CAP Research and technological achievements

The 5 years of the research project have been extremely productive, culminating in more than 50 publications at top-tier conferences and in journals as well as one book which was published by Cambridge Press.

We developed models and algorithms that **enable automated systems to argue when negotiating and interacting with people**. The goal of the arguer agent is that the outcome of the negotiation should be according to its own preferences. We studied game theory-based methods and knowledge-based methods in complex settings and we considered both. We used machine learning methods to predict human decision-making. We presented a methodology of how to integrate human models into the agents' optimization problems in order to generate agents' strategies for interaction proficiently with people. We will present our main findings here:

1. We developed the NegoChat agent, the first negotiation agent that successfully negotiates with people in natural language. The agents' strategies address partial agreements and issue-by-issue interactions. NegoChat's negotiation algorithm is based on bounded rationality, specifically Aspiration Adaptation Theory. We performed a rigorous evaluation of NegoChat, showing its effectiveness.
2. Continuing working under the assumption that automated agents should be able to persuade people in the same way people persuade each other - via dialogs - we presented a novel methodology for persuading people through argumentative dialogs. Our methodology combines theoretical argumentation modeling, machine learning and Markovian optimization techniques that together result in an innovative agent named SPA. Two extensive field experiments, with more than 100 human subjects, show that SPA is able to persuade people significantly more often than a baseline agent and no worse than people are able to persuade each other. To the best of our knowledge this is the first automated argumentative agent that was tested by actually arguing with people.
3. People's cultural background has been shown to affect their negotiation behavior. We considered bilateral negotiations with an alternating-offer protocol and settings that allow parties to choose the extent to which they keep each of their agreements. Our methodology combines a decision theoretic model with classic machine learning techniques in order to predict people's behavior. In extensive experiments conducted in a Colored Trails game, the agent was able to outperform people in the USA, Israel and Lebanon. We then moved to the more

complex setting of a 3-player game on Colored Trails, which is analogous to a market setting in which players need to reach an agreement and commit to or renege from contracts over time. We identified equilibrium strategies for the players. Our experiments showed that the agent using equilibrium strategies for the customer role was able to outperform people playing the same role in Israel and the USA. In contrast, the agent playing the role of the service provider was not able to outperform people. Analysis reveals that this difference in performance is due to the equilibrium proposals being significantly beneficial to the customer players as well as the irrational behavior of human customer players in the game. We then developed a new agent extending the equilibrium agent in two ways in order to handle the uncertainty that characterizes human play in negotiation settings. First, it used a risk adverse strategy using a convex utility function. Second, it reasoned about a possibly bounded rational customer taking into account the country in which it plays. The new agent plays better than people in all countries.

4. We proposed the introduction of an agent to assist in detecting and incriminating a deceptive participant in a chat-room. We designed a game where deception occurs in a text based discussion environment. In this game several participants attempt to collectively detect a deceptive member. We composed an automated agent which participates in this game as a regular player. The goal of the agent is to detect the deceptive participant and alert other members, without raising suspicion itself. We used machine learning on the data collected from human players to design this agent. Extensive evaluation of our agent showed that it succeeds in raising the players' collective success rate in catching the deceptive player.
5. We addressed two main challenges of generating persuasion strategies: How to reveal information to people, and what actions to propose to them. We utilized the framework of a persuasion game, but integrated bounded rational learned models of people. We ran experiments with people in two games and showed the advantages of our approach.
6. We also considered the problem of an automated agent which needs to influence the decision of a group of self-interested agents that must reach an agreement on a joint action. For example, consider an automated agent which aims to reduce the energy consumption of a nonresidential building by convincing a group of people who share an office to agree on an economy mode of air-conditioning and low light intensity. We presented four problems that address issues of minimality and safety in the persuasion process.
7. We also considered a logic-based approach. We studied a variation of Boolean games and assumed that there is some set of variables, the values of which are not directly accessible to players. The players have their own beliefs about these variables, and make decisions about what actions to perform based on these beliefs. The communication takes the form of announcements about the values of

some environmental variables; the effect of an announcement is the modification of the beliefs of the players who hear the announcement so that they accurately reflect the values of the announced variables. By choosing announcements appropriately, the agent can steer the situation away from certain outcomes and towards others.

8. We contributed to the use of automated agents for training people in negotiations and interviewing.

    a. First we studied whether agents can be used for training. There was inconclusive evidence whether practicing tasks with computer agents improves people's performance on these tasks. We studied this question empirically using extensive experiments involving bilateral negotiation and three player coordination tasks played by hundreds of human subjects. We used different training methods for subjects, including practice interactions with other human participants, interacting with agents from the literature, and asking participants to design an automated agent to serve as their proxy in the task. Following training, we compared the performance of subjects when playing against state-of-the-art agents from the literature. The results revealed that in the negotiation settings, in most cases, training with computer agents increased people's performance as compared to interacting with people. In the three player coordination game, training with computer agents increased people's performance when matched with the state-of-the-art agent. These results demonstrate the efficacy of using computer agents as tools for improving people's skills when interacting in strategic settings, saving considerable effort and providing better performance than when interacting with human counterparts.
    b. We developed a Virtual-Suspect system designed for use in police interrogation training simulations. The system allows users to preconfigure various scenarios, as well as different suspect histories and personality types. The responses given by the Virtual-Suspect during the interrogation are selected based on context and the suspect's psychological state, which changes in response to each interrogator's statement. Experiments have shown that the Virtual-Suspect's responses in an interrogation scenario are similar to those of a human respondent. We also developed a virtual body for the VS.

9. In most of the agents we designed models to predict people's decision-making since humans are not necessarily rational decision-makers. The main approach that we developed is the integration of machine learning methods with social sciences methods in order to overcome the relatively small number of training examples one can collect about people's decision-making. We survey here a few of our proposed approaches for predicting and modeling people's behavior.

a. We presented the Linear Weighted-Utility Quantal response human model, which relies on the following two assumptions: 1) Linear Weighted-Utility, i.e. people's subjective utility is a linear combination of attributes, and 2) Logit quantal response: The probability that people will choose an action is proportional to the action's subjective utility.
b. For Cliff-Edge (CE) situations we proposed a new meta algorithm, deviated virtual learning (DVL), which extends existing methods to efficiently cope with environments comprising a large number of alternative decisions at each decision point and where the agent interacts repeatedly with people with similar behavior.
c. We considered exponential decay time series for applications where the time series behavior can change over time due to a change in the user's preferences or a change of environment. We presented an innovative online learning algorithm, which we named Exponentron, for the online prediction of exponential decay time series.
d. We presented a novel synthesis between Collaborative Filtering methods and machine learning classification algorithms to create a fast learning algorithm, CRISP. CRISP exploits the similarities between users in order to apply data from known users to new users, therefore requiring less information on each person. Results from user studies in interruption management indicate that the algorithm significantly improves users' performances in completing the task and their perception of how long it took to complete each task.
e. We proposed a method for predicting people's strategic decisions based on their facial expressions when they have an incentive to hide their decision. We showed that our method outperforms standard SVM as well as humans in predicting subjects' strategic decisions. To the best of our knowledge, this is the first study to present a methodology for predicting people's strategic decisions when there is an incentive to hide facial expressions.