Speculative Game Design of Asymmetric Cooperative Games to Study Human-Machine Teaming

Erik Harpstead
harpstead@cmu.edu
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA

Kimberly Stowers
kim.stwrs@gmail.com
Tim Fletcher Co.
Windsor, Ontario, Canada

Lane Lawley
llawley3@gatech.edu
Georgia Institute of Technology
Atlanta, Georgia, USA

Qiao Zhang
qzhang490@gatech.edu
Georgia Institute of Technology
Atlanta, Georgia, USA

Christopher MacLellan
cmaclellan3@gatech.edu
Georgia Institute of Technology
Atlanta, Georgia, USA

ABSTRACT

While recent advances in Artificial Intelligence and Machine Learning have demonstrated the potential for AI systems to outperform human experts in many domains, including games, AI systems still generally lack the ability to team with humans on complex tasks. One of the barriers to addressing this challenge is a lack of shared task domains in which to do basic research to study Human-Machine Teaming strategies. In our work, we employ speculative game design to create asymmetric cooperative games that can serve as test beds to study human-machine teaming challenges. In this paper, we will describe our general approach and detail the current state of our development efforts.

KEYWORDS

Human-Machine Teaming, Speculative Game Design, Asymmetric Cooperative Games

ACM Reference Format:
Erik Harpstead, Kimberly Stowers, Lane Lawley, Qiao Zhang, and Christopher MacLellan. 2023. Speculative Game Design of Asymmetric Cooperative Games to Study Human-Machine Teaming. In *Foundations of Digital Games 2023 (FDG 2023), April 12–14, 2023, Lisbon, Portugal.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3582437.3587200

1 INTRODUCTION

There have been incredible developments in the areas of Artificial Intelligence and Machine Learning (AI/ML) in recent years, with the advent of computational systems capable of outperforming the greatest human experts at both games, such as *Go* [10] and *StarCraft* [13], as well as real-world tasks such as those simulated within visual range air combat [2]. Despite these advancements, AI agents still generally lack fundamental capabilities necessary for effectively cooperating and coordinating with humans. Basic research is still needed to explore what the future of Human-Machine Teaming (HMT) will look like and what role computational systems will play in supporting human collaborators in the future. Additionally, foundational research is needed to develop human-compatible interactive AI agents that will enable fluid, individualized, and symbiotic adaptation between humans and agents and produce enhanced team performance over time and across tasks [11].

In our work, we aim to address these shortcomings by conducting research along three interconnected research strands: (1) speculative game design to create a battery of HMT challenge tasks, (2) human studies to better understand the semantics and meanings that human collaborators attach to their interactions with computational teammates, and (3) AI engineering work to create human-centered AI agents that can participate with humans in the HMT challenge tasks. In this paper we will describe our approach to the first strand of speculative game design along with some details on initial progress.

2 SPECULATIVE GAME DESIGN TO EXPLORE POSSIBLE HMT FUTURES

A central challenge for HMT research is thinking beyond what is possible with today's technology and envisioning the future. Further, we lack sharable benchmark tasks that could be used to evaluate current HMT progress and test new HMT capabilities. To overcome these barriers, we apply an approach of speculative game design [1] to create a suite of engaging game-based HMT tasks. Figure 1 demonstrates the core concept underlying our speculative design approach. Through the process of designing a collection of games, we will probe the space of possible, plausible, and probable futures to identify those that are preferable.

Within the broader landscape of game design we focus specifically on the design of asymmetric cooperative games [5], which involve a team of players working together in multiple asymmetric roles to accomplish a shared task. For example, a group of "field agent" players might coordinate with an "overwatch" player who has access to more information but cannot directly impact the game world. We take this approach for two reasons. First, it will allow us to use each game to explore multiple HMT scenarios by varying which role(s) the human(s) or AI(s) are fulfilling in the team. Following on from the prior example, we hypothesize that the nature of the HMT dynamic will be different if human players are in the "field agent" role working with an AI "overwatch", versus if they are playing "overwatch" to a team of AI controlled "field agents".

FDG 2023, April 12–14, 2023, Lisbon, Portugal

ACM ISBN 978-1-4503-9855-8/23/04.
https://doi.org/10.1145/3582437.3587200

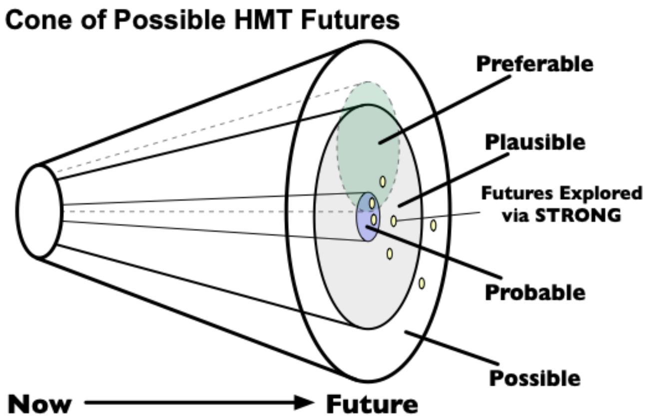

Figure 1: A conceptual representation of our strategy to Speculative Design to explore Human-Machine Teaming futures.

Second, the intentional inclusion of multiple roles gives us affordances for the games to vary along multiple dimensions of teaming interaction, such as naturalism, complexity, number of agents, and embodiment, as well as asymmetric vs. symmetric information and capabilities between agents.

To ensure practical relevance of each game scenario, we pull from existing research in team science and HMT. Prior research suggests future AI agents will need to improve their capabilities in communication, coordination, and adaptation to maximize performance in HMT [11]. In line with this, we are exploring multiple key dimensions between humans and AI agents, including team structure, team roles, and knowledge distribution [3, 8]. We will also examine heterogeneity (team makeup of humans and machines) and task complexity (contextual changes).

In line with Saavedra and colleagues [8], we conceptualize team structure in terms of interdependence. Team structure may be considered pooled, sequential, reciprocal, or fully team interdependent, with each of these structures including varying levels of interaction in the team task. Team roles, on the other hand, we conceptualize across 3 dimensions [3]: dominance (amount of control within team), sociability (amount of interaction with team members), and task orientation (reliability as a team member). Finally, knowledge distribution refers to the specific knowledge, skills, abilities, and other characteristics (KSAOs) of each team member [3].

Concretely, for each game, we follow an approach inspired by Tandem Transformational Game Design [12], to blend the processes of basic research and game design. This process starts with a Game-Driven Goal Delineation cycle, where research members of our team work with the game design team to review existing literature in team science and consult with outside experts to identify fruitful dimensions of the HMT space to target with a new game scenario (e.g., team roles, transparency, etc.). The process then shifts to a Goal-Driven Game Design loop where the game design team will explore multiple prototypes around the target scenario and iteratively playtest them into a full game. Cycling between the loops we will consider the existing prototype games to identify new aspects of team science that they might embody and then proceed to create variations of the games to highlight these dynamics.

3 CURRENT PROGRESS AND FUTURE WORK

To date we have run multiple ideation workshops with HMT experts as well as created one full game and an initial prototype for a second game. Additionally, we have begun development of an AI interface to enable automated interaction with our work going forward.

3.1 Game 1: *Dice Adventure*

Dice Adventure (see Figure 2) is a turn-based dungeon crawling adventure game, in which players take on the role of a dwarf, giant, and human who must collaborate to navigate through maze-like dungeons to reach a goal. Each of the dungeons is filled with challenges such as traps and rocks that have to be broken and monsters that have to be fought by rolling dice. The different characters have asymmetry in their abilities to address each of these challenges (e.g. the giant's available die cannot roll values capable of disarming traps, while the dwarf is similarly unable to break rocks) as well as their movement speed and ability to see through fog of war. Players must leverage their complementary abilities to navigate the dungeon and compensate for challenges that they are individually ill-suited for.

Dice Adventure was designed to be relatively approachable for current AI research approaches (being turn-based and restricted to a gridworld), while leaving open interesting teaming dynamics in the use of fog of war and interdependent task planning. For example, team roles can be explored according to which character roles players take on (e.g., dwarf, giant, or human). Knowledge distribution is reflected in the varying abilities included in each character and their restricted vision of the map. Team structure arises out of the coordination of these abilities between players [8]. For example, *Dice Adventure* can reflect a reciprocal team structure in the way players coordinate their complementary abilities. Within this team structure, we can explore the impact of humans versus machines taking on each role, and explore how that affects performance outcomes such as task score/completion time and team adaptation/learning. We can also examine some of the more nuanced components of team performance, such as symmetry in interaction ("Do the team members participate equally" [3, 4]); and team role adaptation ("Are members able to adapt their role structures" [6]).

Dice Adventure is implemented as a network multiplayer game in the Unity game engine with the eventual goal of being hosted publicly to provide a venue for human players to interact with experimental AI agents. Principle development has been completed with work on the AI interface currently underway. Source code for the game is publicly available at [Link Anonymized].

3.2 Game 2: *What's Cooking*

What's Cooking is a concept for a distributed cooperative cooking game, not unlike *Overcooked* or *PlateUp!*, in which players work together to fill customer orders, however, players are confined to their own individual kitchen spaces, which are connected by conveyor belts. Both raw ingredients and in-progress dishes have to be passed between kitchens in order to accomplish full recipes. Part of the coordination challenge in this concept comes from the inability for individual cooks to know the status of their collaborators and thus coordinate on a complex recipe together.

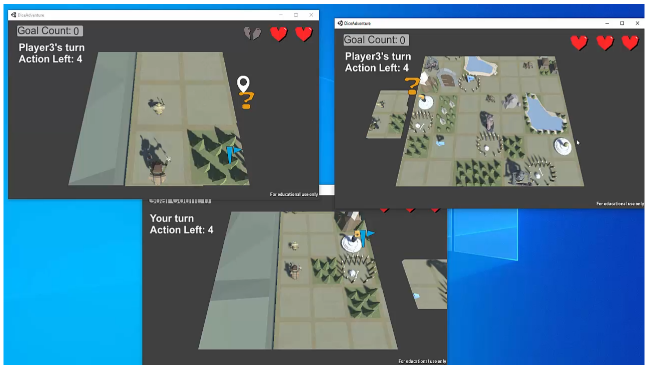

Figure 2: A screenshot of three players instances playing *Dice Adventure*. The dwarf and human (upper left and lower windows) players are exploring together, while the giant (upper right window) has gone off on their own.

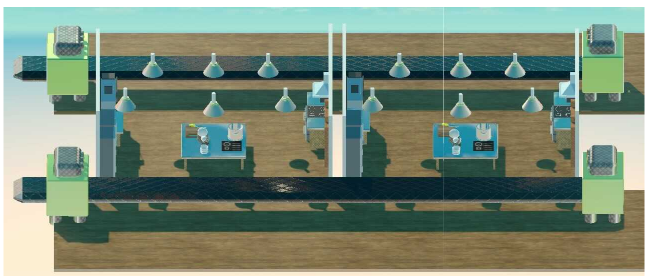

Figure 3: An initial mock up of a sequential kitchen layout in *What's Cooking*. Two player kitchens are connected by conveyor belts on either side that could be used to pass recipe ingredients to each other.

Additionally, the manner in which individual kitchens are connected to each other, provides affordances for the same underlying game system to embody diverse team interdependency structures such as pooled, sequential, reciprocal, and fully interdependent teams [8]. For example, each individual kitchen could be constructed to allow for a pooled form of work where order comes in such that each kitchen is able to make the entire recipe without interactions with other teammates, or they could be linked in a sequential pattern where each kitchen must take its actions in order. Alternatively, more complex relationships could be used, where components of recipes need to be sent back and forth between kitchens to accomplish component steps in an ad hoc manner. Additionally, the structural hierarchy of team roles can be explored in terms of the dominance or control of certain roles [3]. For example, a higher-dominance "quartermaster" role could make ingredient routing decisions between the individual chefs in order to support both their cooking work as well as communication.

What's Cooking is currently in it's section Goal-Driven Game Design loop in which the game development team will be broadly exploring game and level design variations along the team interdependency themes of Saavedra et al [8]. Similar to *Dice Adventure*, *What's Cooking* will be made publicly available as a resource for teaching novel HMT AI strategies.

3.3 Agent Development

The AI agents currently in development for these games are built around the Hierarchical Task Network (HTN) framework [9], in which representations of tasks contain preconditions and effects, as well as compositions of other tasks to be performed in sequence. The leaves of the HTN composition tree are atomic *actions*, which correspond to actions in the game world such as moving or pressing a key.

The interpretability of the HTN model allows it to be modified by the kinds of natural interactions that people naturally use to communicate. Components of the model's strategy, such as the preconditions for performing a certain action or the presence of higher level goals, can be described, discussed, and selectively modified if necessary. This contrasts with the "black box" nature of a neural network-based model, which demands new data and a round of backpropagation to even attempt to modify.

Our current work on interactive construction and modification of HTNs is rooted in the Natural Training Interaction (NTI) framework [7], in which the agent's *knowledge*—e.g., goals, beliefs, and skills—is modifiable by any of several *types* of interaction—e.g., informing, spotlighting, and demonstrating—through a number of *modalities*—e.g. GUIs, gestures, and speech—all according to an overall *pattern* of instruction—e.g., operant conditioning, direct instruction, and apprentice learning.

We are studying how the patterns and interaction types of the NTI framework correspond to the manipulation of HTN models. Work on generalizing demonstrations into re-usable HTNs has shown early promise, but the other interaction types, such as "clarify" and "inform", require verbal interaction. We are preparing to study the correspondence between natural language interactions and NTI interaction types in a wizard-of-Oz (WoZ) format, where participants will instruct an "agent", controlled behind the scenes by another person, how to perform specific tasks in one of our game environments. These instructions will be categorized according to the framework. We will then apply techniques from natural language understanding to construct a parser responsible for converting verbal interactions into HTN modifications.

ACKNOWLEDGMENTS

This work was funded by the Army Research Lab STRONG Program Award #W911NF2120101. The opinions expressed are those of the authors and do not represent the views of the sponsoring agency.

REFERENCES

[1] Paul Coulton, Dan Burnett, and Adrian Ioan Gradinar. 2016. Games as speculative design: allowing players to consider alternate presents and plausible futures. (2016).
[2] DARPA. 2020. AlphaDogfight Trials Final Event. https://www.youtube.com/watch?v=NzdhIA2S35w
[3] Tripp Driskell, James E Driskell, C Shawn Burke, and Eduardo Salas. 2017. Team roles: A review and integration. *Small Group Research* 48, 4 (2017), 482–511.
[4] Ute Fischer, Lori McDonnell, and Judith Orasanu. 2007. Linguistic correlates of team performance: Toward a tool for monitoring team functioning during space missions. *Aviation, space, and environmental medicine* 78, 5 (2007), B86–B95.
[5] John Harris, Mark Hancock, and Stacey D Scott. 2016. Leveraging asymmetries in multiplayer games: Investigating design elements of interdependent play. In

Proceedings of the 2016 Annual Symposium on computer-human interaction in play. 350–361.
[6] Jeffrey A LePine. 2003. Team adaptation and postchange performance: Effects of team composition in terms of members' cognitive ability and personality. *Journal of applied psychology* 88, 1 (2003), 27.
[7] Christopher J MacLellan, Erik Harpstead, Robert P Marinier III, and Kenneth R Koedinger. 2018. A framework for natural cognitive system training interactions. *Advances in Cognitive Systems* 6 (2018), 1–16.
[8] Richard Saavedra, P Christopher Earley, and Linn Van Dyne. 1993. Complex interdependence in task-performing groups. *Journal of applied psychology* 78, 1 (1993), 61.
[9] Earl D Sacerdoti. 1975. *The nonlinear nature of plans.* Technical Report. STANFORD RESEARCH INST MENLO PARK CA.
[10] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. 2016. Mastering the game of Go with deep neural networks and tree search. *nature* 529, 7587 (2016), 484–489.
[11] Kimberly Stowers, Lisa L Brady, Christopher MacLellan, Ryan Wohleber, and Eduardo Salas. 2021. Improving teamwork competencies in human-machine teams: Perspectives from team science. *Frontiers in Psychology* (2021), 1669.
[12] Alexandra To, Elaine Fath, Eda Zhang, Safinah Ali, Catherine Kildunne, Anny Fan, Jessica Hammer, and Geoff Kaufman. 2016. Tandem Transformational Game Design: A Game Design Process Case Study. In *Proceedings of the International Academic Conference on Meaningful Play.*
[13] Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. 2019. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature* 575, 7782 (2019), 350–354.