

TAG: Terraforming Mars

Raluca D. Gaina, James Goodman, Diego Perez-Liebana

Game AI Group
Queen Mary University of London, UK

Abstract

Games and Artificial Intelligence (AI) have had a tight relationship for many years. A multitude of games have been used as environments in which AI players can learn to act and interact with others or the game mechanics directly; used as optimisation problems; used as generators of large amounts of data which can be analysed to learn about the game, or about the players; or used as containers of content which can be automatically generated by AI methods. Yet many of these environments have been very simple and limited in scope. We propose here a much more complex environment based on the boardgame Terraforming Mars, implemented as part of the Tabletop Games Framework: a very large and dynamic action space, hidden information, large amounts of content, resource management and high variability make this problem domain stand out in the current landscape and a very interesting problem for AI methods of multiple domains. We include results of baseline AI game-players in this game and in-depth analysis of the game itself, together with an exploration of problem complexity, challenges and opportunities.

1 Introduction

We present a new game environment for Artificial Intelligence methods based on the boardgame “Terraforming Mars”. This is a game for 1-5 players published by FryxGames in 2016 and currently ranked as one of the best board games in the world¹, with several expansions and a multitude of fan-made content. In this game, players take on the role of CEO of a giant corporation, all aiming to terraform Mars, while competing to obtain the most points at the end. We summarise the main challenges in this game as follows: large, diverse and dynamic action space; large state space; stochasticity; hidden information; wide variety of point-earning and conflicting strategies; large variations in game set-up options; large game parameter space.

As such, we consider this game a very important and intriguing problem for AI players, with parallels that can be drawn to Strategy Games of wider interest to the community, such as “StarCraft”. We aim to study not only the performance of AI techniques in this game, but also carry out an

in-depth analysis of the game itself, to better inform human gameplay. Similar research into complex tabletop games has been carried out in several board and card games. Walton-Rivers et al. (2019), and later Bard et al. (2020), presented “Hanabi” as an interesting challenge for AI with its unique aspect of reverse hidden information: each player can see the others’ cards in hand, but not their own. In this cooperative game, AI players manage to surpass the main challenge of effective communication and reach near-perfect scores. Justensen et al. (2019) introduce “Blood Bowl”, a fantasy football board game, as an AI environment and competition, with its main challenge being a large action space: players need to manage several units across a large board and coordinate their actions for success. Later, “Pandemic” was explored as a complex AI challenge by Chacon et al. (2019) and Sfikas et al. (2020): this cooperative board game combines different types of components (a large board, cards and tokens), with the challenge of managing and coordinating resources between players under time pressure. We consider “Pandemic” to be most similar to the problem domain proposed here; yet, “Terraforming Mars” includes a larger state space, more noise in AI simulations, many players competing for similar goals, and (with some exceptions) a larger action space.

We see applications of search methods which have dominated other domains such as General Game Playing (Genesereth, Love, and Pell 2005) in several board games as well. Monte Carlo Tree Search (Browne et al. 2012) was successfully applied in “Carcassonne” (Heyden 2009), achieving good performance and beating an intermediate beginner human player, but not advanced players; “Settlers of Catan” (Szita, Chaslot, and Spronck 2009), outperforming baseline AI players, but falling short in tests against human players; and “7 Wonders” (Robilliard, Fonlupt, and Teytaud 2014), winning against baseline AI and human players. These games are all diverse and present different challenges, yet we note in particular results of in-depth parameter analysis in “7 Wonders”: the authors find that MCTS theory, results and best parameters established in classical abstract games hold true in their test domain as well. Although “Terraforming Mars” presents different challenges to AI players and a setting which is not completely adversarial or zero-sum, we take these findings forward and apply MCTS in our work.

Table 1: Features of latest TAG games. Average values reported for 100 runs with 2 random players. Speed in milliseconds, for calls to the *next* (N) and *copy* (C) functions of the FM. AS = action space size, HI = hidden information (percentage of game components hidden from the player), SS = state space size (total number of game components), Length = total number of decisions per game.

Feature	Diamant	Dominion	Terraforming Mars
#Players	3-8	2-4	1-5
Speed(N)	0.09	2.50	26.8
Speed(C)	0.19	4.99	125
AS	2	5.04	7.10
HI	0.83	0.51	0.42
SS	35	65	400
Length	19.3	237	460

Eger and Martens (2019) use AI players using a system based on Dynamic Epistemic Logic, with different strategies, to reason about and play “Ultimate Werewolf” against human players, yet found their method to not be regarded as very skillful by the humans. Strategies are evaluated and analysed in-depth in “Ticket to Ride” as well, where de Mesentier Silva et al. (2017) use a series of AI players with different play-styles to test variants of the game with different maps, rules or number of players. Their results show that although preferred play-styles differ per game setting, some cities remain consistently more desirable as the focus of play. Further, the automatic players are able to find game states not covered by the game rules, suggesting automatic play-testing of tabletop games as a ripe area for research. Later, Witter and Lyford (2020) use probabilistic and graph theory concepts to recommend simpler ways of winning “Ticket to Ride” than one might expect. We adopt here a similar manner of breaking down the game analysis and presenting a first analysis of game strategies in “Terraforming Mars”.

“Terraforming Mars” was implemented within the Tabletop Games framework (TAG) (Gaina et al. 2020b)², which aims to study general AI for modern complex tabletop games, and offers a variety of ready-made game components, actions, and rules, as well as analysis and visualisation tools to aid in implementing digital version of tabletop games, as well as AI players to play them. The framework includes 10 other games at the time of writing, ranging from simple abstract domains such as “Tic Tac Toe” to complex environments such as “Pandemic”, “Colt Express” or “Dominion”. We show in Table 1 (full details on Github³) a summary comparison of our implementation to the most recent games added into the framework. We highlight a larger average action space, which increases throughout the game;

²<https://gaigresearch.github.io/projects/TAG>

³<https://github.com/GAIGResearch/TabletopGames/wiki/Game-Statistics>. We note that there may be differences to the physical board games, due to implementation simplifications for the digital versions in TAG.

a similar level of hidden information to ‘Dominion’, and a much larger state space size (due to the hundreds of cards included with the game) and game length. “Terraforming Mars” is slower, due to the large state space and complex mechanics split over different game phases. Additionally, this is the first game in the framework which allows for solo play, to test the abilities of an AI game-player individually, without external influences.

We summarise our contributions as follows: we present and analyse in-depth both the problem domain and the results of 3 AI players compared to publicly available human game statistics, to gain insights into successful strategies and behaviours in this environment.

2 Terraforming Mars

This description of the game refers to the *base game* only, disregarding components, actions and rules added by expansions. Full game rules can be found online⁴.

2.1 Rules Overview

The goal of the game is to terraform Mars, by increasing each of the 3 Global Parameters (Oceans, Temperature and Oxygen) to maximum. The game ends when the goal is achieved, and the player with the highest number of points wins. Points are all summed up at the end of the game and belong to 6 different categories: Terraforming Rating (TR, obtained by increasing Global Parameters or by playing certain cards), board points (obtained by placing City and Greenery tiles on a hexagonal board), milestones (obtained by meeting a certain condition first), awards (obtained by having the highest amount of a certain game component), and cards (obtained by playing cards which award points). The game is played over several *Generations*, each split into 3 phases: Research, Actions and Production. In the Production phase, each player earns resources according to the corresponding production level (with two exception: Energy resources turn into Heat just before production happens; and TR adds up to the Mega Credit production); the Research and Actions phases are described below.

The game has many expansions which add new components or new mechanics, and several variants for the base game as well: Corporate Era includes new project and corporation cards, and enforces all starting resource production to 0, instead of 1. The 1-player (solo) variant becomes a race against time: the player begins with 14 TR instead of 20, and must finish terraforming Mars by the end of generation 14.

2.2 Components

The game is played on a **hexagonal board**, where tiles of different types can be placed. Several counters keep track of the state of the **3 Global Parameters** and each player’s **TR**. Each player has 12 more counters which keep track of their **resources**, and their **production** (or income) of each resource (see Figure 1).

There are 12 **Corporation Cards**. Players choose 1 out of a random selection of 2 at the start of the game, which offers

⁴<https://www.fryxgames.se/TerraformingMars/TMRULESFINAL.pdf>



Figure 1: Resources (top) and production (bottom), from left to right: Mega Credit, Steel, Titanium, Plant, Energy, Heat.

them starting resources and, possibly, persisting effects, or new actions that player can perform (see Figure 2 left).

There are 208 total **Project Cards** (see Figure 2 right), out of which players get dealt 4 each in the Research phase of each Generation (with the exception of the first generation, when each player receives 10 cards instead). Cards have a cost, usually paid with Mega Credit resources (top-left corner), requirements (top-middle; conditions for the card to be playable, if any), tags (top-right corner), a name (top banner), effects applied when played (middle), and points (bottom-right corner). There are 3 types of cards: Automated (immediate effects and tags only), Active (tags, persisting effects and/or unlocking new actions) and Event (immediate effects only; tags only count when the card is played, but not for any other subsequent effects referring to tags played).

2.3 Actions

In the *Research phase*, players may buy cards to their hand from a random selection, using Mega Credit resources.

In the *Actions phase*, players take turns performing 1 or 2 actions, with the following options: Play a card from hand. Actions from their Corporation or Active cards in play. Buy one of 6 Standard Projects (e.g. pay 25 Mega Credits to place a City tile on the board and increase Mega Credit production by 1). Claim a Milestone. Fund an Award. Trade 8 Plant resources for a Greenery tile, or 8 Heat resources to increase the Temperature Global Parameter. Pass (can no longer perform actions this phase; when all players pass, the Actions phase is over).

Terraforming Mars in TAG The version implemented in the TAG framework includes the base game and Corporate Era components, as described above. Expansions are designed in a modular way, such that they can be added in to enable the bonus mechanics, setup or game components they bring. The expansions Hellas, Elysium (which add new boards, milestones and awards) and Venus (which adds a new global parameter, standard project and cards) are implemented as well, but not fully parsed at the time of writing. Most actions on the cards are fully implemented and functional, with the exception of 4 cards featuring:

1. Temporary effects (e.g. discount for next card played).
2. Conditional alternative effects (e.g. “receive X plant production, or Y plant production instead if you have at least Z plant tags in play”).

Actions are highly parameterised and implemented using a sequential model, to avoid extremely large combinatorial action spaces (e.g. a “Play a card” action could be followed by decisions on which resources to use to pay for the card and how to apply the card’s effects, before the action is



Figure 2: Example cards. Left: Corporation Card (Phoblog); the player starts with 10 Titanium and 23 Mega Credit, and, as an effect, Titanium resources are worth 1 extra Mega Credit (4 instead of 3). Right: Project Card (Methane from Titan); it is Automated, costs 28, has Space and Jovian tags, requires Oxygen to be minimum 2, increases Heat and Plant production by 2 each, and gives 2 points at the end.

fully executed). We make use of the new *IExtendedSequence* framework included in TAG for this functionality. What this means for AI players is that there is a delay between executing an action and observing its full effect, which introduces additional uncertainty depending on the player’s horizon (how long the simulated action sequences are).

The modular implementation allows for new cards or full expansions to easily be defined and parsed from JSON, or new action classes created to add new mechanics. The game’s parameters are fully exposed as well, with a total of 30 high-level adjustable options, such as the base exchange rate for Steel and Titanium resources to Mega Credits.

API AI players have access to the complete list of actions currently available and legal, as well as those that are not legal due to lack of resources or missing requirements, through the *Forward Model* object provided. This distinction is made in this game as other actions may become playable depending on a player’s decisions, thus having access to the full list of possibilities (which would also be available to human players in this environment) allows for more informed decision-making.

The *Game State* object received every time a decision is needed encapsulates the entirety of the state that can be observed by the player, including all components (cards, boards, Milestones and Awards, resource counters, Global Parameter counters, etc.) and other game information (current generation, player active effects and actions). The hidden information in this game is the face-down draw decks and the opponent hands (all of which are randomly shuffled in the observation received by the player).

3 Challenges and Opportunities

Overall, the game presents players with a large, dynamic and diverse action space. We highlight the difference to other domains with large action spaces, where most actions are of the same type (e.g. move left/right, up/down; play a card; place a token on a board) and many actions may result in similar effects with slight variations. Here, the game state is made up of hundreds of components, and each action will have a

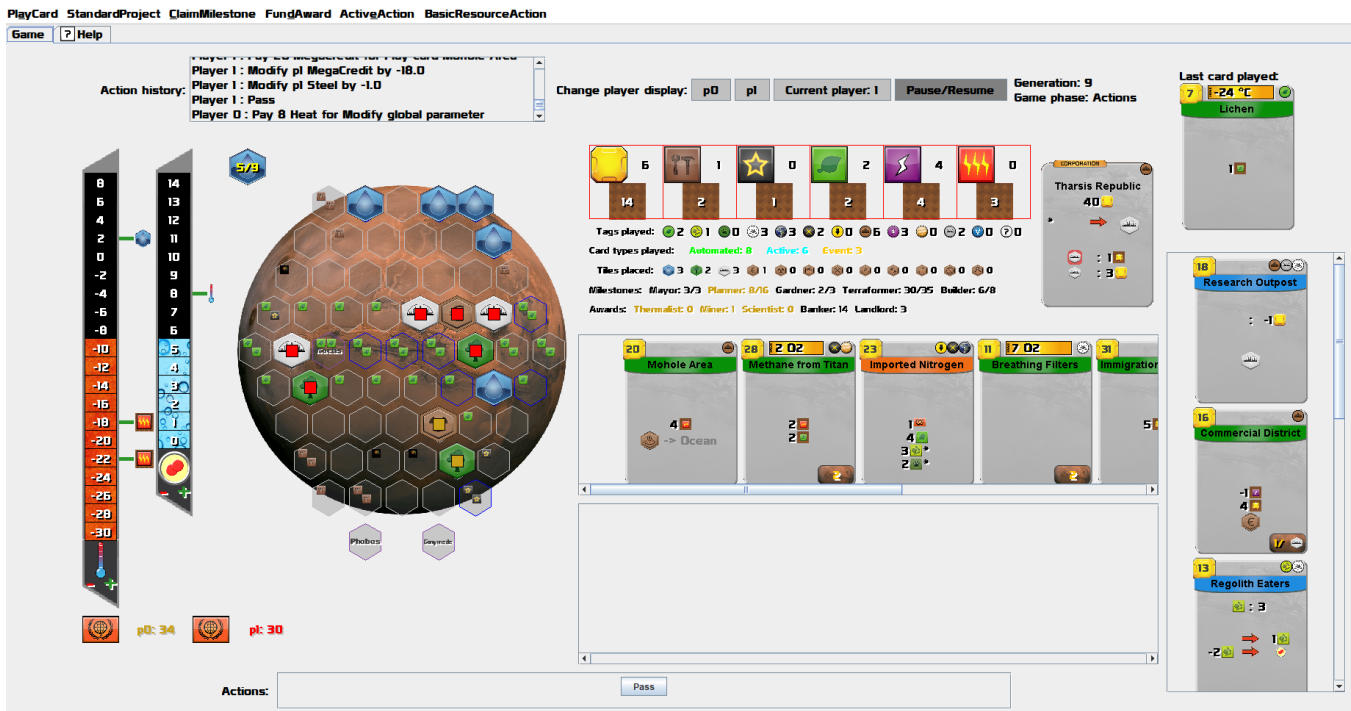


Figure 3: Graphical User Interface for Terraforming Mars in TAG: a mid-game game state between MCTS and a human player.

vastly different effect on one or more of these components. Further, once an action is selected, more may become available, while others would turn illegal. Table 1 shows the average action space size in ‘Terraforming Mars’ to be 7.10, although the maximum observed in 2-player random games reached 46. With expert players, the action space becomes even larger as the game progresses, with the order of actions important as well in key moments in the game (e.g. when to claim Milestones or fund Awards; to claim bonuses from Global Parameters etc.).

Further, the game presents diverse and dynamic strategies which can be used to win. These can be used in isolation, combined or changed throughout the game depending on resources available, in order to maximise performance, with no one strategy identified as dominant. We identify the following large categories of strategies: resource card engine, action card engine (focus on increasing action space), automated card engine (focus on points and resource production), effect card engine (focus on active cards with persisting effects, such as discounts), building engine (focus on populating the board), and terraforming (focus on finishing the game quickly). Although a lot of the knowledge and expertise in the game can be expressed as general advice, the actual game situation heavily influences one’s decisions (in particular, which cards and which actions are actually available to the player), making it a very difficult challenge to encode this knowledge in automatic players.

Lastly, given the sheer amount of cards available, future game states are very difficult to approximate, making AI simulations extremely noisy. As this is also the key in predicting the quality of each decision made, with every game

and best strategy to adopt heavily influenced by the order of cards in the deck, we consider the hidden information here another of the biggest challenges the environment proposes for AI players.

4 Experiments

We use the publicly-available dataset by *ssimeonoff*⁵ to study the game and compare AI players to human players. At the time of writing, there are a total of 9364 games recorded for human players. However, only 93 match the setting tested for this paper (using base game only, Corporate Era, and playing on the base map Tharsis, with base milestones and awards), split as follows: 63 2-player, 26 3-player, 2 4-player and 2 5-player. Given this low number, we acknowledge that human averages are not statistically meaningful (even more so due to the lack of information on the players, we do not know their level of experience with the game), but we keep them as a baseline of comparison.

We run 3 automatic players on the game in their default configuration in TAG: random (randomly chooses 1 legal action), One Step Look Ahead (OSLA; greedily chooses the legal action which leads to the next best state) and Monte Carlo Tree Search (MCTS; iteratively builds a statistical tree by simulating several steps into the future and choosing the best action at the end; for more details on the algorithm, the reader is referred to (Browne et al. 2012), and to (Gaina et al. 2020a) for details on its implementation in TAG). MCTS uses a rollout depth of 10, a UCT K constant of $\sqrt{2}$, and 1000 calls to the Forward Model *next()* and *copy()* functions

⁵<https://ssimeonoff.github.io/>

Table 2: Average final generation reached and score for human, random, OSLA and MCTS players in 2-5-player mirrored games. Average rounded to nearest integer, standard error in brackets.

#players	Generation				Score			
	Human	Random	OSLA	MCTS	Human	Random	OSLA	MCTS
2	14 (2.48)	20 (0.21)	22 (0.21)	18 (0.35)	106 (2.48)	93 (1.44)	114 (1.46)	115 (2.77)
3	12 (0.26)	18 (0.18)	17 (0.19)	15 (1.17)	84 (1.80)	69 (0.85)	88 (0.99)	95 (4.74)
4	10 (1.00)	15 (0.19)	15 (0.13)	12 (0.85)	67 (4.87)	62 (0.75)	78 (0.72)	84 (2.59)
5	12 (0.00)	13 (0.17)	13 (0.12)	10 (0.33)	62 (4.75)	52 (0.48)	70 (0.54)	76 (3.00)

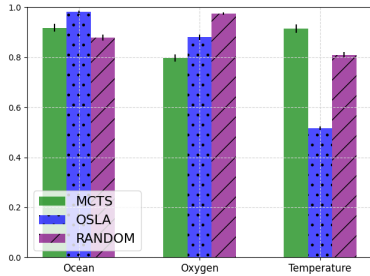


Figure 4: Generation where each Global Parameter was terraformed (normalized based on final generation) for MCTS, OSLA and random in mirrored 2-player games.

as its budget. Both MCTS and OSLA use the same heuristic function to evaluate game states: the number of points they would get if the game were to end in that game state.

We refer to games in which all players are copies of the same agent as *mirrored* games. Each test configuration was run 100 times: mirrored games for all AI players with 1-5 players, and 2-player match-ups of all AI player pairs (MCTS-random, MCTS-OSLA and OSLA-random). We use the logging system available in the framework to record statistics on all games played in line with those recorded for human players in our dataset, along with additional information that we can extract for more in-depth analysis of the game and the performance of the AI players. Examples are the generation when the game finished and the score at each generation (including final score), broken down into each category: TR, milestones and awards, board and cards, the game result (winning player), milestones and awards in play and player corporations.

5 Discussion

In Table 2 we show the average generation and average final score obtained by each of the players under test, including a summary of the human data available, for each configuration of 2, 3, 4 and 5 players. Overall we observe a trend where the numbers of generations required to finish the game is inversely proportional with the number of players, for both humans and AI. Since more players can contribute to the terraforming process, it makes sense that this would be achieved in less generations. MCTS achieves a number of generations closest to human players, while random and OSLA are comparable. In Figure 4 we show more detailed information on the AI players regarding how fast each of the

parameters are terraformed, information which is not available for human players. We notice here that there is a large difference in OSLA’s preference to terraform the Temperature parameter in the first half of the game. We hypothesize that this is due to the standard project (which is always available and can be repeated in a generation) immediately earning a point, and being cheap enough to repeatedly appear in the list of legal actions. OSLA’s short-sightedness considers this action highly beneficial as a result. MCTS and random show a fairly even split across the 3 Parameters. On average, the Ocean parameter is terraformed last, even though it has the smallest amount of steps required; the standard project that places Oceans is more expensive, and only rewards 1 point, whereas the standard project that places Greenery tiles (and therefore raising Oxygen levels) gives a minimum of 2 points, as Greenery tiles are counted separately at the end.

We observe a similar trend for score, with more points being obtained in games with less players, for much the same reasons as mentioned above (similar number of points will be available each game, to be split between more players). Interestingly, we observe both OSLA and MCTS to obtain higher points than human players, with MCTS gaining the most (despite its shorter games shown by lower generation numbers). We note that this is not a direct comparison (but rather a comparison of humans playing against other humans, and MCTS against a copy of itself). Random players do score below humans, showing that even the simple heuristic of maximising points, with no other knowledge of strategies, is quite effective. This finding opens the door for studies on more advanced heuristics. Figure 5 shows the score progression on each generation for MCTS mirrored 2-player games, broken down into the different categories. We observe that cards and board points are fairly similar between the winning and the losing player, and the biggest difference is made by TR and Milestones and Awards: in these categories, the winning player takes the lead early on. In human games, board and card points can often skew the TR or Milestone and Award imbalance: however, the AI players in these experiments do not have policies for placing tiles correctly on the board or choosing and playing the right cards - we hypothesize that heuristic functions focusing on these aspects in particular will see a large boost in performance.

We can extract further interesting statistics from the games played by AIs. In Figure 6 we show the final resource production obtained by each of the players in 2-player mirrored games. We notice that all players favour MegaCredit production, and that the winning player had the highest overall production in all resources. From the ran-

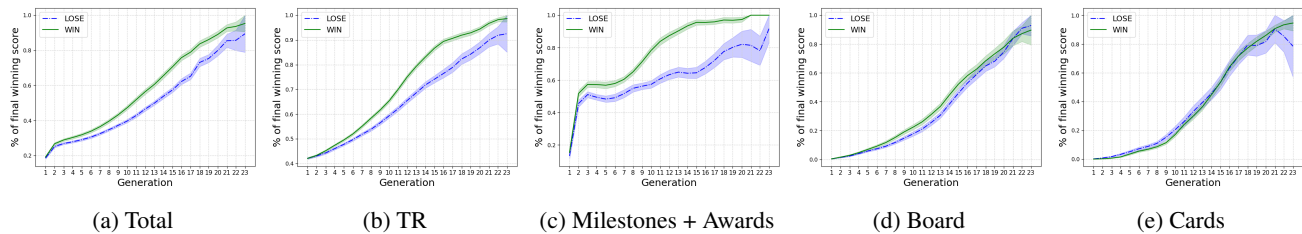


Figure 5: Average score per generation obtained by the winning and losing player, respectively (MCTS 2-player mirrored games). Showing percentage of the final winning score in each category for a normalized view of the score differences. Generation on the X-axis and score percentage on the Y-axis.

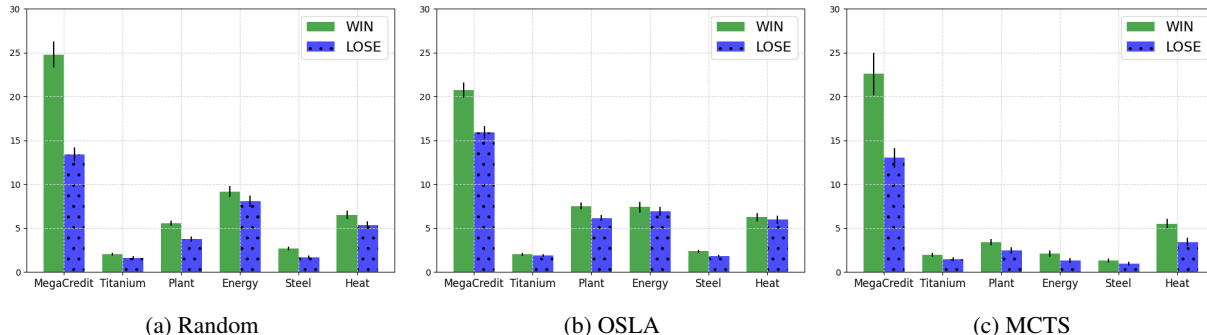


Figure 6: Final production obtained in 2-player mirrored games. Resources on the X-axis, average amount on the Y-axis (including standard error). Production, in order: MegaCredit, Titanium, Plant, Energy, Steel, Heat.

dom player’s performance we can deduce that MegaCredit production (and Energy second, both available through standard projects, unlike the others) is one of the easiest to increase. However, MCTS, the best AI player, ends up with overall lower production in most resources than the other agents, and a focus on Plant and Heat productions, both of which allow for executing more of the basic resource actions and therefore earning more points; Steel and Titanium resources are beneficial for playing cards (and Energy production often a requirement), showing the agent’s lack of focus on strategies involving card engine building. On the one hand, this could suggest that the best way to play the game is to choose the simplest way of gaining points. However, this is yet to be tested against more complex heuristics, which we aim to explore in future work.

MCTS players do show the advantage in long-term planning through award funding, where the player who funded the award ends up winning the points over 70% of the time, as opposed to OSLA and random that see a success rate of 50% or, often, lower. Further, when tested head-to-head in 2-player games, MCTS wins 75% (0.42) of its games against OSLA and 98% (0.01) against random, while OSLA wins 91% (0.03) of its games against random. Overall, these tests suggest MCTS to be the strongest agent.

The *ssimeonoff* database includes 5705 solo games that match the setting used in this study, seeing humans obtain an average score of 75(0.50), with uncertainty of the winning percentage. However, no AI player is able to beat the solo mode of the game, and they obtain a much lower score:

58(0.44) for MCTS, 39(0.40) for OSLA, 30(0.43) for random. This is not surprising given the long games observed, most over the limit of 14 generations imposed in the solo mode, yet it remains an interesting open challenge: how can AI players be taught to efficiently terraform Mars?

6 Conclusions

This paper describes a new game environment for AI methods based on the game “Terraforming Mars” and initial experiments to analyse the game itself, as well as the performance of baseline players included in the Tabletop Games framework (TAG): random, One Step Look Ahead (greedy short-sighted search) and Monte Carlo Tree Search (MCTS). We test these agents against publicly available data recorded from human games. Overall we observe the MCTS player to be on par with human performance, or even better - however, this is not a direct comparison (but rather a comparison of humans playing against other humans, and MCTS against a copy of itself). With the inclusion of “Terraforming Mars” in TAG with an intuitive graphical user interface, we hope to encourage humans to play against the AI directly, to also be able to record more in-depth data about the games and better analyse the AI performance.

Moreover, the parameters used for the MCTS player were as default in the TAG framework; given the multiple decisions which may need to be taken in order to fully play out an action, it is likely that the algorithm could perform much better if given a longer search horizon. The framework includes a more advanced version of MCTS, with many

tunable parameters; modifying these so that the algorithm achieves its full potential in the game could paint a clearer picture of its playstyle, as well as its true strength when compared with human players.

In our work we maintained a simple generic heuristic function regardless of game phase. However, strategies for specific game moments are considered in “Risk” by Gibson et al. (2010), who use MCTS for drafting territories, with a heuristic function learned via linear regression from a set of manually-defined features; this is shown to significantly improve the performance of AI players against the strongest techniques previously tried in the game. We consider this work important for future development of AI players in games with several different phases, such as “Terraforming Mars”.

We can further dive into other topics as well, such as game balance, or categorisation of specific corporation and project cards: with such a vast amount of content as is available in “Terraforming Mars” (even disregarding the multitude of fan-made content), it is a given that every single card is not fully play-tested and ensured to be balanced. The implementation of the game within the TAG framework can offer exactly this missing aspect, allowing for use of TAG tools to analyse and balance the game and its components to maximise player enjoyment. Further, procedural content generation techniques can easily be applied to the structured data included in the game, to create new components or rules for “Terraforming Mars”, building on work by de Mesentier Silva et al. (2018) in “Ticket to Ride” for map generation, or Summerville and Mateas (2016) in “Magic the Gathering” for full card generation from partial specifications.

Acknowledgements

Many thanks to Bob Dillon for playtesting the environment!

References

- Bard, N.; et al. 2020. The Hanabi Challenge: A New Frontier for AI Research. *Artificial Intelligence* 280: 103216.
- Browne, C. B.; et al. 2012. A Survey of Monte Carlo Tree Search Methods. *IEEE Transactions on Computational Intelligence and AI in games* 4(1): 1–43.
- 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*.
- de Mesentier Silva, F.; Lee, S.; Togelius, J.; and Nealen, A. 2017. AI-Based Playtesting of Contemporary Board Games. In *Proceedings of the International Conference on the Foundations of Digital Games - FDG'17*. ACM Press. doi: 10.1145/3102071.3102105. URL <https://doi.org/10.1145/3102071.3102105>.
- de Mesentier Silva, F.; Lee, S.; Togelius, J.; and Nealen, A. 2018. Evolving maps and decks for ticket to ride. In *Proceedings of the 13th International Conference on the Foundations of Digital Games*, 1–7.
- Eger, M.; and Martens, C. 2019. A Study of AI Agent Commitment in One Night Ultimate Werewolf with Human Players. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 15, 139–145.
- Gaina, R. D.; Balla, M.; Dockhorn, A.; Montoliu, R.; and Perez-Liebana, D. 2020a. Design and Implementation of TAG: A Tabletop Games Framework. *arxiv:2009.12065*.
- Gaina, R. D.; Balla, M.; Dockhorn, A.; Montoliu, R.; and Perez-Liebana, D. 2020b. TAG: A Tabletop Games Framework. In *Experimental AI in Games (EXAG), AIIDE 2020 Workshop*.
- Genesereth, M.; Love, N.; and Pell, B. 2005. General game playing: Overview of the AAAI competition. *AI magazine* 26(2): 62–62.
- Gibson, R.; Desai, N.; and Zhao, R. 2010. An automated technique for drafting territories in the board game Risk. In *Sixth Artificial Intelligence and Interactive Digital Entertainment Conference*.
- Heyden, C. 2009. Implementing a computer player for Carcassonne. *Master's thesis, Department of Knowledge Engineering, Maastricht University*.
- Justesen, N.; Uth, L. M.; Jakobsen, C.; Moore, P. D.; Togelius, J.; and Risi, S. 2019. Blood Bowl: A New Board Game Challenge And Competition For AI. In *2019 IEEE Conference on Games (CoG)*, 1–8. IEEE.
- Robilliard, D.; Fonlupt, C.; and Teytaud, F. 2014. Monte-carlo tree search for the game of “7 wonders”. In *Workshop on Computer Games*, 64–77. Springer.
- Sfikas, K.; and Liapis, A. 2020. Collaborative Agent Gameplay in the Pandemic Board Game. In *International Conference on the Foundations of Digital Games*, 1–11.
- Summerville, A. J.; and Mateas, M. 2016. Mystical Tutor: A Magic: The Gathering Design Assistant Via Denoising Sequence-To-Sequence Learning. In *Twelfth artificial intelligence and interactive digital entertainment conference*.
- Szita, I.; Chaslot, G.; and Spronck, P. 2009. Monte-Carlo Tree Search in Settlers of Catan. In *Advances in Computer Games*, 21–32. Springer.
- Walton-Rivers, J.; Williams, P. R.; and Bartle, R. 2019. The 2018 Hanabi Competition. In *2019 IEEE Conference on Games (CoG)*, 1–8. IEEE.
- Witter, R. T.; and Lyford, A. 2020. Applications of Graph Theory and Probability in the Board Game Ticket to Ride. In *International Conference on the Foundations of Digital Games*, 1–4.