

‘Journeys in the Dark’ - Towards Game Master AI in Complex Board Games

Toby Best, Simon Lucas, Raluca Gaina

Game AI Group, Queen Mary University of London, UK
t.j.best@qmul.ac.uk, simon.lucas@qmul.ac.uk, r.d.gaina@qmul.ac.uk

Abstract

The Game Master is a player role synonymous with many tabletop games. The asymmetric gameplay of the role provides different opportunities compared to other players, and can be both cooperative and competitive with the other players in the same game. Though complex environments for exploring human and Artificial Intelligence collaboration exist, few focus on the Game Master role’s semi-cooperative play. Here, we propose a new complex environment based on the board game ‘Descent: Journeys in the Dark (Second Edition)’, as part of the Tabletop Games Framework, showcasing one-versus-many play, tactical combat, and large, dynamic action and state spaces. We include baseline AI player performance of Monte Carlo Tree Search agents in this game, finding them to be well-adept at considering multiple possible end-game conditions compared to the greedy One Step Look Ahead agents. In-depth analysis reveals interesting behaviours and Hero synergies, with the aim of informing the design of games and AI models to enhance human experience in semi-cooperative environments.

1 Introduction

Within many board games and video games, participants are expected to play equal roles. In these symmetrical environments, players have the same opportunities as one another, whether that be in competition by eliminating opponents from play, or collaborating to achieve a common goal.

A subset of tabletop games, known as tabletop role-playing games (TTRPGs), are by contrast one-versus-many asymmetric environments, driven by the power of collective storytelling and interactive narrative formed by the players (Riedl and Bulitko 2013). One player, in the role of Game Master (GM), acts as both narrator and rules referee to guide the other players throughout the game (Tychsen et al. 2005). A GM may have to balance playing both sides of the game simultaneously, taking control of the antagonistic forces as well as any helpful allies who may assist the other players’ characters during play. There is a level of cooperation expected between players, particularly those on the same side, whilst competition still remains between the opposing sides.

Though there have been prior explorations of the use of Artificial Intelligence (AI) agents to play TTRPGs (Martin,

Sood, and Riedl 2018), the research on virtual environments purpose-built for such agents appears rather limited. Additionally, though many TTRPGs may appear to be easy to learn, there is a detailed level of forward planning required to succeed. The complexity and interconnectivity of the rules and the player’s available actions (Shyne 2023) makes it less straightforward for agents to plan their turns when compared to other games, which may lead to unpredictable game outcomes and higher computational costs as a result.

From this, we propose a new TTRPG-like environment for Game AI methods, based on the game ‘*Descent: Journeys in the Dark (Second Edition)*’ (D2e) (Fantasy Flight Publishing, Inc. 2012). Although ‘Descent’ is closer to a competitive board game than a traditional TTRPG such as ‘*Dungeons & Dragons*’ (D&D) (Gygax and Arneson 1974) or ‘*Call of Cthulhu*’ (Press 2017), it possesses more of a minimalistic, strategic and highly structured view of a TTRPG environment. A team of players take the role of adventuring *Heroes*, battling against the *Overlord*, a solo player who controls an army of *Monsters* to directly compete against the *Heroes*.

As such, we believe it has suitable ludic elements, rules and themes that would make the implementation of such an environment worth pursuing. Thus, we construct D2e within the Tabletop Games Framework (TAG) (Gaina et al. 2020a,b), combining several challenges for AI players: one-versus-many play, tactical combat, large and complex state and action spaces, strategic multi-layered decisions through different abilities and card effects, and Game Master AI.

The contributions of this paper are as follows: we detail the properties of the new digital D2e environment implemented within TAG, and initial experimental results measuring the baseline performance of two AI players fighting to win: a greedy short One Step Look Ahead and Monte Carlo Tree Search (Browne et al. 2012), compared against publicly available human player results. Furthermore, we perform a detailed analysis of AI player behaviour and emergent Hero synergies. We discuss challenges and opportunities created, with the aim of informing the design choices for the game and AI models used to play it, and enhance human experience in semi-cooperative environments.

We intend for this to serve as a basis for exploring AI agent behaviours within a more TTRPG-like virtual environment for future research, with a potential roadmap laid

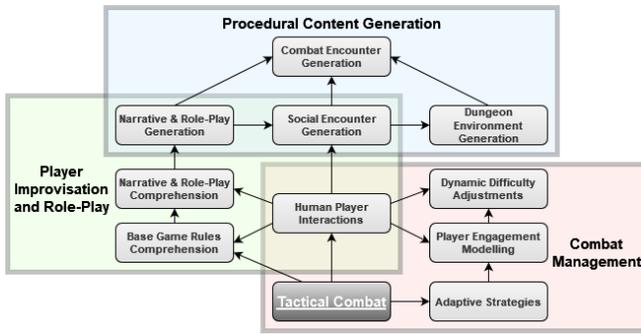


Figure 1: Our intended roadmap for future developments towards Game Master AI. There are three overall sets that encompass each stage - combat management, player improvisation and role-play, and procedural content generation. The current implementation of D2e allows for explorations of AI agents partaking in tactical combat (highlighted).

out in Figure 1. These development intentions include: allowing for the Game Master agent to model players’ performances and adjust difficulty based on player engagement, generating narrative information to reflect players’ actions and the state of the game (Zhu et al. 2023), exploring believable behaviours for non-player characters (Park et al. 2023), and the potential of procedural content generation (as previously explored within ‘Gloomhaven’ (Tijben 2023; Gerhold and Tijben 2023)). We hope these developments will progress towards improving gameplay accessibility in TTRPGs, by form of either assistance tools to help new or inexperienced GMs with running their own games (Santiago III et al. 2023), or to substitute for a human GM if one is unavailable (such as in solo play or in small player groups).

2 Literature Review

There has been much research into both competitive and cooperative complex game environments, using Artificial Intelligence (AI) agents within board, card, tabletop and video games to optimise tactics and strategies, and help expand upon players’ understanding of the game. For example, this can be seen with DeepMind’s ‘AlphaGo’ (Silver et al. 2016), ‘AlphaStar’ (Vinyals et al. 2019) and ‘AlphaZero’ (Silver et al. 2017) agents, which could compete at the same level as professional human players. Refinements of human player strategies were found as a result of adaptations to these AI agents (Egri-Nagy and Törmänen 2020; Shin, Kim, and Kim 2021), with ‘StarCraft II’ (Vinyals et al. 2019; Samvelyan et al. 2019) and ‘MicroRTS’ (Ling et al. 2022) in particular standing out for their tactical play aspect and multi-unit control. The emergence of AI players and assistance tools that directly work within the game engine have allowed for these games to become better understood, and therefore have had an impact upon the play experience. However, this has also lead to some unexpected, and potentially negative, developments to occur, such as cheating in games becoming more accessible (Shin, Kim, and Kim 2020).

AI implementations of modern board games have also helped with analysing game mechanics, such as for creat-

ing tutorials on how to play (Perez-Liebana et al. 2023), or providing a better understanding of common ludemes found in board games (Sousa 2023) and role-playing games (Riel and Monahan 2024). The implementation of ‘Terraforming Mars’ within TAG (Gaina, Goodman, and Perez-Liebana 2021) serves as a recent example of AI agents being implemented to examine a modern board game, in this case used to study the performance of various agents compared to human players within the game’s expansive action space. The high quantity of available actions and the uncertainty of their consequences is acknowledged to prove difficult to predict future game states, a challenge that is shared with TTRPGs (Martin, Sood, and Riedl 2018) due to the collaborative improvisation involved, especially when considering how to construct player agents.

For AI agent implementations within cooperative board games, there has been recent interest in particular towards human-AI cooperation (Ashktorab et al. 2020; Dafoe et al. 2020; Chacón and Eger 2019; Sfikas and Liapis 2020; Gaina and Balla 2022), as well as studies into multi-agent systems for semi-cooperative environments (Amini and Afsharchi 2014; Kitchen, McGroarty, and Aris 2023; van den Bos and Stoelinga 2023), such as one-versus-many games. These games may involve one player fighting against a cooperative team of others, or players may have hidden goals that influence if they win or lose individually. The hide-and-search asymmetric game ‘Scotland Yard’ (Dash et al. 2020) is one such case, with research comparing how adversarial neural network agents adapt to their opposition’s actions. Whereas ‘Scotland Yard’ pits a team of detective players against a single adversary, many TTRPGs often give the Game Master (GM) just as many figures to control as the other players, sometimes more, resulting in many moving parts during play to keep track of.

Regarding tabletop games, the role of GM can be seen as semi-cooperative. TTRPGs at their core are interactive narratives (Riedl and Bulitko 2013), with the GM tasked with guiding the other players through the game’s narrative as well as refereeing the rules and providing obstacles to challenge and oppose them, making the GM just as active a participant as the other players. The asymmetric gameplay and open-ended creative nature of TTRPGs make them prime candidates as challenges for AI to solve (Ellis and Hendler 2017). Previous research into AI agents within TTRPGs include the Discord bot ‘Calypso’ (Zhu et al. 2023), the hypothetical GM-assistant ‘Avalon’ (Santiago III et al. 2023), and large language models trained on actual play (Callison-Burch et al. 2022). Meanwhile, a study on generative agents simulating believable behaviours within non-player characters (Park et al. 2023) highlights the possibility of a single agent governing actions for multiple characters in TTRPGs, for players to interact with in combat and social encounters.

3 Descent

‘Descent: Journeys in the Dark’ (Fantasy Flight Publishing, Inc. 2012) is a fantasy dungeon-crawler board game published by Fantasy Flight Games for 2 to 5 players. Released originally in 2005, with various expansions, campaigns and new quests added over the years, a simplified *Second Edition*



Figure 2: Figure cards for the Warrior Hero ‘Syndrael’ (left) and the Monster ‘Barghest’ (right). Syndrael’s Attributes, Hero Ability and Heroic Feat are listed alongside her stats. The Barghest Minion’s stats and abilities are listed at the top, whilst the Barghest Master’s are listed at the bottom.

(2e) was released in 2012, which this study focuses on. Here, one player becomes the *Overlord*, a Game Master role who takes control of an army of *Monsters* to oppose the other players’ team of *Heroes*. The game is asymmetrically one-versus-many in format - the Heroes work as a team, whereas the Overlord plays individually, and both sides have their own victory conditions in order to win. We only discuss elements of the base game and the quests from its campaign *Heirs of Blood*¹ in this paper. Full rules can be found online².

The playable Heroes and Monsters are represented by cards with the figure’s personal abilities and stats (see Figure 2), which are as follows: *Speed*, how many spaces they can move per action; *Health*, how much damage they can take before they are *defeated*; *Stamina*, how many times they can spend *Fatigue* to power certain actions (for Heroes only); *Defense*, the dice used for a defence roll; and *Attack*, the dice used for an attack roll (for Heroes, this is determined by their equipped weapon). A Hero’s class abilities and equipment are represented via a personal deck of cards, and each card can be exhausted (either temporarily or permanently) to use as actions, and some usable item cards, like armour cards, can grant stat bonuses whilst active.

Turns alternate between the Heroes and the Overlord. When a Figure is *activated* on its turn, it gains two action points to spend on any combination of legal actions. These include: *Move* up to its Speed to an unoccupied space; make an *Attack* against an enemy (either Melee or Ranged, determined by the Hero’s weapon or the Monster’s Attack Type); use a specific *Skill* or *Monster Action*; *Open/Close a Door*, or any other *Special* actions defined by the quest. With the exception of *Move*, which can be split up by another action and continued afterwards, all actions taken must be fully executed before the next can be chosen, although another figure’s abilities can interrupt its results. We discuss this further in Section 4.1.

¹<https://kupdf.net/download/heirs-of-blood-5912ca4fdc0d60a36f959ecc.pdf>

²https://images-cdn.fantasyflightgames.com/ffg/content/descent-second-ed/support/DJ01_Rulebook_ENG.pdf

3.1 Playing as the Heroes

At the start of play, each Hero player chooses an *Archetype*, *Class* and *Hero*. There are four Archetypes: Healer (colour-coded as blue in-game), Mage (yellow), Scout (green) and Warrior (red). Each class determines the Hero’s starting equipment and skills, and a Hero can be any class of their Archetype. An Archetype may only be represented once in the same game. In the base edition used in this work, there are eight classes and Heroes, each Archetype matching two of each. A minimum of two Heroes must be present in a game, so if there are only two players, one player controls two Heroes, yet treats them as acting individually.

In addition to their stats, Heroes have four *Attribute* values: *Might*, *Knowledge*, *Willpower* and *Awareness*, which are used in Attribute tests to determine outcomes of a certain skill, action or card activation. A test is passed if the Hero’s defence roll is lower or equal to that Attribute. Each individual Hero also has their own unique Hero Ability: a special clause that can change how certain actions resolve (e.g. allowing a once-per-turn reroll of a die when you attack, or converting damage suffered from attacks into Fatigue); and a unique Heroic Feat: a once-per-quest action that can turn the tide of battle in their favour (e.g. a free Attack action).

On the Heroes’ turn, players may choose to activate their Hero in any order, which they may freely change each turn. The class skill cards provided to the Hero add to their available options on their turn. Some skills are performed by exhausting the card or by spending Fatigue (or both), and some do not cost action points. Heroes can also spend action points on the following actions: *Rest* to recover all Fatigue at the end of their turn, *Search* an adjacent token, or *Revive* a defeated adjacent Hero. Defeated Heroes are not removed from play, but can only take the *Stand Up* action on their turn, ending their turn immediately. Also, when only two Heroes are present, each Hero also gains a free bonus action, to use only for an extra Attack or to recover 2 hit points.

3.2 Playing as the Overlord

The Overlord takes control of an army of *Monsters*, which come in three variants: *Minion*, *Master* and *Lieutenant* (unique boss-type Monsters with Attributes). Unlike Heroes, a Monster’s Attack or Defense stats are unaffected by equipment, and they do not have a Stamina stat. Minions and masters lack Attributes, and will fail any Attribute test received. A Monster who is defeated is removed from the board.

Minions and masters make up the most common Monsters. Masters often have better stats than minions, and many have skills and abilities that minions cannot access. A *Monster Group* composes all the minions and the master of the same monster type: for example, a *Goblin Archers* group may contain four minions and the master, whereas an *Elementals* group might contain only the master. The number of minions in a group depends on how many Heroes are in play, though only one master can typically exist at a time.

On the Overlord’s turn, they can activate each Monster Group and its Monsters in any order, however all Monsters in a group must end turn before the next group can activate. Monsters can only make one Attack or Monster Action

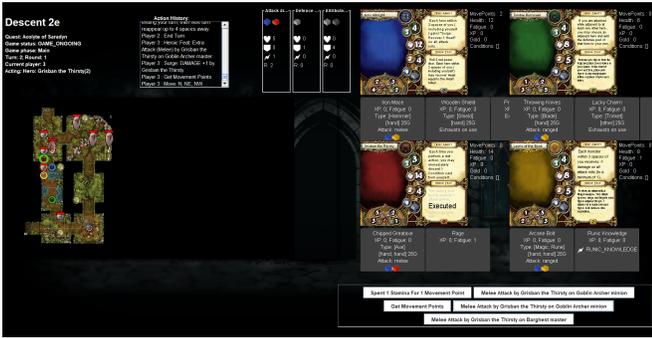


Figure 3: The current Graphical User Interface (GUI) for ‘Descent’ in the TAG Framework. A human player is choosing the action of the Warrior Hero ‘Grisban the Thirsty’.

each turn, and cannot pick both in the same turn. The unique Monster Actions for each Monster are defined on their card, such as the Barghest’s Howl causing a Willpower test for all Heroes within three spaces, suffering 1 Fatigue on a failure.

4 Methods

4.1 Implementation of Descent

The current version of ‘Descent’ implemented in TAG (see Figure 3) fully supports the first quest of the *Heirs of Blood* campaign, *Acolyte of Saradyn*. The Monsters encountered in this quest are Goblin Archers and Barghests. The Overlord can respawn up to two Goblins at the end of their turn to replace any that have been defeated, placing them in Tile ‘4A’ at the top-right of the board (see Figure 4). There are three win conditions available for this quest: the Heroes defeat all the Barghests (Heroes victory); the Overlord defeats all the Heroes (Overlord victory); or the Overlord achieves 7 Fatigue tokens (Overlord victory). The Overlord can gain one Fatigue token at the start and at the end of each of their turns in which they have at least one Goblin in any position within the board tile ‘9A’.

Currently, the game state is fully observable to all players. The estimated state space within the current implementation of D2e lies roughly around 10^{30} to 10^{56} , depending upon the number of players present. Our player agents experienced a mean action space of 6.93 across all games, with a maximum of 67 actions available at one state and a minimum of 1. The framework took on average $841\mu s$ to create a copy of the state, $560\mu s$ to advance the forward model to the next state, and $936\mu s$ to compute the player’s available actions.

Similar to the Ludii framework (Piette et al. 2020) and the Options framework within reinforcement learning (Sutton, Precup, and Singh 1999), the TAG framework allows us to parameterise the actions and use a sequential model in order to keep the immediate action space relatively small and avoid extremely large combinatorial action spaces. Notably, many actions, such as making attacks and Attribute tests, can be interrupted by other figures’ abilities, which can introduce additional uncertainties to future game state predictions. For example, when deciding to make an attack action, the acting figure declares their target and their weapon (if



Figure 4: The game board’s layout for the first quest, Acolyte of Saradyn. The Heroes start at the Entrance ‘E’ (bottom-left), the Goblin Archers at ‘5A’ (top-left), and the Barghests at ‘4A’ (top-right). Goblins can earn Fatigue for the Overlord at ‘9A’ (bottom-middle) in order to win. Up to two defeated Goblins can also respawn anywhere within ‘4A’ each turn.

a Hero). The attack flows through several stages of progression, from rolling the attack dice, to rolling the defence dice, to applying damage, all the while checking for any actions from the other players (including the attacker and defender), which could interrupt and affect the attack. These other actions may include forcing a reroll of a die, adding or subtracting to the rolls to influence their results, or spend *Surges* - a key element of combat that allows the attacker to bolster their attack with special bonus effects. If these actions and decisions were combinatorial instead of sequential, then our already large action space would become even greater, and would negatively affect efficiency when playing the game.

4.2 Physical and Implementation Differences

Due to the nature of the physical version of ‘Descent’, there are some notable components that we chose to omit or otherwise alter, in order to simplify our implementation to be more in-line with our desired end goal of exploring AI models in semi-cooperative asymmetric environments similar to those found in most TTRPGs. As these will impact the play experience of our virtual environment compared to the original game version, we detail our changes as follows.

For one, the TAG implementation does not include the deck of *Overlord Cards*. These are special actions that can help the Overlord or hinder the Heroes, such as immediately ending a Hero’s movement on a failed Awareness test, grant a bonus action to a specific Monster, or force a Hero to attack another Hero (or themselves). We chose to omit this to streamline the environment, as otherwise the Overlord was constantly interrupting other players’ turns to decide if they wanted to use a card. This also meant that the only *Con-*

dition present in the first quest is *Stunned* (afflicted figures only gain 1 action point that turn instead of 2), and can only be inflicted upon Monsters via Healer Archetype Heroes.

Also, when making a Ranged Attack, we calculate a figure’s line of sight from the centre of the attacker’s position to the centre of the target’s position, whereas the original rules calculate line of sight by comparing the corners of the attacker and the target’s spaces. We also impose a hard maximum range of 8 spaces away for the attack to be legal at this stage, as attempting from any greater distance with the available starting equipment would be impossible to succeed.

4.3 AI Players

To examine potential player behaviours that might occur in our implementation of D2e, we use three automatic player agents in this study’s experiments: *Random*, which selects a random available legal action, *One Step Look Ahead (OSLA)*, which greedily selects the action available that leads it to the next best state, and *Monte Carlo Tree Search (MCTS)* (Chaslot et al. 2008; Browne et al. 2012), which constructs a tree simulation of several future actions and selects the best possible predicted outcome. The MCTS agent’s parameters, tuned from preliminary experiments to maximise performance, are set as follows: it has a rollout length of 10, an Upper Confidence Bounds exploration constant K of 0.01, an explore- ϵ of 0.3, and a budget of 300 milliseconds to make decisions at each state. The MCTS decision tree is set to open-loop, due to the random nature of dice rolls involved, and the tree also ignores other players’ potential actions during construction. The Most Average Sampling Technique (MAST) (Bjornsson and Finnsson 2009) is set to consider both actions from the rollout phase and tree phase in our decisions, with a MAST gamma decay of 0. For more on MCTS agent parameters in TAG, please see the Wiki page³.

A weighted sum of seven heuristic features is used to evaluate the game states for OSLA and MCTS agents. The Heroes and Overlord use the same heuristic function, though negated for opposing sides. The features are listed in Table 1, along with the weights used for all games run. Three features are considered beneficial for the Heroes, whilst the other four are beneficial for the Overlord. If a feature is not beneficial for an agent’s team, it is negated during summation instead. Both **HEROES THREAT** and **OVERLORD THREAT** use a breadth-first-search for calculating the precise distance between positions. The full methods used for calculating each heuristic score can be found in the source code⁴. This includes experiment configurations for ease of reproducing results.

5 Experiments and Results

Our data is generated from 7150 games run within TAG. The games run follow the features described in Section 4.1.

As factors such as the number of figures in play or available actions are affected by player numbers, we initially ran

³<http://www.tabletopgames.ai/wiki/agents/MCTS>

⁴<https://github.com/GAIGResearch/TabletopGames/tree/descent>

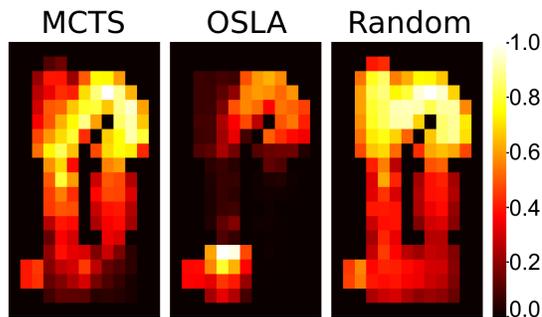


Figure 5: Heat maps of all figures’ movement for 5-Player self-play games. The brighter the space, the more often a figure ended its turn there, from most to least visited. For MCTS, Tiles ‘4A’ and ‘5A’ at the top contain the most frequently visited spaces. For OSLA, these are in the bottom-left, near the Heroes’ start and the score zone of Tile ‘9A’. For Random, most movement occurs in the top board tiles, as this is where the Goblin Archers and Barghests spawn.

6000 games of self-play (with all players using the same agent), in batches of 500 at 2-5 players for each agent, for 2000 games total per agent. Then, we ran 1150 games in a round-robin tournament, where teams of each agent type were pitted against each other in 5-Player games. Though we consider all games within our analyses, the following figures, tables and statistics primarily reference the 5-Player games; similar results are observed for other player counts.

Each game was set up so Hero players were randomly assigned their character, with randomised initial figure placements (and reinforcement placements) as well. A fixed turn order was used: Hero 1 always acted first, then 2, 3 and 4, then the Goblin Archer Master and Minions, then lastly the Barghest Master and Minions. Additionally, a timeout would be inflicted if 20 rounds passed with no winner, assigning a penalty loss to both players, to avoid excessive resource usage in seemingly stalemated situations. Most games lasted no more than an average of 11 rounds for completion, so this felt like an appropriate limit.

One notable challenge encountered was agents deciding to end turns immediately with no other action taken, which often led to frequent timeouts. We resolved this by removing End Turn from the available action space until no other actions were legal. In actual tabletop play, players may want to tactically take no action, or purposefully refuse certain actions, rather than being forced to do so by the environment and losing player agency as a result. We propose the exploration of tactical turn length choices as future work.

In the 5-Player self-play games, the MCTS agent took on average 224 seconds per game, averaging 11 rounds per game at 20 seconds per round. The OSLA agent on the other hand took only 3 seconds per game, averaging 8 rounds per game at 380 milliseconds per round.

5.1 AI Agent Performance

Tables 2 and 3 both show how often each end game condition occurs across all games. Immediately we notice the

Heuristic Feature	Weight	Description
HEROES HP	+ 0.8	Ratio of Heroes’ total current HP to their total maximum HP.
MONSTERS DEFEATED	+ 0.7	Ratio of defeated Barghests (0 HP) to the total number of Barghests.
HEROES THREAT	+ 0.2	How close the Heroes are to the Barghests.
HEROES DEFEATED	- 0.7	Ratio of defeated Heroes (0 HP) to the total number of Heroes.
MONSTERS THREAT	- 0.8	Ratio of Barghests’ total current HP to their total maximum HP.
OVERLORD FATIGUE	- 0.7	Ratio of Overlord’s current Fatigue to their maximum Fatigue (7).
OVERLORD THREAT	- 0.2	How close the Goblin Archers are to the score zone of ‘9A’.

Table 1: The seven heuristic features used to evaluate game states for MCTS and OSLA agents. The weight values listed for each heuristic is taken from the Heroes’ perspective, and are negated when evaluating game states for the Overlord.

Players (Agent)	Game End Condition Met			
	Barghests	Fatigue	Hero Defeat	Timeout
5 (MCTS)	30.2%	65.2%	0.0%	4.6%
4 (MCTS)	21.4%	76.4%	0.0%	2.2%
3 (MCTS)	45.0%	55.0%	0.6%	0.4%
2 (MCTS)	49.2%	50.0%	0.0%	0.8%
5 (OSLA)	10.2%	89.8%	0.0%	0.0%
4 (OSLA)	14.2%	85.8%	0.0%	0.0%
3 (OSLA)	45.6%	54.4%	0.0%	0.0%
2 (OSLA)	48.8%	51.2%	0.0%	0.0%
5 (Random)	0.6%	1.8%	0.0%	97.6%
4 (Random)	2.0%	2.8%	0.0%	95.2%
3 (Random)	8.6%	2.0%	0.6%	89.4%
2 (Random)	9.0%	1.6%	0.0%	89.4%

Table 2: Occurrences of Game End Conditions for the 6000 self-play games, split by 500 games per agent type and number. These are: the Heroes defeat all the Barghests (Heroes win); the Overlord achieves maximum Fatigue (Overlord win); the Overlord defeats all the Heroes (Overlord win); or 20 rounds pass with no declared winner (Timeout).

Overlord Agent	Hero Agents					
	MCTS		OSLA		Random	
	Over.	Time	Over.	Time	Over.	Time
MCTS	68.4%	4.3%	66.7%	0%	88.5%	11.5%
OSLA	100%	0%	88.2%	0%	99.4%	0.6%
Random	0.8%	12.3%	4.5%	16.4%	1.9%	98.1%

Table 3: Results of 1150 5-Player round-robin games between MCTS, OSLA and Random agents, and the frequency of the Overlord winning or a Timeout occurring after 20 rounds passed. Each pair played roughly 130 games.

high frequency of timeouts inflicted upon Random self-play games. This suggests meaningful progress cannot be obtained in ‘Descent’ through passive play - players must actively participate in order to achieve victory. A few timeouts were incurred for MCTS self-play games as well, though its low occurrence, contrasted by reasonably consistent win-rates for both sides, might hint that MCTS is capable of purposefully prolonging games to avoid losing, invoking timeouts via survival rather than an inadequacy to progress.

From this, we can see the asymmetric nature of ‘Descent’ in play in the MCTS versus OSLA matchups of Table 3.

The MCTS agent had a 66.7% win rate when playing as the Overlord against a team of OSLA Heroes. However, when OSLA played the Overlord, the MCTS Heroes team did not win a single match. With how few attacks and kills relate to the Goblins in Table 4 in OSLA self-play games, combined with the high footfall near Tile ‘9A’ in Figure 5, it is plausible that, due to the more short-sightedness of the OSLA agent, which only looks at the next best move available, the OSLA Overlord tends to rush the Goblins to the score zone for a quick win, akin to a ‘Zerg Rush’ strategy within ‘Star-Craft II’. This appears to be something that neither MCTS nor OSLA Heroes can keep up with whilst pursuing their own goals, as their attention may be split towards attacking the Barghests, and thus may not have enough Attack actions available to defeat the overwhelming number of Goblins. As MCTS is capable of planning further ahead, it may be considering various strategies available and planning out its best options, representing a more experienced player, whereas OSLA may represent more of a beginner who is taking the easiest path to victory. There is a need for plan recognition for players to win consistently against such strategies, something previously experienced with human-AI player cooperation within ‘Pandemic’ (Chacón and Eger 2019). An AI agent in such an environment would need to adapt its own tactics, balancing its immediate needs and long-term end goals, whilst predicting other players’ intended action plans.

Though we cannot make a direct comparison, we can look at these results against those of human players, provided by the community-made ‘Unofficial Campaign Tracker’ for ‘Descent’⁵. As of the time of this study, across 1439 games of *Acolyte of Saradyn*, the Heroes were reported to have won 871 games (60.5%), whereas the Overlord had won 568 games (39.5%). As the website is self-reported by the players, and neither clarifies how many are involved per game nor the achieved victory conditions (for the Overlord), it is worth noting that this is only a rough view of the quest’s outcomes by comparison. However, Table 2 shows a higher likelihood of an Overlord victory by contrast. Though the differences in our virtual environment at this stage may be a determining factor, another possible hypothesis is that the AI agents have an easier time controlling the Overlord player, as there are two victory conditions as opposed to the Heroes’ one, and the Overlord has control of all of their units together and thus can plan its moves through all of its units, whereas

⁵<http://d2etracker.com/stats.quests.php>

Hero Name	Archetype / Monster	Attacks Per Game		Targeted Per Game		Kills Per Game		Defeats Per Game	
		MCTS	OSLA	MCTS	OSLA	MCTS	OSLA	MCTS	OSLA
Ashrian	Healer	5.88	2.20	9.90	1.50	1.72	0.30	0.72	0.02
Avric	Healer	4.60	1.51	8.76	1.38	1.44	0.28	0.39	0.00
Leoric	Mage	10.48	2.70	7.95	1.84	3.11	0.40	0.37	0.04
Tarha	Mage	10.06	1.98	7.70	2.39	2.88	0.24	0.70	0.09
Jain	Scout	10.45	2.69	11.57	1.90	2.69	0.88	1.89	0.13
Tomble	Scout	10.36	2.67	7.38	0.90	2.82	0.48	0.81	0.03
Grisban	Warrior	5.90	2.96	9.90	2.57	2.98	0.95	0.91	0.04
Syndrael	Warrior	5.56	2.75	12.05	2.93	2.47	0.83	1.57	0.07
Goblin	Monster	6.24	0.55	3.53	0.14	0.49	0.01	1.48	0.08
Barghest	Monster	1.43	1.25	3.52	2.23	0.31	0.50	0.68	0.45

Table 4: Attacks made, attacks received, defeating blows dealt and times defeated per game for each Figure across the 500 5-Player self-play games for MCTS and OSLA agents, rounded to 2 decimal places. Monster Figures take their group average.

Hero Name	Used Heroic Feat		Average Turn Used	
	MCTS	OSLA	MCTS	OSLA
Ashrian	97.3%	41.6%	2	2
Avric	99.6%	96.5%	2	1
Leoric	62.1%	47.7%	5	2
Tarha	89.8%	27.9%	4	2
Jain	99.2%	100%	1	0
Tomble	98.8%	18.4%	0	2
Grisban	98.0%	97.3%	3	1
Syndrael	99.6%	99.6%	0	0

Table 5: Heroic Feat usage per game for each Hero across the 500 5-Player self-play games for both MCTS and OSLA agents. Turns are rounded to nearest integer, with the first round of the game acting as ‘Turn 0’.

the Hero agents need to coordinate their actions together, a behaviour which may not be fully facilitated by the current infrastructure. Again, as shown in ‘Pandemic’ (Gaina and Balla 2022), coordination amongst agents is crucial yet difficult to achieve, which Table 2 suggests, as the Heroes’ win rates across the 6000 self-play games decreased as the number of participating players increased, as more Hero players were required to be considered when planning that turn, and the Overlord gained more Monsters to challenge the Heroes with, which both would have impacted agent performances.

Further research could look more in-depth towards improving Hero play, more specifically examining any obstacles these agents are facing with cooperative play. This may involve exploring other quests present in the *Heirs of Blood* campaign, and how each agent type performs in comparison to the tracker website’s human player results, when faced with different board layouts and victory conditions. For example, the tracker lists the quest *Blood Will Tell* as favouring the Overlord, who won 127 games (65.1%) compared to the Heroes’ 68 games (34.9%), and so may yield different performances for agents when compared to the first quest.

5.2 Behaviour Analysis

For the performances of the individual components themselves, Table 4 lists the average attacks made and enemies

defeated per game per figure, and Table 5 lists the Heroic Feat usage rates. We note that OSLA performs fewer attacks per game; however, this may be a consequence of the faster pace of OSLA games (the average game length was approximately 8 full rounds for OSLA, as opposed to 11 full rounds for MCTS). Many of the Heroes’ favourite target to attack for OSLA Heroes are the Barghests, whilst there was a mixture of Goblins and Barghests instead for MCTS Heroes. All five Goblins are attacked the most by Widow Tarha for MCTS, which is reasonable as she can make Ranged attacks from a distance, and her Heroic Feat allows her to target two enemies at once. Meanwhile, the Goblins are attacked the most by Grisban in OSLA games - there are three key game elements that might help explain this decision-making: one, as a Warrior, he can only make Melee Attacks; two, his Speed stat of 3 is the worst of all Heroes; and three, no mechanics are currently present that allow Heroes extra movement, other than his counterpart Warrior Syndrael’s Heroic Feat, who cannot be deployed in the same quest as him. Thus, we can infer that, as the OSLA players prefer to rush straight for the Barghests, Grisban cannot keep up with the other Heroes’ pace, and thus is left with the Goblins similarly rushing in the opposite direction towards him. This contrasts with the high usage of Jain and Syndrael’s Heroic Feats, both of which grant extra movement, and therefore considered highly valuable for both agents (with Jain having a 100% usage rate in the very first turn for OSLA games).

5.3 Hero Synergies

As each Hero has their own unique strengths and weaknesses, these can all have an impact on their player’s overall game experience, particularly if they feel too strong or weak; players may not enjoy having to play a weaker Hero as much as one who can more easily hold their own, and likewise the Overlord may find it frustrating if their Monsters and obstacles are easily disposed of by a notably strong Hero. Balance is just as important an aspect of running a TTRPG as narrative and role-play are, and the type of analysis from AI gameplay presented here can help adjust the difficulty of a game to better suit all of the players’ needs.

The most successful four-Hero team for MCTS in our

self-play results is Ashrian, Tarha, Jain and Syndrael, who won 48.5% of their games, whilst the worst is Avric, Tarha, Tomble and Grisban at 10.3%. For OSLA, the best team is Ashrian, Leoric, Jain and Syndrael at 24.0%, whereas Avric, Tarha, Tomble and Syndrael were the only team with a win rate of 0%, losing every game they played together. For individual Heroes, Ashrian performed the best in 5-Player games for both MCTS (35.3% win rate) and OSLA (11.9%), whereas the worst was Avric for both (24.8% and 8.9%).

We can theorise these success rates thus: Ashrian, as a Healer, has access to healing skills that can keep her allies alive. However, she has the highest Speed of all heroes (5, joint with Jain Fairwood), and her Hero Ability can inflict Stunned upon any adjacent Minions for free, effectively cutting their actions per turn in half. Avric Albright's Hero Ability and Feat focus on healing allies, rather than denying Monsters' actions like his counterpart Ashrian can.

Meanwhile, both Jain Fairwood and Syndrael grant extra movement for free with their Heroic Feats, and thus would be highly valued within the agents' heuristic scores to swiftly reduce the distance between Heroes and Barghests - Jain notably uses her Feat in her first turn in every game for OSLA agents. Their Hero Abilities are also useful, as Jain can convert some of the damage she takes from attacks into Fatigue, and Syndrael recovers 2 Fatigue at the end of her turn if she ends her turn without moving, which means she can use both action points to attack and benefit as a result.

We have already noted the flaws with Grisban the Thirsty and his opportunity cost compared to Syndrael, but it is worth acknowledging that his Hero Ability - remove a Condition when taking the Rest action - is not effective in our experimental setup, as the only Condition available in the game is applied by his allies onto the Monsters. Tomble Burrowell, though not as flawed as Grisban, is not as fast as his counterpart Jain (4 Speed versus her 5), has a purely defensive Heroic Feat (vanish from the map where he can't be targeted on one turn, then reappear up to 4 spaces away on the next turn), and though he too has a defensive Hero Ability, he can only use it if he is adjacent to an ally and requires strategic planning to take advantage of, rather than Jain simply having Fatigue available. Lastly, Widow Tarha appears in the worst teams for both OSLA and MCTS, suggesting more advanced play is required to play this Hero efficiently.

6 Challenges and Opportunities

The introduction of the digital 'Descent' TAG environment presents several challenges and opportunities of interest.

Game-playing AI. The complex game with large state and action spaces, asymmetric gameplay and many dynamic parts offers direct challenge for AI game-playing research within tabletop role-playing games. As our focus in this study has been on interesting behaviours in the newly implemented D2e environment, rather than obtaining the best possible behaviours to win, there is still the potential of future research examining the creation of good AI players. The free movement, wide range of actions or abilities the players can enact, and dynamic and conflicting goals, as well as the various strategies or tactics that can be discovered, allow for further study of AI behaviours in complex environments.

Adaptive strategies. As discussed previously in Section 2, human players will often adapt strategies to defeat their opponents, and improve them over time. This element of tactical combat and multi-layered strategy can be realised by the AI agents within D2e similar to the approaches outlined in Holmgård et al.'s study (Holmgård et al. 2018). When considering longer play sessions composed of multiple quests (or a full campaign), or repeated plays of a quest with the same agents, it would be possible to explore weight adjustments for agents to reflect responses to other players' tactics. For example, the Overlord could put more emphasis on defeating the Heroes over pursuing alternative win conditions (e.g. maximising their Fatigue), or react to other players' behaviours, such as rushing Heroes who would rather keep their distance. This would be beneficial to keeping the game engaging for any human players involved. These techniques in changing strategies to adapt to other players' actions could be applied within other strategic asymmetrical and semi-cooperative games, for example the wargames '*OGRE*' (Jackson 1977) or '*Fortress America*' (Gray 1986).

GM player experience modelling. As the core tactical combat aspect demonstrated by our experiments plays is but a minimum part of playing a TTRPG, it is important to remember that a GM in human play often acts in the interest of the group and maintaining player engagement. As our roadmap (see Figure 1) shows, we recognise the opportunity here of modelling varied behaviours for the Overlord player, with the aim of maximising engagement for all players more as a 'Villain', rather than simply defeating the Heroes as an 'Adversary' (Treanor et al. 2015).

Narrative and role-play. Another opportunity raised by this work that we propose for future research is the exploration of naturally integrating role-play and narrative into players' chosen actions. The narrative in tabletop role-playing games is directly tied to both players' decision-making and the consequences of their actions - the action that leads closer to a specific victory condition might not be the action the players want to take for the sake of role-play, and it is up to the Game Master to ensure the players' experience and enjoyment are upheld. Players may also desire to perform more stylish actions that hold more narrative purpose than mechanical, and although not rules-as-written legal, are still within the spirit of the game, and thus allowed with the GM's approval. These 'Rule of Cool' moments for players to attempt something outside the game's expected scope are frequently spontaneous, and although they can certainly try (Mercer 2016), such actions would be unfeasible within virtual game environments without being hard-coded in. A nascent paper by Gallotta et al. (Gallotta et al. 2024) delves into the potential usages of large language models (LLMs) (Kasneci et al. 2023) to benefit gameplay, such as their integration as GM roles and their use as core game mechanics, which may be a possible solution, whilst Värtinen et al. (Värtinen, Hämäläinen, and Guckelsberger 2024) recently explored the usage of ChatGPT-2 for generating game quests within RPGs. Fully automating role-play may risk an undesirable loss of creative expression from the collaborative storytelling aspects of TTRPGs, which is a challenge worth keeping in mind for further developments. However,

the investigation of LLMs working alongside digital environments is a potential opportunity that may help to realise this interpretation of described actions as legal mechanics, allowing GM agents to communicate directly with players to ask them how do they want to do this (Mercer 2016).

Procedural content generation. Lastly, as laid out in Figure 1, procedurally generating quests and combat encounters for ‘Descent’ may be an interesting avenue to explore, much like those within studies of ‘Gloomhaven’ (Gerhold and Tijben 2023; Tijben 2023) before. The full implementation of the environment allows simulation of the generated quests to bring additional value to the evaluation process and increase the quality of the results produced. Our D2e environment already parses quest data from JSON files, so it is easily extensible with new quest setups, although features like new Monster Actions or Heroic Feats will require implementation first.

7 Conclusions and Future Work

This paper described a new game environment adapted from the game ‘Descent: Journeys in the Dark (Second Edition)’ (D2e) in the Tabletop Games Framework (TAG) which presents many challenges in a complex one-versus-many environment, and initial experiments on the performance of Random, One Step Look Ahead (OSLA, a short-sighted greedy search) and Monte Carlo Tree Search (MCTS, a heuristic search of a game tree of available actions) agents. We tested these agents in 500 self-play games of 2-5 players, and 1150 games of robin-robin between the three agents, for a dataset of 7150 games total.

Overall, we observed both MCTS and OSLA to be capable in both roles of Hero and Overlord in one-versus-many play, with MCTS in particular able to make tactical decisions in a combat-based environment, and consider multiple possible end game conditions with multi-layered decision-making. We observed the OSLA player able to make progress towards winning for both sides, showing superiority as the Overlord against MCTS Heroes. Hero synergy analysis revealed Ashrian, Jain Fairwood and Syndrael to feature in most winning combinations, suggesting superior power or ease of play for general AI agents with minimal game-specific knowledge. Widow Tarha appears to be the Hero that requires most advanced strategic play, as the shorter look-ahead agent struggled to make efficient use of her abilities.

Our intended future work (see Figure 1) includes further exploration of modelling Game Masters for AI agents, making the gameplay experience more suitable for more ‘Dungeons & Dragons’-like environments. As we extracted many elements to construct the core tactical combat aspect as the base of our work, it would make sense to build more features on top of our current shell to bring the experience closer to a traditional TTRPG, such as additional quests, levelling up with experience points and unlocking wider class options, and downtime activities that focus more on player experience and game enjoyment over simply defeating opponents, such as travelling between quests or shopping for new equipment. We also hope by integrating more narrative and role-play elements into the chosen actions, we can strive towards

an AI Game Master that can provide a similar game experience to that of a human host. We aim for this to help better inform the design of human-AI semi-cooperative environments and agents in tabletop role-playing games.

Acknowledgements

For the purpose of open access, the author(s) has applied a Creative Commons Attribution (CC BY) license to any Accepted Manuscript version arising.

This work was supported by the EPSRC IGGI CDT (EP/S022325/1).

References

- Amini, S.; and Afsharchi, M. 2014. Finding Better Teammates in a Semi-cooperative Multi-agent System. In *2014 IEEE/WIC/ACM International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT)*, volume 3, 143–150. IEEE.
- Ashktorab, Z.; Liao, Q. V.; Dugan, C.; Johnson, J.; Pan, Q.; Zhang, W.; Kumaravel, S.; and Campbell, M. 2020. Human-ai collaboration in a cooperative game setting: Measuring social perception and outcomes. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW2): 1–20.
- Bjornsson, Y.; and Finnsson, H. 2009. Cadiaplayer: A simulation-based general game player. *IEEE Transactions on Computational Intelligence and AI in Games*, 1(1): 4–15.
- Browne, C. B.; Powley, E.; Whitehouse, D.; Lucas, S. M.; Cowling, P. I.; Rohlfshagen, P.; Tavener, S.; Perez, D.; Samothrakis, S.; and Colton, S. 2012. A survey of monte carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in games*, 4(1): 1–43.
- Callison-Burch, C.; Ippolito, D.; Reitter, D.; Tomar, G. S.; and Martin, L. 2022. Dungeons and Dragons as a Challenge Problem for Artificial Intelligence. In *NAACL Wordplay Workshop*.
- Chacón, P. S.; and Eger, M. 2019. Pandemic as a challenge for human-AI cooperation. In *Proceedings of the AIIDE workshop on Experimental AI in Games*, volume 3, 2.
- Chaslot, G.; Bakkes, S.; Szita, I.; and Spronck, P. 2008. Monte-Carlo tree search: A new framework for game AI. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 4, 216–217.
- Dafoe, A.; Hughes, E.; Bachrach, Y.; Collins, T.; McKee, K. R.; Leibo, J. Z.; Larson, K.; and Graepel, T. 2020. Open problems in cooperative AI. *arXiv preprint arXiv:2012.08630*.
- Dash, T.; Dambekodi, S. N.; Reddy, P. N.; and Abraham, A. 2020. Adversarial neural networks for playing hide-and-search board game Scotland Yard. *Neural Computing and Applications*, 32: 3149–3164.
- Egri-Nagy, A.; and Törmänen, A. 2020. The game is not over yet — Go in the post-AlphaGo era. *Philosophies*, 5(4): 37.
- Ellis, S.; and Hendler, J. 2017. Computers Play Chess, Computers Play Go... Humans Play Dungeons & Dragons. *IEEE Intelligent Systems*, 32(4): 31–34.

- Fantasy Flight Publishing, Inc. 2012. *Descent: Journeys in the Dark Second Edition*. Diamond Comic Distributors.
- Gaina, R.; Balla, M.; Dockhorn, A.; Montoliu Colás, R.; and Perez, D. 2020a. TAG: A tabletop games framework. In *AIIDE Workshops*. CEUR Workshop Proc.
- Gaina, R. D.; and Balla, M. 2022. TAG: Pandemic Competition. In *2022 IEEE Conference on Games (CoG)*, 552–559. IEEE.
- Gaina, R. D.; Balla, M.; Dockhorn, A.; Montoliu, R.; and Perez-Liebana, D. 2020b. Design and implementation of TAG: A tabletop games framework. *arXiv preprint arXiv:2009.12065*.
- Gaina, R. D.; Goodman, J.; and Perez-Liebana, D. 2021. TAG: Terraforming Mars. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 17, 148–155.
- Gallotta, R.; Todd, G.; Zammit, M.; Earle, S.; Liapis, A.; Togelius, J.; and Yannakakis, G. N. 2024. Large Language Models and Games: A Survey and Roadmap. *arXiv:2402.18659*.
- Gerhold, M.; and Tijben, K. 2023. Computer Aided Content Generation—A Gloomhaven Case Study. In *Proceedings of the 18th International Conference on the Foundations of Digital Games*, 1–10.
- Gray, M. 1986. *Fortress America*. Milton Bradley.
- Gygax, G.; and Arneson, D. 1974. *Dungeons & Dragons*. Tactical Studies Rules, Inc.
- Holmgård, C.; Green, M. C.; Liapis, A.; and Togelius, J. 2018. Automated playtesting with procedural personas through MCTS with evolved heuristics. *IEEE Transactions on Games*, 11(4): 352–362.
- Jackson, S. 1977. *OGRE*. Metagaming Concepts.
- Kasneji, E.; Seßler, K.; Küchemann, S.; Bannert, M.; Dementieva, D.; Fischer, F.; Gasser, U.; Groh, G.; Günemann, S.; Hüllermeier, E.; et al. 2023. ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and individual differences*, 103: 102274.
- Kitchen, S.; McGroarty, C.; and Aris, T. 2023. Model Representation Considerations for Artificial Intelligence Opponent Development in Combat Games. In *The International FLAIRS Conference Proceedings*, volume 36.
- Ling, B.; Liu, X.; Jiang, J.; Wu, W.; Wang, W.; Lyu, Y.; and Xu, X. 2022. Master multiple real-time strategy games with a unified learning model using multi-agent reinforcement learning. In *International Conference on Neural Computing for Advanced Applications*, 27–39. Springer.
- Martin, L. J.; Sood, S.; and Riedl, M. O. 2018. Dungeons and DQNs: Toward Reinforcement Learning Agents that Play Tabletop Roleplaying Games. In *INT/WICED@ AIIDE*.
- Mercer, M. 2016. The Rule of Cool! (Game Master Tips). Geek & Sundry, YouTube. Uploaded 16 February 2016. <https://youtu.be/fWZDuFIYkf0>.
- Park, J. S.; O’Brien, J.; Cai, C. J.; Morris, M. R.; Liang, P.; and Bernstein, M. S. 2023. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, 1–22.
- Perez-Liebana, D.; Cakmak, D.; Maghsudi, S.; Spronck, P.; and Thompson, T. 2023. 3.9 The Tabletop Board Games AI Tutor. *Human-Game AI Interaction*, 60–65.
- Piette, E.; Soemers, D. J.; Stephenson, M.; Sironi, C. F.; Winands, M. H.; and Browne, C. 2020. Ludii—the ludemic general game system. In *ECAI 2020*, 411–418. IOS Press.
- Press, P. 2017. *Cthulhu Confidential*. Pelgrane Press.
- Riedl, M. O.; and Bulitko, V. 2013. Interactive narrative: An intelligent systems approach. *Ai Magazine*, 34(1): 67–77.
- Riel, J.; and Monahan, R. 2024. Learning from Ludemes: An Inventory of Common Player Actions within Tabletop Role-Playing Games (TRPGs) to Inform Principled Design of Game-Based Learning Experiences. *International Journal of Role-Playing*, 15: 178–210.
- Samvelyan, M.; Rashid, T.; Schroeder de Witt, C.; Farquhar, G.; Nardelli, N.; Rudner, T. G.; Hung, C.-M.; Torr, P. H.; Foerster, J.; and Whiteson, S. 2019. The StarCraft Multi-Agent Challenge. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, 2186–2188.
- Santiago III, J. M.; Parayno, R. L.; Aiko Deja, J.; and Samson, B. P. V. 2023. Rolling the Dice: Imagining Generative AI as a Dungeons & Dragons Storytelling Companion. *arXiv e-prints*, arXiv-2304.
- Sfikas, K.; and Liapis, A. 2020. Collaborative agent gameplay in the pandemic board game. In *Proceedings of the 15th International Conference on the Foundations of Digital Games*, 1–11.
- Shin, M.; Kim, J.; and Kim, M. 2020. Measuring Human Adaptation to AI in Decision Making: Application to Evaluate Changes after AlphaGo. *arXiv e-prints*, arXiv-2012.
- Shin, M.; Kim, J.; and Kim, M. 2021. Human Learning from Artificial Intelligence: Evidence from Human Go Players’ Decisions after AlphaGo. In *43rd Annual Meeting of the Cognitive Science Society (CogSci 2021): Comparative Cognition: Animal Minds*, 1795–1801. The Cognitive Science Society.
- Shyne, F. 2023. *Automatic Play-testing of Dungeons and Dragons Combat Encounters*. Honors in the Department of Computer Science, Union College.
- Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; Van Den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; et al. 2016. Mastering the game of Go with deep neural networks and tree search. *nature*, 529(7587): 484–489.
- Silver, D.; Hubert, T.; Schrittwieser, J.; Antonoglou, I.; Lai, M.; Guez, A.; Lanctot, M.; et al. 2017. Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm. *CoRR*, abs/1712.01815.
- Sousa, M. 2023. Finding Ludemes in Modern Board Games: Analyzing the Top Number One Games of Board Game Geek. *Board Game Studies Journal*, 17(1): 205–229.

- Sutton, R. S.; Precup, D.; and Singh, S. 1999. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence*, 112(1-2): 181–211.
- Tijben, C. 2023. *Generating Gloomhaven dungeons through evolutionary game design*. B.S. thesis, University of Twente.
- Treanor, M.; Zook, A.; Eladhari, M. P.; Togelius, J.; Smith, G.; Cook, M.; Thompson, T.; Magerko, B.; Levine, J.; and Smith, A. 2015. AI-based game design patterns. In *Proceedings of the 10th International Conference on the Foundations of Digital Games 2015 (FDG 2015)*. USA: Society for the Advancement of Digital Games. ISBN 9780991398249.
- Tychsen, A.; Hitchens, M.; Brolund, T.; and Kavakli, M. 2005. The Game Master. In *ACM International Conference Proceeding Series*, volume 123, 215–222.
- van den Bos, P.; and Stoelinga, M. 2023. With a little help from your friends: semi-cooperative games via Joker moves. In *International Conference on Formal Techniques for Distributed Objects, Components, and Systems*, 155–172. Springer.
- Vinyals, O.; Babuschkin, I.; Chung, J.; Mathieu, M.; Jaderberg, M.; Czarnecki, W. M.; Dudzik, A.; Huang, A.; Georgiev, P.; Powell, R.; et al. 2019. Alphastar: Mastering the Real-Time Strategy Game StarCraft II. *DeepMind blog*, 2: 20.
- Värtinen, S.; Hämäläinen, P.; and Guckelsberger, C. 2024. Generating Role-Playing Game Quests With GPT Language Models. *IEEE Transactions on Games*, 16(1): 127–139.
- Zhu, A.; Martin, L.; Head, A.; and Callison-Burch, C. 2023. CALYPSO: LLMs as Dungeon Master’s Assistants. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 19, 380–390.