Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI 2015)

Automated Agents for Advice Provision

Ariel Rosenfeld
Bar-Ilan University
Ramat-Gan, Israel
rosenfa5@cs.biu.ac.il

Abstract

In this thesis, we focus on automated advising agents. The advice given is a form of relating recommendations or guidance from an automated agent to its human user. Providing the right advice at the right time is extremely complex, and requires a good adaptation to human desires and changing environments. We propose a novel methodology for designing automated advising agents and evaluate it in three real world environments. Our intelligent advising agents were evaluated through extensive field trials, with hundreds of human subjects. A significant increase in human performance as well as a high level of user satisfaction was recorded when they were equipped with our agents.

1 Introduction

It is well known that people's cognitive abilities are limited. Consequently, people have been shown to seek advice in order to improve the outcomes of their (imperfect) decisions [Yaniv, 2004]. That is, a person maintains her decision autonomy but includes others in her decision process.

In this thesis, we focus on automated advising agents. The advice, in our setting, is a form of relating recommendations or guidance from an automated agent to its human user. We focus our attention on agents that offer their advice based solely on the observed behavior of the users and the environment. This setting is referred to as *implicit interaction*. That is, the agent refrains from requesting explicitly defined information from the user.

However, providing beneficial advice in an automated fashion can be very challenging. This complex human-agent interaction involves the modeling of human behavior in changing environments and in human-agent interplay itself.

We consider three real world settings in which we believe automated advising agents can enhance people's performance and capabilities;

1. Advice in Argumentative Dialogs: The agent offers advice to its user while she engages in an argumentative dialog with another human.
2. Advice in Automotive Climate Control Systems (CCSs): The agent offers advice to a human in order to minimize her CSS's energy consumption while keeping the driver comfortable.
3. Advice in Robot Team Collaborations: The agent offers advice to a human operator who operates a large team of semi-autonomous, low-cost robots.

Per setting, we use the following methodology for designing and implementing automated advising agents. First, we build a formal model which is phrased as an optimization problem for the advising agent. Then, we collect contextual data on real world activity using human subjects and physical machines (when needed). Using the obtained data, and based on literature from social science, psychology, and human decision-making studies, we model the behavior of the humans and machines operating in the agent's environment by constructing prediction models using machine learning techniques. Then, based on the prediction models and the optimization problem, the agent finds the most suitable advising policy to achieve its goal: improving the user's performance.

The advising agents are designed, implemented and evaluated in extensive human experiments in *realistic* real-world scenarios using hundreds of human subjects, electric vehicles and simulated and physical robots.

2 Advice in Argumentative Dialogs

An automated agent can help a human when engaging in an argumentative dialog by utilizing its knowledge and computational advantage to provide her with arguments.

First, we examined the well-established Argumentation Theory (see [Walton, 2009]) and its abilities to predict people's arguments. In [Rosenfeld and Kraus, 2014] we examined three experimental settings, with over 130 human conversations and 140 questionnaires, varying in complexity, which show the lack of predictive power of the existing Argumentation Theory. Second, we use Machine Learning (ML) techniques to provide a probability distribution over all known arguments given a partial deliberation. That is, our ML techniques provide the probability of each argument to be used next in a given dialog. Our model achieves 76% accuracy when predicting people's top three argument choices given a partial deliberation. Last, using the prediction model and the newly introduced heuristics of relevance, we designed and evaluated the Predictive and Relevance-based Heuristic agent (PRH). Through an extensive human study with over 200 human subjects, we show that the PRH agent outperforms other agents that propose arguments based on Argumentation Theory, predicted arguments without heuristics or only the heuristics on both people's satisfaction from agents and people's use of the suggested arguments.

Our full results are reported in [Rosenfeld and Kraus, 2014; 2015].

3 Advice in Automotive Climate Control Systems

There is a broad consensus that modifying a driver's behavior may allow a reduction in vehicle fuel consumption [Gonder *et al.*, 2012]. Inspired by these results, we propose an automated agent that advises a driver how to set the car's CCS such that energy consumption is reduced while keeping the driver comfortable. In a repeated interaction environment, the agent offers the driver *adaptive* advice.

While the agent is mainly concerned with the car's energy consumption, the driver is usually more interested in her own comfort level, which changes from one driver to another. This partial conflict, which may occur, is a challenge in the advice provision process—the agent must consider the driver's comfort in order to provide reasonable advice. Furthermore, the agent should take into account the long term effect of each piece of advice.

Using machine learning techniques we were able to identify predicative features which describe the driver's preferences, the environmental factors and a representation of the trust between the driver and the agent, all of which affect her decision-making. This allowed us to construct a personalized driver model for predicting drivers' reactions to advice depending on its content, the context in which it was given and the previous agent-driver interactions.

We used Markov Decision Process modeling, which utilizes the aforementioned models to account for the dynamic and changing environment in which the agent operates, and solved it to calculate an advising policy. Our experiments, with 83 human drivers, show that subjects equipped with our agent significantly saved energy compared to both competing agents. Subjects equipped with MACS reduced their CCS energy consumption by 33% (on average) compared to subjects using a non-advising agent (Silent).

Our full results are reported in [Azaria *et al.*, 2015; Rosenfeld *et al.*, 2015b].

4 Advice in Robot Team Collaborations

The number of multi-robot systems deployed in field applications has risen dramatically over the years. Nevertheless, supervising and operating multiple robots at once is a difficult task for a single operator to execute.

We present a novel methodology that enhances operators' performance by using an intelligent advising agent. The agent provides advice to the operator regarding which actions she should take and acts as a smart filter between the robots' requests and the human operator. Our methodology is not restricted to any certain hardware or algorithm used by the robots, and we consider these factors constants. We first present the *Myopic Advice Optimization Problem*. The optimization problem models the maximization of the operator's performance by selecting when and which advice to provide in a greedy fashion. Then, we evaluate our approach using the Search And Rescue (SAR) task (see [Liu and Nejat, 2013]).

We have extensively tested the performance of our advising agent in both simulated environments (using the Gazebo robot simulation toolbox) and physical deployment (using 10 Hamster robots) with 44 non-expert operators. Experimental results show that our advising agent was able to enhance the operators' performance when managing a team of 10 mobile robots in terms of the task's goals (finding and correctly classifying green objects in a clustered terrain) and reduced the cognitive workload reported by the operators in two distinct SAR environments: an office building floor and a simulated urban open terrain. On average, while equipped with our agent, an operator covered 19.5% more terrain, detected and correctly classified 16.7% more desired objects, reported 7% less workload and reduced 22.3% of the robots' idle time, compared to his benchmark performance without our agent while operating 10 mobile robots. Despite training with simple simulated environments, our agent has shown that it is able to provide solid advice in both complex simulated environments and physical deployment (a short video is available at *http://vimeo.com/119434903*).

Our full results are reported in [Rosenfeld *et al.*, 2015a].

References

[Azaria *et al.*, 2015] Amos Azaria, Ariel Rosenfeld, Sarit Kraus, Claudia V. Goldman, and Omer Tsimhoni. Advice provision for energy saving in automobile climate control systems. In *AI-Magazine*, 2015.

[Gonder *et al.*, 2012] Jeffrey Gonder, Matthew Earleywine, and Witt Sparks. Analyzing vehicle fuel saving opportunities through intelligent driver feedback. Technical report, SAE, 2012.

[Liu and Nejat, 2013] Yugang Liu and Goldie Nejat. Robotic urban search and rescue: A survey from the control perspective. *Journal of Intelligent & Robotic Systems*, 72(2):147–165, 2013.

[Rosenfeld and Kraus, 2014] Ariel Rosenfeld and Sarit Kraus. Argumentation theory in the field; an empirical study of fundamental notions. In *ArgNLP*, 2014.

[Rosenfeld and Kraus, 2015] Ariel Rosenfeld and Sarit Kraus. Providing arguments in discussions based on the prediction of human argumentative behavior. In *AAAI*, 2015.

[Rosenfeld *et al.*, 2015a] Ariel Rosenfeld, Noa Agmon, Oleg Maksimov, Amos Azaria, and Sarit Kraus. Intelligent agent supporting human-multi-robot team collaboration. In *IJCAI*, 2015.

[Rosenfeld *et al.*, 2015b] Ariel Rosenfeld, Amos Azaria, Sarit Kraus, Claudia V. Goldman, and Omer Tsimhoni. Adaptive advice in automobile climate control systems. In *AAMAS*, 2015.

[Walton, 2009] Douglas Walton. Argumentation theory: A very short introduction. In *Argumentation in artificial intelligence*, pages 1–22. Springer, 2009.

[Yaniv, 2004] Ilan Yaniv. Receiving other peoples advice: Influence and benefit. *Organizational Behavior and Human Decision Processes*, 93(1):1–13, 2004.