

Preference Handling for Artificial Intelligence

*Judy Goldsmith
and Ulrich Junker*

■ This editorial explains the benefits of preferences for AI systems and draws a picture of current AI research on preference handling. It thus provides an introduction to the topics covered by this special issue on preference handling.

Intelligent systems as studied and developed in artificial intelligence either support and assist humans acting in the world or act in the world to fulfill tasks like human agents. To act autonomously, the systems must choose among different actions and means of expression; to intelligently support humans' actions, they must understand and respond to the humans' choices. It is a natural assumption that agents who are acting in the world are experiencing the consequences of their actions and are not indifferent with respect to those experiences. An autonomous system such as a Mars rover can experience the consequences of moving along a path by measuring the energy consumption after the move, which will depend on the difficulty of the chosen path, and then judge whether the path was good or bad. To make decisions in a desired way, agents not only need to understand the consequences of their acts but also need a policy, whether pre-computed or not, for choosing. That policy may take into account short- or long-term effects of the choices, but it must have a means of evaluating and of comparing those effects. Preferences achieve this and are the key for agents to make decisions in a desirable (and rational) way. Consequently, preferences are also the key for assistants to understand and support the decisions of human users.

We illustrate the impact and importance of preferences by the following

example. Max Protimesis is trying to choose textbooks for his course on preferences in AI by using an online shopping system. He first asks for books published in the last five years and gets results from psychology, business, and other human-centered disciplines, which do not cover the fundamental normative approaches to decision making. He therefore searches for texts on normative approaches, but they are all from a time before computer science was an academic discipline. Protimesis grows agitated. His students want a "real book," not his brilliant lecture notes. So he eases his constraints and replaces them with preferences: he prefers more recent to older texts, and he prefers theoretical texts to more empirical ones. Given these criteria, and the trade-off between them, he is able to make a good enough, if not ideal, choice such as the one displayed in the third table of figure 1. This story shows that knowledge about user's preferences may be critical to make satisfactory recommendations to a user.

As preferences are so fundamental for decision making, they have extensively been studied in economics, psychology, operations research, and other human-centered disciplines. Classic models are utility functions that map the possible outcomes of the decisions to numeric values and thus allow the comparison and sorting of those outcomes. Another model is that of weak preference orders that describe which

Books (search limited to recent books)	Year	Approach
S. Lichtenstein, P. Slovic: <i>The Construction of Preference</i>	2006	descriptive
A. Tversky, E. Shafir: <i>Preference, Belief, and Similarity: Selected Writings</i>	2003	descriptive
Books (search limited to normative approaches)	Year	Approach
J. von Neumann, O. Morgenstern: <i>Theories of Games and Economic Behaviour</i>	1947	normative
K. J. Arrow: <i>Social Choice and Individual Values</i>	1951	normative
Books (search preferring recent books and normative over prescriptive over descriptive approaches)	Year	Approach
R. L. Keeney, H. Raiffa: <i>Decisions with Multiple Objectives: Preferences and Value Trade-Offs</i>	1976	prescriptive

Figure 1. Search for Books on Preferences.

outcome is at least as preferred as which other outcome. These preference orders can easily be elicited from a user if the set of outcomes is small.

Artificial intelligence brings new application fields to these classic preference models. Preferences are important in recommender systems, e-commerce, multiagent systems, planning and scheduling, configuration and design, and other tasks concerning intelligent decision support or autonomous decision making. However, the cross-fertilization between preference handling and AI also goes in the other direction. Existing AI methods for learning, knowledge representation, and problem solving have led to new preference handling methods. A good example is conditional preferences, which were discussed in a recent article published in *AI Magazine* (Walsh 2007). Another topic is preference-based problem solving as defined in Junker (2008). Moreover, the AI effort to construct models of intelligent behavior has identified new roles of preferences such as reasoning control or the constitution of belief sets (Doyle 2004).

This discussion shows that there are many questions that can be asked about preferences in AI. The primary ones are the following: Which kinds of preference models are of interest? How do we represent them? How do we ob-

tain those preferences? How can we use them in reasoning? How do we actually compute with them? These questions arise for applications of preferences to AI problems as well as for new preference handling methods developed in AI. This special issue seeks to give an idea of the diversity of current AI research in preference handling. It discusses exemplary application fields and some central methods, while pointing out proven and potential benefits of preferences.

As preference handling is becoming important for many fields of AI, we have not been able to cover the full range of preference-related applications in this issue but have had to focus on some of the hot topics. Domains where preference-aware systems are a reality today are recommender systems, personal assistants, and personalized user interfaces. Preference-aware agents help us shop, choose music, arrange photo albums, pick news stories, and perform many other daily functions. The article by Bart Peintner, Neil Yorke-Smith, and Paolo Viappiani (“Preferences in Interactive Systems: Technical Challenges and Case Studies”) surveys the major approaches while discussing questions such as preference learning, system design, and user interaction. There is also a thriving field of intelligent tutoring systems (Sleeman and Brown 1982)

and a rich literature in the field of education on learning styles. We expect that these two studies will soon overlap, and personalized tutors that take into account preferred learning styles will become available.

Autonomous agents such as a Mars rover need to find a best route to a given target while minimizing energy consumption and maximizing the set of tasks that are fulfilled on this route. They must find a suitable sequence of actions that reach the goals and maximize their preferences. In other words, those autonomous systems face a planning problem under preferences. Although planning is one of the central problems in AI, the problem of finding preferred plans or high quality plans has only recently drawn interest, as explained in the article by Jorge Baier and Sheila McIlraith called “Planning with Preferences.” This article nevertheless lists an impressive number of preference-based planning systems. Preference languages are of particular importance in this domain as preferences can be expressed on goals, on actions, and on temporal formulas. McIlraith and Baier further discuss planning algorithms that optimize for preferences or utility, or even, in the case of stochastic domains, expected utility.

Multiagent systems are a rich application field for preference handling.

For example, combinatorial auctions (Sandholm 2007) deal with preferences of the auctioneers, which are represented in the form of bids. This leads to particular questions concerning the aggregation and elicitation of preferences. In any competitive setting, the question of revealing preferences is mated with the question of hiding preferences. The field of mechanism design studies the problem of how to choose the outcome of an auction depending on the reported agents' preferences, while penalizing manipulations. Preference aggregation arises in competitive settings, such as auctions or elections, and in cooperative settings as well. It might seem that a collection of cooperative agents would be able to balance their individual preferences and come up with a collectively acceptable decision, but that is not as easy as it ought to be. The problem of preference aggregation—of gathering individuals' preferences and defining a group decision—is surveyed here in "Preference Handling in Combinatorial Domains: From AI to Social Choice" by Yann Chevaleyre, Ulle Endriss, Jérôme Lang, and Nicolas Maudet. The authors discuss collective decision making for problems with a combinatorial domain such as the selection of a steering committee or the allocation of satellite time among the funding agencies.

Preferences are interesting for many other tasks such as robotics, active vision, and natural language processing. For example, word choice—by a human or a computer agent—and the complexity level of grammatical structures are guided by preferences. If a speech-understanding program recognizes the speaker, that gives it an edge in understanding the speech, in part because it understands the speaker's preferences. We believe that future research will bring new insights in the benefits of preferences for cognitive tasks.

The second part of this special issue surveys methods for preference handling from artificial intelligence and operations research. The articles focus on problems where the set of possible decisions is too large to be described explicitly and a constraint-based or

logical formulation is used instead. We start with simple but rather incomplete preference models, where elicitation is easy but the computation of the best, that is, optimal, outcomes is complicated. This case arises if preferences are specified on different attributes independently of each other. Each of these attributes constitutes a different viewpoint for comparing the decisions. In this case, the preferences can be combined into a partial order, the famous Pareto-dominance partial order, which means that one outcome is preferred to another outcome if it is strictly preferred to the other outcome for all of the attributes on which the two outcomes differ. This partial order permits different trade-offs, and the field of multiobjective optimization has developed algorithms for computing the tradeoffs systematically. Matthias Ehrgott's article, "Multiobjective Optimization," gives a thorough survey of these methods.

In "Preferences in Constraint Satisfaction and Optimization," Francesca Rossi, Brent Venable, and Toby Walsh discuss soft constraints and conditional-preference networks (CP-nets). Soft constraints are more general than hard constraints and associate a degree of satisfaction with each of the decisions. CP-nets express preferential dependencies between attributes. A book buyer may prefer hardcover to paperback for mathematical books, which he or she has to read many times, and paperback to hardcover for introductory books, which he or she will read once. Here, the preference on the book cover depends on the book's type. Rossi, Venable, and Walsh discuss the relationships between soft constraints and CP-nets and mention topics such as abstraction, explanation, learning, and elicitation. More about CP-nets can be found in Ronen Brafman's and Carmel Domshlak's "Preference Handling—An Introductory Tutorial."

As previously mentioned, preferences may have other roles than guiding choices. For example, they may help to construct belief sets and deal with incomplete knowledge. This was discovered in the attempt to formalize commonsense reasoning, and the pioneering work of John McCarthy on circumscription and Jon Doyle on rea-

son maintenance systems are two instructive examples of this endeavor. Reason maintenance systems strongly influenced the newer field of answer-set programming (ASP), and there is a close correspondence between the concepts of those fields. In "Preferences and Nonmonotonic Reasoning," Gerhard Brewka, Ilkka Niemelä, and Mirosław (Mirek) Truszczynski give a nice introduction to ASP and explain the difference between beliefs and desires. Although ASP is designed to compute preferred belief sets, it can also be extended to represent preferences between choices that an agent has to make.

The preference representations considered so far have been limited in the sense that they could not model an arbitrary preference order or utility function over a combinatorial domain. The classic work on multiattribute utility theory (MAUT) (Keeney and Raiffa 1976, Fishburn 1964) introduces (generalized) additive utility functions, which are the sum of local utility functions over subsets of attributes. These factored utility functions can model preferential dependencies between attributes and decisions with nondeterministic effects. However, the elicitation of those factored utility functions is a major bottleneck in applying this model. Dariusz Braziunus and Craig Boutilier give a survey on elicitation methods in their article "Elicitation of Factored Utilities." They also show how to make good decisions if the utility function is not completely known and the gains for acquiring additional preference information might not outweigh the elicitation costs.

Preference elicitation is also the topic of the final article "User-Involved Preference Elicitation for Product Search and Recommender Systems" by Pearl Pu and Li Chen. The article returns to interactive systems and gives practical guidelines for eliciting preferences in those systems. The purpose is to prompt the user to give preference information by means of examples and targeted information. The central ideas are to show diverse examples to stimulate user critiques and to avoid extreme examples that are not of interest to the user. There is much more

Symposium on Abstraction, Reformulation and Approximation

SARA

<http://tidel.mie.utoronto.ca/sara09/>

The SoCS and SARA symposia will be co-located together with IJCAI in Southern California.

July 6-10, 2009

Lake Arrowhead, California
Submission Deadline: March 1

SoCS

**Symposium
on
Combinatorial
Search**

<http://www.search-conference.org/>

work on learning human preferences, and much of that falls under the rubric of data mining. We recommend the surveys on preference learning given in Fürnkranz and Hüllermeier (2005) and the Netflix Prize.¹

In this issue, we focused on preference handling for artificial intelligence. Preferences are of interest for many other branches of computer science, from very large scale integration (VLSI) design to databases and human-computer interfaces. And preferences are, of course, a subject of active research in decision analysis, operations research, and related fields. Simon French and colleagues provide a modern introduction to decision making and analysis (French, Maule, and Papamichail 2009) and Alexis Tsoukiàs

(2008) discusses the history of decision analysis and its interaction with artificial intelligence.

This special issue only scratches the surface of preferences in AI. We encourage you to continue your investigation into the articles here and those mentioned in the references. Furthermore, we invite you to submit articles and to attend events on advances in preference handling at upcoming conferences. The last AAAI-sponsored workshop on preference handling in AI was held at the July, 2008 AAAI conference in Chicago, Illinois. In 2009, preference handling will be the topic of the multidisciplinary conference on Algorithmic Decision Theory, which is planned to be held in Venice, October 2009.

Note

1. www.netflixprize.com/

References

- Brafman, R. I., and Domshlak, C. 2008. Preference Handling: An Introductory Tutorial. *AI Magazine* 30(1).
- Doyle, J. 2004. Prospects for Preferences. *Computational Intelligence* 20(2): 11–136.
- Fishburn, P. C. 1964. *Decision and Value Theory*. New York: Wiley.
- French, S.; Maule, J.; and Papamichail, N. 2009. *Decision Behaviour, Analysis, and Support: Decision Making and How Analysis and Information Systems May Support This*. Cambridge, UK: Cambridge University Press.
- Fürnkranz, J., and Hüllermeier, E. 2005. Preference Learning. *Künstliche Intelligenz* 19(1): 60–61.
- Junker, U. 2008. Preference-Based Problem Solving for Constraint Programming. In *Recent Advances in Constraints: CSCP 2007, LNAI 5129*, ed. F. Fages, F. Rossi, and S. Soliman, 94–111. Berlin: Springer.
- Keeney, R. L., and Raiffa, H. 1976. *Decisions with Multiple Objectives*. New York: John Wiley and Sons.
- Sandholm, T. 2007. Expressive Commerce and Its Application to Sourcing: How We Conducted \$35 Billion of Generalized Combinatorial Auctions. *AI Magazine* 28(3): 45–58.
- Sleeman, D. H., and Brown, J. 1982. *Intelligent Tutoring Systems*. Boston: Academic Press.
- Tsoukiàs, A. 2008. From Decision Theory to Decision Aiding Methodology. *European Journal of Operational Research* 187(1): 138–161.
- Walsh, T. 2007. Representing and Reasoning with Preferences. *AI Magazine* 28(4): 59–70.

Judy Goldsmith received her Ph.D. from the University of Wisconsin-Madison in 1988 and is a professor of computer science at the University of Kentucky, where she is affiliated with the cognitive science programs. Her current research is on preferences, reasoning with uncertainty, communicating about policies in a stochastic world, and computational complexity.

Ulrich Junker received his Ph.D. in computer science from the University of Kaiserslautern in 1992. His long-term research interests include the study of preferences for automated problem solving. Junker is a distinguished scientist at ILOG.