Effectiveness of Teamwork-Level Interventions through Decision-Theoretic Reasoning in a Minecraft Search-and-Rescue Task

Extended Abstract

David V. Pynadath
University of Southern California

Nikolos Gurney
University of Southern California

Sarah Kenny
University of Southern California

Rajay Kumar
University of Southern California

Stacy C. Marsella
Northeastern University

Haley Matuszak
Northeastern University

Hala Mostafa
Bitsight Technologies

Pedro Sequeira
SRI

Volkan Ustun
University of Southern California

Peggy Wu
RTRC

KEYWORDS

Human-agent teaming, POMDPs, Theory of Mind

ACM Reference Format:

David V. Pynadath, Nikolos Gurney, Sarah Kenny, Rajay Kumar, Stacy C. Marsella, Haley Matuszak, Hala Mostafa, Pedro Sequeira, Volkan Ustun, and Peggy Wu. 2023. Effectiveness of Teamwork-Level Interventions through Decision-Theoretic Reasoning in a Minecraft Search-and-Rescue Task: Extended Abstract. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 3 pages.

Autonomous agents offer the promise of improved human teamwork through automated assessment and assistance during task performance [15, 16, 18]. Studies of human teamwork have identified various processes that underlie joint task performance, while abstracting away the specifics of the task [7, 11, 13, 17]. We present here an agent that focuses exclusively on teamwork-level variables in deciding what interventions to use in assisting a human team. Our agent does not directly observe or model the 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 these teamwork variables and updates its beliefs over them using a Bayesian Theory of Mind [1], applying Partially Observable Markov Decision Processes (POMDPs) [9] in a recursive manner to assess the state of the team it is currently observing and to choose interventions to best assist them.

Our work makes use of the testbed and ACs developed by performers on DARPA's Artificial Social Intelligence for Successful Teams (ASIST) program. Experimentation in the ASIST program has used a simulated urban search and rescue (USAR) task that involves clearing and avoiding hazards while rescuing victims of a disaster [4, 5]. The experiments described in this investigation bring distributed teams of three participants together with an agent as an advisor. Our agent uses the following ACs as sensors:

Carnegie-Mellon University's (CMU's) TED (Team Effectiveness Diagnostic) generates four measures: collective effort, appropriate skill use, appropriate use of strategies, and communications[7].

CMU's BEARD (Background of Experience, Affect, and Resources Diagnostic) generates individual/team profiles of the following: anger, anxiety, gaming experience, Reading the Mind in the Eyes (RMIE)[2], and the Santa Barbara Sense of Direction (SBSOD)[8].

Cornell's Trust AC measures players' trust in their teammates, including the players' compliance behaviors.

GELP (Gallup Emergent Leadership Prediction) provides a measure of emergent leadership for each participant, derived from audio data, NLP, task competency scores, the intake survey data, etc. [10].

IHMC Joint Activity (by The Florida Institute for Human and Machine Cognition) uses a graph representation of task dependencies to report on changes in the team members' activity status.

Rutgers Utility Agent provides a set of measures related to the needs of the team members.

University of Central Florida's (UCF's) Player Profiler predicts team members' 'potential' for teamwork and task performance [3].

In this initial work, we use the following team process variables [11] whose measures were recently validated [12]:

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

Coordination: "synchronizing or aligning the members' actions."

Motivating and confidence building: "activities that develop and maintain members' motivation and confidence"

Systems monitoring: "tracking team resources…and factors in the team environment…to ensure that the team has what it needs to accomplish its goals and objectives".

Team monitoring: " assisting others in the performance of their tasks (by providing feedback or coaching…)".

Our agent is built within PsychSim. PsychSim provides reusable AI technology for generating multiagent systems capable of populating game environments [14]. These agents represent a decision-making model, similar to interactive POMDPs [6], that generates behavior by reasoning from a declarative representation of their

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom.

goals and beliefs. We manually constructed a dynamic influence diagram model that captures various hypothesized dependencies among team-level variables. The edges in this diagram capture correlations between team-process characteristics and AC-provided profiles of individual players and the team as a whole (e.g., a participant who scores high on CMU's BEARD anxiety scale would contribute to a lower expected affect management capability of the team). These also include effects of team-process characteristics on observed behaviors (e.g., a team with good coordination would have a higher likelihood of advancing their joint activities, as monitored by IHMC's AC). There are also trigger conditions for possible interventions as a function of AC variables. Finally, there are edges to capture the effects of our agent's candidate interventions on both team-process and AC variables (e.g., cheerleading is likely to increase the team's motivating process variable and CMU TED's measure of collective effort). The dependency structure and the weights on the links within it were selected based on background knowledge and significant guesswork.

The agent considers 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 early still, so exploration should be valued more.

Notify phase (late): The agent reminds the team that it is getting late, 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.

The agent's goal (as reflected in its reward function) is to increase the team-process variables. To incentivize our agent to choose its interventions, we introduce a positive effect on the team-process variables into our influence diagram. Our agent can then compute the expected reward of any applicable intervention using existing domain-independent algorithms [9].

Participant recruitment and data collection for thirteen teams were conducted at Arizona State University. One surprising result from the program's complete set of experiments was that there was no significant improvement in task performance with a human advisor over the no-advisor condition. However, it was still unfortunate to observe that teams working with our agent did the worst in terms of mission score (utility) and completing tasks started (error rate). Some of this may be attributable to the teams paired with our agent generally scoring lower on pre-trial game skill metrics than those paired with other teams. However, it is also the case that our agent performed the fewest interventions out of all of the advisors. Furthermore, it chose its interventions without any explicit modeling of the task. All of these factors no doubt contributed to poor scores on these two metrics. On the other hand, our agent had the best score for collective effort, even when compared against the human advisor.

We limited our agent to perform each intervention type only once per trial, allowing us to divide any trial where the intervention was performed into pre- and post-intervention phases. We use the trigger conditions for each intervention as its target behavior. For example, the activity completion frequency after cheerleading interventions was higher than that before for all trials, suggesting a positive effect of the intervention.

The "Report drop" intervention became applicable if there was a team member who had not satisfied any outstanding requests from another (when at least two such requests had been made). For the teams where our agent intervened, there were 46 nonresponsive events on average per trial, while for those where it did not, there were only 26. So on the whole, our agent chose to intervene for teams where nonresponsiveness was a more prevalent issue. The frequency of nonresponsiveness decreased after the intervention in 12 of the 18 trials in which our agent performed it, suggesting another positive effect.

The "Remind of best practices" intervention was applicable when a requestor had at least five outstanding requests total, regardless of requestee. It is thus an analogue of the "Report drop" intervention, which identifies a requestee with outstanding requests (although limited to a specific requestor). The frequency of the unfulfilled request trigger condition decreased after the intervention in only two of the 20 trials in which it was chosen, suggesting a negative result. One possible reason is that there was almost no correlation between the number of occurrences of this trigger condition with final score. In contrast, the "Report drop" condition did have a noticeable correlation with score. It is thus possible that this intervention was simply irrelevant to the team's goals and plans.

Given that the current model was created through manual input (and with a great deal of guesswork, despite the combined expertise of the team) and with no information about the task, it is pleasantly surprising that it was able to perform comparably to (and sometimes better than) the human and other agent advisors on some of the experimental metrics. This result suggests that the AC variables are indeed useful in informing advice to a human team, and that a domain-independent agent advisor is indeed feasible. Furthermore, our methodology provides a means to identify the relative value of different AC variables in triggering interventions (e.g., using the number of requests of A that it has not fulfilled was more useful in these experiments than the number of unfulfilled requests that B has made of others).

Of course, the relatively small number of data points makes it hard to know how much those findings will generalize. The more important contribution of the data described in the previous section will be in the refinement of the model. In particular, we can use the data to refine the dependencies captured in the current PsychSim dynamic influence diagram. As a result, our current agent represents a novel agent for incorporating team-process variables, social-science analyses from outside researchers, and human behavior data into a unified model capable of autonomously assisting a human team. Furthermore, it provides a platform for evaluating candidate variations on that model in terms of their strengths and weaknesses across different teams and tasks.

ACKNOWLEDGMENTS

The authors thank the research teams at CMU, Cornell, Gallup, IHMC, Rutgers, and UCF who developed the ACs that serve as input to our agent. Thanks also go to Nancy Cooke and her team at Arizona State University for their ongoing work on collecting data for the evaluation of the agent. The experiment was approved by all of the performers' Institutional Review Boards (IRBs) and the US Army Human Research Protections Office (HRPO). This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. W911NF-20-1-0011. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Defense Advanced Research Projects Agency (DARPA).

REFERENCES

[1] Chris Baker, Rebecca Saxe, and Joshua Tenenbaum. 2011. Bayesian theory of mind: Modeling joint belief-desire attribution. In *Proceedings of the Annual Meeting of the Cognitive Science Society*.

[2] Simon Baron-Cohen, Sally Wheelwright, Jacqueline Hill, Yogini Raste, and Ian Plumb. 2001. The "Reading the Mind in the Eyes" Test revised version: A study with normal adults, and adults with Asperger syndrome or high-functioning autism. *The Journal of Child Psychology and Psychiatry and Allied Disciplines* 42, 2 (2001), 241–251.

[3] Rhyse Bendell, Jessica Williams, Stephen M Fiore, and Florian Jentsch. 2021. Towards artificial social intelligence: Inherent features, individual differences, mental models, and theory of mind. In *International Conference on Applied Human Factors and Ergonomics*. 20–28.

[4] Christopher C Corral, Keerthi Shrikar Tatapudi, Verica Buchanan, Lixiao Huang, and Nancy J Cooke. 2021. Building a synthetic task environment to support artificial social intelligence research. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 65, 1 (2021), 660–664.

[5] Jared T Freeman, Lixiao Huang, Matt Woods, and Stephen J Cauffman. 2021. Evaluating artificial social intelligence in an urban search and rescue task environment. In *AAAI Fall Symposium on Computational Theory of Mind for Human-Machine Teams*.

[6] Piotr J Gmytrasiewicz and Prashant Doshi. 2004. Interactive POMDPs: Properties and preliminary results. In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems*, Vol. 3. 1374–1375.

[7] J Richard Hackman and Ruth Wageman. 2005. A theory of team coaching. *Academy of Management Review* 30, 2 (2005), 269–287.

[8] Mary Hegarty, Anthony E Richardson, Daniel R Montello, Kristin Lovelace, and Ilavanil Subbiah. 2002. Development of a self-report measure of environmental spatial ability. *Intelligence* 30, 5 (2002), 425–447.

[9] Leslie Pack Kaelbling, Michael L Littman, and Anthony R Cassandra. 1998. Planning and acting in partially observable stochastic domains. *Artificial intelligence* 101, 1-2 (1998), 99–134.

[10] Ellyn Maese, Pablo Diego-Rosell, Les DeBusk-Lane, and Nathan Kress. 2021. Development of emergent leadership measurement: Implications for human-machine teams. In *AAAI Fall Symposium on Computational Theory of Mind for Human-Machine Teams*.

[11] Michelle A Marks, John E Mathieu, and Stephen J Zaccaro. 2001. A temporally based framework and taxonomy of team processes. *Academy of Management Review* 26, 3 (2001), 356–376.

[12] John E Mathieu, Margaret M Luciano, Lauren D'Innocenzo, Elizabeth A Klock, and Jeffery A LePine. 2020. The development and construct validity of a team processes survey measure. *Organizational Research Methods* 23, 3 (2020), 399–431.

[13] Philip M Podsakoff, Scott B MacKenzie, Robert H Moorman, and Richard Fetter. 1990. Transformational leader behaviors and their effects on followers' trust in leader, satisfaction, and organizational citizenship behaviors. *The Leadership Quarterly* 1, 2 (1990), 107–142.

[14] David V Pynadath and Stacy C Marsella. 2005. PsychSim: Modeling theory of mind with decision-theoretic agents. In *Proceedings of the International Joint Conference on Artificial Intelligence*. 1181–1186.

[15] Sangwon Seo, Lauren R Kennedy-Metz, Marco A Zenati, Julie A Shah, Roger D Dias, and Vaibhav V Unhelkar. 2021. Towards an AI coach to infer team mental model alignment in healthcare. In *Proceedings of the IEEE Conference on Cognitive and Computational Aspects of Situation Management*. 39–44.

[16] Gita Sukthankar, Katia Sycara, Joseph A Giampapa, Chris Burnett, and Alun Preece. 2007. Towards a model of agent-assisted team search. In *Procceedings of the First Annual Conference of the International Technology Alliance in Network and Information Science*.

[17] Bruce W Tuckman and Mary Ann C Jensen. 1977. Stages of small-group development revisited. *Group & Organization Studies* 2, 4 (1977), 419–427.

[18] Sheila Simsarian Webber, Jodi Detjen, Tammy L MacLean, and Dominic Thomas. 2019. Team challenges: Is artificial intelligence the solution? *Business Horizons* 62, 6 (2019), 741–750.