

Satisfaction-Maximizing Autonomous Agents

Scientific Abstract (PIs Sarne and Aumann 3007/24)

The design of autonomous agents that work alongside human users, or carry out tasks on their behalf, is a major field of study, both in academia (primarily within the Autonomous agents and multi-agent systems community (AAMAS)) and industry. Academics and industry experts alike are continually probing new paradigms and architectures to augment the capabilities of these agents, emphasizing their efficiency, adaptability, and contextual applicability. A crucial aspect in the design of agent-based systems is the clear definition of their goals and the criteria for measuring success. This foundational step establishes the agent's operational framework and guides its decision-making processes. In much of the existing systems, these goals are often quantified using some "objective" success criteria, which are measurable and directly tied to the agent's performance, e.g., task completion duration, net profit achieved, quality of the outcome obtained. Alas, prior literature has commonly demonstrated that actual user satisfaction with the agent-based system is not directly correlated with measures of its final outcome. In particular, in settings of *repeated interaction*, where the task of interest involves a sequence of decisions to be made or multiple interactions with the user, it has been shown that user satisfaction is greatly affected by the *dynamics* of the interaction.

The proposed research focuses on designing effective methods for enhancing user-satisfaction in agent and multi-agent systems in settings of repeated interaction. In particular, the two main objectives are: (i) study and develop prediction models for user satisfaction from autonomous agents in collaborative human-agent settings, and (ii) develop general and domain-specific mechanisms, methods and protocols (termed "agent strategies") for MAS designers for increasing satisfaction (both aggregate and individual) of users in such environments. We plan to focus in two primary archetypal models: (1) When the agent is serving, working alongside or acting on behalf of a single user; (2) When the agent is serving or collaborating with two or more users. Both models can be considered a form of a multi-agent system (MAS), as the user/s and the autonomous agent can be seen as distinct agents within the system.

This shift from traditional performance-centric approaches to one that prioritizes user satisfaction in collaborative interactions with autonomous agents has the potential to greatly promote user-agent collaborations for more intuitive and rewarding interactions across diverse domains. The insights to be gained will not only benefit end-users but also provide practical guidance to industries involved in agent development, contributing to the creation of more successful and widely accepted products. Additionally, the research establishes a systematic, data-driven approach to understanding and forecasting user contentment, empowering AI researchers and data scientists with a strong foundation for deeper insights into factors influencing user satisfaction.