Artificial Intelligence Applied to Assistive Technologies and Smart Environments: Papers from the 2015 AAAI Workshop

I Prefer to Eat . . .

Thomas E. Allen
University of Kentucky
www.cs.uky.edu/∼teal223/
thomas.allen@uky.edu

Abstract

In this challenge paper, we consider the importance of preferences in smart homes and assistive environments and discuss the potential application of models and algorithms developed within the computational preferences community. We suggest the value of future research collaborations.

The Importance of Preferences in Smart Environments

A chime sounds and Joe's wristwatch vibrates, reminding him that it is time to eat. Joe was recently diagnosed with early-stage Alzheimer's disease and also suffers from depression. In his case these interfere with executive function, making it difficult for him to make decisions, keep track of the time, and make and carry out plans. However, with the help of an intelligent home environment, Joe is still able to live independently. The house knows what is in Joe's refrigerator and pantry. Furthermore, the kitchen has already compiled a list of recipes that are "feasible": Joe can assemble those dishes and has actually been known to eat them. They have also been screened for ingredients that are incompatible with his medications or violate other dietary constraints. With the help of a system similar to that described by Bouchard et al. (2014), Joe can prepare a suitable meal. Still, he has difficulty making decisions. How should the choice be narrowed down?

It is well established that assistive environments should be personalized (see, for example, Augusto et al. [2013]). They should learn, and then know, their inhabitants' preferences. Suppose that Joe's kitchen has some sort of representation of his preferences. It would know whether Joe wants a single option or the sense of control that comes from choosing from a list of suggestions. It may even be able to model Joe's mood and cognitive state at a given point in time and be able to adjust to his "good" and "bad" days.

What is involved in representing preferences? The house must be able to recognize context—his good and bad days, the activities he is engaged in, whether guests are present, the season, day of the week, and so on—and to reason with that context. Context can include the combinations of dishes, or even what Joe had for lunch the previous day. He might prefer the same thing every morning for breakfast, but a greater variety for his evening meal. He may prefer some sort of "comfort food" on cold, rainy days, but something relatively light after working in the garden on a hot summer day. He may always prefer that certain foods, such as peas and carrots, occur together. On particularly challenging days he may be altogether incapable of preparing anything more complicated than a prepackaged meal.

Preference Representations

Preferences often have a combinatorial structure. It is conceivable that Joe could choose from dozens, if not hundreds, of recipes and other consumable food items (frozen meals, canned items, condiments, etc.). Moreover, these could be combined in almost any imaginable fashion to compose a given meal. The number of possible meals is thus exponential in the number of consumable items. Furthermore, if context—previous meals, the state of the house and its inhabitant, and so on—is included in the preference representation, the number of possible outcomes could be vast indeed. It is unlikely that we could even list all of the possibilities explicitly. We thus need some sort of factored representation of the possibilities that avoids the problem of having to treat every possibility as an atom. That is, we need to represent outcomes according to their *features*.

However, even with a factored representation, it would be impossible to elicit from Joe a ranking over all possible meal outcomes. An alternative approach is to use some sort of rating system—for example, one in which each meal with a rating of $\star\star\star$ is preferred to each with a rating of $\star\star$; however, this approach has its difficulties also. First, there is the problem of consistency of ratings (from day to day and person to person). Conceivably, when he is despondent, Joe may give *everything* the lowest possible rating. Furthermore, we are likely to be left with a lot of ties for the top rating, making it difficult to recommend a most preferred choice or list of top choices of manageable size. Finally, ratings do not necessarily reflect rankings. Consider that Joe could give a top rating to meals that he knows would win high approval from his doctor, family, and friends, but nonetheless be happier most days with poutine or even the humble peanut butter and jelly sandwich—dishes to which Joe awards only two stars. For similar reasons, we understand that Netflix is no longer relying primarily on ratings for its recommendations.

Issues such as these have led to the development of *compact formalisms* for representing preferences. *General additive independence (GAI) value functions* (Bacchus and Grove 1995; Fishburn 1999) assign numerical value to the anticipated *utility* of each outcome to the subject. However, not all utility functions are additive. Moreover, it is not always clear how to assign numerical values to human preferences. For example, Joe may know that he prefers tomato bisque soup to clam chowder without being able to quantify such a preference. Preferences can also be modeled compactly as constraint satisfaction problems (CSPs) through methods that employ *soft constraints* (Bistarelli, Montanari, and Rossi 1997; Bistarelli et al. 1999; Rossi, Venable, and Walsh 2011). An advantage of this approach is that it allows for the use of constraint solvers. However, as with GAI value functions, it introduces the problem of quantifying preferences that are more naturally expressed in a qualitative way.

Conditional preference networks (CP-nets) (Boutilier et al. 2004) offer a qualitative approach. CP-nets consist of a node for each feature in the model—e.g., a binary variable indicating inclusion of a menu item or multivalued variable representing a category (e.g., beverage) or state of the world (e.g., the day of the week or Joe's mood). An arc from one node to another indicates that the preference over the latter features depends in general on the value of the former. Such preferences are encoded for each node in *conditional preference tables* specifying the preference for some or all assignments to the parent variables. Such preferences over the values of features take the form of rules such as, "If it is a cold day, I prefer hot tea to iced tea, all else being equal." CP-nets have garnered significant interest within the computational preferences community and have been proposed for applications including cybersecurity (Bistarelli, Fioravanti, and Peretti 2007), automated negotiation (Aydoğan et al. 2013), and interest-matching in social networks (Wicker and Doyle 2007). A number of extensions have been proposed, including *PCP-nets*, a probabilistic variant closely related to Bayesian networks (Cornelio et al. 2013; Bigot et al. 2013). Many problems involving CP-nets and their variants, including learning a CP-net that is consistent with comparison data and using the resulting CP-net to determine which of two arbitrary outcomes is preferred, are known to be computationally hard in the worst case (Goldsmith et al. 2008; Lang and Mengin 2009). However, the problems of finding the *most preferred* and k-*best* outcomes can be solved in polynomial time in the number of features when domains are constrained—for example, when certain items are unavailable (Boutilier et al. 2004). This is precisely what is needed for us to recommend one or more best choices to Joe.

Conclusions

Compact preference models such as CP-nets and soft constraints offer considerable potential for customizing assistive technologies and smart environments. We believe this potential should be explored in collaborative research in which preference models and algorithms are integrated in prototype environments and evaluated, first with simulated data (or data repurposed from previous experiments) and ultimately with human subjects.

References

Augusto, J. C.; Callaghan, V.; Cook, D.; Kameas, A.; and Satoh, I. 2013. Intelligent environments: a manifesto. *Human-Centric Computing and Information Sciences* 3(1):1–18.

Aydoğan, R.; Baarslag, T.; Hindriks, K. V.; Jonker, C. M.; and Yolum, P. 2013. Heuristic-based approaches for CP-nets in negotiation. In *Complex Automated Negotiations: Theories, Models, and Software Competitions*. Springer. 113–123.

Bacchus, F., and Grove, A. 1995. Graphical models for preference and utility. In *Proceedings of the Eleventh conference on Uncertainty in artificial intelligence*, 3–10. Morgan Kaufmann Publishers Inc.

Bigot, D.; Fargier, H.; Mengin, J.; and Zanuttini, B. 2013. Probabilistic conditional preference networks. *Septièmes Journées d'Intelligence Artificielle Fondamentale* 57.

Bistarelli, S.; Montanari, U.; Rossi, F.; Schiex, T.; Verfaillie, G.; and Fargier, H. 1999. Semiring-based CSPs and valued CSPs: Frameworks, properties, and comparison. *Constraints* 4(3):199–240.

Bistarelli, S.; Fioravanti, F.; and Peretti, P. 2007. Using CP-nets as a guide for countermeasure selection. In *Proc. 2007 ACM Symposium on Applied Computing*, 300–304. ACM.

Bistarelli, S.; Montanari, U.; and Rossi, F. 1997. Semiring-based constraint satisfaction and optimization. *Journal of the ACM (JACM)* 44(2):201–236.

Bouchard, B.; Bouchard, K.; and Bouzouane, A. 2014. A smart range helping cognitively-impaired persons cooking. In *Twenty-Sixth IAAI Conference*.

Boutilier, C.; Brafman, R.; Domshlak, C.; Hoos, H.; and Poole, D. 2004. CP-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *Journal of Artificial Intelligence Research* 21:135–191.

Cornelio, C.; Goldsmith, J.; Mattei, N.; Rossi, F.; and Venable, K. B. 2013. Updates and uncertainty in CP-nets. In *26th Australasian Joint Conference on Artificial Intelligence*.

Fishburn, P. 1999. Preference structures and their numerical representations. *Theoretical Computer Science* 217(2):359–383.

Goldsmith, J.; Lang, J.; Truszczynski, M.; and Wilson, N. 2008. The computational complexity of dominance and consistency in CP-nets. *J. Artif. Intell. Res.(JAIR)* 33:403–432.

Lang, J., and Mengin, J. 2009. The complexity of learning separable ceteris paribus preferences. In *IJCAI-09*, 848–853. San Francisco, CA, USA: Morgan Kaufmann.

Rossi, F.; Venable, K. B.; and Walsh, T. 2011. A Short Introduction to Preferences: Between Artificial Intelligence and Social Choice. *Synthesis Lectures on Artificial Intelligence and Machine Learning* 5(4):1–102.

Wicker, A. W., and Doyle, J. 2007. Interest-matching comparisons using CP-nets. In *Proc. AAAI*.