobjective optimization [111]),
- the special issue of the journal *Annals of Operations Research* (vol. 163) in 2008 (see [245]),
- the Algorithmic Decision Theory conferences held in 2009, 2011 and 2013 (see [247, 61, 226]),
- the special issue (vol. 175) of the journal *Artificial Intelligence* in 2011, (see [94]),
- the Dagstuhl seminar² (14101) on “Preference Learning” in 2014, and

¹ <http://www.dagstuhl.de/en/program/calendar/semhp/?semnr=04271>

² <http://www.dagstuhl.de/en/program/calendar/semhp/?semnr=14101>

- the establishment of the EURO Working Group on Preference Handling which regroups the international community concerned by this subject.³

This article makes an overview of the recent advances in this area, with a focus on reasoning, argumentation, deontic reasoning, representation, learning, and non-classical models (preference representations based on fuzzy sets and beyond). The reader only needs to be acquainted with basic notions of discrete mathematics and logic.

Traditionally preferences have always been modelled as binary relations applied to a set which we will denote as A . For the purpose of this article we will only consider finite or enumerable infinite sets which can be of three types:

- enumeration of objects;
- subsets of vector spaces (typically $A \subseteq \mathbb{R}^n$, A being an enumerable subset);
- subsets of the product space of n attributes X_i ($A \subseteq \prod_i X_i$).

Hence preferences over composed objects such as trees or graphs are not considered here.

The use of binary relations includes the cases where preferences are expressed either as binary comparisons of objects among them (relative comparison) or as binary comparisons of objects to “norms” or “standards” (absolute comparison), see [83].

Basic references on this issue can be considered: [115, 178, 243, 117, 251, 169, 269, 275, 229, 119, 4, 219]. In this article we adopt the notation introduced in [251]. The usual definitions of (a)symmetric, (ir)reflexive, transitive, Ferrers⁴, etc. relations apply. We use a generic preference relation (being just reflexive) denoted \succeq to be read “at least as good as”, from which we can get an asymmetric part, denoted \succ (and usually called strict preference) and a symmetric one denoted \sim . The symmetric part can be distinguished in indifference (denoted \approx) and incomparability (denoted \bowtie), the later being irreflexive (while indifference is reflexive). In Section 7 we will see that also the asymmetric part can be further decomposed in several relations. Besides using an explicit representation (in terms of sets), preferences are usually represented using graphs, directly representing the binary relation, matrices (a non graphical representation of a graph) and, under precise conditions, numerically (so that the ordering resulting on the set A can be expressed using the natural ordering of numbers).

We adopt the term “preference structure” in order to denote collections of preference relations which establish ordered partitions of A and fulfil a number of properties. Such preference structures are named as “weak orders”, “interval orders”, “*PQI* interval orders”, etc. (for definitions see [219]). Preference

³ The reader is invited to check the website <http://preferencehandling.free.fr> for an account of all the activities related to this domain.

⁴ A relation $R \subseteq A \times A$ is Ferrers if $\forall x, y, z, w \in A, (xRy \wedge zRw) \rightarrow (xRw \vee zRy)$.

structures can be “characterised”: showing the necessary and sufficient conditions for which actual preferences upon a set A happen to be one of these structures. Such representation theorems can be of three types:

- direct conditions upon the binary relations and their combinations;
- forbidden configurations of the associated graph structure;
- specific conditions admitting numerical presentations.

A typical example is the interval order preference structure. The first type of characterisation states that \succeq is an interval order iff $\forall x, y, z, w \in A (x \succ y \wedge y \sim z \wedge z \succ w) \rightarrow x \succ w$. The second type of characterisation states that \succeq is an interval order iff no subgraph on 4 vertices x, y, z, w can just consist of the arcs (x, y) , (z, w) and neither (x, w) nor (z, y) . The third type of characterisation states that \succeq is an interval order iff $\forall x, y \exists l, r : A \mapsto \mathbb{R} l(x) < r(x) : x \succ y \leftrightarrow l(x) > r(y)$.

We already introduced the difference between “relative” and “absolute” comparisons (or preferences). In the first case $\succeq \subseteq A \times A$. In the second case $\succeq \subseteq A \times N \cup N \times A$ where N is the set of norms or standards to which elements of A have to compare in order to make an assessment (for instance when we say that $x \in A$ is “good” we assume that $x \succ y$, $y \in N$ being the norm or standard of “good”). Another difference can be established between “direct” and “extended” preferences. The first ones are as usual represented by $\succeq \subseteq A \times A$. The second ones are represented by $\succcurlyeq \subseteq 2^A \times 2^A$. Generally a certain coherence between \succeq and \succcurlyeq is expected when they concern the same set A . In other terms, extended preferences concern the comparison of whole subsets of A among them and not just among single elements (the reader can see [25, 212, 211] for the multiple semantics of this type of preferences). A third distinction can be introduced between “first” and “second” order preferences. First order preferences are the usual orderings upon a set A . Second order preferences instead concern any potential order among the orderings of the set A . Consider the case where we have n orderings \succeq_j upon a set A representing preferences holding under different scenarios (let us call J the set of scenarios; and $j \in J$). The existence of an ordering relation $\triangleright \subseteq J \times J$ represents a second order preference among the orderings obtained for each single scenario. Typical cases are likelihood comparisons (scenario i is likely to occur not less than scenario j) or importance comparisons (the order according to dimension i is at least as important as the order according to dimension j : $\succeq_i \triangleright \succeq_j$). The reader should note that in most cases second order preferences are not independent from first order preferences (cfr. [197, 52]) and that despite their intuitive appealing they do not constitute “primitive information” for the construction of decision models (see [83]).

The types of preferences we discussed until now introduce purely ordinal information. In the case such preferences admit a numerical representation (for instance of the type: $x \succ y \leftrightarrow f(x) > f(y)$) this is not unique: all monotonic increasing transformations of the numerical representation are admissible. In other terms we only consider the ordinal information numbers convey. The values (numbers) we can associate to elements of A in order to respect the

order induced by some preference model do not admit any “quantitative” interpretation. Should we be interested in more “quantitative” information we need to be able to model “differences of preferences”: the difference of preference between x and y is at least as large as the difference of preference between z and w ($xy \succsim zw$). The reader can see more in [52] and [56].

Preferences are conveyed through “preference statements”: *I like x , y is better than z , I do not like w , x combined to y is worse than z combined to w etc.* As such they can also be modeled as logical sentences (using some appropriate language). For this purpose we will make reference to some basic logical notation (for a basic reference see [285]). Moreover we draw the attention of the reader to issues related to the semantics of logical inference (model theory) since preferences can and have been used in order to extend reasoning (for an introduction see [95]).

The article is organized as follows. Section 2 reviews the literature about the use of preferences in extending reasoning models. Section 3 presents the use of preferences in argumentation theory. Section 4 discusses deontic logic, an approach to model preference statements in specially tailored languages. Section 5 discusses the problem of compact (in computational terms) representations of preferences. Section 6 presents the literature about preference learning. Section 7 reviews non conventional preference models, mainly established using logical approaches.

2 Preferences in reasoning

I have an appointment for the first time with Björn, a Swedish man. When I arrive, I am introduced to a short man with dark hair and dark eyes. He is Björn. My surprise (and maybe yours as well) is due to the fact that the prototypical Swedish person is tall, blond and blue-eyed. Everyday we reason with incomplete information, we assume the world is as normal as possible and jump to conclusions, that may be later given up upon learning new information. Such kind of reasoning is called *nonmonotonic*, to distinguish it from the traditional deductive inference in which the set of conclusions grows proportionally to the set of available information. So, following a famous syllogisms, from “All men are mortal” and “Socrates is a man”, we can derive that “Socrates is mortal” and this conclusion will remain no matter what other information we may add later. In other words, in classical logic, we cannot retract a previously obtained conclusion. If something was derivable at some point, it will still be derivable if we add more premises. This can be expressed formally by the *Monotonicity* property for the classical consequence relation:

$$\text{If } \Gamma \vdash \alpha, \text{ then } \Gamma \cup \beta \vdash \alpha \tag{1}$$

where Γ is a finite set of formulas and α and β are formulas of a propositional language \mathcal{L}^5 , built up from a finite set \mathcal{P} of propositional symbols and

⁵ For simplicity here we consider classical propositional logic.

the usual connectives ($\neg, \wedge, \vee, \rightarrow, \leftrightarrow$). An *interpretation* is a total function $\mathcal{P} \rightarrow \{0, 1\}$ that assigns a truth value (0 or 1 or, equivalently, false or true) to any propositional letter. An interpretation w is said to be a *model* of a formula α (denoted by $w \models \alpha$) if and only if w makes α true in the usual truth functional way. The notion of model captures the *semantics* of the logical connectives. The syntactic counterpart $\Gamma \vdash \alpha$ means that α is *deducible* from Γ (where Γ may be empty).

Our everyday reasoning does not satisfy monotonicity. The most famous bird in computer science is Tweety: if I know that Tweety is a bird, I will assume that Tweety flies. However, upon learning that Tweety is a penguin, I will withdraw the previous conclusion. So, we have that $\Gamma \vdash \alpha$, but $\Gamma \cup \beta \not\vdash \alpha$, where β is the information that Tweety is a penguin.

As Robert Koons notes [176], defeasible reasoning has been object of philosophical investigations since Aristotle’s *Topics* and *Posterior Analytic*, but the subject received particular interest from researchers in artificial intelligence during the last forty years. The need to investigate and formalize nonmonotonicity emerged in the 1970s, when artificial intelligence researchers were facing knowledge representation problems [239, 200]. To John McCarthy, nonmonotonicity is what characterizes common sense reasoning. The 1980s saw a great development of formalisms to capture nonmonotonic reasoning: circumscription [200, 201, 189], default logics [240, 112], modal nonmonotonic logics [203, 204], an epistemic reformulation of McDermott’s logic, i.e. autoepistemic logic [210] which was in its turn modified and investigated by Halpern and Moses in [150], and extended logic programming [133, 134], a fragment of Reiter’s default logic. One may view all major formalisms for nonmonotonic reasoning as different approaches to the problem of identifying belief sets preferred for reasoning, once the world is assumed to be as normal as possible [67].

Yoav Shoham [261, 260, 262] observed that, in spite of the differences between the various formalisms for nonmonotonic reasoning, it is possible to provide a unifying semantical framework for nonmonotonic logics by generalizing the notion of minimal models introduced in circumscription.⁶ Shoham considers any standard logic \mathcal{L} , that is, any logic with the usual model-theoretic semantics, such as propositional logic, first-order predicate logic, and modal logic. A *preference logic* is obtained by associating \mathcal{L} with a strict partial preference order \succ on interpretations. If in classical logic the meaning of a formula is the set of models that satisfy it, Shoham shows that nonmonotonic logics are obtained by adding a preference ordering, which allows the logic to focus only on a subset of these interpretations. This captures the idea behind all nonmonotonic logics, that is to assume the world is as normal as possible and use this assumption to identify those models that are ‘preferable’ in a certain respect. So, when waiting for Björn, I’m justified to expect a blond, tall and blue-eyed man as the prototypical Swedish person is blond, tall and with blue eyes. An inference in the preferential framework can be seen as a selection of those conclusions that hold in all maximally preferred interpretations. Thus,

⁶ A similar but less general proposal was made by Bossu and Siegel [45].

from having that $\Delta \models \alpha$ if α is true in all models of Δ (as in classical logic), we now have that $\Delta \models_{\succ} \alpha$ if α is true in all preferred models of Δ . Also, from the previous it does not automatically follow that $\Delta \cup \beta \models \alpha$. This is so because the set of preferred models of $\Delta \cup \beta$ may not be a subset of the set of preferred models of Δ . Shoham needs to modify the usual notions of satisfaction and entailment to take the ordering on interpretations into account. In [260] he defines a preferred model and preferential entailment as follows:

Definition 1 (Preferred model) An interpretation m preferentially satisfies α (written $m \models_{\succ} \alpha$) iff m satisfies α and there is no other interpretation m' (with $m' \succ m$) such that m' satisfies α . Then, m is said to be a preferred model of α .

Definition 2 (Preferential entailment) α preferentially entails β (written $\alpha \models_{\succ} \beta$) iff the models of β are a superset of the preferred models of α .

Shoham claims that some nonmonotonic logics are special cases of his general framework while the connections to others is not clear. The system that is closest to Shoham’s framework is the family of logics based on circumscription. In general terms, circumscribing a predicate is reducing the individuals that satisfy that predicate in the theory to only those that are necessary in view of the theory. Predicate circumscription assumes that the objects that can be shown to have a certain property by reasoning from certain facts are all the objects that satisfy such given predicate, as in the famous missionaries and cannibals puzzle.⁷ Thus, circumscription means preferring those models of the theory that have minimal extensions of the predicates in question. Circumscription can then be captured by Shoham’s preferential model semantics. That is not surprising since, as Shoham himself acknowledges, “the notion of preferred models was implicit in McCarthy’s work from the start” [260, p. 237].

Unlike circumscription and the minimal knowledge logic of Halpern and Moses, the relation of Reiter’s default logic with Shoham’s framework is less clear. Default logics are consistency-based logics as they privilege the syntactic approach in the definition of nonmonotonic inference. A default theory is a

⁷ The puzzle asks how three missionaries and three cannibals can cross a river. The information given is that there is a rowboat for two people and that, if the cannibals outnumber the missionaries, the missionaries will be eaten. As McCarthy observed, if someone would suggest to use the bridge or the helicopter, we would be irritated. This is because the assumption is that only the elements explicitly provided in the puzzle are assumed to exist. If we want to be able to avoid the excessive qualification (for example, saying that the capacity of the boat does not change in the course of the action, that no bridge or helicopter exist, etc.), circumscription is a candidate to achieve this:

[Circumscription] will allow us to conjecture that no relevant objects exist in certain categories except those whose existence follows from the statement of the problem and common sense knowledge. When we circumscribe the first order logic statement of the problem together with the common sense facts about boats etc., we will be able to conclude that there is no bridge or helicopter. [200, p.30]

first-order classical logic theory to which nonmonotonic inference rules are added. These are *default rules* and have the form:

$$\alpha : \beta_1, \dots, \beta_n / \gamma \quad (2)$$

where $\alpha, \beta_1, \dots, \beta_n$ and γ are closed predicate logic formulae. The meaning of Equation 2 is the following: if α is known, and if it is consistent to assume β_1, \dots, β_n , then infer γ . The crucial notion of default logics is that of extensions (conclusion sets), which are obtained by applying as many default rules as possible without running into an inconsistency. A default theory can have one, several or no extensions. Shoham translates default theories into a modal one and adds a preference relation on the Kripke structures. He then considers examples where the default theory and its translation into a modal theory have a different number of extensions. Introducing a preference criterion does not lead to capture Reiter's definition of extensions. Shoham, then, concludes that more work needs to be done to cast light on that relationship. This was finally achieved by David Makinson [195], who showed that Reiter's default logic cannot be captured by the preferential entailment framework. The reason is that, unlike preferential logics, default logics are not cumulative.⁸

The variety of nonmonotonic formalisms raised the question of whether it was possible to provide a systematic approach that could classify, distinguish and clarify the relations between each formalism. Dov Gabbay was the first to suggest to study the different consequence relations defined by the different nonmonotonic systems [127]. The step undertaken was not an obvious one, as not all nonmonotonic formalisms assumed a consequence relation [155, 156]. In his seminal paper, Gabbay worked on Gentzen-style consequence relations to single out the minimal conditions a nonmonotonic consequence relation \vdash should satisfy in order to represent a nonmonotonic logic. In addition to *Monotonicity* (1), a classical consequence relation satisfies other two properties, *Reflexivity*:

$$\Gamma \cup \alpha \vdash \alpha \quad (3)$$

and *Cut*:

$$\text{If } \Gamma \vdash \alpha \text{ and } \Gamma \cup \alpha \vdash \beta, \text{ then } \Gamma \vdash \beta \quad (4)$$

Assuming that α and β are formulas, $\alpha \vdash \beta$ should be read as “ β is a plausible consequence of α ”. If we substitute \vdash to the classical consequence relation, we obtain the defeasible versions of Reflexivity ($\alpha \vdash \alpha$) and Cut, to which Gabbay added a weak version of Monotonicity (which obviously cannot hold in a nonmonotonic system), called *Cautious Monotony*:

⁸ For the same reason, Reiter's default logic does not satisfy Cautious Monotonicity (5).

$$\text{If } \Gamma \vdash \alpha \text{ and } \Gamma \vdash \beta, \text{ then } \Gamma \cup \alpha \vdash \beta \quad (5)$$

Cautious Monotony is the converse of Cut, and it is also called *Cumulative Monotony* because it says that it is safe to draw consequences and then use them as additional premises.

If Gabbay did not provide a semantics for his properties, five years later, Kraus, Lehmann and Magidor [179] developed Gabbay's and Shoham's works and characterized nonmonotonic consequence relations both proof-theoretically and semantically:

[N]one of the nonmonotonic systems defined so far in the literature [...] may represent all nonmonotonic inference systems that may be defined by preferential models. The framework of preferential models, therefore, has an expressive power that cannot be captured by negation as failure, circumscription, default logic or autoepistemic logic. [...] The main point of this work, therefore, is to characterize the consequence relations that can be defined by models similar to Shoham's in terms of proof-theoretic properties. To this end, Gabbay's conditions have to be augmented. [179, p. 168-169]

The weakest logical system introduced by Kraus, Lehmann and Magidor, *system C* (for *cumulative*), adds to the properties of Reflexivity, Cut and Cautious Monotonicity proposed by Gabbay, the inference rules of *Left Logical Equivalence*:

$$\text{If } \models \alpha \leftrightarrow \beta \text{ and } \alpha \vdash \gamma, \text{ then } \beta \vdash \gamma \quad (6)$$

and *Right Weakening*:

$$\text{If } \models \alpha \rightarrow \beta \text{ and } \gamma \vdash \alpha, \text{ then } \gamma \vdash \beta \quad (7)$$

Among the rules that can be derived in *C*, we should mention the *And* rule:

$$\text{If } \alpha \vdash \beta \text{ and } \alpha \vdash \gamma, \text{ then } \alpha \vdash \beta \wedge \gamma \quad (8)$$

This rule guarantees that plausible consequences can be accumulated via conjunction, and so can be seen as a principle of cumulativity that all basic nonmonotonic logics satisfy. Kraus, Lehmann and Magidor develop a semantic account for *C*, and provide a representation theorem.

Even though system *C* satisfies the minimal requirement for a nonmonotonic logic, Kraus, Lehmann and Magidor consider it to be too weak. A stronger and best fitted for nonmonotonic inference systems is *system P* (for *preferential*), which generalizes Shoham's preferential semantics and so it is of particular interest here. This system is equivalent to the one proposed by

Ernest Adams in the context of conditional logic [2] and to the ‘conservative core’ of Judea Pearl and Hector Geffner’ probabilistic system for default reasoning [222]. The system P consists of all the rules of C with the addition of the Or rule:

$$\text{If } \alpha \sim \gamma \text{ and } \beta \sim \gamma, \text{ then } \alpha \vee \beta \sim \gamma \quad (9)$$

An important derived rule of P is D , originally suggested by Makinson in a personal communication to the authors:

$$\text{If } \alpha \wedge \neg\beta \sim \gamma \text{ and } \alpha \wedge \beta \sim \gamma, \text{ then } \alpha \sim \gamma \quad (10)$$

The relevance of D resides in the fact that it allows the principle of reasoning by cases, which is a problem in nonmonotonic reasoning, as Pearl’s famous example [220] shows. Suppose that we know that male birds fly and, as a separate default rule, that female birds fly. The default theory in which, in addition to the previous two default rules, we know that Tweety is a bird (without having information about its gender) has as only extension the original information that Tweety is a bird. This is so because the prerequisites of the two rules cannot be verified. This means that the conclusion that Tweety flies is not obtained, even though it is an intuitively desirable conclusion.

The semantics of P is based on the notion of preferential model. Preferential models are cumulative ordered models in which the agent has a preference over worlds (instead of set of worlds, as in Shoham’s version). The preference relation \succ is a strict partial order (it was a well-order⁹ in Shoham’s framework), satisfying the *smoothness condition*, a technical condition ensuring the existence of a minimal element when we deal with infinite sets of formulas.

Cumulative reasoning systems provide expected results for the inheritance of defaults in taxonomies but cannot support transitive reasoning. Kraus, Lehmann and Magidor show that in the presence of the rules of C , transitivity and monotonicity are equivalent [179, Lemma 3.4]. Hence, frameworks like system P are too general as they do not allow chaining of defaults. In order to overcome such limitation, other approaches introduce preference statements between defaults in order to constrain the possible preference relations. Examples for such approaches are prioritized circumscription [189,201], pointwise circumscription [190,135], and prioritized default logic [46,64,22,90].

Approaches using preferences among default rules avoid also to be too weak in presence of too many conflicting defaults, as it is the case of systems based on preferential models. Preference information can be explicit (the information

⁹ A well-order relation on a set Q is a total order on Q with the property that every non-empty subset of Q has a minimal element. A total ordering of a finite set is trivially a well-ordering, while this is not the case for infinite sets. Intuitively, such ordering on possible worlds represents the plausibility order that an agent assigns to a world. An agent’s current beliefs are the minimal worlds at the bottom of the ranking. The higher a world is, the less plausible it is for the agent.

has to be specified by the knowledge engineer, as in [189,63]) or implicit. Approaches based on an implicit preference information often give higher priority to more specific rules [274,179,273,131]. But, as observed in [65,67], there may be other reasons than specificity to prefer one rule over another. In legal reasoning, for example, the criteria of recency or authority (for instance, federal law overrides state law) may also be used in case of conflict between laws.

Systems like Brewka's [65] include a flexible treatment of preferences over defaults (unlike in the explicit preference ordering approaches, where the preference is fixed) and, at the same time, by representing the priority information within the logical language, handle other criteria than specificity. Starting from [64] which added a partial order on defaults to Reiter's default logic [240], the idea is to extend the language to make reasoning about default priorities possible (via the naming of default rules), to generate default extensions and, finally, to keep only those extensions compatible with the priority information.

Example 1 [65] Suppose that we have the following default rules and facts:

$d_1 : bird \rightarrow flies$
 $d_2 : penguin \rightarrow \neg flies$
penguin
bird

Reiter's default theory would give two extensions, namely $E_1 = Th\{penguin, bird, flies\}$ and $E_2 = Th\{penguin, bird, \neg flies\}$.¹⁰ Suppose now that we also have the information that d_2 has priority over d_1 (denoted as $d_2 \prec d_1$). Clearly, only the second extension is compatible with this information. The system proposed in [65] has a mechanism that ensures that E_1 is eliminated.

Delgrande and Schaub [90,89] also consider a default theory where preference information is part of the theory and preferences among default rules can be part of a default rule. The novelty is that they show that such a generalised default theory can be translated into a standard default theory without preferences, but in which defaults are applied in the appropriate order. This means that all default logic theorem provers can be used for the prioritised version.

Frameworks for preferential reasoning have not been limited to the propositional context [179,187]. Among the proposals to capture defeasible reasoning in logics other than propositional one, we recall Lehmann and Magidor' extension to preferential predicate logics [186], defeasible deontic logics [214] (cf. Section 4), preferential extensions of description logics [70,138,73], and semantics for preferential reasoning in modal logics [71].

After this short *excursus*, we may ask ourselves whether a universal logic of nonmonotonic reasoning is possible. Doyle and Wellman [96] answered negatively to this question. Each formalism for nonmonotonic reasoning can be seen as a special theory of preferential or rational inference (*i.e.* how to select the maximally preferred states). If one could combine the different rational choices made by the several nonmonotonic formalisms into a single choice, then one

¹⁰ *Th* denotes that an extension is a theory (*i.e.*, it is a deductively closed set).

could claim that a universal logic of nonmonotonic reasoning exists. Yet, by adapting the framework for the aggregation of individual preferences of social choice theory, they show that a negative result similar to Arrow’s impossibility theorem [20] can be obtained for preferential nonmonotonic logics.

3 Preferences in argumentation theory

An argumentation system contains alternative arguments for or against some conclusions. Argument-based systems were mostly developed within artificial intelligence to study defeasible reasoning [230, 231, 192, 213, 233, 43].¹¹ The non-monotonicity lies in the fact that an argument may be defeated by another argument, which in turns may support the opposite claim. Formal argumentation can thus be seen as a generalized way of nonmonotonic reasoning [42], and indeed several nonmonotonic formalisms including Nute’s Defeasible Logic [144], Simari’s DeLP [129], logic programming Default Logic [107] have been shown to conform to the standard semantics of argumentation theory.

Abstract argumentation theory studies the positions that a rational agent can take in presence of a given set of arguments, where some arguments are in conflict with others. Arguments and the conflict relations are considered as generally as possible. Arguments are abstract entities, *i.e.* their internal structure is disregarded. Likewise, the conflict relation is left unspecified. If two arguments are in conflict, this roughly means that they cannot both hold. Abstracting away from the internal structure of the arguments and from the precise meaning of the conflict relation allows us to study how to reason in presence of a conflicting set of arguments in the most general way. Some approaches assume a particular logic [235] while others do not specify the underlying logic [43, 107].

Argumentation theory is central within artificial intelligence [30] as it provides a logic-based formalism for the treatment of defeasible reasoning and conflict resolution [263, 6, 171, 208, 36], negotiation [267, 180, 11], and argumentation-based dialogues [13, 234, 175].

An *argumentation framework* is simply a set of arguments and a binary relation among them. Dung (who can be considered to be the father of abstract argumentation) identified the binary relation with an *attack* relation [107]. Given an argumentation framework, the aim of argumentation theory is to identify and characterize the sets of arguments (*extensions*) that can reasonably survive the conflicts expressed in the framework. In general, given an argumentation framework, there are several possible extensions [107].

In order to simplify the discussion, we only consider finite argumentation frameworks.

Definition 3 (Argumentation framework) An *argumentation framework* AF is a tuple (Ar, \mathcal{R}) , where Ar is a set of arguments and \mathcal{R} is a binary

¹¹ However, John L. Pollock’s work, who developed Roderick Chisholm’s ideas [81, 82] into a theory of *prima facie reasons* and their *defeaters*, was rather motivated by epistemological questions in philosophy of science.




Fig. 1 An argumentation framework.

relation on Ar (i.e., $\mathcal{R} \subseteq Ar \times Ar$). An argument A *attacks* an argument B iff $(A, B) \in \mathcal{R}$.

An argumentation framework can be represented as a directed graph in which the arguments are represented as nodes and the attack relations as arrows. For instance, the argumentation framework (Ar, \mathcal{R}) where $Ar = \{A, B, C, D\}$ and $\mathcal{R} = \{(A, B), (B, A), (A, C), (B, C), (C, D)\}$ is represented in Figure 1. There we have that A and B attack each other, both A and B attack C , and C attacks D .

Different semantics have been proposed to define the acceptability of arguments in an argumentation framework. In Dung’s original extension approach [107], an extension is a subset of Ar that represents the set of arguments that can be accepted. Dung’s semantics are based on the notion of *conflict-freeness*, namely a set should not be self-contradictory nor include arguments that attack each other [107]. This ensures that no extension will support contradictory conclusions.

Definition 4 (Conflict-free / defence) Let (Ar, \mathcal{R}) be an argumentation framework. The set $S \subseteq Ar$ is *conflict-free* if and only if there are no $A, B \in S$ such that $(A, B) \in \mathcal{R}$.

We also say that S *defends* A (or, the argument A is *acceptable with respect to* S) if, $\forall B \in Ar$ such that $(B, A) \in \mathcal{R}$, $\exists C \in S$ such that $(C, B) \in \mathcal{R}$.

As we have seen, conflict-freeness is the minimal requirement for an extension. The most common acceptability semantics used in the literature are the following:

Definition 5 (Acceptability semantics) Let $AF := (Ar, \mathcal{R})$ be an argumentation framework and set $S \subseteq Ar$.

- S is an *admissible* extension if and only if it is conflict-free and defends all its elements.
- S is a *complete* extension if and only if it is conflict-free and contains precisely all the elements it defends, i.e., $S = \{A \mid S \text{ defends } A\}$.
- S is a *grounded* extension if and only if S is the smallest (w.r.t. set inclusion) complete extension of AF .
- S is a *preferred* extension if and only if S is maximal (w.r.t. set inclusion) among admissible extensions of AF .

- S is a *stable* extension if and only if S is conflict-free and $\forall B \notin S, \exists A \in S$ such that $(A, B) \in \mathcal{R}$.

It is known that for every argumentation framework, there exists at least one admissible set (the empty set), exactly one grounded extension, one or more complete extensions, one or more preferred extensions, and zero or more stable extensions.

The admissible extensions of the argumentation framework in Figure 1 are $\emptyset, \{A\}, \{B\}, \{A, D\}$ and $\{B, D\}$. The preferred and stable extensions are $\{A, D\}$ and $\{B, D\}$, the complete extensions are $\emptyset, \{A, D\}$ and $\{B, D\}$ and, finally, the grounded extension is \emptyset .

Though the generality of the framework we have briefly reviewed is part of its attractiveness, it has been argued that arguments do not have the same strengths [263, 32, 75]. Preferences can be added and taken into account in order to evaluate arguments [10, 235, 7, 12, 209]. Consider the following example due to Amgoud and Cayrol.

Example 2 [9] Let (Ar, \mathcal{R}) be an argumentation framework with $Ar = \{A, B, C\}$ and $\mathcal{R} = \{(A, B), (B, C)\}$. The set of acceptable arguments is $\{A, C\}$. However, suppose that argument B is preferred to A and to C . How can the preference over arguments and the attack relation be combined when deciding which arguments to accept? One natural way is to say that, since B is preferred to A , it can defend itself from the attack of A . This would lead to accept B and reject C .

Adding preference relations allows for more expressivity. Not only we can express that some arguments are in conflict, but also that some arguments are preferable to others, for example, because they express more probable beliefs or promote more important values [31]. Dung’s framework has then been extended by introducing preference relations into argumentation systems [10, 235, 7, 8, 209].

Simari and Loui [263] introduced preferences over arguments, and [32] considered arguments from prioritised beliefs in inconsistent knowledge bases. A natural domain of application of argument-based systems has been the modelling of legal disagreement [232, 257, 143]. Indeed, one of the first extensions of Dung’s argumentation framework was inspired by legal reasoning [235]. In order to increase the potential for implementation and following [106], arguments were expressed in a logic-programming language and their conflicts decided with the help of (defeasible) priorities over rules.

A more direct extension of Dung’s framework is the one proposed in [9], where a preference ordering enriches Dung’s framework. In the literature on preference-based argumentation, the attack relation in a preference-based argumentation framework is called *defeat*, and is denoted by Def .

Definition 6 (Preference-based argumentation framework (PAF)) A preference-based argumentation framework (PAF) is a triplet (Ar, Def, \succeq) , where Ar is a set of arguments, Def is the *defeat* binary relation on Ar , and \succeq is a (partial or total) preorder defined on $Ar \times Ar$.

Thus, $A \succeq B$ means that argument A is at least as preferred as B and the relation \succ is the strict counterpart of \succeq . As illustrated in Example 2, one idea to combine preference and attack relations is that if an argument B is preferred to its attacker A , then A 's attack against B is not successful and B is accepted. Informally, the idea in [9] is to remove those attacks that conflict with preferences and calculate Dung's classical semantics on the resulting argumentation framework. So, what is the relation between the two frameworks? Kaci and van der Torre [166] show that a preference-based argumentation framework can *represent* an argumentation framework:

Definition 7 (PAF representing an AF) A preference-based argumentation framework (Ar, Def, \succeq) represents an argumentation framework (Ar, \mathcal{R}) iff $\forall A, B \in Ar$, it is the case that $(A, B) \in \mathcal{R}$ iff $(A, B) \in Def$ and it is not the case that $B \succ A$.

It should be easy to see that each preference-based argumentation framework represents one argumentation framework, whereas each argumentation framework can be represented by various preference-based argumentation frameworks [166].

Inspired by Perelman's work on persuasion [225, 224], Bench-Capon [31] extends standard argumentation framework to take into account values promoted by arguments and defines *value-based argumentation frameworks (VAF)*. The idea is that in practical reasoning, two individuals may agree on the fact that an argument attacks another argument, but may disagree on whether that attack is successful because the two arguments promote different values and the two individuals disagree on the preference over those values. So, preferences over arguments are determined by the values those arguments support, like human life, world security and good world relations can be values promoted by some arguments in a political debate over whether invading Iraq or not [21]. Dung's framework is thus enriched by adding a non-empty set of values \mathcal{V} , a function *val* that assigns a value to each argument, and a partial order $>$ over values. The idea is that an argument A defeats (successfully attacks) an argument B iff $(A, B) \in \mathcal{R}$ and the value promoted by B is not more important than the value promoted by A (i.e. not $val(B) > val(A)$).

Kaci and van der Torre [166] extend Bench-Capon's value-based argumentation frameworks in two directions. They take into account the possibility that arguments support multiple values, and consider various types of preferences over values. In [31], if values v_1 is preferred to value v_2 , then each argument supporting v_1 is preferred to each argument supporting v_2 . Kaci and van der Torre claim that real-world situations are more complex and consider two additional preferences: that v_1 is preferred to value v_2 according to a first preference relation if and only if at least one argument supporting v_1 is preferred to each argument supporting v_2 , and that v_1 is preferred to value v_2 according to a second preference relation if and only if each argument supporting v_1 is preferred to at least one argument supporting v_2 . Similarly as for the relation between *PAF* and *AF*, Kaci and van der Torre show that a

VAF represents a *PAF* if and only if, for any two arguments A, B , it is the case that $A \succeq B$ if and only if $val(A) > val(B)$ or $val(A) = val(B)$.

Instead of having a pre-specified preference relation among arguments or values, one may consider the case in which arguments can express preferences between other arguments. This is the route explored by Modgil with the introduction of *extended argumentation framework (EAF)* [209]. Suppose, for example, that the BBC and the CNN disagree on today’s weather forecast in London and that argument C says that BBC are more trustworthy than CNN. Argument C is expressing a preference for the argument that today it will be dry in London since BBC said so (argument A) over argument B which claims the opposite as CNN forecasted a rainy day. Dung’s argumentation framework is then extended by adding a second attack relation \mathcal{D} to Ar and the standard binary attack relation \mathcal{R} . \mathcal{D} ranges from an argument to an element of \mathcal{R} : if A attacks (B, C) (denoted by $(A, (B, C))$), then A claims that C is preferred to B . In an extended argumentation framework, the success of an attack is relative to the set of arguments S one is currently considering. Thus, the notion of defeat is parametrised w.r.t. S : let $S \subseteq Ar$, A is said to *defeats_S* B ¹² iff $(A, B) \in \mathcal{R}$ and $\nexists C \in S$ such that $(C, (A, B)) \in \mathcal{D}$.

In [209] and [15] it was observed that ignoring those attacks where the attacked argument is stronger than the attacker does not always give intuitive results. It can happen that the resulting extension violates the basic condition imposed on acceptability semantics, namely the conflict-freeness of extensions. This is problematic as, in turns, it may lead to violate the rationality postulates put forward in [72].

Example 3 Suppose that an argumentation framework contains only two arguments A and B , and that A attacks B . Assume also that $B \succ A$. Preference-based argumentation frameworks run into troubles because, since B is preferred to A , the attack against B fails, i.e. no defeat relation holds between the two arguments. The problem is that, by removing an attack, we remove an important piece of information from the graph, namely that there is a conflict between two arguments. By doing so, the two arguments may end up in the same acceptable extension, violating the basic requirement of conflict-freeness that grounds the idea of Dung’s extensions as representing *coherent* positions.

Amgoud and Vesic [15, 16] propose a new preference-based argumentation framework that guarantees conflict-free extensions to deal with the above problem. However, according to [167, 164, 165], the source of the problem concerning removing attacks resides in a misunderstanding of Dung’s framework. The abstract nature of Dung’s framework imposes a careful instantiation of it and, in particular, a suitable choice of the appropriate defeat relation. A preference-based argumentation framework with a symmetric conflict relation guarantees to prevent the undesirable result [164]. The preference relation is then used to decide the direction of the defeat relation between two arguments.

¹² The notation *defeats_S* makes it explicit that defeat is parametrised w.r.t. S .

4 Deontic logics

The Swedish man we encountered in Section 2 is in fact a serial killer and my life is now in danger. I know that nobody should kill. However, in order to save my own life, I have to consider the possibility of killing Björn. If I have the choice between a bloody murder and a bloodless one, I should go for the second option. Philosophers, logicians and computer scientists have reasoned about such a situation, which is one of the many paradoxes (or puzzles) arising in deontic logic, a logic to reason about concepts like obligations and permissions. The gentle murderer paradox says that one should not kill but, if one does, he has the obligation to kill gently [124]. The paradox arises from the fact that in one of the most familiar systems of deontic logic, the so-called *Standard Deontic Logic* (SDL), from those premises one can derive that it is obligatory to kill *tout court*, a hardly defensible conclusion.

Deontic logic studies concepts that have a clear practical relevance for law, ethics, human and artificial institutions, security systems etc. But, as it happened for modal logic (which strongly influenced deontic logic [205]¹³), contributions have also been made on a more theoretical level. Even though the first natural applications were to formalise legal reasoning [265], deontic logic has been increasingly used in computer applications. As observed by Thorne McCarty [202], one of the primary applications of deontic logic is to detect a violation of an obligation and to trigger the appropriate sanction to such violation.

The use of deontic logic is especially useful as a knowledge representation language in all those situations in which a system designer wants to take into account the violation of some obligations and the appropriate action [161,162]. The study of logical systems of deontic logic, the formal analysis of normative systems, the formal representation of legal knowledge, the specification of aspects of norm-governed multi-agent systems and autonomous agents, as well as normative aspects of protocols for communication, negotiation and multi-agent decision making are all among the topics of DEON, a biennial international conference on deontic logic in computer science. Recently, an extension of multi-agent systems with concepts traditionally studied in deontic logic gave rise to a new area called *Normative Multi-agent Systems* [41] with the satellite NorMAS workshops¹⁴.

The beginning of deontic logic can be traced back to the Thirties, when the Danish philosopher Jørgensen discussed the logical character of imperatives [163]:

¹³ An exception was the work by Ernst Mally [196], an early pioneer of deontic logic and the first one to have used the term *Deontik*. His work was not influenced by modal logic but his impact on the discipline was undermined by technical problems [122].

¹⁴ According to the definition of the first workshop on normative multiagent systems in 2005, “Normative Multi-Agent Systems are multi-agent systems with normative systems in which agents can decide whether to follow the explicitly represented norms, and the normative systems specify how and in which extent the agents can modify the norms” [41].

[A]ccording to a generally accepted definition of logical inference only sentences which are capable of being true or false can function as premises or conclusions in an inference; nevertheless it seems evident that a conclusion in the imperative mood may be drawn from two premises one of which or both of which are in the imperative mood. [163, p. 290]

The point is that imperatives, legal statutes, moral standards etc. are usually not viewed as being true or false. Expressions like “Mark, leave the room!” command a specific behaviour and are not descriptive. Being nondescriptive, they cannot be termed true or false. Thus, they cannot be premise or conclusion of a logical inference. The Jørgensen’s dilemma expresses the fact that, though there certainly exists a logical study of normative concepts, it seems difficult to have a logic of normative concepts. The logic of imperatives is tightly connected to deontic logic [151], and some authors claim they are essentially the same discipline.

The discussion of whether norms have truth values continued, but the first formal system of deontic logic was given only in 1951 by the Finnish philosopher Georg Henrik von Wright [293,292]. It is consensus to fix in von Wright’s work the beginning of deontic logic. Subsequent deontic logic systems built on his work, though essentially any aspect of von Wright’s logic has been criticised and von Wright himself proposed several systems to overcome difficulties he encountered.

Many deontic logic systems have been proposed in the literature, like Standard Deontic Logic (SDL), von Wright’s Old System (OS) [292], Chellas’ Minimal Deontic Logic (MDL) [77], Hansson’s Preference-based Deontic Logic (PDL) [152] and variants of these logics. By taking obligation to be the primary concept and represent it by the operator O (so $O\alpha$ reads as “It is obligatory that α ”)¹⁵, van der Torre classified those logics on the basis of the following three properties [271]:

(Weakening) $O(\alpha \wedge \beta) \rightarrow O\alpha$
 (And) $(O\alpha \wedge O\beta) \rightarrow O(\alpha \wedge \beta)$
 (Violations) wff: $\alpha \wedge O\neg\alpha$

The first two properties should be clear. The third property gives the well-formed formulas (wff) to express violations in the language, by saying that it may happen that $\neg\alpha$ was obligatory and nevertheless we have that α is the case. So, if a certain logic (like OS) does not have Violations, it means that it lacks the possibility to express violations in its language. SDL satisfies all three properties. MDL satisfies Weakening and Violations, and PDL is the only system in which Weakening does not hold¹⁶, but it satisfies And and Violations.

¹⁵ Taking O as the primary concept means that other notions, like permissible, can be defined from O . So, for example, if we denote ‘permissible’ by P , we have that $P\alpha \leftrightarrow \neg O\neg\alpha$.

¹⁶ For explanations of why Weakening cannot hold in preference-based deontic logics, see [157, 139, 140, 152].

We give here the syntax and the semantic of the most cited of such systems, that is Standard Deontic Logic (SDL), a monadic deontic logic that builds upon propositional logic.

Definition 8 (SDL) Let \mathcal{L} be the language built upon a denumerable set \mathbf{P} of propositional variables, the usual connectives \neg and \rightarrow and the operator O . Axioms of SDL are the following:

- (*Taut*) All tautologies of \mathcal{L} .
- (*K*) $O(\alpha \rightarrow \beta) \rightarrow (O\alpha \rightarrow O\beta)$
- (*D*) $\neg O\perp$

SDL is closed under the following rules of inference:

- (Modus Ponens) If $\vdash \alpha$ and $\vdash \alpha \rightarrow \beta$, then $\vdash \beta$
- (Necessitation) If $\vdash \alpha$ then $\vdash O\alpha$

To the reader familiar with modal logic, it will be clear that SDL is just the normal modal logic KD with a reinterpretation of the \Box symbol as O for obligation. Since its birth deontic logic was seen as a branch of modal logic, thanks to the similarities between modal notions of necessity and possibility and deontic notions of obligations and permissions. Dissimilarities between the two fields were, however, noticed by von Wright [295].

The semantics of SDL is a possible worlds (Kripke) semantics.

Definition 9 (Kripke semantics) A possible world (Kripke) model for a deontic theory in SDL is a tuple $M = \langle W, R_O, V \rangle$ that consists of a nonempty set of worlds W , a binary serial deontic accessibility relation R_O between worlds¹⁷, and a valuation function V that assigns a truth value to atomic propositions in each world $w \in W$. Intuitively, $R_O(w, w')$ means that w' is an ideal deontic alternative to the actual world w and that in w' it holds everything obligatory in world w . A formula $O\alpha$ is true in w in M (denoted $M, w \models O\alpha$) iff $M, w' \models \alpha$ for all w' with $R_O(w, w')$.

A particular class of obligations that was highly relevant for the development of formal deontic logic systems are *contrary-to-duty* ones. A contrary-to-duty obligation expresses what one should do when obligations have been violated, like Gabriella's violation of the obligation to abstain from killing or, to take a less prosaic example, what Saint Paul said in the first letter to the Corinthians:

It is good for a man not to touch a woman. But if they cannot contain, let them marry: for it is better to marry than to burn. (Cited in [287].)

Many of the deontic logic paradoxes are related to contrary-to-duty paradoxes. Probably the most famous one is the gentle murderer paradox [124], which we informally encountered at the beginning of this section.

¹⁷ An accessibility relation is serial if, for every world w , there is at least one world accessible to w .

Example 4 (The gentle murderer paradox) Suppose that the following are wff's in a SDL theory:

1. Gabriella should not kill Björn: $O\neg k$
2. If Gabriella kills Björn, then she should do it gently: $k \rightarrow O(k \wedge g)$
3. Gabriella kills Björn: k

$k \rightarrow O(k \wedge g)$ is a contrary-to-duty obligation of $O\neg k$, as it says what one should do if one violates an obligation. The problem with the gentle murderer is that in a SDL theory with these sentences, one can derive Ok . By applying Modus Ponens to (2) and (3) we obtain $O(k \wedge g)$, thus Ok by Weakening. Our theory is thus inconsistent as it contains both Ok and $O\neg k$.

As we have seen in Section 2, default logic formalises the reasoning based on default assumptions. The world is assumed as normal as possible (Swedish men are assumed to be tall, blond and with blue eyes, birds are assumed to be able to fly etc.) unless evidence to the contrary. If this is the case, previously obtained conclusions may be not any longer derivable. The idea behind Shoham's proposal was to add preferences to nonmonotonic logics, where preferences represent different degrees of normality. In the most normal case, Swedish men are tall and blond.

Deontic logic can also be a preference-based logic. As preferences served to represent and treat exceptions in a preference-based nonmonotonic logic, so they can be used to represent and treat violations in a preference-based deontic logic. In other words, similar to degrees of normality in preference-based default logics, preferences can be seen as degrees of ideality in a logic for normative reasoning. Indeed, many preference-based deontic logics are default logics (see also Nute's collection on defeasible deontic logic [214]).

This was the way taken by Bengt Hansson to treat paradoxes in deontic logic. He realised that paradoxes arose because the semantics used in deontic logic was too rigid [151]. As seen in Definition 9, the truth condition for deontic statements is defined by considering only the actual world and the worlds that are accessible from the actual one (in which everything that is obligatory in the actual world holds):

In SDL, norms are assumed to refer exclusively to what obtains in the best possible alternatives. [...] In SDL, only what is compatible with the best is not wrong. [152, p. 83]

So, for example, the gentle murderer leads to a paradox because we obtain two conflicting obligations, Ok and $O\neg k$. No world can possibly satisfy both. The solution prospected by Hansson was to move from the "best possible alternatives" to a hierarchy of alternatives by defining a preference ordering \succeq over alternative worlds. Depending on the properties of \succeq , different logics can be obtained. Hansson, for example, at the beginning assumed only reflexivity though in general, the preference ordering over the worlds is any partial pre-order. The gentle murderer paradox can be solved by distinguishing accessible worlds in which $\neg k$ is true (and k is false) and worlds in which k is true (and

$\neg k$ is false) and say that the first kind of worlds are better than the second one.

A dyadic deontic logic (or logic of conditional obligation) is a logic in which an obligation is relative to some circumstances. An obligation $O(\alpha|\beta)$ means that “ α is obligatory, if β is the case”. Clearly, monadic obligations of type $O\alpha$ are a special kind of dyadic obligations where the antecedent is a tautology $O(\alpha|\top)$.

Hansson gave a semantics of dyadic obligations in which an “ideality ordering” \succeq over possible worlds is added [151]. When a preference relation over worlds is added, the intuition is that $O(\alpha|\beta)$ means that worlds in which $\alpha \wedge \beta$ is true are preferred to those in which $\neg\alpha \wedge \beta$ is true. Dyadic deontic logics with a preference ordering were also proposed by Danielsson [87], though it is Hansson’s work that is the most cited since it is the most easily accessible [1].

Properties like Weakening and And can be reformulated for dyadic logics:

$$\begin{aligned} \text{(Weakening)} \quad & O(\alpha_1 \wedge \alpha_2|\beta) \rightarrow O(\alpha_1|\beta) \\ \text{(And)} \quad & O(\alpha_1|\beta) \wedge O(\alpha_2|\beta) \rightarrow O(\alpha_1 \wedge \alpha_2|\beta) \end{aligned}$$

Deontic logics in a dyadic form were previously introduced by von Wright [292] and Rescher [241]. Among the several dyadic logics that have been proposed in the literature we recall those by Chellas [77], Alchourrón [3], Lewis [188] (an extension of Hansson’s [151]), Føllesdal and Hilpinen [122], and van Fraassen [286]. One difference is that, unlike [188,151,122,286], the logics of [77,3] satisfy the property called strengthening of the antecedent:

$$\text{(Strengthening of the Antecedent)} \quad O(\alpha|\beta_1) \rightarrow O(\alpha|\beta_1 \wedge \beta_2)$$

van der Torre [271] shows that logics that satisfy strengthening of the antecedent cannot formalise contrary-to-duty reasoning because strengthening of the antecedent is one of the key properties that lead to contrary-to-duty paradoxes. On the other hand, logics that do not have strengthening of the antecedent can formalise the contrary-to-duty paradoxes. However, Alchourrón [3], Castañeda [74], and Tan and van der Torre [268] criticise the lack of strengthening of the antecedent as - in the words of Castañeda - it is a “negative solution that looks like overkill” (though such criticism is not undisputed - see, for example, van Benthem et al [34]). A solution to this problem proposed by van der Torre is a two-phase deontic logic [271]. The idea is to have a logic that allows the combination of two desirable properties for a dyadic deontic logic, that is strengthening of the antecedent and weakening of the consequent, by forbidding the application of strengthening of the antecedent after weakening of the consequent, a sequence that leads to paradoxical results.

5 Compact representations of preferences

As already mentioned in the introduction, preferences are traditionally viewed as binary relations applied to a set A , such that $\succeq \subseteq A \times A$. If the size of

A is reasonable then an explicit representation of \succeq is feasible both from a cognitive (human agent) and a computational (artificial agent) point of view.

Consider now the case of choosing a digital camera out of an e-commerce site. Digital cameras are described by tens of features (each with different possible values: size of the memory, type of the lens, batteries, brand, price, etc.). If we denote by X_i the values of each possible attribute, the whole space of possible digital cameras is a large subset of $\prod_i X_i$. Considering only binary attributes it is easy to see that the size of A rapidly approaches 2^n , n being the number of attributes. There is no way to handle such a set for several different reasons:

- a human agent cannot compare the 2^n potential options in order to compile his preferences;
- there is no space in the memory of many artificial agents allowing to store such a set;
- supposing that we manage to get the whole set of comparisons and that we managed to store it somewhere, there is no way to compute something operational out of it, such as verifying if it is at least partially ordered and (if this is the case) identifying the maximal elements of A ($\text{Max}(A)=\{x : \nexists y \in A : y \succ x\}$).

Problems of the same type arise when the set A on which preferences are expected to apply results from other combinatorial manipulations such as being constructed as the power set of some set Ω of elementary actions ($A \subseteq 2^\Omega$). It is the case when a “portfolio” of actions (or candidates) needs to be chosen out of a list of alternatives.

An “obvious” solution to this problem could be to use some numerical representation for the different attributes and/or elementary actions and then aggregate them appropriately. In a typical case, each attribute is associated to a value function (u_i) and then these are summed to an overall value (U). However, there are several problems:

1. Suppose we have $\succeq_i \subseteq X_i \times X_i$. The conditions under which $\exists u_i : X_i \mapsto \mathbb{R}$ such that $x \succeq_i y \Leftrightarrow u_i(x) \geq u_i(y)$ are restrictive (\succeq_i needs to be a weak order) and do not hold a priori (see also [51]).
2. Suppose that the conditions for having such value functions are satisfied. The conditions under which $\exists U : A \mapsto \mathbb{R}$ such that $U = \sum_i u_i$ represents the preference \succeq over A are even more restrictive:
 - the u_i need to be more than ordinal measures (to be more precise: they need to be interval measures such that differences of preferences can be explicitly considered);
 - the u_i need to be commensurable among them (to be more precise: the differences of preferences along one attribute need to be comparable to the differences of preferences along any other attribute, such comparison establishing the trade-offs among the attributes);
 - the \succeq_i (and thus u_i) need to be preferentially independent among them (to be more precise: $x \succeq_i y$ needs to hold independently from how x

and y compare along all other attributes and this condition should hold among any subset of attributes).

Once again such conditions are not naturally met (see [51]).

Generally speaking we are looking how to compare vectors of the $\prod_i X_i$ space among them (and then possibly find an appropriate way to “measure” such vectors in some appropriate value scale). Under a conjoint measurement theory perspective, given two vectors (from the vector space $\prod_i X_i$) $x_1 = \langle x_1^1 \cdots x_1^n \rangle$ and $x_2 = \langle x_2^1 \cdots x_2^n \rangle$, the most general model representing a global preference should be

$$\langle x_1^1 \cdots x_1^n \rangle \succeq \langle x_2^1 \cdots x_2^n \rangle \text{ iff } \exists F : X_i^{2n} \mapsto \mathbb{R} \text{ such that } F(x_1^1 \cdots x_1^n, x_2^1 \cdots x_2^n) \geq 0$$

The function F is characterisable in many different ways: decomposable, transitive, skew symmetric, additive etc. (for more details see [178, 243, 53, 55, 54, 52]).

However, this is of little practical interest since it does not tell us how to handle the cases where the simple additive model does not hold (as in the case of preferential dependencies). The basic idea, proposed in the recent years has been to develop appropriate languages for compact representations (see [184]). Before presenting some of such languages we should mention that these are classified in the literature according to a number of criteria:

- expressiveness;
- concision;
- cognitive relevance;
- computational complexity.

For a discussion about these issues the reader can see [78, 199, 301].

Compact representations of preference models have been considered mainly in order to take into account preferential dependencies and conditioning which impede the use of simple additive conjoint measurement functions. The two main languages developed for this purpose are CP-nets (accounting for conditional preference statements) and GAI-networks (accounting for generalised additivity). We mention (without discussing it) the use of lexicographic preferences, a specific case of totally ordered attributes space studied in [116] and recently used for both aggregation and learning algorithms facing conditional preferences (see [302], [185]). We will also briefly discuss the approach consisting in directly modelling preference statements as logical sentences. We will not discuss the problem of choosing a “portfolio” of actions out of a set of alternatives, since there is no specific language developed for this purpose (although the reader can have a look at [183]). The literature rather focusses on how to “extend” preferences expressed upon a set Ω to the power set 2^Ω and how to handle efficiently the optimisation problems deriving from such an operation. The reader can see more details in [25, 212, 255, 283].

CP-nets CP-nets have been conceived in order to allow the representation (and efficient computation) of situations where somebody wants to claim that

$\langle xw \rangle \succeq \langle yw \rangle$, but $\langle yz \rangle \succeq \langle xz \rangle$ (the preference between x and y depends on what comes together to x and to y). A typical example concerns choosing between red and white wine: *I prefer red wine to white wine if we eat meat, but I prefer white wine to red wine if we eat fish*. We can distinguish between:

- unconditional “Ceteris Paribus” (CP) preferences along some of the attributes of X . For instance, if we say that x_1^3 (the first value of attribute 3) is preferred to x_2^3 (the second possible value for the same attribute) all other things being the same (*ceteris paribus*), expressed as $x_1^3 \succeq_{CP} x_2^3$, we can deduce (assuming that there are three attributes in our domain) that, for example $\langle x_1^1 x_1^2 x_1^3 \rangle \succeq \langle x_1^1 x_1^2 x_2^3 \rangle$, and in general that $\langle x_i^1 x_j^2 x_1^3 \rangle \succeq \langle x_i^1 x_j^2 x_2^3 \rangle$ with i, j being the indices associated to X_i and X_j .
- conditional preferences, where preferences along a certain attribute (let us say attribute k) are conditioned by preferences expressed on another attribute (let us say attribute l).

Technically speaking, a CP-net is represented through a directed graph among the variables (attributes), where the maximal elements of the graph are the ones where preferences are unconditioned. It is easy to see that when the graph is acyclic, computing within a CP-net is very efficient (while it is not the case otherwise). The reader can see more about CP-nets in [49, 59, 93, 299]. CP-nets have been extended to TCP-nets [58, 57] in order to take into account the “relative importance” that some attributes may have with respect to other ones and from that to more general CP-theories (see [300]), and to UCP-nets (see [48]) in order to consider conditional utilities. The case of multiple agents has been discussed in [248]. CP-nets have been applied in several configuration problems as well as in planning.

GAI-networks The intuition behind additive utility (such that $U(x) = \sum_j u_j(x^j)$) is that the contribution of each attribute to the overall utility is independent from all other attributes (and subsets of attributes). However, decomposable and additive conjoint measurement functions hold only under very strong conditions. Under less restrictive conditions [23, 115]) we can instead have utility functions based on a “generalised additive independence” such that:

$$U(x) = \sum_{i=1}^k u_i(x^{C_i})$$

where:

- C_i are subsets of the set of attributes X (and there are k of these factors);
- $D_{C_i} = \prod_{j \in C_i} X^j$;
- $\exists u_i : D_{C_i} \mapsto \mathbb{R}$.

The general idea within such functions is to measure both the utility contribution of each single attribute and the utility contribution of each subset of attributes (as if they were single ones). The reason for this is due to the positive or negative “synergies” two or more attributes may present: *for instance*

having white wine (first attribute) with fish (second attribute) has a positive synergy (which implies giving a positive reward in the function), while having red wine (first attribute) with fish (second attribute) has a negative synergy (which implies giving a penalisation in the function). Such generalised additivity includes many of the existing utility aggregation procedures (additive, k -additive etc.). However, what is important with such functions is the possibility to represent them in suitable graphical models, named GAI networks allowing for nice compact representations and efficient computation even for rather complex dependencies. The reader can see more details in [105, 141, 142].

Logical representations The basic idea here is to work directly on a set of sentences (which express preferences, values, desires and their opposite) instead of the binary relations. We can distinguish two approaches.

The first one consists in associating to each sentence $\varphi \in \mathcal{L}$ (where \mathcal{L} is a set of preferential sentences) a numerical value representing the contribution of φ (when logically satisfied) to the overall utility of the agent owing \mathcal{L} . Such values can also represent priorities in satisfying sentences or “distances” from some “target state of the agent” (see for instance [32, 66, 182, 221]). It is easy to show that in most cases such an approach boils down to use possibility theory (and possibility logic) as representation language (see for instance [33, 108]). A slightly different approach has been developed starting from a Constraint Satisfaction Programming perspective: the idea here is to see preference statements as compact constraints and use the power of constraint programming in order to solve preference aggregation and recommendation problems. The reader can see more in [246, 249, 250, 132]. For other approaches the reader can also see [68, 44].

The second approach consists in establishing a “logic of preferences”, an idea going back to the 60s [294, 296]. We have discussed several aspects of this approach in Section 4. We just recall here that the principle consists in elaborating a language allowing to express sentences of the type $\varphi \wedge \neg \psi \triangleright \neg \varphi \wedge \psi$: it is preferred a situation where φ holds and ψ does not hold to one where φ does not hold and ψ holds. For more details the reader can see [37, 168, 284].

6 Analytics of preferences: learning and eliciting

We now take the point of view of Björn, who has invited Gabriella for dinner. He needs to decide the menu (consisting of a first course, a second course, a beverage and a dessert) for the evening’s dinner based on his very limited information about her preferences. In addition, he might be able to use stereotype information (e.g. it is known that many researchers in logic-based artificial intelligence have a preference for spicy foods), or he may try to extrapolate preferences by observations of past behaviors (she has been noticed at the conference dinner choosing chicken and not beef). Finally, he might call her on the phone in order to ask her very specific questions about her preferences,

as for instance comparison queries (*would you prefer pasta or risotto as a first course?*). Queries of specific semantics need to be asked for attributes that are not preferentially independent; for example if the preference over beverages depends on the instantiation of the second course, one cannot simply ask to compare between red and white wine, but Björn may ask: *given that the second course is fish, do you prefer red or white wine ?* It will however be impossible for him to ask Gabriella to reveal her complete preference relation over the (combinatorial) combinations, so he will do better asking questions that are informative (for instance, if he knows for sure that she is allergic to nuts, he will not ask preference questions involving this item; and neither will ask about caviar, if this item is too expensive given the available budget).

The topic of preference learning or eliciting has recently raised substantial interest in the communities of operations research and artificial intelligence. Note that in the former, the term preference elicitation is more frequent, while in the latter (especially in the subfield of machine learning) the term preference learning is common (due to different emphasis on the process for the former and on the data for the latter). Indeed there is a convergence between these communities in addressing these problems; a stream on preference learning has been organized at the *European conference on Operations Research (EURO)* for a number of years; we also mention the recent initiative of establishing the workshop *From Multi-criteria Decision Aid to Preference Learning (DA2PL)*. In order to emphasize this convergence we propose to adopt the term of *preference analytics* as more general term.

Learning or eliciting preferences means to acquire preference information in either direct or indirect way, from preference statements, critiques to examples, observations of user's clicking behaviour, etc. The study of the assessment of the preferences of a decision maker goes back to several decades; particular emphasis has been given to the elicitation of utility functions for multi-attribute and multi-criteria settings [173]. Classic approaches for utility elicitation focus on high risk decision and aim at assessing the decision maker's utility very precisely. The decision maker is asked a number of questions in order to assess precisely the parameters of the utility function, with the exact questions to be asked depending on the adopted *protocol*.

Decision makers are asked questions that can be *local*, focusing on attributes in isolation, or *global*, aimed at comparing complete outcomes. In particular, *standard gamble queries (SGQ)* ask the following: *“Choose between option x_0 for sure or a lottery $\langle x^\top, l, x_\perp, 1-l \rangle$ ”* (where the best option x^\top is obtained with probability l , the worst option x_\perp is obtained with probability $1-l$). An answer to a SGQ gives a constraint on the utility of x_0 . Assume, without loss of generality that $u(x^\top)=1$ and $u(x_\perp)=0$; the expected utility of the lottery $\langle x^\top, l, x_\perp, 1-l \rangle$ is l . Consider, for instance, the standard elicitation procedure for additive models; the decision analyst would typically ask to consider an attribute (for example, dessert), ask for the best (say, the Italian dessert tiramisù) and worst value (apple pie). He will then ask *local standard gamble queries* for each remaining color to assess its local utility value (*value function*). The subsequent problem is that of refining the intervals on local

utility values, which can be done with the so-called *bound queries*. The next step is then of assessing the “scaling” factors that relate the weighting attributes to each other. In order to do this, a “reference” outcome is fixed and typical questions are based on the notion of *indifference swaps* (asking the decision maker to assess the required changes to make two alternatives equally preferred) in order to assess the relative importance of different features or criteria.

The classic approach to elicitation suffers from a number of drawbacks. Gamble queries (and similar questions) are difficult to respond. The precision attained with classical elicitation methods is often unnecessary and the cognitive cost might just not be worth the effort. While classic elicitation protocols are well-founded, and can lead to being able to rank alternatives from the best to the worst, their applicability has been questioned. Starting from a couple of decades ago, in the operations research community, several researchers [158, 84] started to deal with the problem of eliciting a utility function when only incomplete information is available. Instead of fully eliciting a utility function, indirect elicitation methods assess a utility function from assignment examples (for instance, examples are assigned to classes). More recently, researchers in artificial intelligence [47, 76, 298, 50, 288, 289] have developed elicitation techniques for utility elicitation, with the goal of mitigating the cognitive cost for the user. Indeed in AI, the aim is that of developing agents that act regionally on behalf of the user; eliciting the preferences of the user in an effective way is therefore crucial.

Preference disaggregation methods The UTA method [158] is an assessment procedure for a set of utility functions based on linear programming. Local utilities are piecewise linear; the interval is divided in subintervals and linear interpolation is used to approximate the utility contribution of a given feature. Pairwise preferences are expressed with linear inequalities; slack variables are added to allow inconsistencies in the preference information. Since several utility functions may in general be feasible, a typical approach is the minimization of the sum of such slack variables, yielding a utility function that fits as good as possible the available preference information. An alternative method is MACBETH¹⁸ [84], also based on a system of linear inequalities, that asks the user to give some reference levels for each feature/criterion and information about the difference of satisfaction between the values of a given feature. Utility elicitation methods goes beyond additive models; several researchers have considered models based on the Choquet integral; for instance the TOMASO decision support system [198] and the MYRIAD software tools [181]. A review of methods for *preference disaggregation* (the general term for this kind of assessment; this is in contrast to *preference aggregation* that combines preference information given a specified model) can be found in [159].

The idea that, when presented with limited information about the user preferences, there is not just one but many consistent utility functions, gives

¹⁸ Measuring Attractiveness by a Categorical Based Evaluation Technique.

rise to robustness concerns. Let us assume that at a given point in the interaction \mathcal{P} represents the set of preference statements available and $\mathcal{U}_{\mathcal{P}}$ the set of possible overall utility functions consistent with those. In robust ordinal regression [147], a choice x is necessarily preferred to y , written \succeq^N , if, for all feasible utility functions u , it holds that $u(y) \geq u(x)$ (with strict necessary preference $x \succ^N y$ if $u(y) > u(x)$ for all feasible utility functions u); x is possibly preferred to y , written \succeq^P , if there exists at least one utility function u such that $u(x) \geq u(y)$ with $u \in \mathcal{U}_{\mathcal{P}}$. The properties of the necessary preference \succeq^N and of possible preference relation \succeq^P are analyzed: it is easy to see that $\succeq^N \subseteq \succeq^P$ (if something is necessary, it is also possible); moreover it holds that if something is necessarily preferred to something else, the latter cannot be possibly preferred to the former¹⁹: if $x \succ^N y$ then it cannot hold $y \succ^P x$. The method UTA^{GMS} provides linear programming formulations in order to compute necessary and possible rankings; notice that when additional preference statements are added, \succeq^P might only decrease and \succeq^N increases. The method $UTADIS^{GMS}$ (UTilités Additives DIScriminantes) is the analogous for sorting problems [148] (alternatives need to be assigned to classes, that are ordered from the most to the worst preferred), given a set of examples of assignments and assuming an underlying additive model, the method associates to each alternative the set of necessary and possible assignments.

Analytic Hierarchy Process (AHP) A number of methods base the evaluation of alternatives under conflicting criteria on subjective judgments of the decision maker about pairwise comparisons of alternatives [191]. The most famous approach of this family is the *Analytic Hierarchy Process (AHP)* [253], a method widely used for decision aid. The decision maker provides pairwise information about the intensity of the different criteria that are mapped in a numerical scale, constructing an evaluation matrix. The principal eigenvector of the matrix is used as a weighting vector. The items are also evaluated in a pairwise fashion with respect to each of the criteria; then, item evaluations are multiplied by the weighting vector in order to assess the ranking. Despite its popularity, AHP has been criticized by several authors. In particular, [85] criticizes the method with respect to the semantics of the priority vector derived from the principal eigenvalue method, as the method derived violates a condition of order preservation between importance weights. More problematic is that AHP, in its classic formulation, may display the (usually undesired) phenomenon of *rank reversal*: adding, as additional choice, an alternative that is a copy of a dominated item, can impact the final ranks (obtained with AHP) of other alternatives, and even change the item that is ranked first (in other words aggregation with AHP violates the independence of irrelevant alternatives). Belton and Gear [28] give a numeric example for which AHP gives $B \succ A \succ C$ but, after adding D (a clone of B) AHP gives $A \succ B \sim D \succ C$. Dyer [110] shows how rank reversal of this kind might happen even when D is dominated by B (and not just a clone). These observations sparked a vigor-

¹⁹ There is an obvious similarity to approaches using modal logic, however no formal logical treatment is made in robust ordinal regression.

ous debate in the community, with several replies to Dyer's papers [153,254] and followups [109], with the opinions differing on the interpretation of rank reversal and its acceptability, and the applicability of variants of AHP that circumvents the problem.

Leaving aside technical problems, we note that the procedure underlying AHP is quite cognitive demanding: the construction of the comparison matrices involves asking the decision maker a number of questions that is quadratic in the number of constraints and features; therefore this method can only be applied to decision problems involving a small number of items and features. Indeed, consider the problem of choosing a camera out of n possible choices and m criteria (quality of lens, memory, software,...), the user will be requested to answer $m^2 + mn^2$ questions (for example, if there are 20 different models of cameras and 10 criteria, the user will be asked to answer 520 questions, something considered unacceptable in typical electronic commerce applications).

Reasoning about similar decision problems Departing from traditional axiomatic approaches in economics, [136] analyzes (from a formal point of view) the situation of a decision maker who makes use of previous choices in memory in order to estimate the utility of a choice in a new problem. This is similar to what happens in case-based reasoning (a subfield of artificial intelligence) where solutions to previous similar problems are adapted in order to solve the problem at hand. This idea has been also considered in preference-based systems, in particular in approaches to recommendation based on case-based reasoning [206,266].

Adaptive Utility Elicitation Ideally, a system for *automated elicitation and recommendation* will only consider cognitive plausible forms of interactions, focusing on the available alternatives of the current decision problem. A number of researchers [47,76,298,50,288] have proposed the idea of an interactive utility-based recommender system. It is assumed that the user has a latent utility function that dictates his preferences; the system maintains a "belief" (whose nature will be clearer in a moment) about such utility function u . The general schema is as follows (bearing some similarity to *active learning*):

1. Some initial user preferences \mathcal{P}_0 are given; initialize belief
2. **Repeat** until the *belief* meets some termination condition
 - (a) Ask user a query q
 - (b) Observe the user response r
 - (c) Update the *belief* given r
3. Recommend the item optimal according to the current *belief*

A number of alternative proposals have been made with respect to 1) how preference uncertainty is represented in a belief, 2) which criterion is used to make a recommendation, and 3) how to select the question that is asked next. In the following Table, we outline some possibilities.

	minimax-regret approach	maximin-utility approach	Bayesian approach
knowledge representation	constraints	constraints	prob. distribution
which option to recommend?	minimax regret	maximin utility	expected utility
which query to ask next?	worst-case regret reduction	worst-case maximin improvement	expected value of information

A possibility for representing the current belief about the utility is to encode user responses with constraints (as in UTA) and reason about all possible consistent utility functions (as in robust ordinal regression) making use of a robust decision criterion to select the item to recommend. While maximin is a possibility [290], Boutilier et al. [50] suggest to adopt minimax regret, that is a less conservative robust criterion for decision making under uncertainty [258, 177]. The intuition behind the approach of *minimax regret* is that of an adversarial game; the recommender selects the item reducing the “regret” with respect to the “best” item when the unknown parameters are chosen by the adversary. As before, \mathcal{P} represents the set of preference statements available and $\mathcal{U}_{\mathcal{P}}$ the set of possible overall utility functions consistent with those. The *max regret* of an option x is the maximum difference between the utility of the best item and the utility of x when an adversary is choosing the utility function $u \in \mathcal{U}_{\mathcal{P}}$:

$$\text{MR}(x; \mathcal{U}_{\mathcal{P}}) = \max_{y \in X} \max_{u \in \mathcal{U}_{\mathcal{P}}} u(y) - u(x) \quad (11)$$

with X being the set of possible choices. The *minimax regret* $\text{MMR}(W)$ of $\mathcal{U}_{\mathcal{P}}$ and the *minimax optimal item* $x_{\mathcal{P}}^*$ are then found by finding the item associated with the smallest maximum regret: $\text{MMR}(\mathcal{U}_{\mathcal{P}}) = \min_{x \in X} \text{MR}(x, \mathcal{U}_{\mathcal{P}})$ and $x_{\mathcal{P}}^* = \arg \min_{x \in X} \text{MR}(x, \mathcal{U}_{\mathcal{P}})$. The advantages of regret-based approach are threefold: 1) it is easy to update our knowledge about the user: whenever a query is answered, we treat the answer r as a new preference and derive a new set $\mathcal{U}_{\mathcal{P} \cup r}$ ²⁰ 2) simple “priors” can be encoded with constraints on $\mathcal{U}_{\mathcal{P}}$ and 3) there are efficient heuristics that directly use the computation of minimax regret to choose the queries to ask next to the user: the *current solution strategy* [298, 50] asks the user to compare $x_{\mathcal{P}}^*$ and its adversarial choice $y_{\mathcal{P}}^*$ associated with maximum regret. The approach comes with limitations too, as it cannot deal with noisy responses and the formulation of the optimization depends on the assumption about the utility.

A principled idea is that of asking queries with high (a posteriori) improvement in decision quality. Assuming a query q that can have \mathcal{R} responses, for each response $r \in \mathcal{R}$ we can consider the updated utility space $\mathcal{U}_{\mathcal{P} \cup r}$ assuming

²⁰ For instance $\mathcal{U}_{\mathcal{P}}$ is a polytope in the case of a linear utility model defined by a vector of weights w ; $\mathcal{U}_{\mathcal{P} \cup r}$ will be then a subset of $\mathcal{U}_{\mathcal{P}}$, consisting of all instantiations of w that satisfies r in addition to \mathcal{P} .

that the answer to q is r ; the associated minimax regret is $\text{MMR}(\mathcal{U}_{\mathcal{P} \cup r})$ (that can only be lower or equal to $\text{MMR}(\mathcal{U}_{\mathcal{P}})$). A good query will significantly reduce minimax regret in each of the scenarios. A (non probabilistic) notion of myopic value of information [288] can be defined considering the *worst-case regret reduction* of query q when the utility function must lie in $\mathcal{U}_{\mathcal{P}}$:

$$\text{RR}(q; \mathcal{U}_{\mathcal{P}}) = \min_{r \in \mathcal{R}} [\text{MMR}(\mathcal{U}_{\mathcal{P}}) - \text{MMR}(\mathcal{U}_{\mathcal{P} \cup r})] = \text{MMR}(\mathcal{U}_{\mathcal{P}}) - \max_{r \in \mathcal{R}} \text{MMR}(\mathcal{U}_{\mathcal{P} \cup r}). \quad (12)$$

For example, consider the question $q = \text{“Do you prefer pizza or pasta?”}$. The possible responses are $\mathcal{R} = \{\text{pizza} \succ \text{pasta}, \text{pasta} \succ \text{pizza}, \text{pasta} \sim \text{pizza}\}$. In order to compute the regret reduction $\text{RR}(q; \mathcal{U}_{\mathcal{P}})$ (with respect to a set of already stated preferences \mathcal{P}) we need to consider $\mathcal{U}_{\mathcal{P} \cup \{\text{pizza} \succ \text{pasta}\}}$ (the set of utility functions satisfying $\text{pizza} \succ \text{pasta}$ in addition to \mathcal{P}), $\mathcal{U}_{\mathcal{P} \cup \{\text{pizza} \prec \text{pasta}\}}$ and $\mathcal{U}_{\mathcal{P} \cup \{\text{pizza} \sim \text{pasta}\}}$. In each of these utility spaces we compute minimax regret, consider the response associated with the highest value (in order to be robust with respect to the less favorable response), and subtract the current minimax regret value.

We can now use RR to numerically compare queries, as for example establishing that asking to compare pizza with pasta is less informative than asking to compare the shepherd’s pie with risotto. In particular asking to compare *pasta alla carbonara* with *pasta in bianco* (pasta with butter) when the former is known to be necessarily preferred to the latter (according to the \succ^N introduced before) gives no regret reduction at all; nor there is value in repeating questions in this model. The “best” query according to this measure is then the query with maximal regret reduction; $q^* = \arg \max_q \text{RR}(q; \mathcal{U}_{\mathcal{P}})$. Notice that a query can have significantly different RR values depending on the $\mathcal{U}_{\mathcal{P}}$ (the currently known preferences may have a strong impact on the value of the query). The straightforward approach for query selection would be to consider all candidate queries, evaluate RR and pick the one with highest value; however this is impractical for large outcomes spaces. Practical methods for framing query selection as an optimization problems for comparison queries (and choice queries, that extend comparison queries to a set of elements) are thoroughly discussed in [288].

Alternatively, one could assume a Bayesian standpoint: this has the advantage of handling noisy information, can exploit prior information (if available) and can be used with different assumption about the *choice model* of the user. We assume distributional information about the parameters involved in the user’s utility function; the belief $\theta(w)$ is a probability distribution over the parameters of the utility function that encodes the knowledge of the system about the user’s preferences; expected utility of a given item x is given by $EU_{\theta}(x) = \int u(x; w) \theta(w) dw$. The recommendation x_{θ}^* is the one associated with maximal expected utility under the current probabilistic belief: $EU_{\theta}^* = \max_{x \in A} EU_{\theta}(x)$; and $x_{\theta}^* = \arg \max_{x \in A} EU_{\theta}(x)$. When a new preference is acquired (for instance, the user states that he prefers apples over orange), the distribution is updated according to Bayes, using Monte Carlo methods, or in-

ference scheme based on expectation-propagation [207], in particular Trueskill [154] can be adapted to preference elicitation [149].

The problem of deciding which questions to ask could be formulated as a *Partially Observable Markov Decision Process (POMDP)* [47], however it is impractical to solve for non trivial cases. A more tractable approach is to consider (myopic) *Expected Value Of Information (EVOI)*, the difference between the expected posterior utility (of the best recommendation in the updated belief) $\text{EPU}_\theta^*(q)$ associated to a query q and the current EU_θ^* :

$$\text{EVOI}_\theta(q) = \text{EPU}_\theta^*(q) - \text{EU}_\theta^* = \sum_{r \in \mathcal{R}} P_\theta(r) \text{EU}_{\theta|r}^* - \text{EU}_\theta^* \quad (13)$$

where \mathcal{R} is the set of possible responses (answers), θ is the current belief distribution, $\theta|r$ the posterior and $P_\theta(r)$ the probability of a given response according to θ , whose value depends on the assumptions made about the *choice model*. Returning to our example about the dinner, a possible query could be to choose the preferred dish among pasta, risotto and shepherd’s pie (these three dishes constitute \mathcal{R} in this example); the posterior distributions would be $\theta|risotto$, $\theta|pasta$ and $\theta|shepherd$ respectively (the distribution of the parameters conditioned to the selected dish *being the preferred one among those mentioned*); $P_\theta(risotto)$ would weigh numerically (let’s say, for instance, 75% of probability) the scenario of preferring risotto to both pasta and to the shepherd’s pie; and similarly we can assess $P_\theta(pasta)$ and $P_\theta(shepherd)$. In order to ask the most informative question, we then ask the query $q^* = \arg \max \text{EVOI}_\theta(q)$ with highest EVOI or, equivalently, the query with highest EPU (since the current EU^* value can be considered as a constant when choosing the query maximizing EVOI). For choice queries (“*Among the following options, which one do you prefer?*”), Viappiani and Boutilier [289] showed that the problem of finding the optimal query is tightly connected to the problem of finding an optimal recommendation set (the generation of shortlisted alternatives from which the user makes a selection [236], as in the display of search engine results) and near-optimal queries can be computed efficiently with worst-case guarantees.

Learning preferences from data A number of methods in the machine learning community have been developed in order to assess a set of parameters consistent with what is known about the user; these include methods based on support vector machines, such as SVM-rank [160]. The common ground of these approaches is that they fit an assumed model using the available data, and use the learned model to make predictions. For the problem of learning a utility function, this can mean that each user’s preference is viewed as a constraint (an hyperplane in case of linear models) on the parameter space, and *max-margin* learners aim at identifying the set of parameters that maximize the minimum distance from the nearest hyperplane. These methods are “pointwise” in the sense that a single best guess of the user’s utility function is provided as output. While these methods work well for tasks such as prediction (as this is the setting they were designed to), they are not readily apt for interactive systems, when one needs to assess which question the system

should ask next (the focus of machine learning is most often of learning from available data).

Some approaches do not make the assumption that a latent numeric utility function exists. These include models from ranking, such as *Mallows* models [86]; for an approach for learning Mallows models see for instance Lu and Boutilier [193]. At the extreme *model-free* approaches do not make any specific model assumption at all, but rely on local estimation techniques. Common approaches are based on clustering and on using the information about the nearest neighbors in order to make estimations about the preferences of the user (implicitly there is an assumption of regularity, meaning that preferences of neighbors are similar). These approaches are popular for *label ranking* (a type of preference problem where the input data is based on preference rankings that are assigned to a specific label), in particular for the problem of predicting preferences based on demographic information, where the notion of similarity between user is defined naturally. A case-based nearest neighbors approach is proposed by Brinker and Hüllermeier [69], while Yu, Wan and Le [303] propose a method based on decision trees.

If one knows specific information about the underlying preference model, it is advantageous to exploit this information. A number of works focused on algorithms for learning preferences under specific model assumptions. In particular efforts have been made to learn lexicographic models [302] (the lexicographic assumption greatly simplifies the learning task, as only very few rankings are consistent with a lexicographic model) and preferences over sets [297]. As mentioned before in Section 5, CP-nets are a compact language representation for preferences in multi-attribute domains. If we know that a user has preferences that are consistent with a CP net, what is the best way to learn them? Chevaleyre et al. [80] address this problem from a theoretical point of view, providing learnability results. If we want to reason about possible CP networks consistent with current information, we may use *Probabilistic CP networks (PCP-nets)* [38], a compact representation of a probabilistic distribution over CP networks; PCP-nets can be used for learning by conditioning on available information. Finally, [44] describes how to learn conditionally lexicographic preference relations.

7 Non classical preference models

Until now we have always considered preferences as model of “certain information”. Indeed they have been considered as binary relations and the language used in order to formalise any theory about them has been (obviously) classic logic; the reader can check this in all basic texts about preference modelling: [115, 119, 251].

However, it is reasonable that if asked about a preference between any two x and y one could reply “*I do not know*” or more generally hesitate replying partially and/or ambiguously. The problem of representing values and uncertainty is not really new: Ramsey [238] and De Finetti [88] have addressed

the issue already in the 30s and it has been formalised in decision analysis both in prescriptive terms [291] and normative ones [258]. The problem with these approaches is that they are limited by the way uncertainty and hesitation are modelled: practically only probability (although subjective) is considered along with economically rational preferences ([115]). Such limitations gave rise to alternative approaches either within the decision theory community [118, 136, 174, 256] or within the artificial intelligence community (see for instance [99]). In the following we are going to survey the results of the latter approach.

7.1 Fuzzy sets

The basic idea here is the one to create and/or use languages specifically tailored to uncertainty modelling purposes, and more specifically the language of fuzzy sets. The first attempts to use fuzzy sets for preference modelling purposes date back to the late 70s (see [215], [216], [252]). The major challenge was to translate the theoretical structures already used in order to characterise and work with preferences into the new language: what is a fuzzy partition? How to define a fuzzy transitive binary relation? What should be a fuzzy preference structure? There is a wide literature in this area (partially surveyed in [211]) for which we cite some classic references: [100], [121], [145], [169], [227].

If, on the one hand, fuzzy (or *valued* as they are often are called) preference relations allowed to introduce a more nuanced and realistic preference modelling language, on the other hand they opened a certain number of problems.

- Preference structures and representation theorems require the introduction of complex logical sentences. For instance the definition of transitivity ($\forall x, y, z \in A \ x \succeq y \wedge y \succeq z \rightarrow x \succeq z$) requires to translate the universal quantifier as well as the connectives “and” and “imply” under form of appropriate functions. This problem has been addressed through the use of T-norms and conorms [101, 259] possibly satisfying the De Morgan principle (such that $T(x, y) = N(S(N(x), N(y)))$), where T , S and N are the functions representing T-norms, T-conorms and negations respectively). It has been soon proved that the system of functional equations which results when usual preference structures need to be characterised does not admit a unique solution (see [5, 121]). This is not surprising knowing the truth functionality problems of fuzzy reasoning and introduces a degree of freedom which needs to be managed during the modelling process.
- Most of the times valued preference relations are practically sentences where preferences are associated to some “measure of uncertainty” (it is typically the case of expected utility). Under such a perspective when several different valued preferences need to be aggregated what practically we get is a problem of aggregating the associated measures (as in most ordered statistics problems), which in this case are supposed to be fuzzy values. The problem has been addressed extending well known aggregation procedures through the introduction of the Choquet and Sugeno integrals which are the more general ordered statistics we can conceive (see [146]).

The reader should note that this approach allows to include probability measures above preference statements within the same theoretical framework. Actually the concept of fuzzy measure is nothing more than the one of “capacity”: a function f over the power set of a set Ω ($f : 2^\Omega \mapsto [0, 1]$) such that $f(\emptyset) = 0$ and $A \subseteq B \subseteq \Omega \rightarrow f(A) \leq f(B)$. Probabilities are just additive capacities. All that said, such tools implicitly introduce a commensurability hypothesis among these different measures, which is far from being true in practice.

- A specific way to address the problem of measuring the uncertainty associated to preference statements has been the use of possibility distributions (see [102]): these replace the additive property characterising probabilities (seen as capacities) with a pure ordinal sum ($\pi(A \cup B) = \max(\pi(A), \pi(B))$), π being a possibility distribution). If we now consider purely ordinal preference statements and purely ordinal likelihoods (such as described by possibilities) we get a possibilistic version of “Qualitative Decision Theory” an attempt to establish an ordinal version of classic Decision Theory (see [103], [98]). The problem here is that the resulting decision rules are either overconfident or not decisive and thus operationally of little interest (see [99], [114]).

7.2 Beyond fuzzy sets

As already mentioned it might be the case that it is not possible to establish precisely whether a certain relation holds or not. However, hesitation can be due either to incomplete information (missing values, unknown replies, unwillingness to reply etc.) or to contradictory information (conflicting evaluation dimensions, conflicting reasons for and against the relation, inconsistent replies etc.). More generally speaking while we try to assess the belief of a sentence or the value of something we may face the, rather common, situation where both positive information (reasons, values) and negative information (reasons, values) are available. Typical cases include positive and negative witnesses, majorities for and vetoes against (a preference or a statement), arguments for and against, gains and losses etc. Such situations have been considered in argumentation theory ([272]), value theory ([242]), cognitive studies about decision under risk and uncertainty ([170], as well as in philosophy and formal logic (see [97]), [26] and [27]). The common idea behind these approaches is that the negative information (reasons, values) is not just the complement of the positive one, but needs to be considered explicitly and formalised appropriately. In formal logic this idea has been further developing multi-valued logics and more precisely four-valued logics (see in [17, 18, 35, 113, 123, 120, 137, 172, 270, 277]).

In the case of preference modelling, the use of such logics was first suggested in [276] and [92]. Such logics extend the semantics of classical logic through two hypotheses:

- the complement of a first order formula does not necessarily coincide with its negation;
- truth values are only partially ordered (in a bilattice), thus allowing the definition of a boolean algebra on the set of truth values.

The result is that using such logics, it is possible to formally characterise different states of hesitation when preferences are modelled (see [279], [280]). Furthermore, using such a formalism, it becomes possible to generalise the concordance/discordance principle (used in several decision aiding methods) as shown in [278] and several characterisation problems can be solved (see for instance [281]). More recently (see [228, 125, 19, 91, 217, 218, 282]) it has been suggested to use the extension of such logics for continuous valuations.

Among others, this research allowed to show that such continuous valuations correspond to the logical counterpart of the concept of bi-capacity introduced as measure of “bipolar” preference measurement. The issue of bipolarity returned of interest in the recent years (see [104]) through different contributions where the presence of clearly distinct positive and negative reasons are considered in representing preferences and supporting decisions (see [14, 39, 40]).

8 Conclusions

Preferences are a key element of decision making and a basic concept for several research fields such as economics, decision theory, game theory, artificial intelligence, classification, databases, etc. This article presented the state of the art of preferences in artificial intelligence; it introduced techniques for reasoning about, arguing, representing and learning preferences. Note that, due to the width of the topic, we did not cover some important domains where handling preferences have also been considered as planning [24], personalized user interfaces [128], and the recently very active field of preference-based reinforcement learning [126].

Preferences play a role in almost all sub-areas of artificial intelligence, as witnessed by the diversity of the formalisms employed in the different sections of this article. We prospect that in the future preferences will play an even increasing role (some prominent research directions are mentioned in the different sections of this article). Indeed, artificial intelligence aims at producing computational artifacts that can help humans in a number of problems acting on their behalf; reasoning, explaining, learning and in general handling preferences are central issues to be tackled in any non trivial artificial intelligence system. There are several applications [223] of this research area, some already deployed in practice, including personalized and location-aware recommendation systems [62] and interactive personalized configuration systems [264].

We envision that the research issues covered by this survey will be more and more interconnected: for instance, we can foresee the development of preference elicitation strategies for non classical preference models (following the recent

works of [29], where adaptive elicitation techniques are proposed for assessing preferences dictated by a Choquet integral), or the development of richer languages for representation of complex preferences, that can then be used for argumentation. A critical point is that of providing explanations (generated automatically) to the user on how his preferences are treated (aggregated, assessed,...) and the reasons behind a particular action taken by the system. This issue is crucial in the research field of recommender systems [130]; a recommender system should be able to explain why a particular product is suggested.

Finally, preferences constitute the central element of negotiations and social choice problems, that arise from the fact that different agents (or organizations) have conflicting objectives, expressed in terms of preferences. Voting systems are ways to aggregate the preferences of different users (or agents) in order to make a collective choice (as in elections); however for novel application domains, such as internet-based decision support tools for groups of users, new frameworks need to be developed (for instance, voting systems with incomplete preference profiles and incremental elicitation of votes, following [194]; alternatively autonomous agents might engage in online argumentations in order to choose the best candidate for a job). Computational social choice is the field that studies algorithmic methods to reason about collective choices, as such preferences are a central element.

We expect that our general survey about preferences in artificial intelligence can be of interest to the researchers of various sub-disciplines of artificial intelligence, contributing in the widespread adoption of preference handling methods and fostering the development of new research directions at the intersection of different fields.

Acknowledgments

We thank the three anonymous referees for their valuable comments and suggestions that helped us improving the content and readability of the paper.

Gabriella Pigozzi benefited from the support of the AMANDE project of the French National Research Agency (ANR-13-BS02-0004); Paolo Viappiani is supported by the ELICIT project funded by the French National Research Agency through the IDEX Sorbonne Universités under grant ANR-11-IDEX-0004-02.

References

1. Åqvist, L.: Deontic logic. In: D. Gabbay, F. Guenther (eds.) *Handbook of Philosophical Logic*, pp. 147–264. Kluwer Academic, Dordrecht (1984)
2. Adams, E.: *The Logic of Conditionals*. Reidel, Dordrecht (1975)
3. Alchourrón, C.: Philosophical foundations of deontic logic and the logic of defeasible conditionals. In: J.J. Meyer, R. Wieringa (eds.) *Deontic Logic in Computer Science: Normative System Specification*, pp. 43–84. John Wiley, New York (1993)

4. Aleskerov, F., Bouyssou, D., Monjardet, B.: Utility maximization, choice and preference. Springer-Verlag, Berlin (2007). 2nd edition
5. Alsina, C.: On a family of connectives for fuzzy sets. *Fuzzy sets and Systems* **16**, 231–235 (1985)
6. Amgoud, L.: A formal framework for handling conflicting desires. In: Proceedings of ECSQARU'03, vol. 2711, pp. 552–563 (2003)
7. Amgoud, L., Cayrol, C.: Integrating preference orderings into argument-based reasoning. In: Proceedings of ECSQARU'97, pp. 159 – 170 (1997)
8. Amgoud, L., Cayrol, C.: On the acceptability of arguments in preference-based argumentation framework. In: Proceedings of UAI'98, pp. 1 – 7 (1998)
9. Amgoud, L., Cayrol, C.: A reasoning model based on the production of acceptable arguments. *Annals of Mathematics and Artificial Intelligence* **34**, 197–216 (2002)
10. Amgoud, L., Cayrol, C., LeBerre, D.: Comparing arguments using preference orderings for argument-based reasoning. In: Proceedings of the 8th International Conference on Tools with Artificial Intelligence, pp. 400 – 403 (1996)
11. Amgoud, L., Dimopoulos, Y., Moraitis, P.: A unified and general framework for argumentation-based negotiation. In: Proceedings of AAMAS'07, pp. 158:1–158:8 (2007)
12. Amgoud, L., Dimopoulos, Y., Moraitis, P.: Making decisions through preference-based argumentation. In: Proceedings of KR'08, pp. 113 – 123 (2008)
13. Amgoud, L., Maudet, N., Parsons, S.: Modeling dialogues using argumentation. In: Proceedings of ICMAS'00, pp. 31–38 (2000)
14. Amgoud, L., Prade, H.: Using arguments for making and explaining decisions. *Artificial Intelligence* **173**, 413–436 (2009)
15. Amgoud, L., Vesic, S.: Generalizing stable semantics by preferences. In: COMMA, pp. 39–50 (2010)
16. Amgoud, L., Vesic, S.: A new approach for preference-based argumentation frameworks. *Annals of Mathematics and Artificial Intelligence* **63**, 149–183 (2011)
17. Arieli, O., Avron, A.: The value of the four values. *Artificial Intelligence* **102**, 97–141 (1998)
18. Arieli, O., Avron, A., Zamansky, A.: Ideal paraconsistent logics. *Studia Logica* **99**, 31–60 (2011)
19. Arieli, O., Cornelis, C., Deschrijver, G.: Preference modeling by rectangular bilattices. In: Proceedings of MDAI 2006, LNAI 3885, pp. 22–33. Springer Verlag, Berlin (2006)
20. Arrow, K.: Social choice and individual values. J. Wiley, New York (1951). 2nd edition, 1963
21. Atkinson, K., Bench-Capon, T., McBurney, P.: Persuasive political argument. In: Computational Models of Natural Argument, IJCAI'05 workshop, pp. 44–51 (2005)
22. Baader, F., Hollunder, B.: Priorities on defaults with prerequisites, and their application in treating specificity in terminological default logic. *Journal of Automated Reasoning* **15**(1), 41–68 (1995)
23. Bacchus, F., Grove, A.: Graphical models for preference and utility. In: Proceedings of UAI'95, pp. 3–10 (1995)
24. Baier, J.A., McIlraith, S.A.: Planning with preferences. *AI Magazine* **29**(4), 25–36 (2008)
25. Barberà, S., W., W.B., Pattanaik, P.: Ranking Sets of Objects. In: S. Barberà, P. Hammond, C. Seidl (eds.) *Handbook of Utility Theory, Vol 2: Extensions*, pp. 893–977. Springer Verlag, Berlin (2004)
26. Belnap, N.: How a computer should think. In: Proceedings of the Oxford International Symposium on Contemporary Aspects of Philosophy, pp. 30–56. Oxford (1976)
27. Belnap, N.: A useful four-valued logic. In: G. Epstein, J. Dunn (eds.) *Modern uses of multiple valued logics*, pp. 8–37. D. Reidel, Dordrecht (1977)
28. Belton, V., Gear, T.: On a short-coming of saaty's method of analytic hierarchies. *Omega* **11**(3), 228–230 (1983)
29. Benabbou, N., Perny, P., Viappiani, P.: Incremental elicitation of choquet capacities for multicriteria decision making. In: ECAI 2014 - 21st European Conference on Artificial Intelligence, 18-22 August 2014, Prague, Czech Republic, pp. 87–92 (2014)
30. Bench-Capon, T., Dunne, P.: Argumentation in artificial intelligence. *Artificial Intelligence* **171**, 619–641 (2007)

31. Bench-Capon, T.J.M.: Persuasion in practical argument using value-based argumentation frameworks. *Journal of Logic and Computation* **13**, 429–448 (2003)
32. Benferhat, S., Dubois, D., Prade, H.: Argumentative inference in uncertain and inconsistent knowledge bases. In: *Proceedings of UAI'93*, pp. 411–419 (1993)
33. Benferhat, S., Dubois, D., Prade, H.: Towards a possibilistic logic handling of preferences. *Applied Intelligence* **14**, 303–317 (2001)
34. Benthem, J., Grossi, D., Liu, F.: Deontics = betterness + priority. In: G. Governatori, G. Sartor (eds.) *Deontic Logic in Computer Science, LNCS*, vol. 6181, pp. 50–65. Springer Verlag, Berlin (2010)
35. Bergstra, J., Bethke, I., Rodenburg, P.: A propositional logic with four values: true, false, divergent and meaningless. *Journal of Applied Non-Classical Logics* **5**, 199–217 (1995)
36. Besnard, P., Hunter, A.: *Elements of Argumentation*. MIT Press (2008)
37. Bienvenu, M., Lang, J., Wilson, N.: From preference logics to preference languages, and back. In: *Proceedings of KR 10*, pp. 214–224 (2010)
38. Bigot, D., Zanuttini, B., Fargier, H., Mengin, J.: Probabilistic conditional preference networks. *CoRR* **abs/1309.6817** (2013)
39. Bistarelli, S., Pini, M., Rossi, F., Venable, K.: From soft constraints to bipolar preferences: modelling framework and solving issues. *Journal of Experimental and Theoretical Artificial Intelligence* **22**, 135 – 158 (2010)
40. Bistarelli, S., Pini, M., Rossi, F., Venable, K.: Uncertainty in bipolar preference problems. *Journal of Experimental and Theoretical Artificial Intelligence* **23**, 545 – 575 (2011)
41. Boella, G., van der Torre, L., Verhagen, H.: Introduction to normative multiagent systems. *Computation and Mathematical Organizational Theory, Special issue on Normative Multiagent Systems* **12**(2-3), 71–79 (2006)
42. Bondarenko, A., Dung, P., Kowalski, R., Toni, F.: An abstract, argumentation-theoretic approach to default reasoning. *Artificial Intelligence* **93**, 63–101 (1997)
43. Bondarenko, A., Toni, F., Kowalski, R.: An assumption-based framework for non-monotonic reasoning. In: *Proceedings of the 2nd International Workshop on Logic Programming and Nonmonotonic Reasoning*, pp. 171–189 (1993)
44. Booth, R., Chevaleyre, Y., Lang, J., Mengin, J., Sombattheera, C.: Learning conditionally lexicographic preference relations. In: *Proceedings of ECAI'10*, pp. 269–274 (2010)
45. Bossu, G., Siegel, P.: Saturation, nonmonotonic reasoning and the closed-world assumption. *Artificial Intelligence* **25**, 13–65 (1985)
46. Boutilier, C.: What is a default priority. In: *Proceedings of Canadian Society for Computational Studies of Intelligence Conference*, pp. 140–147 (1992)
47. Boutilier, C.: A POMDP formulation of preference elicitation problems. In: *Proceedings of AAAI'02*, pp. 239–246 (2002)
48. Boutilier, C., Bacchus, F., Brafman, R.: UCP-networks: A directed graphical representation of conditional utilities. In: *Proceedings of UAI'01*, pp. 56–64 (2001)
49. Boutilier, C., Brafman, R., Hoos, H., Poole, D.: Reasoning with conditional ceteris paribus preference statements. In: *Proceedings of UAI'99*, pp. 71–80 (1999)
50. Boutilier, C., Patrascu, R., Poupart, P., Schuurmans, D.: Constraint-based Optimization and Utility Elicitation using the Minimax Decision Criterion. *Artificial Intelligence* **170**, 686–713 (2006)
51. Bouyssou, D., Marchant, T., Pirlot, M., Perny, P., Tsoukiàs, A., Vincke, P.: *Evaluation and decision models: a critical perspective*. Kluwer Academic, Dordrecht (2000)
52. Bouyssou, D., Marchant, T., Pirlot, M., Tsoukiàs, A., Vincke, P.: *Evaluation and decision models with multiple criteria: Stepping stones for the analyst*. Springer Verlag, Boston (2006)
53. Bouyssou, D., Pirlot, M.: Preferences for multiattributed alternatives: Traces, dominance, and numerical representations. *Journal of Mathematical Psychology* **48**, 167–185 (2004)
54. Bouyssou, D., Pirlot, M.: Conjoint measurement tools for MCDM. In: J. Figueira, S. Greco, M. Ehrgott (eds.) *Multiple Criteria Decision Analysis: State of the Art Surveys*, pp. 73–132. Springer Verlag, Boston (2005)

55. Bouyssou, D., Pirlot, M.: Following the traces: - an introduction to conjoint measurement without transitivity and additivity. *European Journal of Operational Research* **163**, 287–337 (2005)
56. Bouyssou, D., Pirlot, M.: Conjoint measurement models for preference relations. In: D. Bouyssou, D. Dubois, M. Pirlot, H. Prade (eds.) *Decision Making Process*, pp. 617–672. J. Wiley, New York (2009)
57. Brafman, R., Dimopoulos, Y.: Extended semantics and optimization algorithms for CP-networks. *Computational Intelligence* **20**, 219–245 (2004)
58. Brafman, R., Domshlak, C.: Introducing variable importance tradeoffs into CP-Nets. In: *Proceedings of UAI'02*, pp. 69–76 (2002)
59. Brafman, R., Domshlak, C.: Graphically structured value-function compilation. *Artificial Intelligence* **172**, 325–349 (2008)
60. Brafman, R., Domshlak, C.: Preference handling: An introductory tutorial. *AI Magazine* **30**, 58–86 (2008)
61. Brafman, R., Roberts, F., Tsoukiàs, A.: *Proceedings of ADT 2011*. LNAI 6992, Springer Verlag, Berlin (2011)
62. Braunhofer, M., Kaminskas, M., Ricci, F.: Location-aware music recommendation. *IJMIR* **2**(1), 31–44 (2013)
63. Brewka, G.: Preferred subtheories: An extended logical framework for default reasoning. In: *Proceedings of the 11th International Joint Conference on Artificial Intelligence - Volume 2, IJCAI'89*, pp. 1043–1048. Morgan Kaufmann Publishers Inc. (1989)
64. Brewka, G.: Adding priorities and specificity to default logic. In: C. MacNish, D. Pearce, L.M. Pereira (eds.) *Logics in Artificial Intelligence, Lecture Notes in Computer Science*, vol. 838, pp. 247–260. Springer Berlin Heidelberg (1994)
65. Brewka, G.: Reasoning about priorities in default logic. In: *Proceedings of the 12th National Conference on Artificial Intelligence (MIT)*, pp. 940–945 (1994)
66. Brewka, G.: A rank-based description language for qualitative preferences. In: *Proceedings of ECAI 04*, p. 303307 (2004)
67. Brewka, G., Niemela, I., Truszczynski, M.: Preferences and nonmonotonic reasoning. *AI magazine* **29**, 69–78 (2008)
68. Brewka, G., Niemel, I., Truszczynski, M.: Answer set optimization. In: *Proceedings of IJCAI'03*, pp. 867–872 (2003)
69. Brinker, K., Hüllermeier, E.: Case-based label ranking. In: J. Fürnkranz, T. Scheffer, M. Spiliopoulou (eds.) *ECML, Lecture Notes in Computer Science*, vol. 4212, pp. 566–573. Springer (2006)
70. Britz, K., Heidema, J., Meyer, T.A.: Semantic preferential subsumption. In: *Proceedings of KR'08*, pp. 476–484 (2008)
71. Britz, K., Meyer, T., Varzinczak, I.: Preferential Reasoning for Modal Logics. *Electronic Notes in Theoretical Computer Science* **278**, 55–69 (2011)
72. Caminada, M., Amgoud, L.: On the evaluation of argumentation formalisms. *Artificial Intelligence* **171**, 286–310 (2007)
73. Casini, G., Straccia, U.: Rational closure for defeasible description logics. In: *Proceedings of JELIA'10*, pp. 77–90 (2010)
74. Castañeda, H.: The paradoxes of deontic logic: The simplest solution to all of them in one fell swoop. In: R. Hilpinen (ed.) *New Studies in Deontic Logic: Norms, Actions and the Foundations of Ethics*, pp. 37–85. D.Reidel, Dordrecht (1981)
75. Cayrol, C., Royer, V., Saurel, C.: Management of preferences in assumption-based reasoning. In: *Proceedings of IPMU'92*, pp. 13–22 (1993)
76. Chajewska, U., Koller, D., Parr, R.: Making rational decisions using adaptive utility elicitation. In: *Proceedings of AAAI'00*, pp. 363–369 (2000)
77. Chellas, B.: Conditional obligation. In: S. Stunland (ed.) *Logical Theory and Semantical Analysis*, pp. 23–33. D. Reidel, Dordrecht (1974)
78. Chevaleyre, Y., Endriss, U., Lang, J.: Expressive power of weighted propositional formulas for cardinal preference modeling. In: *Proceedings of KR'06*, pp. 145–152 (2006)
79. Chevaleyre, Y., Endriss, U., Lang, J., Maudet, N.: Preference handling in combinatorial domains: From AI to social choice. *AI Magazine* **29**(4), 37–46 (2008). URL <http://www.aaai.org/ojs/index.php/aimagazine/article/view/2201>

80. Chevaleyre, Y., Koriche, F., Lang, J., Mengin, J., Zanuttini, B.: Learning ordinal preferences on multiattribute domains: The case of cp-nets. In: J. Fürnkranz, E. Hüllermeier (eds.) *Preference Learning*, pp. 273–296. Springer Verlag, Berlin (2011)
81. Chisholm, R.: *Perceiving*. Princeton University Press, Princeton (1957)
82. Chisholm, R.: *Theory of Knowledge*. Prentice-Hall, Englewood Cliffs (1966)
83. Colorni, A., Tsoukiàs, A.: What is a decision problem? preliminary statements. In: *Proceedings of ADT'13, LNAI 8176*, pp. 139 – 153. Springer Verlag, Berlin (2013)
84. Bana e Costa, C.A., Vansnick, J.C.: MACBETH - an interactive path towards the construction of cardinal value functions. *International transactions in operational Research* **1**, 489–500 (1994)
85. Bana e Costa, C.A., Vansnick, J.C.: A critical analysis of the eigenvalue method used to derive priorities in AHP. *European Journal of Operational Research* **187**, 1422–1428 (2008)
86. Critchlow, D.E., Fligner, M.A., Verducci, J.S.: Probability models on rankings. *Journal of Mathematical Psychology* **35**, 294 – 318 (1991)
87. Danielsson, S.: *Preference and obligation*. Studies in the logic of ethics. Filosofiska föreningen, Uppsala (1968)
88. de Finetti, B.: La prevision: Ses lois logiques, ses sources subjectives. In: *Annales de l'Institut Henri Poincaré* **7**, pp. 1–68. Paris (1937). Translated into English by Henry E. Kyburg Jr., *Foresight: Its Logical Laws, its Subjective Sources*. In Henry E. Kyburg Jr. and Howard E. Smokler (1964, Eds.), *Studies in Subjective Probability*, 53–118, Wiley, New York
89. Delgrande, J.P., Schaub, T.: Expressing preferences in default logic. *Artificial Intelligence* **123**(1-2), 41–87 (2000)
90. Delgrande, J.P., Schaub, T.H.: Compiling reasoning with and about preferences into default logic. In: *Proceedings of the 15th International Joint Conference on Artificial Intelligence - Volume 1, IJCAI'97*, pp. 168–174. Morgan Kaufmann Publishers Inc. (1997)
91. Deschrijver, G., Arieli, O., Cornelis, C., Kerre, E.: A bilattice-based framework for handling graded truth and imprecision. *Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* **15**, 13–41 (2007)
92. Doherty, P., Driankov, D., Tsoukiàs, A.: Partial logics and partial preferences. In: *Proceedings of CEMIT'92*, pp. 525–528 (1992)
93. Domshlak, C., Brafman, R.: CP-nets - reasoning and consistency testing. In: *Proceedings of KR'02*, pp. 121–132 (2002)
94. Domshlak, C., Hüllermeier, E., Kaci, S., Prade, H.: Preferences in AI: an overview. *Artificial Intelligence* **175**, 1037–1052 (2011)
95. Doyle, J.: Prospects for preferences. *Computational Intelligence* **20**, 111–136 (2004)
96. Doyle, J., Wellman, M.: Impediments to universal preference-based default theories. *Artificial Intelligence* **49**, 97–128 (1991)
97. Dubarle, D.: *Essai sur la généralisation naturelle de la logique usuelle*. *Mathématique, Informatique, Sciences Humaines* N° **107**, 17–73 (1989). 1963 manuscript, published posthumously
98. Dubois, D., Fargier, H., Perny, P.: Qualitative decision theory with preference relations and comparative uncertainty: An axiomatic approach. *Artificial Intelligence* **148**, 219–260 (2003)
99. Dubois, D., Fargier, H., Perny, P., Prade, H.: Qualitative decision theory: from Savage's axioms to non-monotonic reasoning. *Journal of the ACM* **49**, 455–495 (2002)
100. Dubois, D., Grabisch, M., Modave, F., Prade, H.: Relating decision under uncertainty and multicriteria decision making models. *International Journal of Intelligent System* **15**, 967–979 (2000)
101. Dubois, D., Prade, H.: A class of fuzzy measures based on triangular norms. *International Journal of General Systems* **8**, 43–61 (1982)
102. Dubois, D., Prade, H.: *Possibility theory*. Plenum Press, New-York (1988)
103. Dubois, D., Prade, H.: Possibility theory as a basis for qualitative decision theory. In: *Proceedings of IJCAI'95*, pp. 1924–1930 (1995)
104. Dubois, D., Prade, H.: An introduction to bipolar representations of information and preference. *International Journal of Intelligent Systems* **23**, 866–877 (2008)

105. Dubus, J., Gonzales, C., Perny, P.: Multiobjective optimization using GAI models. In: Proceedings of IJCAI'09, pp. 1902–1907 (2009)
106. Dung, P.: An argumentation semantics for logic programming with explicit negation. In: Proceedings of the 10th Logic Programming Conference, pp. 616–630 (1993)
107. Dung, P.M.: On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n -person games. *Artificial Intelligence* **77**, 321–357 (1995)
108. Dupin de Saint-Cyr, F., Lang, J., Schiex, T.: Penalty logic and its link with Dempster-Shafer theory. In: Proceedings of UAI'94, pp. 204–211 (1994)
109. Dyer, J.S.: A clarification of "remarks on the analytic hierarchy process". *Management Science* **36**(3), pp. 274–275 (1990). URL <http://www.jstor.org/stable/2631949>
110. Dyer, J.S.: Remarks on the analytic hierarchy process. *Management Science* **36**(3), pp. 249–258 (1990). URL <http://www.jstor.org/stable/2631946>
111. Ehrgott, M.: Multiobjective optimization. *AI Magazine* **29**(4), 47–57 (2008). URL <http://www.aaai.org/ojs/index.php/aimagazine/article/view/2198>
112. Etherington, D.: Reasoning with incomplete information. Pitman, London (1988)
113. Fages, F., Ruet, P.: Combining explicit negation and negation by failure via belnap's logic. *Theoretical Computer Science* **171**, 61–75 (1997)
114. Fargier, H., Sabadin, R.: Qualitative decision under uncertainty: back to expected utility. *Artificial Intelligence* **164**, 245–280 (2005)
115. Fishburn, P.: *Utility Theory for Decision Making*. J. Wiley, New York (1970)
116. Fishburn, P.: Lexicographic orders, utilities and decision rules: a survey. *Management Science* **20**, 1442–1471 (1974)
117. Fishburn, P.: *Interval Orders and Interval Graphs*. J. Wiley, New York (1985)
118. Fishburn, P.: *Nonlinear preference and utility theory*. Johns Hopkins University Press, Baltimore (1988)
119. Fishburn, P.: Preference structures and their numerical representations. *Theoretical Computer Science* **217**, 359–383 (1999)
120. Fitting, M.: Bilattices and the semantics of logic programming. *Journal of Logic Programming* **11**, 91–116 (1991)
121. Fodor, J., Roubens, M.: *Fuzzy preference modelling and multicriteria decision support*. Kluwer Academic, Dordrecht (1994)
122. Føllesdal, D., Hilpinen, R.: Deontic logic: An introduction. In: R. Hilpinen (ed.) *Deontic Logic: Introductory and Systematic Readings*. D. Reidel, Dordrecht (1971)
123. Font, J., Moussavi, M.: Note on a six valued extension of three valued logics. *Journal of Applied Non-Classical Logics* **3**, 173–187 (1993)
124. Forrester, J.W.: Gentle Murder, or the Adverbial Samaritan. *Journal of Philosophy* **81**, 193–196 (1984)
125. Fortemps, P., Słowiński, R.: A graded quadrivalent logic for ordinal preference modelling : Loyola-like approach. *Fuzzy Optimization and Decision Making* **1**, 93–111 (2002)
126. Fürnkranz, J., Hüllermeier, E., Cheng, W., Park, S.H.: Preference-based reinforcement learning: a formal framework and a policy iteration algorithm. *Machine Learning* **89**, 123–156 (2012)
127. Gabbay, D.: Theoretical foundations for nonmonotonic reasoning in expert systems. In: Proceedings Nato Advanced Study Institute on Logic and Models of Concurrent Systems, pp. 439–457. Springer Verlag, Berlin (1985)
128. Gajos, K., Weld, D.S.: Preference elicitation for interface optimization. In: Proceedings of UIST'05, pp. 173–182 (2005)
129. García, A., Simari, G.: Defeasible logic programming: an argumentative approach. *Theory and Practice of Logic Programming* **4**, 95–138 (2004)
130. Gedikli, F., Jannach, D., Ge, M.: How should I explain? A comparison of different explanation types for recommender systems. *Int. J. Hum.-Comput. Stud.* **72**(4), 367–382 (2014)
131. Geffner, H., Pearl, J.: Conditional entailment: Bridging two approaches to default reasoning. *Artificial Intelligence* **53**(2-3), 209–244 (1992)
132. Gelain, M., Pini, M., Rossi, F., Venable, K., Wilson, N.: Interval-valued soft constraint problems. *Annals of Mathematics and Artificial Intelligence* **58**, 261–298 (2010)

133. Gelfond, M., Lifschitz, V.: Logic programs with classical negation. In: D.H. Warren (ed.) *Logic programming*, pp. 579–597. MIT Press, Cambridge, MA (1990)
134. Gelfond, M., Lifschitz, V.: Classical negation in logic programs and disjunctive databases. *New Generation Computing* **9**, 365–385 (1991)
135. Gelfond, M., Przymusinska, H., Przymusinski, T.: On the Relationship Between Circumscription and Negation as Failure. *AIJ* **38**, 75–94 (1989)
136. Gilboa, I., Schmeidler, D., Wakker, P.: Utility in case-based decision theory. *Journal of Economic Theory* **105**, 483–502 (2002)
137. Ginsberg, M.: Multivalued logics: a uniform approach to reasoning in Artificial Intelligence. *Computational Intelligence* **4**, 265–316 (1988)
138. Giordano, L., Olivetti, N., Gliozzi, V., Pozzato, G.L.: Alc + t: a preferential extension of description logics. *Fundamenta Informaticae* **96**, 341–372 (2009)
139. Goble, L.: A logic of good, would and should. part 1. *Journal of Philosophical Logic* **19**, 169–199 (1990)
140. Goble, L.: A logic of good, would and should. part 2. *Journal of Philosophical Logic* **19**, 253–276 (1990)
141. Gonzales, C., Perny, P.: GAI networks for utility elicitation. In: *Proceedings KR’04*, pp. 224 – 234 (2004)
142. Gonzales, C., Perny, P., Queiroz, S.: Preference aggregation with graphical utility models. In: *Proceedings of AAAI’08*, pp. 1037–1042 (2008)
143. Gordon, T.: *The Pleading Game. An Artificial Intelligence Model of Procedural Justice*. Kluwer, Dordrecht (1995)
144. Governatori, G., Maher, M., Antoniou, G., Billington, D.: Argumentation semantics for defeasible logic. *Journal of Logic and Computation* **14**(5), 675–702 (2004)
145. Grabisch, M.: Fuzzy integral in multicriteria decision making. *Fuzzy Sets and Systems* **69**, 279–298 (1995)
146. Grabisch, M., Labreuche, C.: Fuzzy measures and integrals in MCDA. In: J. Figueira, S. Greco, M. Ehrgott (eds.) *Multiple Criteria Decision Analysis: State of the Art Surveys*, pp. 563–608. Springer Verlag, Boston (2005)
147. Greco, S., Mousseau, V., Slowinski, R.: Ordinal regression revisited: Multiple criteria ranking using a set of additive value functions. *European Journal of Operational Research* **191**, 416–436 (2008)
148. Greco, S., Mousseau, V., Slowinski, R.: Multiple criteria sorting with a set of additive value functions. *European Journal of Operational Research* **207**, 1455–1470 (2010)
149. Guo, S., Sanner, S.: Real-time multiattribute bayesian preference elicitation with pairwise comparison queries. In: *AISTATS*, pp. 289–296 (2010)
150. Halpern, J., Moses, Y.: Towards a theory of knowledge and ignorance: Preliminary report. In: *Proceedings of NMR’84*, pp. 125–143 (1984)
151. Hansson, B.: An analysis of some deontic logics. *Noûs* **3**, 373–398 (1969)
152. Hansson, S.: Preference-Based Deontic Logic (PDL). *Journal of Philosophical Logic* **19**, 75–93 (1990)
153. Harker, P.T., Vargas, L.G.: Reply to ”remarks on the analytic hierarchy process” by j. s. dyer. *Management Science* **36**(3), pp. 269–273 (1990). URL <http://www.jstor.org/stable/2631948>
154. Herbrich, R., Minka, T., Graepel, T.: Trueskilltm: A bayesian skill rating system. In: *Proceedings of NIPS’06*, pp. 569–576 (2006)
155. Israel, D.: What’s wrong with non-monotonic logic? In: *Proceedings of AAAI’80*, pp. 99–101 (1980)
156. Israel, D.: The role(s) of logic in Artificial Intelligence. In: D.M. Gabbay, C.J. Hogger, J.A. Robinson (eds.) *Handbook of Logic in Artificial Intelligence and Logic Programming, Volume I*, pp. 1–31. Oxford University Press, Oxford (1993)
157. Jackson, F.: On the semantics and logic of obligation. *Mind* **94**, 177–196 (1985)
158. Jacquet-Lagrèze, E., Siskos, Y.: Assessing a set of additive utility functions for multicriteria decision making: the UTA method. *European Journal of Operational Research* **10**, 151–164 (1982)
159. Jacquet-Lagrèze, E., Siskos, Y.: Preference disaggregation: 20 years of MCDA experience. *European Journal of Operational Research* **130**, 233–245 (2001)

160. Joachims, T.: Optimizing search engines using clickthrough data. In: Proceedings of KDD'02, pp. 133–142 (2002)
161. Jones, A., Sergot, M.: Deontic logic in the representation of law: Towards a methodology. *Artificial Intelligence and Law* **1**, 45–64 (1992)
162. Jones, A., Sergot, M.: On the characterisation of law and computer systems: The normative systems perspective. In: J.J. Meyer, R. Wieringa (eds.) *Deontic Logic in Computer Science*. 1993 (John Wiley & Sons)
163. Jørgensen, J.: Imperatives and logic. *Erkenntnis* **7**, 288–296 (1938)
164. Kaci, S.: Refined preference-based argumentation frameworks. In: COMMA, pp. 299–310 (2010)
165. Kaci, S.: *Working with Preferences: Less Is More*. Springer Verlag, Berlin (2011)
166. Kaci, S., van der Torre, L.: Preference-based argumentation: Arguments supporting multiple values. *Journal of Approximate Reasoning* **48**(3), 730–751 (2008)
167. Kaci, S., van der Torre, L.W.N., Weydert, E.: Acyclic argumentation: Attack = conflict + preference. In: Proceedings of ECAI'06, pp. 725–726 (2006)
168. Kaci, S., van der Torre, L.: Reasoning with various kinds of preferences: logic, non-monotonicity, and algorithms. *Annals of Operations Research* **163**, 89114 (2008)
169. Kacprzyk, J., Roubens, M.: *Non Conventional Preference Relations in Decision Making*. Springer Verlag, LNAMES n. 301, Berlin (1988)
170. Kahneman, D., Tversky, A.: Prospect theory: An analysis of decision under risk. *Econometrica* **47**, 263–291 (1979)
171. Kakas, A., Moraitis, P.: Argumentation based decision making for autonomous agents. In: Proceedings of AAMAS'03, pp. 883–890 (2003)
172. Kaluzhny, Y., Muravitsky, A.: A knowledge representation based on the Belnap's four valued logic. *Journal of Applied Non-Classical Logics* **3**, 189–203 (1993)
173. Keeney, R.L., Raiffa, H.: *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. John Wiley and Sons, New York (1976)
174. Köbberling, V., Wakker, P.: Preference foundations for nonexpected utility: A generalized and simplified technique. *Mathematics of Operations Research* **28**, 395–423 (2003)
175. Kok, E.M., Meyer, J.J.C., Prakken, H., Vreeswijk, G.: A formal argumentation framework for deliberation dialogues. In: Proceedings of ArgMAS'10, pp. 31–48 (2010)
176. Koons, R.: Defeasible reasoning. In: E.N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy*, on-line. Stanford University, Stanford (2009)
177. Kouvelis, P., Yu, G.: *Robust Discrete Optimization and Its Applications*. Kluwer Academic, Dordrecht, Dordrecht (1997)
178. Krantz, D., Luce, R., Suppes, P., Tversky, A.: *Foundations of measurement*, vol. 1: Additive and polynomial representations. Academic Press, New York (1971)
179. Kraus, S., Lehmann, D., Magidor, M.: Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence* **44**, 167–207 (1990)
180. Kraus, S., Sycara, K., Evenchik, A.: Reaching agreements through argumentation: a logical model and implementation. *Artificial Intelligence* **104**, 1–69 (1998)
181. Labreuche, C., Huédé, F.L.: MYRIAD: a tool suite for MCDA. In: Proceedings of EUSFLAT'05, pp. 204–209 (2005)
182. Lafage, C., Lang, J.: Propositional distances and compact preference representation. *European Journal of Operational Research* **160**, 741–761 (2005)
183. Lang, J.: Logical preference representation and combinatorial vote. *Annals of Mathematics and Artificial Intelligence* **42**, 3771 (2004)
184. Lang, J.: Logical representation of preferences. In: D. Bouyssou, D. Dubois, M. Pirlot, H. Prade (eds.) *Decision-Making Process: Concepts and Methods*, pp. 321–363. J. Wiley, New York (2009)
185. Lang, J., Mengin, J., Xia, L.: Aggregating conditionally lexicographic preferences on multi-issue domains. In: Proceedings of CP 2012, pp. 973 – 987 (2012)
186. Lehmann, D., Magidor, M.: Preferential logics: the predicate calculus case. In: Proceedings of TARK'90, pp. 57–72 (1990)
187. Lehmann, D., Magidor, M.: What does a conditional knowledge base entail? *Artificial Intelligence* **55**, 1–60 (1992)
188. Lewis, D.: Semantic analysis for dyadic deontic logic. In: S. Stunland (ed.) *Logical Theory and Semantical Analysis*, pp. 1–14. D. Reidel, Dordrecht (1974)

189. Lifschitz, V.: Computing circumscription. In: Proceedings of IJCAI'85, pp. 121–127 (1985)
190. Lifschitz, V.: Pointwise circumscription. In: M. Ginsberg (ed.) *Readings in Nonmonotonic Reasoning*, pp. 179–193. Morgan Kaufmann Publishers Inc. San Francisco, CA, USA (1987)
191. Lootsma, F.: *Multi-criteria decision analysis via ratio and difference judgement*. Kluwer Academic, Dordrecht (1999)
192. Loui, R.: Defeat among arguments: a system of defeasible inference. *Computational Intelligence* **2**, 100–106 (1987)
193. Lu, T., Boutilier, C.: Learning mallows models with pairwise preferences. In: Proceedings of ICML'11, pp. 145–152 (2011)
194. Lu, T., Boutilier, C.: Robust approximation and incremental elicitation in voting protocols. In: *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16-22, 2011*, pp. 287–293 (2011)
195. Makinson, D.: General theory of cumulative inference. In: M.G. M. Reinfrank J. de Kleer, R. Sandewall (eds.) *Non-Monotonic Reasoning*, LNAI 346, pp. 55–76. Springer Verlag, Berlin (1989)
196. Mally, E.: *Grundgesetze des Sollens. Elemente der Logik des Willens*. Graz: Leuschner & Leubensky (1926)
197. Marchant, T.: Towards a theory of MCDM: stepping away from social choice theory. *Mathematical Social Sciences* **45**, 343–363 (2003)
198. Marichal, J.L., Meyer, P., Roubens, M.: Sorting multi-attribute alternatives: the TOMASO method. *Computers & Operations Research* **32**, 861–877 (2005)
199. Marquis, S., Lang, J., Liberatore, P., Marquis, P.: Expressive power and succinctness of propositional languages for preference representation. In: Proceedings of KR'04, pp. 203 – 212 (2004)
200. McCarthy, J.: Circumscription: A form of nonmonotonic reasoning. *Artificial Intelligence* **13**, 27–39 (1980)
201. McCarthy, J.: Applications of circumscription to formalizing of commonsense knowledge. *Artificial Intelligence* **28**, 89–116 (1986)
202. McCarty, L.T.: Modalities over actions: 1. model theory. In: Proceedings of (KR'94), pp. 437 – 448. Morgan Kaufmann (1994)
203. McDermott, D.: Non-monotonic logic i. *Artificial Intelligence* **13**, 41–72 (1982)
204. McDermott, D., Doyle, J.: Non-monotonic logic ii. *Journal of ACM* **29**, 33–57 (1982)
205. McNamara, P.: Deontic logic. In: E.N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy*, on-line. Stanford University, Stanford (2010)
206. McSherry, D., Stretch, C.: Automating the discovery of recommendation knowledge. In: Proceedings of IJCAI'05, pp. 9–14 (2005)
207. Minka, T.P.: Expectation propagation for approximate bayesian inference. In: Proceedings of UAI'01, pp. 362–369 (2001)
208. Modgil, S.: Nested argumentation and its application to decision making over actions. In: Proceedings of ArgMAS'05, pp. 57–73. Springer Verlag, Berlin (2006)
209. Modgil, S.: Reasoning about preferences in argumentation framework. *Artificial Intelligence Journal* **173**, 901–934 (2009)
210. Moore, R.: Semantic considerations on nonmonotonic logic. *Artificial Intelligence* **25**, 75–94 (1985)
211. Moretti, S., Öztürk, M., Tsoukiàs, A.: Preference modelling. In: M. Ehrgott, S. Greco, J. Figueira (eds.) *State of the Art in Multiple Criteria Decision Analysis*. New revised version. Springer Verlag, Berlin (to appear)
212. Moretti, S., Tsoukiàs, A.: Ranking sets of possibly interacting objects using Shapley extensions. In: Proceedings of KR 12, pp. 199–209 (2012)
213. Nute, D.: Defeasible reasoning and decision support systems. *Decision Support Systems* **4**, 97–110 (1988)
214. Nute, D. (ed.): *Defeasible Deontic Logic*. Synthese Library 263. Kluwer Academic Publishers (1997)
215. Orlovsky, S.: Decision making with a fuzzy preference relation. *Fuzzy Sets and Systems* **1**, 155–167 (1978)
216. Ovchinnikov, S.: Structure of fuzzy binary relations. *Fuzzy Sets and Systems* **6**, 169–195 (1981)

217. Öztürk, M., Tsoukiàs, A.: Modelling uncertain positive and negative reasons in decision aiding. *Decision Support Systems* **43**, 1512 – 1526 (2007)
218. Öztürk, M., Tsoukiàs, A.: Bipolar preference modelling and aggregation in decision support. *International Journal of Intelligent Systems* **23**, 970–984 (2008)
219. Öztürk, M., Tsoukiàs, A., Vincke, Ph.: Preference modelling. In: M. Ehrgott, S. Greco, J. Figueira (eds.) *State of the Art in Multiple Criteria Decision Analysis*, pp. 27 – 72. Springer Verlag, Berlin (2005)
220. Pearl, J.: *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, San Mateo, CA (1988)
221. Pearl, J.: System Z: A natural ordering of defaults with tractable applications to default reasoning. In: *Proceedings of TARK'90*, pp. 121–135 (1990)
222. Pearl, J., Geffner, H.: Probabilistic semantics for a subset of default reasoning. Technical Report CSD-8700XX, R-93-III, Computer Science Dept., UCLA (1988)
223. Peintner, B., Viappiani, P., Yorke-Smith, N.: Preferences in interactive systems: Technical challenges and case studies. *AI Magazine* **29**(4), 13–24 (2008)
224. Perelman, C.: *Justice, Law and Argument*. Reidel: Dordrecht (1980)
225. Perelman, C., Olbrechts-Tyteca, L.: *The New Rhetoric: A Treatise on Argumentation*. University of Notre Dame Press, Notre Dame (1969)
226. Perny, P., Pirlot, M., Tsoukiàs, A.: *Proceedings of ADT 2013*. LNAI 8176, Springer Verlag, Berlin (2013)
227. Perny, P., Roy, B.: The use of fuzzy outranking relations in preference modelling. *Fuzzy Sets and Systems* **49**, 33–53 (1992)
228. Perny, P., Tsoukiàs, A.: On the continuous extension of a four valued logic for preference modelling. In: *Proceedings of IPMU'98*, pp. 302–309 (1998)
229. Pirlot, M., Vincke, P.: *Semi Orders*. Kluwer Academic, Dordrecht (1997)
230. Pollock, J.: *Knowledge and Justification*. Princeton University Press, Princeton (1974)
231. Pollock, J.: Defeasible reasoning. *Cognitive Science* **11**, 481–518 (1987)
232. Prakken, H.: A tool in modelling disagreement in law: preferring the most specific argument. In: *Proceedings of the 3rd International Conference on Artificial Intelligence and Law*, pp. 165–174 (1991)
233. Prakken, H.: An argumentation framework in default logic. *Annals of Mathematics and Artificial Intelligence* **9**, 91–131 (1993)
234. Prakken, H.: Coherence and flexibility in dialogue games for argumentation. *Journal of Logic and Computation* **15**, 1009–1040 (2005)
235. Prakken, H., Sartor, G.: Argument-based extended logic programming with defeasible priorities. *Journal of Applied Non-Classical Logics* **7**, 25–75 (1997)
236. Price, R., Messinger, P.R.: Optimal recommendation sets: Covering uncertainty over user preferences. In: *Proceedings of AAAI'05*, pp. 541–548 (2005)
237. Pu, P., Chen, L.: User-involved preference elicitation for product search and recommender systems. *AI Magazine* **29**(4), 93–103 (2008). URL <http://www.aaai.org/ojs/index.php/aimagazine/article/view/2200>
238. Ramsey, F.: *Foundations of Mathematics and other Logical Essays*. Routledge & P. Kegan, London (1931). Collection of papers published posthumously, edited by R.B Braithwaite
239. Reiter, R.: Logic and data bases. In: H. Gallaire, J. Minker (eds.) *On closed world data bases*, pp. 55–76. Plenum Press, New York (1978)
240. Reiter, R.: A logic for default reasoning. *Artificial Intelligence* **13**, 81–132 (1980)
241. Rescher, N.: The logic of preference. In: *Topics in Philosophical Logic, Synthese Library*, vol. 17, pp. 287–320. Springer Verlag, Berlin (1968)
242. Rescher, N.: *Introduction to Value Theory*. Prentice Hall, Englewood Cliffs (1969)
243. Roberts, F.: *Measurement theory, with applications to Decision Making, Utility and the Social Sciences*. Addison-Wesley, Boston (1979)
244. Roberts, F.: Computer science and decision theory. *Annals of Operations Research* **163**, 209–253 (2008)
245. Roberts, F., Tsoukiàs, A.: Special issue on computer science and decision theory. *Annals of Operations Research* **163**, 270 (2008)
246. Rossi, F.: Constraints and preferences: Modelling frameworks and multi-agent settings. In: G. Della Riccia, D. Dubois, R. Kruse, H. Lenz (eds.) *Similarities and Preferences*, pp. 305 – 320. CISM series, Springer Verlag, Berlin (2008)

247. Rossi, F., Tsoukiàs, A.: Proceedings of ADT 2009. LNAI 5783, Springer Verlag, Berlin (2009)
248. Rossi, F., Venable, K., Walsh, T.: mCP nets: Representing and reasoning with preferences of multiple agents. In: Proceedings of AAAI'04, pp. 729–734 (2004)
249. Rossi, F., Venable, K., Walsh, T.: Preferences in constraint satisfaction and optimization. *AI Magazine* **29**, 58–68 (2008)
250. Rossi, F., Venable, K., Walsh, T.: A Short Introduction to Preferences. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers (2011)
251. Roubens, M., Vincke, P.: Preference Modeling. LNEMS 250, Springer Verlag, Berlin (1985)
252. Roy, B.: Partial preference analysis and decision aid: The fuzzy outranking relation concept. In: D. Bell, R. Keeney, H. Raiffa (eds.) *Conflicting objectives in Decisions*, pp. 40–75. J. Wiley, New York (1977)
253. Saaty, T.: *The Analytic Hierarchy Process, Planning, Priority Setting, Resource Allocation*. McGraw-Hill, New York (1980)
254. Saaty, T.L.: An exposition on the ahp in reply to the paper "remarks on the analytic hierarchy process". *Management Science* **36**(3), pp. 259–268 (1990). URL <http://www.jstor.org/stable/2631947>
255. Salo, A., Keisler, J., Morton, A.: *Portfolio Management*. Springer Verlag, Berlin (2011)
256. Samuelson, P.: Probability and the attempts to measure utility. *Economic Review* **1**, 117–126 (1950)
257. Sartor, G.: A formal model of legal argumentation. *Ratio Juris* **7**, 212–226 (1994)
258. Savage, L.: *The Foundations of Statistics*. J. Wiley, New York (1954). Second revised edition, 1972
259. Schweizer, B., Sklar, A.: *Probabilistic Metric Spaces*. North Holland, Amsterdam (1983)
260. Shoham, Y.: A semantical approach to nonmonotonic logics. Proceedings of the Symposium on Logic in Computer Science pp. 275–279 (1987)
261. Shoham, Y.: Nonmonotonic logics: Meaning and utility. In: Proceedings of IJCAI'87, pp. 388–393 (1987)
262. Shoham, Y.: *Reasoning about Change*. MIT Press, Boston (1987)
263. Simari, G., Loui, R.: A mathematical treatment of defeasible reasoning and its implementation. *Artificial Intelligence* **53**, 125–157 (1992)
264. Sinz, C., Haag, A., Narodytska, N., Walsh, T., Gelle, E., Sabin, M., Junker, U., O'Sullivan, B., Rabiser, R., Dhungana, D., Grünbacher, P., Lehner, K., Federspiel, C., Naus, D.: Configuration. *IEEE Intelligent Systems* **22**(1), 78–90 (2007)
265. Smith, T.: Legal expert systems: discussion of theoretical assumptions. Ph.D. thesis, University of Utrecht (1994)
266. Smyth, B.: Case-based recommendation. In: *The Adaptive Web*, LNCS 4321, pp. 342–376. Springer Verlag, Berlin (2007)
267. Sycara, K.: Persuasive argumentation in negotiation. *Theory and Decision* **28**, 203–242 (1990)
268. Tan, Y.H., van der Torre, L.: How to combine ordering and minimizing in a deontic logic based on preferences. In: *Deontic Logic, Agency and Normative Systems*. Proceedings of the ?eon'96 Workshop in Computing, pp. 216–232. Springer Verlag, Berlin (1996)
269. Tanguiane, A.S.: *Aggregation and Representation of Preferences*. Springer-Verlag, Berlin (1991)
270. Thomason, R., Horty, J.: Logics for inheritance theory. In: M. Reinfrank, J. de Kleer, M. Ginsberg, E. Sandewall (eds.) *Non-Monotonic Reasoning*, pp. 220–237. Springer Verlag, Berlin (1987). LNAI 346
271. van der Torre, L.: Reasoning about obligations: Defeasibility in preference-based deontic logic. Ph.D. thesis, Erasmus University Rotterdam (1997)
272. Toulmin, S.: *The Uses of Argument*. Cambridge University Press, Cambridge (1958)
273. Touretzky, D.S.: A skeptic's menagerie: conflictors, preemptors, reinstaters, and zombies in nonmonotonic inheritance. In: *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence (IJCAI-91)*, pp. 478–483. Morgan Kaufmann (1991)

274. Touretzky, D.S., Horty, J.F., Thomas, R.H.: A clash of intuitions: The current state of nonmonotonic multiple inheritance systems. In: *IJCAI-87*, pp. 476–482. Morgan Kaufmann (1987)
275. Trotter, W.: *Combinatorics and partially ordered sets*. John Hopkins University Press, Baltimore (1992)
276. Tsoukiàs, A.: Preference modelling as a reasoning process: a new way to face uncertainty in multiple criteria decision support systems. *European Journal of Operational Research* **55**, 309–318 (1991)
277. Tsoukiàs, A.: A first-order, four valued, weakly paraconsistent logic and its relation to rough sets semantics. *Foundations of Computing and Decision Sciences* **12**, 85–108 (2002)
278. Tsoukiàs, A., Perny, P., Vincke, P.: From concordance/discordance to the modelling of positive and negative reasons in decision aiding. In: D. Bouyssou, E. Jacquet-Lagrèze, P. Perny, R. Slowinski, D. Vanderpooten, P. Vincke (eds.) *Aiding Decisions with Multiple Criteria: Essays in Honour of Bernard Roy*, pp. 147–174. Kluwer Academic, Dordrecht (2002)
279. Tsoukiàs, A., Vincke, P.: A new axiomatic foundation of partial comparability. *Theory and Decision* **39**, 79–114 (1995)
280. Tsoukiàs, A., Vincke, P.: Extended preference structures in MCDA. In: J. Climaco (ed.) *Multicriteria Analysis*, pp. 37–50. Springer Verlag, Berlin (1997)
281. Tsoukiàs, A., Vincke, P.: Double threshold orders: A new axiomatization. *Journal of Multi-criteria Decision Analysis* **7**, 285–301 (1998)
282. Turunen, E., Öztürk, M., Tsoukiàs, A.: Paraconsistent semantics for Pavelka style fuzzysentential logic. *Fuzzy Sets and Systems* **161**, 1926–1940 (2010)
283. Uckelman, J.: Alice and Bob will fight: The problem of electing a committee in the presence of candidate interdependence. In: *Proceedings of MPREF'10*, pp. 73–78 (2010)
284. van Benthem, J., Girard, P., Roy, O.: Everything else being equal: A modal logic approach to ceteris paribus preferences. *Journal of Philosophical Logic* **38**, 83125 (2009)
285. van Dalen, D.: *Logic and Structure*. Springer Verlag, Berlin (1983)
286. van Fraassen, B.: The logic of conditional obligation. *Journal of Philosophical Logic* **1**, 417–438 (1972)
287. van Fraassen, B.: Values and the heart's command. *The Journal of Philosophy* **70**, 5–19 (1973)
288. Viappiani, P., Boutilier, C.: Regret-based optimal recommendation sets in conversational recommender systems. In: *Proceedings of the third ACM conference on Recommender systems*, pp. 101–108 (2009)
289. Viappiani, P., Boutilier, C.: Optimal bayesian recommendation sets and myopically optimal choice query sets. In: *Proceedings of NIPS'10*, pp. 2352–2360 (2010)
290. Viappiani, P., Kroer, C.: Robust optimization of recommendation sets with the maximin utility criterion. In: *Proceedings of ADT'13*, pp. 411–424 (2013)
291. von Neumann, J., Morgenstern, O.: *Theory of games and economic behaviour*. Princeton University Press, Princeton (1947). 2nd edition
292. von Wright, G.: Deontic Logic. *Mind* **60**, 1–15 (1951)
293. von Wright, G.: *An Essay in Modal Logic*. North-Holland, Amsterdam (1951)
294. von Wright, G.: *The logic of preference*. Edinburgh University Press, Edinburgh (1963)
295. von Wright, G.: Deontic logic and the theory of conditions. In: R. Hilpinen (ed.) *Deontic Logic: Introductory and Systematic Readings*, pp. 159–177. D. Reidel, Dordrecht (1971)
296. von Wright, G.: The logic of preference reconsidered. *Theory and Decision* **3**, 140–169 (1972)
297. Wagstaff, K.L., desJardins, M., Eaton, E.: Modelling and learning user preferences over sets. *Journal of Experimental and Theoretical Artificial Intelligence* **22**, 237–268 (2010)
298. Wang, T., Boutilier, C.: Incremental utility elicitation with the minimax regret decision criterion. In: *Proceedings of IJCAI'03*, pp. 309–316 (2003)
299. Wilson, N.: Consistency and constrained optimisation for conditional preferences. In: *Proceedings of ECAI'04*, pp. 888 – 894 (2004)

-
300. Wilson, N.: Extending CP-nets with stronger conditional preference statements. In: Proceedings of AAAI'04, pp. 735 – 741 (2004)
 301. Wilson, N.: Efficient inference for expressive comparative preference languages. In: Proceedings of IJCAI'09, pp. 961–966 (2009)
 302. Yaman, F., Walsh, T., Littman, M., desJardins, M.: Learning lexicographic preference models. In: J. Fürnkranz, E. Hüllermeier (eds.) Preference Learning, pp. 251–272. Springer Berlin, Berlin (2011)
 303. Yu, P., Wan, W., Lee, P.: Decision tree modeling for ranking data. In: J. Fürnkranz, E. Hüllermeier (eds.) Preference Learning, pp. 83–106. Springer Verlag, Berlin (2011)