A Multi-Agent Game for Studying Human Decision-making (Demonstration)

Han Yu, Xinjia Yu, Su Fang Lim, Jun Lin, Zhiqi Shen, and Chunyan Miao
Joint NTU-UBC Research Centre of Excellence in Active Living for the Elderly (LILY)
Nanyang Technological University, 50 Nanyang Avenue, Singapore 639798
{han.yu, xyu009, sflim2, jlin7, zqshen, ascymiao}@ntu.edu.sg

ABSTRACT

Understanding how human beings delegate tasks to trustees when presented with reputation information is important for building trust models for human-agent collectives. However, there is a lack of suitable platforms for building large scale datasets on this topic. We describe a demonstration of a multi-agent game for training students in the practice of Agile software engineering. Through interacting with agent competitors in the game, the behavior data related to users' decision-making process under uncertainty and resource constraints are collected in an unobtrusive fashion. These data may provide multi-agent trust researchers with new insight into the human decision-making process, and help them benchmark their proposed models.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence - *Intelligent Agents*

Keywords

Reputation, trust, decision-making, bench-marking, game

1. INTRODUCTION

As agents are increasingly required to operate in socially complex settings, understanding how people delegate tasks to resource constrained trustees when presented with their reputation information is of significant importance to multi-agent systems (MASs) researchers [7, 8, 9]. However, there is a lack of datasets to help MAS researchers gain insight into this topic and benchmark proposed agent models.

In this paper, we describe a demonstration[1] of a multi-agent game based informatics platform - Agile Manager (AM) - to help researchers understand how people delegate tasks to trustee agents with different behavior patterns in an environment characterized by uncertainty. The game is targeted at people who wish to learn how to make efficient task allocation decisions in the Agile software development (ASD) practice [6]. It is capable of unobtrusively collecting the players' in-game behavior data (including their mood, task allocation strategy in response to the changes in trustee agents' reputation and emotional state, mouse movement trajectories, etc.). It can potentially produce a large scale dataset with new sources of information to provide insight into human decision-making behaviors to MAS researchers, and build a benchmark for future agent trust models.

[1]`http://www.agelesslily.org/demo_agilemanager/`

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*

2. DEMONSTRATION CONTENT

The AM game is a simulated environment where a player can experience the challenges involved in managing an ASD team. The game interface and system architecture are shown in Figure 1. The player acts as the manager of an ASD team of ten programmers each controlled by an agent. Each level of the game consists of several Sprints of development. The player needs to allocate a number of tasks to the programmer agents (PAs) in each Sprint [6] to maximize the chance of success of the software project. Each task is characterized by its *value*, *difficulty*, *required effort*, and *deadline*.

Similar to real world ASD teams, each PA has different internal characteristics in terms of *competence*, *capacity*, and *emotion*. Competence is the probability that a PA can complete a given task with high quality. It is the ground truth information used by the game system to generate PA behaviors. It remains invisible to both the player and the artificial intelligence (AI) competitors. Capacity represents the maximum effort a PA can commit to working on assigned tasks during a given Sprint. It is presented to the player in the form of a workload bar as shown in Area C of Figure 1. In odd numbered levels of the game, more competent PAs also have larger capacities; whereas in even numbered levels, more competent PAs have smaller capacities. The emotion of a PA is calculated following the approach in [3]. It stems from the OCC cognitive theory of emotions. The resulting emotions are presented to the player in the form of the facial expression of a PA exemplified by Area A of Figure 1.

Each PA has characteristics which are affected by the interaction with the player. They include *reputation* and *current workload*. A PA's reputation depends on its past performance in completing the given tasks with high quality and on time. The higher the ratio between a PA's competence and the difficulty of the task assigned to it, the more likely it is for the task to be completed with high quality. A PA's reputation is computed following the $BRSEXT$ approach in [8]. It is presented to the player in the form of a star rating as shown in Area B of Figure 1.

A team of ten agents, with the same characteristics as the PAs controlled by the player, allocate the same set of tasks among themselves following the SWORD approach [10]. This

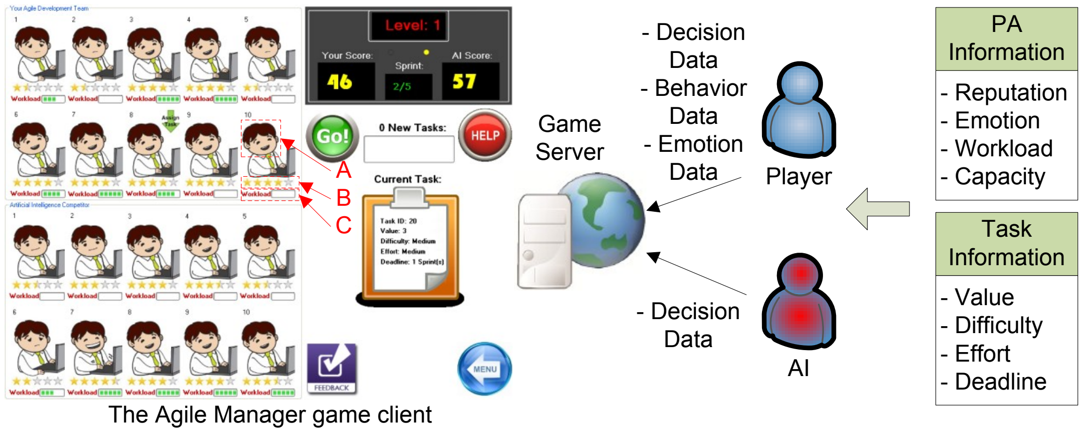

Figure 1: An overview of the *Agile Manager* game platform.

team is the AI competitor against the player. The player can benefit from the game by observing the AI's strategy to improve his/her decision-making.

3. DISCUSSIONS AND FUTURE WORK

As we aim to build a dataset on how people delegate tasks to resource constrained trustees when presented with their reputation information to support future research in multi-agent trust modeling, the AM game collects a wide range of data generated by the players. The types of data include:

1. every task allocation decision from the human and the AI players (i.e., time taken between successive decisions, which task was assigned to which PA, and snapshots of each PA's current situation including its perceived reputation and workload);
2. players' self reported strategies used to allocate tasks in each game session;
3. players' self reported emotions after each game session (i.e., a mix of 6 basic emotions [2]) and the corresponding facial expressions (captured with *AffectButton* [1]);
4. mouse movement and mouse click positions in the game interface during each game session;
5. effectiveness of persuasive computing techniques based on the Elaboration Likelihood Model (ELM) [5];
6. players' self reported profile information.

Beyond the work described here, the game is on track of being used in large scale user studies. The resulting dataset can potentially fill an important gap in existing reputation ratings datasets, such as *Epinions* [4], which focus on reflecting people's opinions on products instead of task allocation decisions. In subsequent research, we will incorporate features enabling the game to study how to optimize the usage of persuasive techniques for people from different backgrounds to influence their decisions.

4. ACKNOWLEDGMENTS

This research is supported in part by Interactive and Digital Media Programme Office, National Research Foundation hosted at Media Development Authority of Singapore (Grant No.: MDA/IDM/2012/8/8-2 VOL 01).

5. REFERENCES

[1] J. Broekens and W.-P. Brinkman. AffectButton: Towards a standard for dynamic affective user feedback. In *Proceedings of the 3rd International Conference on Affective Computing and Intelligent Interaction (ACII'09)*, pages 1–8, 2009.

[2] P. Ekman. An argument for basic emotions. *Cognition and Emotion*, 6(3-4):169–200, 1992.

[3] J. Lin, Ailiya, C. Miao, and Z. Shen. A FCM based approach for emotion prediction in educational game. In *Proceedings of the 7th International Conference on Computing and Convergence Technology (ICCCT'12)*, pages 980–986, 2012.

[4] P. Massa and P. Avesani. Trust-aware bootstrapping of recommender systems. In *Proceedings of ECAI 2006 Workshop on Recommender Systems*, pages 29–33, 2006.

[5] R. E. Petty and J. T. Cacioppo. *Communication and persuasion: central and peripheral routes to attitude change*. Springer-Verlag, 1986.

[6] F. Shull, T. Dyba, and T. Dingsoyr. What do we know about agile software development? *IEEE Software*, 26(5):6–9, 2009.

[7] H. Yu, C. Miao, B. An, C. Leung, and V. R. Lesser. A reputation management approach for resource constrained trustee agents. In *Proceedings of the 23rd International Joint Conference on Artifical Intelligence (IJCAI'13)*, pages 418–424, 2013.

[8] H. Yu, Z. Shen, C. Leung, C. Miao, and V. R. Lesser. A survey of multi-agent trust management systems. *IEEE Access*, 1(1):35–50, 2013.

[9] H. Yu, Z. Shen, C. Miao, and B. An. Challenges and opportunities for trust management in crowdsourcing. In *Proceedings of the IEEE/WIC/ACM International Conference on Intelligent Agent Technology (IAT'12)*, pages 486–493, 2012.

[10] H. Yu, Z. Shen, C. Miao, and B. An. A reputation-aware decision-making approach for improving the efficiency of crowdsourcing systems. In *Proceedings of the 12th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'13)*, pages 1315–1316, 2013.