

Research Statement - Paolo Viappiani

paolo.viappiani@gmail.com
<http://www.cs.toronto.edu/~paolo/>

VISION

Personalization can help people to better deal with the complexities of choices in their lives. Personalized preference-aware interactive systems are in great demand in many domains: for instance, ambient intelligence systems can make our homes more secure, comfortable and better organized, cognitive assistants help the elderly in their daily activities, tutoring systems facilitate learning, adaptive hypermedia alter their content according to the user. The interest in personalization is evident the most in the web, where in recent years there has been a surge in research on personalized recommender systems; these systems include information search, music and shopping recommendation and personalized newspapers.

In the future, these personalized tools will make more and more use of Artificial Intelligence techniques, such as machine learning, constraint processing and probabilistic inference. They will need to adapt to external circumstances, be proactive, and reason about uncertainty. In particular, the quality of these personalized services crucially depends on the underlying model of the users preferences, and the acquisition of such model has to be performed in a discreet and possibly unobtrusively manner, in order to minimize user effort, while maximizing performance given the (limited) acquired information. The task of preference elicitation is made even more challenging by the idiosyncratic nature of human interaction. Behavioral decision theorists have long studied how biases such as *framing* and *anchoring* often arise in many common circumstances, where human behavior deviates from that of a pure rational agent. Effective systems necessarily require the design of quantitative models to account for these phenomena, in accordance with findings from decision theory and behavioral psychology.

Moreover, from an orthogonal prospective, the web is increasingly *social* as user interactions take place in online social networks. This means that new opportunities arise for computational tools that leverage the information sources provided by online communities and agents that provide recommendations for groups of individuals. The development of effective strategies of recommendation in a social contest is an exciting new research direction.

Motivated by these challenges, my research aims at fostering our understanding of the problems related to automatic personalization and at providing both theoretical and practical solutions (models, formalisms, algorithms and prototypes).

RECENT RESEARCH ACHIEVEMENTS

Preference elicitation and conversational recommender systems have been the central focus of my PhD and postdoctoral research. The acquisition of a preference model (preference elicitation) is a challenging problem as users have limited cognitive abilities. It is unrealistic, in the setting of a web-interaction to ask the user to answer a large number of cognitively complex questions (as a classic method of utility elicitation would do). However, limited knowledge about the user is often enough to provide near optimal recommendations: this intuition is behind the idea of *adaptive* utility (or preference) elicitation.

Conversational systems can effectively produce recommendations and help the user in stating his preferences with limited effort. In the last few years, I have considered two main orthogonal and complementary approaches: the first based on critiquing and probabilistic reasoning, the second based on minimax regret optimization. The advantage of the first is that critiquing is easy for users to use and understand, and can prevent behavioral decision biases as *the anchoring effect*. The advantages of the second are the ability of giving worst-case

guarantees to the quality of the recommendations and the possibility of optimizing *query sets* with respect to the *ex post* quality of a recommendation (analogue of the value of the information without the necessity of obtaining probabilistic prior). These approaches are presented in the next sections.

1. Example-critiquing with Adaptive Suggestions

For my doctoral dissertation at the *Ecole Polytechnique Federale de Lausanne* (EPFL), I have explored the topic of preference elicitation in *preference-based search*, defined as the problem of finding a target item from an electronic collection available on line.

Biases typical of human decision making, such as *means-objectives*, *anchoring* and *prominence effect* (studied by behavioral decision theory) can arise when a user is using a web tool to search for specific items in an electronic catalog. With user studies, I showed that common web search interfaces induce the users in stating incorrect preferences due to *means-objective*, leading to poor decision accuracy (~25%).

Better preference models can be acquired when the users are allowed to *critique* real examples presented to them. Indeed, *more accurate preference models are obtained when preferences are expressed on users' own initiative*, supported by behavioral decision theory. According to our experimental study, most of the preferences (79%) emerge from positive critiques that identify an opportunity that the user had not considered before.

To stimulate the expression of correct preferences, we proposed *example-critiquing with model-based suggestions*, a conversational framework that enables users to incrementally construct preference models by critiquing shown examples and that suggests items to stimulate preference expression (based on reasoning about preference uncertainty and Bayes learning), where the uncertainty over the user model is represented by probabilistic distributions over the possible preferences. The use of suggestions results (according to our user studies) in a dramatic improvement of *decision accuracy*, improving the chance of identifying the user's most preferred item.

Our model-based suggestions make the user aware of its true preferences, supporting the psychological process of *preference construction*. Such suggestions can be adapted, learning from the user's past actions (using Bayes inference). User studies showed that interactive tools with suggestions provided by our model achieve higher decision accuracy than traditional product search interfaces (up to 70%). We also discussed efficient implementation of interactive tools for Preference-based search in practice, both in databases and in configurable catalogs .

- P. Viappiani, B. Faltings and P. Pu. [Preference-based Search using Example-Critiquing with Suggestions](#). *JAIR* (Journal of Artificial Intelligence Research), 27, 2006, p. 465-503.
- Paolo Viappiani, Pearl Pu, Boi Faltings: [Preference-based search with adaptive recommendations](#). *AI Communications* 21(2-3): 155-175 (2008)
- P. Viappiani, B. Faltings and P. Pu. [Evaluating Preference-based Search Tools: a Tale of Two Approaches](#). *Proceedings of the Twenty-first National Conference on Artificial Intelligence (AAAI-06)*, Boston, MA, USA, July, 2006, pp. 205-211
- P. Pu, P. Viappiani and B. Faltings. [Increasing User Decision Accuracy using Suggestions](#). *ACM conference on Human factors in computing systems (CHI06)*, Montreal, Canada, April, 2006, pp. 121-130.

2. Regret-based Utility Elicitation and Optimization

My postdoctoral research (with Prof. Craig Boutilier) focused on merging the advantages of utility-based elicitation (i.e. with a precise, sound and quantitative semantics) with *open-*

ended interactions that let the users express their preferences in a way that is natural to them. We considered how to integrate example-critiquing and utility-based elicitation, and how to make recommendation in problems with subjective features. These are discussed next.

2.1 Optimal Recommendation Sets

We designed a regret-based recommendation system and developed several adaptive elicitation strategies that acquire an utility model with limited user effort. Minimax regret optimization is a criterion for "robust" decision-making, that bounds the worst case loss with respect to the actual best recommendation. Moreover, it is an efficient driver of adaptive utility elicitation, as the minimax regret current solution gives insight about which part of the utility space is more critical (and should be elicited first).

While we focused on recommender systems, many of our advances in minimax regret optimization are general and can be applied to different problems (such as regret-based planning). We formulated the problem of minimax regret optimization for database problems as an adversarial search problem with alpha-beta cuts, and our algorithm for online computation of minimax regret allows fast (acceptable in an interactive settings) computation for up to several thousand items.

We considered recommendation strategies for optimal recommendation sets and developed a new decision criterion, *setwise minimax regret* to select recommendations that are both optimal given the current information *and* maximize regret reduction when incorporating user's feedback (equivalent of EVOI for strict uncertainty).

We performed experiments comparing our regret-based model with traditional approaches based on heuristic utility functions and similarity-based navigation, achieving a dramatic improvement versus the state of the art.

- P. Viappiani and C. Boutilier. [Regret-based Optimal Recommendation Sets in Conversational Recommender Systems](#). *The 3rd ACM Conference on Recommender Systems (RecSys 2009)*

2.2 Preference Elicitation with Subjective Features

We considered preference elicitation with subjective features: *user-defined features* that differ from the standard, fixed in advance, catalog attributes. We cast this as a concept learning problem, where the system needs to learn just enough about the concept in order to give a good recommendation.

We give a minimax regret formulation for feature elicitation and develop strategies for both choosing the query to ask ("Does this item belong to the concept?") and meta-strategies to decide when to focus on eliciting the utility parameters and when focus on concept learning. The latter are divided in phased strategies (the system learns first the concept and then the utility) and interleaved elicitation strategies (the system simultaneously learns about the concept and the preferences). Results show better performance of the interleaved strategies; moreover our strategies outperform traditional concept learning techniques as the "halving" algorithm.

- C. Boutilier, K. Regan and P. Viappiani. [Simultaneous Elicitation of Preference Features and Utility](#). *To appear in the proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI-10)*, Atlanta, USA, July, 2010.
- C. Boutilier, K. Regan and P. Viappiani. [Online Feature Elicitation in Interactive Optimization](#). *Proceedings of the Twenty-Sixth International Conference on Machine Learning (ICML 2009)*, Montreal, June, 2009, p.10.

FUTURE DIRECTIONS

Artificial Intelligence aims in general at providing services and make decisions on behalf of an user. However, the desired agent behaviors, as well as the user needs and preferences are rarely, if never, completely specified upfront. In other words, AI agents need to optimize their plans and actions with respect to *uncertain objectives*. This is true for many different domains, as for example, recommender systems, personal agents and cognitive assistants, robotic, electronic commerce applications, personalized marketings and computational advertisement.

The long term goal of my research is to provide personalized tools that can help people to solve complex decision problems, while optimizing user-machine interaction, learning from the environment and adapting to unforeseen circumstances. I am interested in all of the following: developing new machine learning algorithms for emerging applications, providing new algorithms and solutions for user adaptation and personalization, understanding how people interact with intelligent systems, and building working prototypes.

In the future I want to extend my research in several directions. One such direction is the elicitation of preference models in virtual social network (such as Facebook and LinkedIn) for *social-aware* recommendation systems.

A peculiar phenomenon is the *cascading of preferences in social networks* (the effect caused by the influence of a node towards its neighbors). While much work has been dedicated to studying this phenomenon in both qualitative and quantitative ways, only recently researchers in AI and machine learning have started to think about how to actively exploit cascading. The design of effective recommendation strategies that optimally triggers the cascade effect (in order to maximize adoption) is currently an open problem. We are interested in the cascading effect with a fine-grained representation of user preferences (differently from the current literature), where the peers are associated with a well-defined utility function; the system needs to acquire both information about the user and information about the network structure in order to make effective recommendations.

Recommendation systems also need to account for the trade-off between user satisfaction and revenue. Quite surprisingly current systems as collaborative filtering focus only on rating prediction, neglecting that electronic commerce sites are self-interested, and must obey to budget constraints (as well to management preferences). This opens up opportunities for investigation at the intersection between personalization, multi-agent systems and game theoretic analysis. For this project, I am considering collaborations with EPFL, the university where I obtained my PhD.

As discussed before, the performance of a recommender system crucially depends on the user model. Technologies such as eye and gaze tracking can acquire additional information about an user in an unobtrusively manner. I am interested in integrating such tools (that are increasingly becoming inexpensive, thus available to a larger public) in recommender and decision-support systems. The challenge design of effective profiling techniques and machine learning algorithms in order to take advantages of these sensory information to improve the user model. For this research direction, I am looking for partners in HCI labs.

As mentioned in the previous section, minimax regret is a powerful tool for adaptive utility elicitation. I am considering a number of extensions of my previous work on adaptive utility elicitation (continuing to collaborate with the AI group at the University of Toronto):

1. I plan to give a formulation of regret-based optimization for ranking problems and multi-agent systems. For instance, consider computational advertisement: the system must produce an ordered list of ads, where the items in the first positions have more weight (because of prominence). The optimization needs to account for the preferences (utilities) of the advertisers and the available information about the current user to decide which ads to show in which position. Regret-based ranking would produce the allocation that minimizes worst-case loss with respect to uncertainty over the agents' utilities.

-
2. I am currently considering hybrid probabilistic and regret models. The advantages are manifold. We can account for user response errors and biases in a principled way (using probabilistic behavioral models from discrete choice theory), yet we maintain the advantages of elicitation with strict uncertainty (robust guarantees, no prior needed).
 3. However, when prior are available and user responses can be represented probabilistically, elicitation and recommendation can be formulated from a full Bayesian perspective. Current ongoing work on efficient selection of EVOI-optimal query strategies in very large datasets. I am also interested in prototypes and experiments with actual decision problems, to test adaptive utility elicitation in real scenarios.
 4. *Decision-theoretic planning* for autonomous agents and robotic systems (either (either regret-based or Bayesian). Instead of a single, predetermined reward function, the agent needs to optimize its action with respect to an uncertain quantitative. This research will include investigations over reinforcement learning strategies with uncertain rewards and transition probabilities, multi-attribute value functions and multi-objective optimization. I am also interested in practical strategies and user interfaces for eliciting the preferred robot behavior.

Another future interest is the study of contextual preferences and the relations between preferences, beliefs and intentions from a formal (logic-based) point of view. To this end, I expect to conduct research in collaboration with the LAMSADE group of the French University of Paris-Dauphine. The idea is to give a well defined semantics to contextual preferences and a formal description of the decision biases that occur in human decision making.

From a more practical prospective, I envision several applications of personalized artificial intelligence. One domain is that of personalized entertainment: video-games whose difficulty is tailored to the particular user, maximizing user enjoyment instead of match outcome. Ubiquitous personalization in mobile environments and ambient intelligence constitute attractive domains as well.

Attached to Paolo Viappiani's application (paolo.viappiani@gmail.com).

Updated on April, 23, 2010.

SELECTED PUBLICATIONS

These are my 5 most representative publications. *For all publications, see the next page.*

C. Boutilier, K. Regan and P. Viappiani. [Simultaneous Elicitation of Preference Features and Utility](#). *To appear in the proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI-10)*, Atlanta, USA, July, 2010 [acceptance rate: 28%]

- *A new model for preference elicitation, where the user can have preferences on individual "subjective features" (such as a "cool car"), casted as a concept learning problem. The system simultaneously learns the user preferences and the features on where they are defined.*
- *Among the very few papers to be accepted for both oral and poster presentation: according to the conference organizers, "posters are being used to highlight exceptional papers, ones that merit attention in addition to oral presentation"*

P. Viappiani, B. Faltings and P. Pu. [Preference-based Search using Example-Critiquing with Suggestions](#). *JAIR (Journal of Artificial Intelligence Research)*, 27, 2006, p. 465-503. [acceptance rate of JAIR submissions in 2006: 18%]

- *This article shows how personalized suggestions based on probabilistic reasoning can tackle behavioral issues in decision-making; resulting in better decisions.*
- *JAIR is considered among the most influential journals for Artificial Intelligence. As of April, 23, 2010, this article has been cited 47 times according to Google Scholar.*

P. Viappiani, B. Faltings and P. Pu. [Evaluating Preference-based Search Tools: a Tale of Two Approaches](#). *Proceedings of the Twenty-first American Conference on Artificial Intelligence (AAAI-06)*, Boston, MA, USA, July, 2006, pp. 205-211. [acceptance rate 19%]

- *AAAI is a top-tier Computer Science International conference, and one of the leading conferences in Artificial Intelligence. This paper, describing a substantial part of my PhD thesis, has been cited 17 times (as of April, 23, 2010)*

P. Viappiani and C. Boutilier. [Regret-based Optimal Recommendation Sets in Conversational Recommender Systems](#). *The 3rd ACM Conference on Recommender Systems (RecSys 2009)*, pp. 101-108. [acceptance rate for long papers: 17%]

- *This paper describes advances of regret-based preference elicitation, resulting in effective acquisition of the user's utility. We show how, when adopting minimax regret as criterion, the best recommendation set is also the best query set. ACM RecSys is the emerging conference for research in recommendation systems.*

B. Faltings, P. Pu, M. Torrens and P. Viappiani. [Designing Example-critiquing Interaction](#). *ACM International Conference on Intelligent User Interfaces (IUI 2004)*, Madeira, Portugal, January, 2004, pp. 22-29. [acceptance rate 28%]

- *This paper, presented at one of the main UI conferences, has been cited 43 times (as of April, 23, 2010) according to Google Scholar.*

Book chapters

Pearl PU, Boi Faltings, Li Chen, Jiyong Zhang, and Paolo Viappiani. [Usability Guidelines for Product Recommenders based on Example Critiquing Research](#). Handbook on Recommender Systems, Kantor, P.B.; Ricci, F.; Rokach, L.; Shapira, B. (Eds.), Springer. ([TO APPEAR](#))

Referred Journal Papers

Paolo Viappiani, Pearl Pu, Boi Faltings: [Preference-based search with adaptive recommendations](#). *AI Communications* 21(2-3): 155-175 (2008)

B. Peintner, P. Viappiani, and N. Yorke-Smith. [Preferences in Interactive Systems: Technical Challenges and Case Studies](#). *AI Magazine* 29(4), 13-24, Winter 2008.

P. Viappiani, B. Faltings and P. Pu. [Preference-based Search using Example-Critiquing with Suggestions](#). *JAIR* (Journal of Artificial Intelligence Research), 27, 2006, p. 465-503.

Refereed Publications in International Conferences with Proceedings

C. Boutilier, K. Regan and P. Viappiani. [Simultaneous Elicitation of Preference Features and Utility](#). *To appear in the proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI-10)*, Atlanta, USA, July, 2010. [acceptance rate: 28%; among the very few papers to be accepted for oral and poster presentation, 'to highlight exceptional papers']

P. Viappiani and C. Boutilier. [Regret-based Optimal Recommendation Sets in Conversational Recommender Systems](#). *The 3rd ACM Conference on Recommender Systems (RecSys 2009)*, pp. 101-108. [acceptance rate for long papers: 17%]

C. Boutilier, K. Regan and P. Viappiani. [Preference elicitation with Subjective Features](#). *The 3rd ACM Conference on Recommender Systems (RecSys 2009)*, pp. 341-344. [acceptance rate for short papers: 28%]

C. Boutilier, K. Regan and P. Viappiani. [Online Feature Elicitation in Interactive Optimization](#). *Proceedings of the Twenty-Sixth International Conference on Machine Learning (ICML 2009)*, Montreal, June, 2009. [acceptance rate for long papers: 24%]

P. Viappiani, P. Pu and B. Faltings. [Conversational Recommenders with Adaptive Suggestions](#). *Recommender Systems 2007*, Minneapolis, Minnesota, USA, October, 2007, pp. 89-96.

P. Viappiani and B. Faltings. [Design and Implementation of Preference-based Search](#). *The 7th International Conference on Web Information Systems Engineering (WISE)*, LNCS4255, Wuhan, China, October, 2006, pp. 72-83. [acceptance rate 20%]

P. Viappiani, B. Faltings and P. Pu. [Evaluating Preference-based Search Tools: a Tale of Two Approaches](#). *Proceedings of the Twenty-first American Conference on Artificial Intelligence (AAAI-06)*, Boston, MA, USA, July, 2006, pp. 205-211. [acceptance rate 19%]

P. Pu, P. Viappiani and B. Faltings. [Increasing User Decision Accuracy using Suggestions](#). *ACM conference on Human factors in computing systems (CHI06)*, Montreal, Canada, April, 2006, pp. 121-130. [acceptance rate 23%]

P. Viappiani and B. Faltings. [Implementing Example-based Tools for Preference-based Search](#). *The Sixth International Conference on Web Engineering (ICWE 06)*, ACM, Menlo

Park, California, USA, July, 2006, pp. 89-90.

P. Viappiani, B. Faltings and P. Pu. [The Lookahead Principle for Preference Elicitation: Experimental Results](#). *Seventh International Conference on Flexible Query Answering Systems (FQAS)*, Milan, Italy, June, 2006, pp. 378-389.

B. Faltings, P. Pu, M. Torrens and P. Viappiani. [Designing Example-critiquing Interaction](#). *ACM International Conference on Intelligent User Interfaces (IUI 2004)*, Madeira, Portugal, January, 2004, pp. 22-29. [acceptance rate 28%]

Communications in International Workshops and Symposia (refereed)

Paolo Viappiani and C. Boutilier. [Optimal Set Recommendations Based on Regret](#). *IJCAI'09 Workshop on Intelligent Techniques for Web Personalization & Recommender Systems (ITWP'09)*, Pasadena, California, USA, July, 2009

P. Viappiani and B. Faltings. [Preference-based search for Configurable Products](#). *AAAI'07 Configuration Workshop*. Vancouver, Canada, July, 2007.

P. Viappiani and B. Faltings. [Adaptive Suggestions for Example based Tools](#). *ECAI'06 Multidisciplinary Workshop on Advances in Preference Handling*, Riva del Garda, Italy, August, 2006, pp. 134-140.

P. Viappiani, B. Faltings, V. Schickel-Zuber and P. Pu. [Stimulating Preference Expression Using Suggestions](#). *Mixed-Initiative Problem-Solving Assistants, AAAI Fall Symposium Serie*, November, 2005, pp. 128-133.

Viappiani, B. Faltings, V. Schickel-Zuber and P. Pu. [Stimulating Preference Expression Using Suggestions](#). *IJCAI-05 Multidisciplinary Workshop on Advances in Preference*, Edinburgh, UK, August, 2005, pp. 186-191.

P. Viappiani, P. Pu and B. Faltings. [Acquiring User Preferences for Personal Agents](#). *Personalized Agents. AAAI Fall Symposium Serie*, Cape Code, MA, USA, October, 2002, pp. 53-59.

Dissertations

P. Viappiani. [Preference-based search with suggestions](#). PhD Thesis 3876, Ecole Polytechnique Fédérale de Lausanne (EPFL), 2007. (*nominee for the EPFL Best Thesis Award*)

P. Viappiani. Modellizzazione di Agenti Razionali. Tesi di Laurea. Politecnico di Milano, 2002.

Technical Reports

J. Zhang, P. Pu and P. Viappiani. [A Study of User's Online Decision Making Behavior](#). Technical Report, Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland, December, 2006.

P. Viappiani and B. Faltings. [Design and Implementation of Preference-based Search](#). Technical Report LIA-REPORT-2006-004, Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland, May, 2006.

Unpublished

J. Zhang, P. Pu and P. Viappiani. *Supporting Online Decision Making with Anchoring Examples*. 2006.

