Shih, JL. et al. (Eds.) (2023). Proceedings of the 31st International Conference on Computers in Education. Asia-Pacific Society for Computers in Education

Improving Teamwork through a Decision-Theoretic Coach in a Minecraft Search-and-Rescue Game

David V. PYNADATH[a*], Nikolos GURNEY[a], Sarah KENNY[a], Rajay KUMAR[a], Stacy C. MARSELLA[b], Haley MATUSZAK[b], Hala MOSTAFA[c], Pedro SEQUEIRA[d], Vokan USTUN[a] & Peggy WU[e]

[a]*Institute for Creative Technologies, University of Southern California, USA*
[b]*Northeastern University, USA*
[c]*Bitsight Technologies, USA*
[d]*SRI International, USA*
[e]*Raytheon Technologies Research Center , USA*
*pynadath@usc.edu

Game-based learning provides a promising methodology for teaching skills such as teamwork (Martín-Hernández et al., 2021; Riivari et al., 2021; Syynimaa et al., 2022). Studies of human teamwork have identified various team processes that underlie good joint task performance (e.g., Marks et al., 2001). The multiagent research community has operationalized many of these findings into domain-independent algorithms that separate task-specific knowledge from reusable task-independent teamwork knowledge (e.g., Tambe, 1997). These frameworks were successfully deployed in facilitating teamwork in real-world human teams when supplied with task-specific knowledge (e.g., Chalupsky et al., 2001).

We present here an agent that focuses exclusively on teamwork-level variables in deciding what interventions to use in coaching a human team in a game-based search-and-rescue environment in Minecraft (Corral et al., 2021). The game is an urban search and rescue (USAR) task that involves clearing and avoiding hazards while rescuing victims of a disaster. The task is performed by distributed teams of three participants, possibly with a human or agent as a coach. The coaches are not acting team members, but they are able to communicate, via text chat, with the human team members, who themselves can directly engage each other in high-volume, unconstrained voice communication.

Our agent does not directly observe the game environment or the people in it, but instead relies on input from analytic components (ACs) (developed by other research teams) that process environmental information and output only teamwork-relevant measures. Our agent models teamwork variables and updates its beliefs over them using PsychSim, which provides reusable AI technology for agents capable of populating game environments (Pynadath & Marsella, 2005). PsychSim integrates two AI technologies: recursive models (Gmytrasiewicz & Durfee, 1995) and decision-theoretic reasoning in the form of POMDPs (Kaelbling et al., 1998). PsychSim has been used to build interactive training games for a variety of soft skills, such as negotiation (Kim et al., 2009) and avoiding risky behavior under social pressure (Miller et al., 2011). For the current search-and-rescue task, we built an agent that assesses the team it is currently observing and chooses interventions to best assist them. Details on our agent and some objective measures of its purposes are available in an earlier publication (Pynadath et al., 2023). In this paper, we look deeper into some specific observations of human behavior when interacting with our agent in the game.

Our agent used team-process variables (Marks et al., 2001) whose measures were recently validated (Mathieu et al., 2020), but only those relevant to the game environment:

- **Affect management:** "activities that foster emotional balance, togetherness, and effective coping with stressful demands and frustration."

- **Coordination:** "the process of synchronizing or aligning the members' actions."
- **Motivating and confidence building:** "activities that develop and maintain members' motivation and confidence while working toward team goals."
- **Systems monitoring:** "activities such as tracking team resources (e.g., money) and factors in the team environment (e.g., inventories) to ensure that the team has what it needs to accomplish its goals and objectives."
- **Team monitoring:** "members assisting others in the performance of their tasks (by providing feedback or coaching or assisting with the task itself)."

None of these processes are directly observable. We thus model them as hidden variables that our agent seeks to improve through the following candidate interventions:

- **Reflection:** Between trials, the agent prompts the team with a reminder about a situation in the first trial where a player was stuck on a threat plate for an inordinate amount of time.
- **Cheerleading:** The agent congratulates a player on successfully achieving a goal (more specifically, moving a victim to a triage area).
- **Report drop:** The agent reports on a noteworthy lack of performance by one player (more specifically, failure to respond to outstanding requests by a teammate).
- **Notify phase (early):** The agent reminds the team that it is still early in the game, so exploration should be valued more.
- **Notify phase (late):** The agent reminds the team that it is getting late in the game, so exploration should be valued less.
- **Remind of best practices:** The agent suggests that someone help a player who has a number of outstanding, but unaddressed, requests.
- **Prompt activity:** The agent asks about any possible issues upon observing that the team has been predominantly idle for a period of time.

Data for thirteen teams were collected at Arizona State University. One immediate question was whether the teams even noticed our agent's interventions. For one team, five of eight interventions were verbally acknowledged; for two others, none were. For all other teams, one or two interventions were explicitly acknowledged. Breaking down these counts by intervention type, none of the *cheerleading* interventions were acknowledged. The two *notify phase* interventions were the most commonly acknowledged (14 of 38), while the various request/activity based interventions were less frequently acknowledged (8 of 38).

The post-trial survey asked a few questions about the participants' opinions of their coach. Not surprisingly, the human coaches received much higher trust scores than the agent coaches. All of the agent coaches were distrusted, but ours was distrusted the least.

Beyond these aggregate results, we found it necessary to examine the actual behaviors of individual teams, in addition to listening to the dialogue within those teams. The relatively small number of teams makes this in-depth examination feasible, although it also potentially limits the generalizability of the observations made.

One recurring issue was a disconnect between what our agent perceived as requests vs. what the team members actually intended as requests. Participants could mark blocks in Minecraft (e.g., for victims, for obstructed rooms), making an implicit request (e.g., to evacuate/triage the victim inside, to clear the obstruction). However, marking these blocks was sometimes used as implicit communication or even just a mnemonic device, not a request. This difference of perspective led to occasional confusion when our agent intervened to stimulate satisfaction of "requests". For example, in one trial, our agent performed a *remind of best practices* intervention regarding Red's requests, but Red then expressed no awareness of any requests when asked by another team member to clarify.

As a more positive example, there was another team where our agent also chose *remind of best practices* for Red and where a team member also asked for clarification. However, in this case, Red responded by explicitly identifying a victim, its type, and its location. The other team member then saw this victim and triaged it. Our agent's intervention thus triggered the desired shared situational awareness and subsequent task completion.

A more ambiguous example starts off just as these first two, with a team member asking for clarification of Red's requests after an agent intervention. In this case, Red admits to accidentally forgetting to pick up a marker. Despite our agent’s mistaken reaction to a nonexistent request, the communication prompted by its intervention preempted the misunderstanding that would have occurred if the discussion had not taken place.

As another example of the evaluation challenge, our agent performed the late version of its “notify phase” interventions while the team was in the middle of discussing its strategy. There was a pause while the speaker considered the agent's message, but then the team resumed discussing and executing its current plan, which was in conflict with our agent's recommendation. The team was clearly not complying, but it is not clear whether they were better or worse off as a result. Thus, while the lack of compliance is objectively measurable, any wider-scope learning outcomes are harder to evaluate.

Acknowledgements

This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. W911NF-20-1-0011 and by the U.S. Army Research Laboratory (ARL) under contract number W911NF-14-D-0005. Any opinions, findings, or recommendations do not necessarily reflect the views of DARPA or the US Government.

References

Chalupsky, H., Gil, Y., Knoblock, C. A., Lerman, K., Oh, J., Pynadath, D. V., Russ, T. A., Tambe, M., et al. (2001). Electric elves: Applying agent technology to support human organizations. *Proceedings of the Innovative Applications of Artificial Intelligence Conference*, 1, 51–58.

Corral, C. C., Tatapudi, K. S., Buchanan, V., Huang, L., & Cooke, N. J. (2021). Building a synthetic task environment to support artificial social intelligence research. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 65(1), 660–664.

Gmytrasiewicz, P. J., & Durfee, E. H. (1995). A rigorous, operational formalization of recursive modeling. *Proceedings of the International Conference on Multiagent Systems*, 125–132.

Kaelbling, L. P., Littman, M. L., & Cassandra, A. R. (1998). Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1-2), 99–134.

Kim, J. M., Hill Jr, R. W., Durlach, P. J., Lane, H. C., Forbell, E., Core, M., Marsella, S., Pynadath, D.V. & Hart, J. (2009). BiLAT: A game-based environment for practicing negotiation in a cultural context. *International Journal of Artificial Intelligence in Education*, 19(3), 289-308.

Marks, M. A., Mathieu, J. E., & Zaccaro, S. J. (2001). A temporally based framework and taxonomy of team processes. *Academy of Management Review*, 26(3), 356–376.

Martín-Hernández, P., Gil-Lacruz, M., Gil-Lacruz, A. I., Azkue-Beteta, J. L., Lira, E. M., & Cantarero, L. (2021). Fostering University Students’ Engagement in Teamwork and Innovation Behaviors through Game-Based Learning (GBL). *Sustainability*, 13(24), 13573.

Mathieu, J. E., Luciano, M. M., D’Innocenzo, L., Klock, E. A., & LePine, J. A. (2020). The development and construct validity of a team processes survey measure. *Organizational Research Methods*, 23(3), 399–431.

Miller, L. C., Marsella, S., Dey, T., Appleby, P. R., Christensen, J. L., Klatt, J., & Read, S. J. (2011). Socially optimized learning in virtual environments (SOLVE). *In Interactive Storytelling: Fourth International Conference on Interactive Digital Storytelling*, 4, 182-192.

Pynadath, D. V., Gurney, N., Kenny, S., Kumar, R., Marsella, S. C., Mostafa, H., Sequeira, P., Ustun, V., and Wu, P. (2023). “Effectiveness of teamwork-level interventions through decision-theoretic reasoning in a Minecraft search-and-rescue task”. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*.

Pynadath, D. V., & Marsella, S. C. (2004). Fitting and compilation of multiagent models through piecewise linear functions. *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems*, 1197–1204.

Syynimaa, K., Lainema, K., & Lainema, T. (2022). Identifying Challenges in Virtual Teams: A Case Study of Teamwork in a Game-Based Learning Environment. *International Association for Development of the Information Society*.

Riivari, E., Kivijärvi, M., & Lämsä, A. M. (2021). Learning teamwork through a computer game: for the sake of performance or collaborative learning? *Educational Technology Research and Development*, 69, 1753-1771.

Tambe, M. (1997). Towards flexible teamwork. *Journal of Artificial Intelligence Research*, 7, 83–124.