Research Statement - Paolo Viappiani

paolo.viappiani@cs.aau.dk
http://www.cs.aau.dk/~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).

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. During my PhD thesis I have considered an approach based on critiquing and qualitative reasoning, able to give personalized suggestions in order to stimulate preference expression.

More recently, I have considered principled approaches based on an explicit representation of the user's possible utility function, considering two distinct frameworks for utility elicitation: strict uncertainty with *minimax regret* recommendations and *Bayesian* inference. Our seminal result (discussed in point 2 and 5 below) makes it possible to characterize the *optimal query set* (the query that gives most information) for both frameworks.

In the following I review my research achievements with respect to these different approaches.

I also present results from recent works on using machine learning techniques for the optimization of human computation.

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 c*ritique* 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.

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)*

3 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. *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.

4 Bayesian Utility Elicitation: Optimal Queries and Optimal Recommendation Sets

Bayesian approaches to utility elicitation assign priors to the value of utility parameters; these distributions are updated according to observations of user responses. In this setting (myopic) Expected Value of Information (EVOI) is a natural criterion for selecting queries. However, EVOI-optimization is usually computationally prohibitive. We examined EVOI optimization using choice queries, queries in which a user is ask to select her most preferred product from a set. We showed that, under very general assumptions, the optimal choice query w.r.t. EVOI coincides with the optimal recommendation set, that is, a set maximizing the expected utility of the user selection. Since recommendation set optimization is a simpler, *submodular* problem, this can greatly reduce the complexity of both exact and approximate (greedy) computation of optimal choice queries. We also considered the case where user responses to choice queries are e*rror-prone* (using both constant and *mixed multinomial logit noise models*) and provide w*orst-case* guarantees. Finally we developed a local search technique for query optimization that works extremely well with large outcome spaces.

These results have been awarded a prestigious "spotlight presentation" at NIPS 2010, and, togheter with the previous results related to regret-based elicitation, have been accepted for presentation at AAAI 2011 in the "New Scientific and Technical Advances in Research (NECTAR)" track.

- Paolo Viappiani and Craig Boutilier. Optimal Bayesian Recommendation Sets and Myopically Optimal Choice Query Sets. In *Advances in Neural Information Processing Systems 23 (NIPS)*, 2352--2360, Vancouver, BC, Canada, 2010.
- Paolo Viappiani and Craig Boutilier. Recommendation sets and choice queries: There is no exploration/exploitation tradeoff! Procceedings of the Twenty-fifth AAAI Conference on Artificial Intelligence (AAAI-11), NECTAR track San Francisco, CA, USA, 2011.

5 AI-based techniques for "Human Computation"

Crowd-sourcing, or Human Computation, is the act of outsourcing a problem to a group or a community. Human "experts" are used to solve problems that present difficulties for algorithmic methods (examples include Amazon's Mechanical Turk, the ESP game). We considered a Bayesian approach to the problem of interactively learning complex concepts using crowd-sourcing. In our setting, the system needs to select both the queries to ask and the expert to ask them to. We proposed recommendation techniques, inference methods, and query selection strategies to assist a user charged with choosing a configuration that satisfies some (partially known) concept. Our model is able to simultaneously learn the concept definition and the "types" (representing different skill levels) of the experts.

In a way analogous to preference elicitation, learning is most effective when one focuses on the parts of the problem that are most relevant: the system needs to select the experts whose answers are (predicted to be) as informative as possible. However, we note that there might value in asking a query to an expert other than that predicted to be most "knowledgeable" because we may learn more about the types of other experts. Adopting ideas from Reinforcement Learning, we model the *exploration/exploitation* trade-off, where e*xploration* in this setting means learning more about the types of the experts, while *exploitation* means focusing on learning more about the concept.

We have evaluate dour model with simulations, showing that our Bayesian strategies are effective even in large concept spaces with many uninformative experts. In particular, the strategy that explicitly models the exploration/exploitation tradeoff is most effective.

- Paolo Viappiani, Sandra Zilles, Howard J. Hamilton, and Craig Boutilier. Learning complex concepts using crowdsourcing: A Bayesian approach. In *Proceedings of the Second Conference on Algorithmic Decision Th*eory (ADT-11}, Piscataway, NJ, USA, 2011.

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 c*ascading 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 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. 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 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. ***(current collaboration with Jorge Baier, Pontificia Universidad Catolica, Chile)***

5. I am interested in interactive sequential optimization problems such as extensions of multi-armed bandits where one can actively "query" an expert about the value of a given arm (using *comparison queries* of the kind: ``among these two arms, which one do you think is best?''). Since querying is associated with a cost, one needs to choose both when to query and which query to ask (if no query is asked, an arm is pulled). The total reward is cumulative, therefore asking the right query is crucial. I am particularly interested in the case of the presence of several experts (of different reliability levels) simultaneously. A possible application is Reinforcement Learning with an active demonstrator. ***(current collaboration with INRIA Sequel, France)***

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.