The 2022 Ludii AI Competition

Éric Piette[1] *, Dennis J.N.J. Soemers[2], Matthew Stephenson[3] *and* Cameron Browne[2]

[1] *ICTEAM, UCLouvain, Belgium*

[2] *Department of Advanced Computing Sciences, Maastricht University, The Netherlands*

[3] *College of Science and Engineering, Flinders University, Australia*

E-mail: eric.piette@uclouvain.be

Abstract. The Ludii AI Competition involves general game playing events focused on developing agents that can play a wide variety of board games. In the 2022 edition, three competition tracks were proposed: Kilothon, General Game Playing, and Learning. All tracks used the Ludii general game system to provide the necessary games and API. This paper reports the motivation, context, and results of the 2022 Ludii AI Competition.

Keywords: General Game Playing, Board Game, Competition, Ludii, Kilothon

1. INTRODUCTION AND MOTIVATION

General Game Playing (GGP) is a field of research in which the goal is to develop agents that can play general games without human intervention (Pitrat, 1968). The games are typically described in a specific Game Description Language (GDL). For example, the Stanford GDL (Love et al., 2008), in which the rule descriptions are provided in low-level logic, led to the description of many dozens of games and to a new impetus for GGP research (Genesereth and Thielscher, 2014). The development of many innovative and state-of-the-art GGP approaches was motivated by different GGP competitions, such as the International General Game Playing Competition (IGGPC) (Genesereth et al., 2005; Genesereth and Björnsson, 2013) organised from 2005 to 2016, in which the Stanford GDL was used. The majority of successful approaches were based on Monte-Carlo Tree Seach techniques (Browne et al., 2012), such as Ary (Méhat and Cazenave, 2010) and Cadiaplayer (Björnsson and Finnsson, 2009). Other original and successful approaches, such as WoodStock (Koriche et al., 2017b) (the last IGGPC winner in 2016), proposed to combine Monte-Carlo simulations with stochastic constraint-based search propagation techniques, and to filter equivalent positions and moves by the detection of contraint-based symmetry (Koriche et al., 2017a). Other GGP approaches from the different editions of IGGPC can be found in (Świechowski et al., 2015).

IGGPC has not been organised since 2016, and we can observe a slowdown in GGP advancements for board games[1]. One of the reasons is that the Stanford GDL has many issues regarding how it models games. This low-level logic language does not capture the many commonly-used high-level game concepts (such as geometry of the board, common end conditions, or the basic moves). Each of these game concepts needs to be expressed from scratch in any new game description that requires them. This limitation is addressed in the recent Ludii general game system (Browne et al., 2020;

*Corresponding author. E-mail: eric.piette@uclouvain.be.

[1] GGP literature is available at: http://www.general-game-playing.de/literature.html

1389-6911/$35.00 © 0 – IOS Press. All rights reserved.

Piette et al., 2020), which is being developed as part of the European Research Council (ERC) funded Digital Ludeme Project (DLP).[2]

In the Ludii game description language, games are not written in low-level logic, but rather expressed using a large library of ludemes (Browne, 2021). Many commonly-used high-level concepts are summarised in a single ludeme, or combinations of a few ludemes. This makes it significantly easier for humans to write, read, modify, and understand game descriptions, in particular for board games that are traditionally or commonly played around the world (Browne et al., 2019). The Ludii general game system,[3] in its current version 1.3.7, includes 1,146 games and has a community of designers frequently adding new games.

Ludii's object-oriented state representations (Piette et al., 2021) similarly makes concepts such as game boards, graph elements, directions and connections, and other commonly-used game aspects, available for any game to any agent. This, together with some examples,[4] instructions (Soemers et al., 2022b), tutorials,[5] and practical demos (Browne et al., 2020), facilitate the implementation of GGP agents by junior and senior developers and researchers with Ludii. For this reason, and thanks to the possibility to play remotely with Ludii, the DLP team proposed to use this general game system as a new competition platform, particularly to renew the motivation to undertake the GGP challenge (Stephenson et al., 2019).

2. THE 2022 LUDII AI COMPETITION

The 2022 Ludii AI Competition was run by the DLP team members of Maastricht University's Department of Advanced Computing Sciences over July 25–29, 2022.

This competition was a set of GGP events focussed on developing agents that can play a wide variety of board games. The Ludii version used for this competition (1.3.2) includes a substantial number and wide diversity of games. This made it possible to propose tracks requiring agents to play over 1000 games.

All tracks were run online alongside the 25^{th} Computer Olympiad.[6] The communication between the organisers and the participants was realised through a Slack chat, and every participant was required to be present and online 30 min before the beginning of each round. Finally, each game played remotely during the competition could be followed by any spectator thanks to the Graphical User Interface (GUI) and the spectator mode offered by Ludii.

2.1. Competition Games

All games were provided to the agents in the Ludii game description format (*.lud), and implemented and run in Ludii. In this edition, the organisers decided to guarantee up front that all the games would be:

- Adversarial.

[2]http://ludeme.eu
[3]Ludii is publicly available at: https://ludii.games/
[4]Ludii Example AI GitHub Repository is available at: https://github.com/Ludeme/LudiiExampleAI
[5]Ludii Tutorials are available at: https://github.com/Ludeme/LudiiTutorials
[6]25th Computer Olympiad: https://icga.org/?page_id=3468

- Turn-based.
- Sequential.
- Fully observable.

Even with games limited to these types, a sufficiently complex challenge was provided by each of the proposed tracks. However, Ludii also includes games that are outside of these categories, such as games with hidden information, simultaneous-move games, and cooperative games. These games did not take part in this edition of the competition, but could be used for future and more complex GGP challenges.

2.2. Tracks

Three tracks were proposed:

(1) Kilothon: This track challenges competitors to maximise their performance (average utility) in more than 1000 different games, playing against a standard UCT (Kocsis and Szepesvári, 2006) baseline with specific time constraints for each game. This track does not involve interaction between the participants. Participating in this competition was possible at any time between the 1^{st} of February 2022 and the 29^{th} of July 2022.
(2) GGP: The participants had to play against other participants on six different games, unknown to the participants before the competition. The track's requirements were the closest to IGGPC, but using Ludii instead of the Stanford GDL.
(3) Learning: This competition was run under similar conditions to the GGP track, however, the set of six different games used were announced months before the actual competition. This track invited the agents to learn before competing, and allowed the implementation of different types of GGP agents.

The complete rules of each track are described in Section 3.

2.3. Participation

The competition was announced on different channels such as the Ludii forum,[7] mail lists, Twitter, and other social media, from February 2022 onwards. To participate in any of the tracks, a registration mail had to be sent to the DLP team. The only information required was the first and last name of the participant, the name of the agent, and an email address to contact them. Additionally, the participants could provide a short description of their agent if desired.

All agents were run directly from their own machine. All agents should have been implemented according to the Ludii Java AI API[8] or the Ludii Python AI API.[9] However, if a participant was able to implement a wrapper to Ludii in any other languages, this was also allowed. Finally, despite this not being a requirement, the organisers also encouraged the winners to make their source code open source if possible after the competition, or to share it privately with the DLP team for research purposes.

[7] The Ludii Forum is available at: https://ludii.games/forums/
[8] Ludii Java AI API is available at: https://github.com/Ludeme/LudiiAI
[9] Ludii Python AI API is available at: https://github.com/Ludeme/LudiiPythonAI

The organisers did not allow the use of direct copies of the AIs already available in the open source repository of Ludii,[10] but if one of these Ludii AIs was used as a basis, this should be mentioned by the participants in the description of their agents, alongside a description of how it is different form the original code.

Participants could enter all three tracks, but could only enter one agent per track.

2.4. Prizes

From 2005 to 2013, IGGPC was organised in conjunction with the AAAI Conference and offered a prize of $10,000 to the authors of the winning agent (Genesereth et al., 2005). This substantial amount of money attracted more and more participants every year, and was part of the reputation of this competition. Tristan Cazenave, full professor at the LAMSADE in the Université Paris Dauphine-PSL proposed to the organisers of the Ludii AI Competition 2022 to follow the same idea with an attractive prize of money for each track.

A total prize pool of €9000 for the three tracks was funded by his PaRis Artificial Intelligence Research InstitutE chair (PR[AI]RIE)).[11] The winner of each track received €3000.

3. COMPETITION RULES

This Section describes the rules of each track of the competition.

3.1. Kilothon track

In the previous GGP competitions, including IGGPC, an important limitation on the number of possible games was the relatively high level of complexity, from a user perspective, of the logic-based nature of the Stanford GDL. The game library offered by Ludii[12] is significantly larger, containing in its current version 1139 games. In comparison to Stanford GDL, in reality this number is even higher, because Ludii includes different rulesets and a set of options for many of its games (e.g., different number of players, board sizes, movement types, etc.), whereas these would all be described as completely distinct games in the Stanford GDL. 1462 rulesets and 1,929,197 option combinations are implemented. See Figure 1 for an overview of the diversity of games included in Ludii.

To take advantage of this large number of games, the organisers of the Ludii AI competition proposed a new GGP challenge, taking the form of a competition track called the Kilothon. In this track, participants do not interact or directly compete against each other, but rather try to win as many games as possible against a simple UCT (Kocsis and Szepesvári, 2006) baseline agent, in 1094 games from the Ludii library that satisfy the conditions described in Section 2.1.

To make it possible to run this challenge in less than a day, the organisers decided to define some time constraints for each agent playing the Kilothon games. For each game, only one minute of *smart* thinking time is allocated to each agent. After this time, the agent is forced to play randomly until the end of the game. Participants could distribute this minute of thinking time as they wish for their agent.

[10] Ludii Repository is available at: https://github.com/Ludeme/Ludii

[11] PR[AI]RIE: https://prairie-institute.fr/)

[12] The Ludii game library is available at: https://ludii.games/library.php

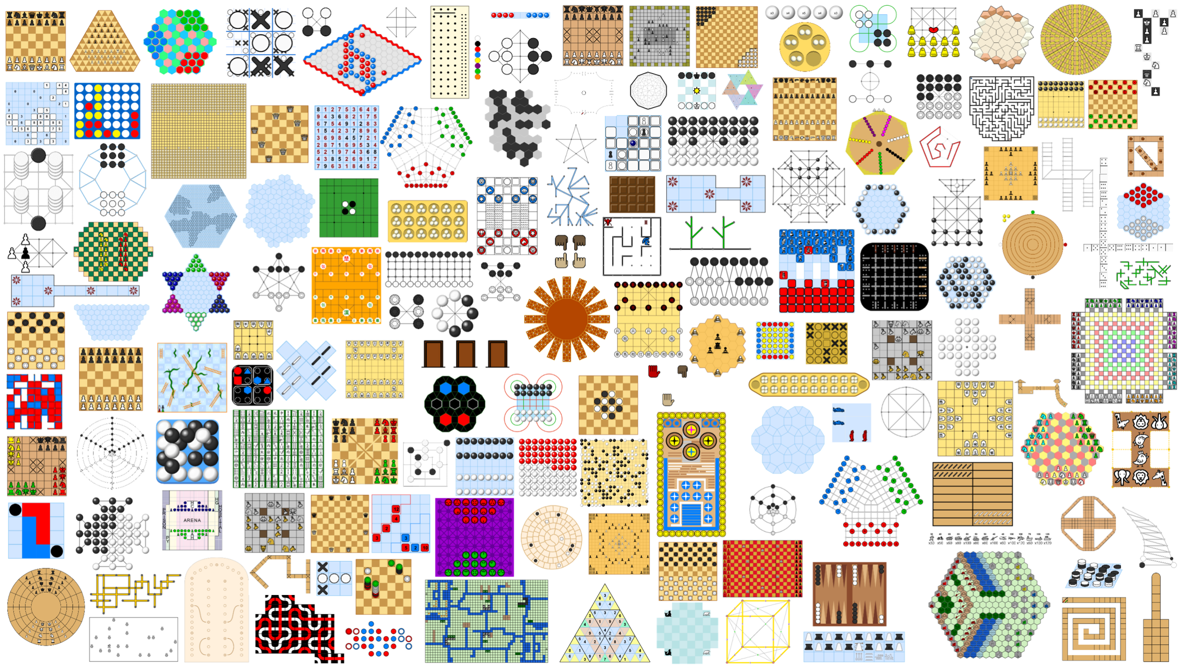

Fig. 1. Overview of a subset of Ludii games.

The UCT baseline used 0.5s of thinking time per move during this *smart* minute, and then played randomly. This allows a full Kilothon to be run in approximately 17 hours.

The participant always took the role of Player 1 (P1) in all games, and UCT took the role of each of the other players (P2 to potentially P16). The organisers implemented a simple piece of code provided to all the challengers through the Java API to run a single Kilothon.[13]

After each Kilothon, the utility score of Player 1 (between -1 and 1) was stored. Utility scores were obtained by linearly scaling the rankings obtaind by players, such that 1 is always the best possible utility, and -1 is always the worst possible utility. For example, obtaining third place in a 7-player game would give a utility of $1 - \left((3-1) \times \frac{2}{7-1}\right) \approx 0.333$. A Kilothon ended when all games were completed, and the results of the agent were sent automatically by mail to the organisers. The Kilothon results were also made available to the participant in a CSV format, for them to decide if they wish to improve their agent and run the Kilothon again. All participants could run any number of Kilothons they wish from February until the last day of the competition on the 29^{th} of July, but only the last Kilothon run by each participant was taken into account by the organisers. The winner of the Kilothon was simply the participant obtaining the highest average utility over all the games. In the unlikely case that two entrants tied on success rate, the total number of moves made by each agent would be used to break the tie.

3.2. General Game Playing track

The competition rules of this track were the closest to the GGP competitions organised from 2005 to 2016 by Stanford University.

[13]Kilothon video instructions are available at: https://youtu.be/bEdJqyUQ1Hg

It involved a Swiss tournament style run over six rounds of two-player games. In each round, a new game was provided, and each participant had to face another participant, playing the game twice (each starting once as Player 1). The games were not named or provided to the agents beforehand.

Each agent had 30 minutes of thinking per match, which can be allocated as desired. If this time has expired before the match ended, then the offending agent lost.

All matches were played over 25-27 July, 2022. The match-ups for each round were announced each morning by the organising team. Two rounds happened every day, one in the morning and one in the afternoon. All agents were required to remain online during the competition, so that the opponents were able to play their matches.

The remote mode of Ludii was used to play all matches. The participants were required to register on the Ludii forum and to log in through Ludii with their forum account.[14]

The games were selected by the competition organisers to cover a variety of categories. Figure 2 shows the six games of the GGP track. The games were played in the following order:

(1) Upper Hand is a 3D placement game played on a Shibumi board. The goal of the game is to have at least 27 marbles of his colour on the board. The average game length is pretty short ($\sim$31 turns) and was ideal to start the GGP track.
(2) ConHex is an abstract strategy game. Both players aim to complete a contiguous chain of connected cells between the 2 sides of their colour. Cells are claimed by conquering them. In order to conquer a cell, players must surround it with pegs played on the vertices. The organisers decided to select that game because it uses a mix of different types of game board elements: the players actually play on vertices, but win by capturing the cells of an irregularly-tiled board.
(3) Shogun is a capturing game played on an 8×8 grid. Each player can remove the other's pieces by moving onto them, as in chess. The game was selected because of its stochastic aspect. The pieces can move at a certain distance, which is randomly modified after their movement.
(4) Tablut is a traditional game belonging to the family of the Tafl games. This game was selected because of the asymmetric setup and victory conditions of each player.
(5) Oware is a two-row mancala-style game, originating from, and very popular in, West Africa. This game was selected to represent the Mancala games because of their originality compared to other common board games.
(6) Shobu is an abstract strategy game played on four boards. This game was selected because of the plurality of boards used and because each player always has to select two moves per turn.

3.3. Learning Track

The Learning track followed the same rules as the GGP track above, with the only exception being that the games used were revealed beforehand over March and April: one every week. This allowed participants to prepare for the competition and train their agents on the games beforehand.

Compared to the GGP track, the games decided by the organisers for the learning track were selected to involve more complex aspects such as a longer game length, a larger branching factor, etc. Figure 3 shows the six games of the Learning track. The games were played in the following order:

[14] A short video tutorial presenting the remote mode is available at: https://youtu.be/vl69mx7p7zM

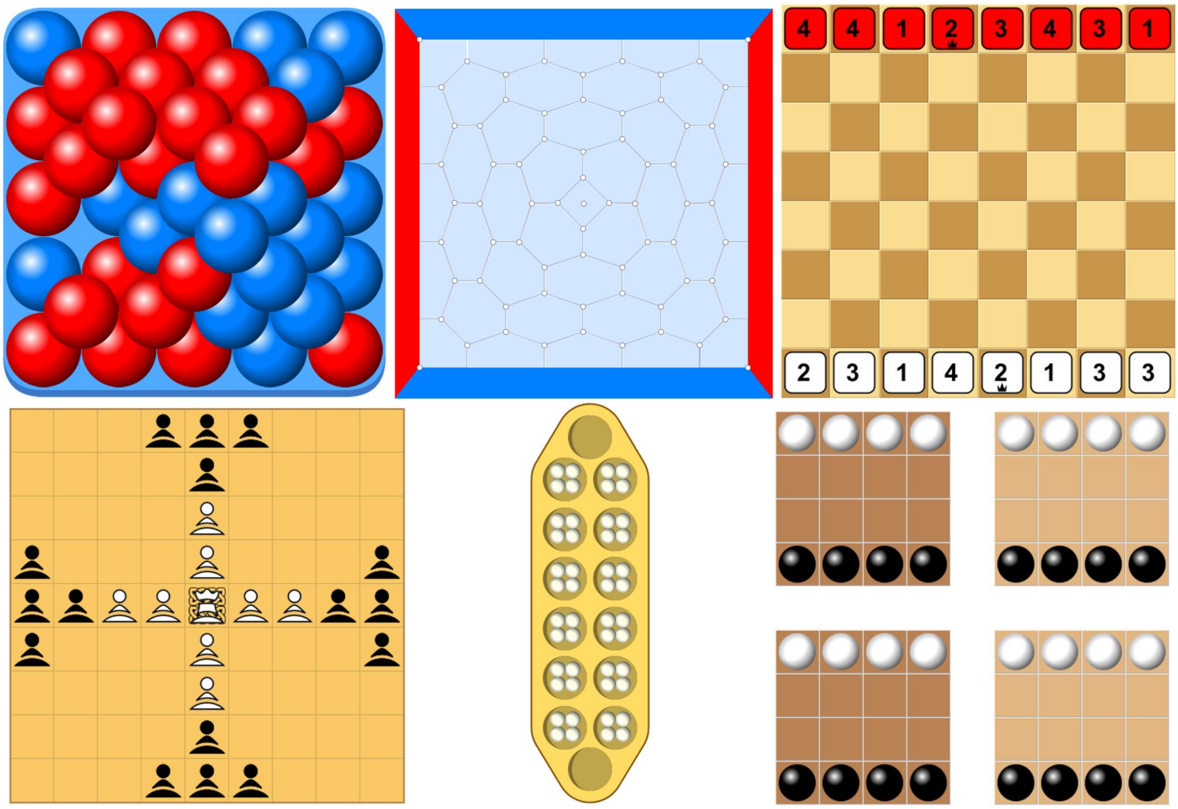

Fig. 2. The six games played during the GGP track. On the top row: Upper Hand, ConHex, and Shogun. On the bottom row: Tablut, Oware, and Shobu.

(1) Bashni is a game with leaping captures played in Russia during the nineteenth century. However, instead of removing pieces from the board when they are captured, they are stacked underneath the capturing piece. This game was selected because of its stacking feature, which is rarely studied.
(2) Ploy is a capturing game, with the goal of capturing the enemy commander or reducing the opposing army to a single commander. Each piece has an indicator which determines in which directions the piece can move. This can be altered by rotating the piece 45 degrees to the left or right. Rotation moves are a rare feature in board games, which makes it challenging to design algorithms using well this move type.
(3) Quoridor is a race game. The object of the game is to advance the pawn to the opposite side of the board. Each turn, the player may either move their pawn to an adjacent cell, or place a wall between two cells to block any movement between the corresponding cells. This game was selected because it involves play on different graph elements (cells as well as edges).
(4) Mini Wars is a simplified tabletop war game designed by one of the organisers. It involves an economic system, and buying an army of different piece types to capture buildings placed on the board. This game was selected for its originality and complexity.
(5) Plakoto is a game related to Backgammon that is typically played as part of Tavli. The game was selected to represent the Backgammon-like games, and because it involves dice.
(6) Lotus is played on a Kensington board. This board is irregularly tiled, which is a property that we rarely see in games used in existing research for learning in games.

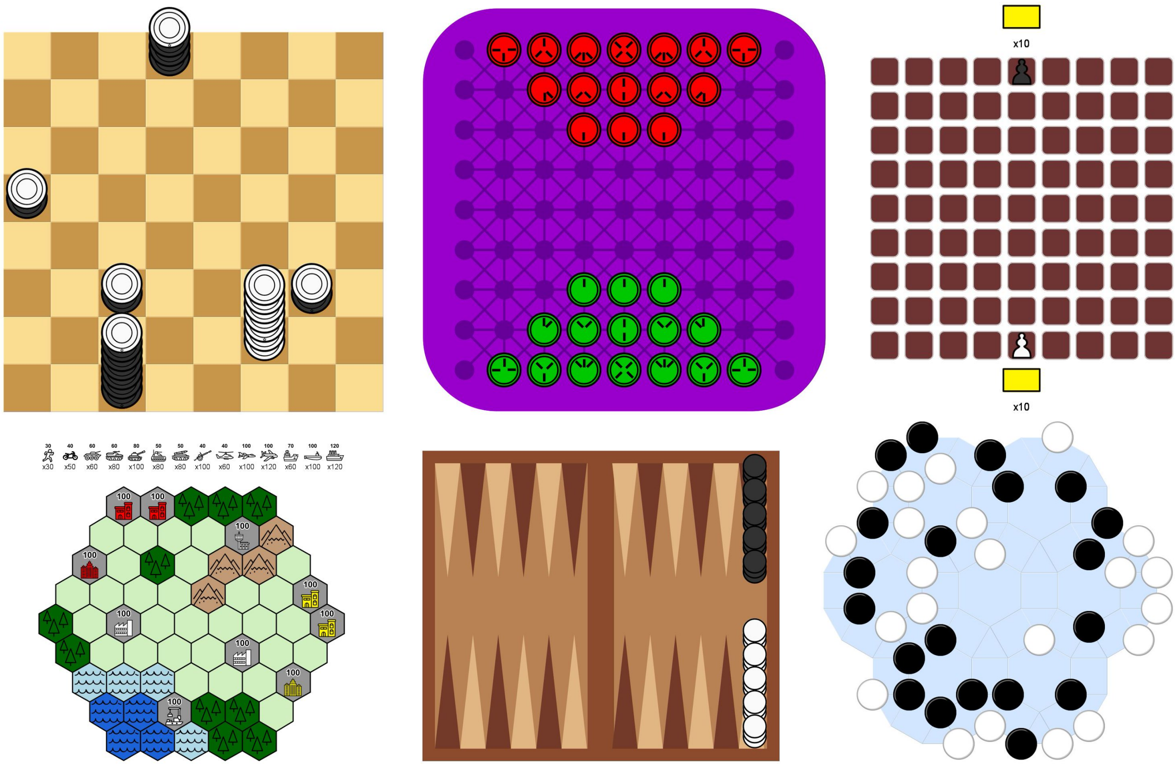

Fig. 3. The six games played during the Learning track. On the top row: Bashni, Ploy, and Quoridor. On the bottom row: Mini Wars, Plakoto, and Lotus.

4. RESULTS

On the 29th of July, the results and winners of the 2022 Ludii AI Competition were announced on the Ludii Forum,[15] on GitHub,[16] and on Twitter.[17]

Each track had a different number of registrations coming from different countries such as Brazil, China, France, the Netherlands, and the United Kingdom. The following subsections describe the results and the winning agents of each track.

4.1. Kilothon results

Seven participants registered for the Kilothon track, but only four of them succeeded to send valid results before the deadline of the competition. As described in Section 3.1, the participants were ranked according to the average utility obtained over all games. The final ranking is the following:

(1) *UBFM Contender* designed by Cyprien Michel–Deletie - utility = 0.231.
(2) *MCTimeS* designed by Ludovic Perrier - utility = 0.031.
(3) *SHOT-br* designed by Victor Putrich - utility = -0.034.
(4) *Zkealinvo* designed by Jingyang Zeng - utility = -0.456.

[15]Results on the Ludii Forum: https://ludii.games/forums/showthread.php?tid=1107
[16]Results on the GitHub page: https://github.com/Ludeme/LudiiAICompetition
[17]Results on Twitter: https://twitter.com/LudiiGames/status/1552930380419928069

Cyprien Michel–Deletie, student at the ENS de Lyon, won the 2022 Ludii Kilothon Competition.

His agent, *UBFM Contender*, was based on Unbounded Best-First Minimax (UBFM) (Korf and Chickering, 1996; Cohen-Solal, 2020). UBFM expands the tree search down the best path of the tree. It tries to follow immediate success until it either succeeds or fails, at which point the tree will grow in a different direction. The main behaviour of his agent is to use random playouts to estimate the number of turns the game will last at the first turn. It used this value each turn to set the time allocated to its decision-making. The estimation of the duration of the game is updated each time terminal states are encountered during the search. Additionally, its agent used two heuristics. (1) A material heuristic to minimise the enemy pieces and maximise its own pieces; (2) A mobility heuristic to maximise its movement and to minimise the movements of the opponent. UBFM is used only for non-stochastic games, for the others, UCT is used.

4.2. GGP Results

The GGP track had 8 official registrations (1x Brazil, 1x China, 2x the Netherlands, and 4x from France).

Only 3 of these registrations joined the Slack chat of the competition, but one of these participants was not able to join the actual competition due to differing timezone complications and was disqualified. Consequently, the GGP track was played between 2 participants *"MCTimeS - Ludovic Perrier"* and *"Zkealinvo - Jingyang Zeng"*, who competed on the 6 games described in Section 3.2. Table 1 shows the winner of each match played. After the first 5 rounds, both agents reached the same number of victories. However, in the final round, *Jingyang Zeng* won on both sides and won the 2022 edition of the Ludii GGP Competition.

His agent, *Zkealinvo*, is based on Thompson Sampling (Thompson, 1933). In other words, the algorithm estimates child nodes by establishing Beta distribution. At first, all the child nodes of the root are explored for the limited times. During the exploration, the distribution parameters are updated according to the rewards obtained on each playout. Then, the algorithm samples values from the distribution of the root's child nodes. The highest value is regarded as the best child which is going to be explored and update the parameters of the distribution. This process is kept in the loop until the time run out or iterations reaches the maximum setting. In the end, the algorithm returns the root's child node, which has the highest mean for its distribution. The author proposes an improvement relative to the updating method of distribution parameters. Some information such as the visit counts of each node is not fully used. A further improvement could be to incorporate the All Moves As First (AMAF) heuristic (Brügmann, 1993) to enhance the performance.

The complete record of the GGP matches played can be downloaded at this address:

https://ludii.games/downloads/Ludii_AI_Competition_2022_Trials/GGP_Track.zip.

4.3. Learning results

The Learning track had 6 official registrations (1x Brazil, 1x China, 1x the Netherlands, 3x France).

Unfortunately, none of the participants joined the Slack chat of the competition or the expected matches to run. To still be able to run the Learning track, the organisers invited the GGP participants to compete on the games proposed on this track. However, no prize money was rewarded for this track.

Table 1

The winner of each match played during the GGP track.

Game	Player 1 Vs. Player 2	Winner
	Day 1 - 25th July	
Upper Hand	*Zkealinvo* Vs. *MCTimeS*	*MCTimeS*
Upper Hand	*MCTimeS* Vs. *Zkealinvo*	*Zkealinvo*
ConHex	*Zkealinvo* Vs. *MCTimeS*	*MCTimeS*
ConHex	*MCTimeS* Vs. *Zkealinvo*	*MCTimeS*
	Day 2 - 26th July	
Shogun	*Zkealinvo* Vs. *MCTimeS*	*Zkealinvo*
Shogun	*MCTimeS* Vs. *Zkealinvo*	*Zkealinvo*
Tablut	*Zkealinvo* Vs. *MCTimeS*	*MCTimeS*
Tablut	*MCTimeS* Vs. *Zkealinvo*	*Zkealinvo*
	Day 3 - 27th July	
Oware	*Zkealinvo* Vs. *MCTimeS*	*MCTimeS*
Oware	*MCTimeS* Vs. *Zkealinvo*	*Zkealinvo*
Shobu	*Zkealinvo* Vs. *MCTimeS*	*Zkealinvo*
Shobu	*MCTimeS* Vs. *Zkealinvo*	*Zkealinvo*

Table 2 shows the winner of each match played. Two matches ended by timeout because an agent used more than 30min. The symbol * means that the match ended by a timeout. After 6 rounds, *Jingyang Zeng* won this track as well.

The complete record of the GGP matches played can be downloaded at this address:

https://ludii.games/downloads/Ludii_AI_Competition_2022_Trials/Learning_Track.zip.

Table 2

The winner of each match played during the Learning track.

Game	Player 1 Vs. Player 2	Winner
	Day 1 - 25th July	
Bashni	*Zkealinvo* Vs. *MCTimeS*	*MCTimeS*
Bashni	*MCTimeS* Vs. *Zkealinvo*	*MCTimeS*
Ploy	*Zkealinvo* Vs. *MCTimeS*	*Zkealinvo*
Ploy	*MCTimeS* Vs. *Zkealinvo*	*MCTimeS*
	Day 2 - 26th July	
Quoridor	*Zkealinvo* Vs. *MCTimeS*	*Zkealinvo*
Quoridor	*MCTimeS* Vs. *Zkealinvo*	*Zkealinvo*
Mini Wars	*Zkealinvo* Vs. *MCTimeS*	*Zkealinvo**
Mini Wars	*MCTimeS* Vs. *Zkealinvo*	*MCTimeS**
	Day 3 - 27th July	
Plakoto	*Zkealinvo* Vs. *MCTimeS*	*Zkealinvo*
Plakoto	*MCTimeS* Vs. *Zkealinvo*	*MCTimeS*
Lotus	*Zkealinvo* Vs. *MCTimeS*	*Zkealinvo*
Lotus	*MCTimeS* Vs. *Zkealinvo*	*Zkealinvo*

5. FUTURE LUDII AI COMPETITIONS

The organising team of the 2022 Ludii AI Competition was able to run successfully 2 of its 3 tracks. The organisers wish to run this competition every year and to increase the number of participants. For example, the Kilothon track, which can be run over a longer period of time, is easily accessible to any researcher or developer. Particularly, the organisers think that this challenge could be attractive for Master or PhD students to experiment with ideas developed during their courses and research. Concerning the GGP and Learning tracks, the organisers wish to announce them earlier next year, and multiple times to attract more people. The GGP track run through the Ludii system has the potential to motivate the design of innovative GGP approaches. The learning track may incentivise the creation of new types of general learning approaches, using for example deep learning techniques (Soemers et al., 2022a) or spatial features to improve the search (Soemers et al., 2022c).

The future editions of the Ludii AI Competition could incorporate tracks dedicated to other game types. In particular, games with hidden information (such as card games) can attract researchers working on search and learning techniques applicable to these games.

The creation of competitions using Ludii is not limited to GGP-related events (Stephenson et al., 2019). In the long term, Ludii could be used to run competitions on many other research topics, such as Puzzle Solving, Procedural Content Generation (PCG), or provided the most interpretable explanation of strategies (XAI).

The GitHub page used to run the Ludii AI Competition 2022 is available at:

https://github.com/Ludeme/LudiiAICompetition.

ACKNOWLEDGEMENTS

The Digital Ludeme Project is funded by the European Research Council through ERC Consolidator Grant #771292) and is run by Cameron Browne at Maastricht University's Department of Advanced Computing Sciences. Special thanks to Tristan Cazenave, Paris Artificial Intelligence Research Institute (PRAIRIE) chair, for securing funding for the Ludii AI Competition 2022 prizes.

REFERENCES

Björnsson, Y. & Finnsson, H. (2009). CadiaPlayer: A simulation-based general game player. *IEEE Transactions on Games*, *1*(1), 4–15.

Browne, C., Powley, E., Whitehouse, D., Lucas, S., Cowling, P.I., Rohlfshagen, P., Tavener, S., Perez, D., Samothrakis, S. & Colton, S. (2012). A Survey of Monte Carlo Tree Search Methods. *IEEE Transactions on Computational Intelligence and AI in Games*, *4*(1), 1–49.

Browne, C., Soemers, D.J.N.J., Piette, É., Stephenson, M., Conrad, M., Crist, W., Depaulis, T., Duggan, E., Horn, F., Kelk, S., Lucas, S.M., Neto, J.P., Parlett, D., Saffidine, A., Schädler, U., Silva, J.N., de Voogt, A. & Winands, M.H.M. (2019). Foundations of Digital Archæoludology. Technical report, Schloss Dagstuhl Research Meeting, Germany.

Browne, C., Stephenson, M., Piette, É. & Soemers, D.J.N.J. (2020). A Practical Introduction to the Ludii General Game System. In T. Cazenave, H.J. van den Herik, A. Saffidine and I.-C. Wu (Eds.), *Advances in Computer Games. ACG 2019*. Lecture Notes in Computer Science (Vol. 12516, pp. 167–179). Springer, Cham.

Browne, C. (2021). Everything's a Ludeme: Well, Almost Everything. In *Proceedings of the XIIIrd Board Game Studies Colloquium (BGS 2021)*. Paris.

Brügmann, B. (1993). Monte Carlo Go. Technical report, Max Planck Institute of Physics, München, Germany.

Cohen-Solal, Q. (2020). Learning to Play Two-Player Perfect-Information Games without Knowledge. https://arxiv.org/abs/2008.01188.

Genesereth, M.R. & Björnsson, Y. (2013). The International General Game Playing Competition. *AI Magazine*, *34*(2), 107–111.

Genesereth, M. & Thielscher, M. (2014). *General Game Playing*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers.

Genesereth, M.R., Love, N. & Pell, B. (2005). General Game Playing: Overview of the AAAI Competition. *AI Magazine*, *26*(2), 62–72. http://www.aaai.org/ojs/index.php/aimagazine/article/view/1813.

Kocsis, L. & Szepesvári, C. (2006). Bandit Based Monte-Carlo Planning. In J. Fürnkranz, T. Scheffer and M. Spiliopoulou (Eds.), *Machine Learning: ECML 2006*. Lecture Notes in Computer Science (Vol. 4212, pp. 282–293). Springer, Berlin, Heidelberg.

Korf, R.E. & Chickering, D.M. (1996). Best-first Minimax Search. *Artificial Intelligence*, *84*(1–2), 299–337.

Koriche, F., Lagrue, S., Piette, É. & Tabary, S. (2017a). Constraint-Based Symmetry Detection in General Game Playing. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17* (pp. 280–287).

Koriche, F., Lagrue, S., Piette, É. & Tabary, S. (2017b). WoodStock : Un programme-joueur générique dirigé par les contraintes stochastiques. *Revue D'Intelligence Artificielle (RIA)*, *31*(3), 307–336.

Love, N., Hinrichs, T., Haley, D., Schkufza, E. & Genesereth, M. (2008). General Game Playing: Game Description Language Specification. Technical report LG-2006-01, Stanford Logic Group.

Méhat, J. & Cazenave, T. (2010). Ary, a general game playing program. *2010 Board Games Studies colloquium*.

Piette, É., Browne, C. & Soemers, D.J.N.J. (2021). Ludii Game Logic Guide. https://arxiv.org/abs/2101.02120.

Piette, É., Soemers, D.J.N.J., Stephenson, M., Sironi, C.F., Winands, M.H.M. & Browne, C. (2020). Ludii – The Ludemic General Game System. In G.D. Giacomo, A. Catala, B. Dilkina, M. Milano, S. Barro, A. Bugarín and J. Lang (Eds.), *Proceedings of the 24th European Conference on Artificial Intelligence (ECAI 2020)*. Frontiers in Artificial Intelligence and Applications (Vol. 325, pp. 411–418). IOS Press.

Pitrat, J. (1968). Realization of a general game-playing program. In *IFIP Congress (2)* (pp. 1570–1574).

Soemers, D.J.N.J., Mella, V., Browne, C. & Teytaud, O. (2022a). Deep Learning for General Game Playing with Ludii and Polygames. *ICGA Journal*, *43*(3), 146–161.

Soemers, D.J.N.J., Piette, É., Stephenson, M. & Browne, C. (2022b). Ludii User Guide. Technical report, Maastricht University. https://ludii.games/downloads/LudiiUserGuide.pdf.

Soemers, D.J.N.J., Piette, É., Stephenson, M. & Browne, C. (2022c). Spatial State-Action Features for General Games. Under review.

Stephenson, M., Piette, É., Soemers, D.J.N.J. & Browne, C. (2019). Ludii as a Competition Platform. In *Proceedings of the 2019 IEEE Conference on Games (COG 2019)* (pp. 634–641). London.

Świechowski, M., Park, H., Mańdziuk, J. & Kim, K.-J. (2015). Recent Advances in General Game Playing. *The Scientific World Journal*, *2015*.

Thompson, w.R. (1933). On the likehood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, *25*(3-4), 285–294.