ICGA Journal 46 (2024) 28–35
DOI 10.3233/ICG-240251
IOS Press

GameTable Working Group 1 meeting report on search, planning, learning, and explainability

Dennis J.N.J. Soemers [a,*], Jakub Kowalski [b], Éric Piette [c], Achille Morenville [c] and Walter Crist [d]

[a] *Department of Advanced Computing Sciences, Maastricht University, the Netherlands*
[b] *Faculty of Mathematics and Computer Science, University of Wrocław, Poland*
[c] *ICTEAM, UCLouvain, Belgium*
[d] *Centre for the Arts in Society, Leiden University, the Netherlands*

Abstract. The inaugural in-person meeting for the "GameTable" COST Action's Working Group 1 (WG1) on Search, Planning, Learning, and Explainability took place at the Leiden Institute of Advanced Computer Science (LIACS) on January 31st, 2024. The primary aims of this meeting were to facilitate talks and discussions on, and connect researchers interested in, three core research goals: (1) human-like game-playing AI, (2) imperfect-information games within a general game playing context, and (3) explainable search and reinforcement learning in games. This report provides a summary of the discussions and talks that took place during the meeting.

1. GAMETABLE – WORKING GROUP 1

GameTable[1] is a network of researchers, funded by a European Cooperation in Science & Technology (COST) grant, studying tabletop games from the perspectives of a variety of disciplines such as artificial intelligence (AI), archaeology, and mathematics. The Action was initiated in October 2023, and it will remain funded up to October 2027.

The network consists of five working groups (WGs), of which Working Group 1 – entitled Search, Planning, Learning, and Explainability – focuses on topics such as search algorithms (Russell and Norvig, 2020) and reinforcement learning (Sutton and Barto, 2018) for automated game playing, general game playing (Pitrat, 1968), and explainability and interpretability of game-playing algorithms and their decisions or decision-making processes. The first official, in-person meeting (with an option for remote participation via Zoom) was organised on January 31, 2024, at the Leiden Institute of Advanced Computer Science in Leiden, the Netherlands. The meeting, which followed an Action-wide kickoff event in Leiden on the preceding two days (Piette et al., 2024), was primarily organised by Dennis Soemers and Jakub Kowalski, as leaders of the working group, with Mike Preuss providing local support as employee in Leiden. This report summarises the discussions and talks that took place in the meeting.

*Corresponding author. E-mail: dennis.soemers@maastrichtuniversity.nl.
[1] https://gametable.network/

2. MEETING REPORT

Working Group 1 (WG1) of the GameTable COST Action has a wide array of different goals and objectives, which we aim to make progress on over the four-year lifetime of the Action. Many of these goals revolve around further development of AI techniques for automated game-playing in general contexts (applicable to many different tabletop games). The aim is to develop techniques that may serve to estimate and analyse how humans in various cultures and periods of time may have played and experienced such games, as originally envisioned and explored in the Digital Ludeme Project (Browne, 2018; Browne et al., 2019; Soemers et al., 2019; Crist and Soemers, 2023). For this first meeting, three specific topics were selected to be discussed in detail: (1) human-like game-playing AI, (2) imperfect information games in a general game playing context, and (3) explainable search and reinforcement learning in games. By no means do these three topics exhaustively cover all topics of interest to the group, but they were identified as important and suitable topics of discussion for the first meeting. The discussions that took place around these topics are presented in more detail in the following three subsections.

When plans for the meeting were first announced to members of the group, in November 2023, members were asked to explicitly express any potential interest in attending. The event was open to be joined by any interested members of the Action in principle, but it was not possible to provide reimbursements for travel and accommodations to everyone. Based on expressions of interest, areas of expertise, and levels of experience, a total of 20 participants were offered reimbursements. We ensured to have a suitable mix of early-career researchers and more experienced ones, as one of the explicit goals of COST Actions in general is to enhance visibility of, and facilitate networking and new opportunities for, young researchers in particular. The total number of participants (including remote participation) was 28.

2.1. Human-like game-playing AI

This session started with one of the younger researchers among the participants, Dominik Jeurissen, presenting recent work (Jeurissen et al., 2024) on using Large Language Models (LLMs) to play the game of NetHack in a zero-shot manner (i.e., without any fine-tuning on the game). Thanks to the vast amounts of "general knowledge" modern LLMs have observed in their training, likely including many pieces of text written about popular games such as NetHack, these models appear to be capable of making meaningful decisions in response to prompts about the current game state, as well as having the ability to produce language-based outputs (e.g., explanations) about their decisions. While the development of human-like AI was not necessarily the primary focus of this research, there are still interesting potential relations to be explored. For example: may the nature of the model, which was originally trained to effectively communicate with humans in human languages, as opposed to being an algorithm designed or trained to play as well as possible, aid in producing a game-player that behaves more human-like?

The remainder of the session on human-like game-playing AI shifted to a more open-ended discussion, among the group of all participants, around questions such as:

- What does it mean for game-playing AI to be human-like?
- Why is human-like game-playing AI for tabletop games considered important?
- What are fruitful initial research directions and possible opportunities for grant proposals within this topic?

In the context of video games, there has been more work on aspects of human-like AI (Milani et al., 2023), for purposes such as creating believable non-player characters or creating bots that can stand in for disconnected players in multi-player games. In the context of board games, there has been some work on e.g. developing cognitively plausible planning or learning algorithms (Mańdziuk, 2011, 2012). However, in GameTable, the focus is more on developing AI that can play games such that, when statistics related to the gameplay experience are gathered, they will hopefully accurately estimate the statistics that would have been collected by observing humans playing those games. These statistics can be almost anything related to the gameplay experience including, but not limited to, win/draw/loss percentages, durations (in turns or moves) of games, distributions of usage of different sites (e.g., cells or vertices) on a board, and estimators of (differences between) skill floors and ceilings.

Where possible, social and cultural contexts in which humans play(ed) games ought to be accounted for. In contrast to the prevailing focus on competitiveness and strong or optimal play in most game AI research, humans often play(ed) in contexts where winning – or delaying or avoiding losses – at all costs is not the only (or even main) objective. Humans play tabletop games together with (and often observed by) other humans (Crist et al., 2016), who they generally still want to be able to get along with after play, or may even wish to impress with their unique or attractive and exciting playing style. Going beyond the formal rules of a game (which themselves are also often already very much in flux), this implies that there are often implicit rules (of etiquette) which human players tend to largely abide by. For example, in many contexts they will not indefinitely stall a game to avoid losing, but instead opt for a more interesting move (potentially leading to a faster loss), or at the very least simply concede or agree to a tie and end the game, rather than continuing play for an infinite amount of time. Humans may also intentionally *not* follow certain rules (i.e., cheat), but also cheating would have to be done in an intelligent and subtle manner.

Another key aspect to human-like AI is that there is no such thing as a single, individual player that can capture and represent the experience that humans playing games would have had. Games are generally played by a wide variety of different players, of different (and changing) levels of skill, attention spans, social expectations, reasons for playing other than winning, and so on. Sometimes, games are played competitively in a formally organised tournament setting or for money bet in a tavern, possibly between roughly equally-skilled players. Other times, games may be played between teachers and students or (grand)parents and children, which may sometimes be for educational purposes and other times simply to pass the time. This implies that we would often want a larger distribution of many different AI players, rather than just a single “human-like” AI player.

The discussions at the meeting raised many of such interesting questions, but they continue to be largely open research questions after our meeting. Some initial suggestions following from the discussion consist of looking at play traces of chess and Go (for which even historic traces of play from centuries ago have been recorded), and seeking out increased collaborations with researchers who could provide relevant insights from the field of psychology, in addition to the cultural and AI research fields that are already well represented in GameTable.

2.2. Imperfect information games (in a general game playing context)

The field of General Game Playing (GGP) (Pitrat, 1968) has produced a wealth of research (Björnsson and Schiffel, 2016), including many competitions (Genesereth et al., 2005; Genesereth and Björnsson, 2013) and game description languages and frameworks (Love et al., 2008; Genesereth and Thielscher, 2014; Kowalski et al., 2019; Piette et al., 2020). This includes support for modelling imperfect-information games – games where parts of the state can be unobservable to some players – in some frameworks (Thielscher, 2010; Soemers et al., 2024), but the development of general agents that can play any arbitrary imperfect-information game is severely lacking in comparison to the work on such agents for perfect-information games.

There have been a few attempts at designing such agents (Schofield et al., 2012; Schofield and Thielscher, 2019), but their playing strength is often still underwhelming in practice. There have been great advances in algorithms for imperfect-information games over the past decade (Moravčík et al., 2017; Brown and Sandholm, 2017, 2019; Brown et al., 2020; Schmid et al., 2023), but they tend to require an extent of domain knowledge (to infer information sets from public states or construct efficient representations of public belief states) that is not available in GGP contexts. A notable exception is Regularised Nash Dynamics (Perolat et al., 2022), which can be implemented with less need for such knowledge (or ability to automatically infer it from game descriptions), although this work is primarily known from results obtained with large neural network and massive compute resources, which also poses a challenge in practice for scaling up to GGP with many and varied games.

Many of the discussions we had and questions that were raised in this session revolved around what kinds of frameworks or infrastructure would need to be built or improved to accelerate research towards truly general game playing in an imperfect-information context, with acceptable levels of playing strength being practically feasible on a large scale in terms of number of games. GDL-II (Thielscher, 2010) supports such research in principle, but its complexity of use and computational inefficiency impede successful developments in practice. Ludii (Piette et al., 2020) already supports the implementation of imperfect-information games in its game description language in theory (Soemers et al., 2024), but in practice it would require substantial additions to the language for such games to be describable in a manner that is as accessible and user-friendly as the language is for perfect-information games (Morenville and Piette, 2024). The question whether it is really necessary to have a single framework in which every (tabletop) imperfect-information imaginable can be implemented in a non-programming language has also been raised. Perhaps we can simply make do with a large array of games being implemented in a general-purpose programming language (as done in, e.g., OpenSpiel (Lanctot et al., 2019)), or with multiple frameworks and game description languages for multiple different subsets of imperfect-information games. Aside from frameworks for modelling games, it is likely we could help accelerate research in this area by organising new competitions (Togelius, 2016) focused on imperfect-information GGP.

In addition to these general discussions, Ondřej Kubíček and David Milec contributed talks on their research. This included an exploration of a sound method to introduce search at test time on top of a policy that was originally trained without search (Kubíček et al., 2024), as well the development of sound and robust methods that can play well against strong opponents, whilst also managing to exploit weaker opponents (Milec et al., 2024).

2.3. Explainable search and reinforcement learning (in games)

The session on explainable search and reinforcement learning (RL) methods (with a focus on applications in tabletop games) was filled with contributed talks on recent as well as planned research. Manuel Eberhardinger explored how synthesising executable programs as policies may enhance explainability, as programs can be considered inherently explainable or interpretable (Eberhardinger et al., 2023a,b). Ekkehard Schnoor contributed a talk with a focus on using concept-based approaches (Kim et al., 2018; Achtibat et al., 2023) to discover human-understandable concepts learned using

AlphaZero-like combinations of deep learning and search (Silver et al., 2018; McGrath et al., 2022; Hammersborg and Strümke, 2023; Schut et al., 2023). Yngvi Björnsson provided arguments as to why the game of Chess should be considered an ideal benchmark for current research on explainable search and RL (Pálsson and Björnsson, 2023). Mark Winands presented a recent study on explaining recommendations made by a search algorithm for SameGame (Sironi et al., 2023), and discussed different types and forms of explanations that could be presented to humans. Finally, Hendrik Baier contributed a talk on the recently launched *hyPEr ExpeRt collaborative AI assistant* (PEER) project,[2] which aims to build artificial intelligence (AI) techniques that are explainable, aware of the need to explain in their own search procedures, trustworthy, and human-centric.

3. CONCLUSION

The first in-person meeting for Working Group 1 of the GameTable COST Action produced fruitful discussions on future research plans within the scope of the group and the Action, and enabled researchers at various career stages to meet and present work to each other. The Action will continue to fund similar meetings, combined meetings and events with other working groups, training schools, short-termin scientific missions, and other networking and dissemination activities until October 2027. We will also investigate opportunities to actively collaborate and write new grant proposals in the areas of research that have been discussed, and continue to highlight and support junior researchers in the field. Researchers who are interested in joining the Action should feel free to apply via the COST website.[3]

ACKNOWLEDGEMENTS

This article is based upon work from COST Action CA22145 – GameTable, supported by COST (European Cooperation in Science and Technology). We thank all the participants who attended and contributed to our meeting.

REFERENCES

Achtibat, R., Dreyer, M., Eisenbraun, I., Bosse, S., Wiegand, T., Samek, W. & Lapuschkin, S. (2023). From attribution maps to humand-understandable explanations through concept relevance propagation. *Nature Machine Intelligence*, *5*(9), 1006–1019. doi:10.1038/s42256-023-00711-8.

Björnsson, Y. & Schiffel, S. (2016). General game playing. In *Handbook of Digital Games and Entertainment Technologies* (pp. 1–23). Singapore: Springer Singapore.

Brown, N., Bakhtin, A., Lerer, A. & Gong, Q. (2020). Combining deep reinforcement learning and search for imperfect-information games. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan and H. Lin (Eds.), *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)* (Vol. 33, pp. 17057–17069). Curran Associates, Inc.

Brown, N. & Sandholm, T. (2017). Superhuman AI for heads-up no-limit poker: Libratus beats top professionals. *Science*, *359*(6374), 418–424. doi:10.1126/science.aao1733.

[2] https://peer-ai.eu/en/

[3] https://www.cost.eu/actions/CA22145/#tabs+Name:Working%20Groups%20and%20Membership

Brown, N. & Sandholm, T. (2019). Superhuman AI for multiplayer poker. *Science*, *365*(6456), 885–890. doi:10.1126/science.aay2400.

Browne, C. (2018). Modern techniques for ancient games. In *IEEE Conference on Computational Intelligence and Games* (pp. 490–497). Maastricht: IEEE Press.

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.

Crist, W., de Voogt, A. & Dunn-Vaturi, A.-E. (2016). Facilitating interaction: Board games as social lubricants in the ancient near East. *Oxford Journal of Archaeology*, *35*(2), 179–196. doi:10.1111/ojoa.12084.

Crist, W. & Soemers, D.J.N.J. (2023). The digital ludeme project: Combining archaeological and computational methods for the study of ancient board games. *Journal of Archaeological Science: Reports*, *49*.

Eberhardinger, M., Maucher, J. & Maghsudi, S. (2023a). Learning of generalizable and interpretable knowledge in grid-based reinforcement learning environments. In *Proceedings of the Nineteenth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment. AIIDE '23*. AAAI Press.

Eberhardinger, M., Maucher, J. & Maghsudi, S. (2023b). Towards explainable decision making with neural program synthesis and library learning. In *17th International Workshop on Neural-Symbolic Learning and Reasoning*.

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

Genesereth, M.R. & Björnsson, Y. (2013). The international general game playing competition. *AI Magazine*, *34*(2), 107–111. doi:10.1609/aimag.v34i2.2475.

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.

Hammersborg, P. & Strümke, I. (2023). Reinforcement learning in an adaptable chess environment for detecting human-understandable concepts. *IFAC-PapersOnLine*, *56*(2), 9050–9055. 22nd IFAC World Congress. doi:10.1016/j.ifacol.2023.10.135.

Jeurissen, D., Perez-Liebana, D., Gow, J., Çakmak, D. & Kwan, J. (2024). Playing NetHack with LLMs: Potential & Limitations as Zero-Shot Agents. arXiv:2403.00690.

Kim, B., Wattenberg, M., Gilmer, J., Cai, C., Wexler, J., Viegas, F. & Sayres, R. (2018). Interpretability beyond feature attribution: Quantitative Testing with Concept Activation Vectors (TCAV). In J. Dy and A. Krause (Eds.), *Proceedings of the 35th International Conference on Machine Learning*. Proceedings of Machine Learning Research (Vol. 80, pp. 2668–2677). PMLR.

Kowalski, J., Maksymilian, M., Sutowicz, J. & Szykuła, M. (2019). Regular boardgames. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence* (Vol. 33, pp. 1699–1706). AAAI Press.

Kubíček, O., Burch, N. & Lisý, V. (2024). Look-ahead search on top of policy networks in imperfect information games. In *Proceedings of the International Joint Conference on Artificial Intelligence*. Accepted.

Lanctot, M., Lockhart, E., Lespiau, J.-B., Zambaldi, V., Upadhyay, S., Pérolat, J., Srinivasan, S., Timbers, F., Tuyls, K., Omidshafiei, S., Hennes, D., Morrill, D., Muller, P., Ewalds, T., Faulkner, R., Kramár, J., Vylder, B.D., Saeta, B., Bradbury, J., Ding, D., Borgeaud, S., Lai, M., Schrittwieser, J., Anthony, T., Hughes, E., Danihelka, I. & Ryan-Davis, J. (2019). OpenSpiel: A Framework for Reinforcement Learning in Games. arXiv:1908.09453.

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.

Mańdziuk, J. (2011). Towards cognitively-plausible game playing systems. *IEEE Computational Intelligence Magazine*, *6*(2), 38–51. doi:10.1109/MCI.2011.940626.

Mańdziuk, J. (2012). Human-like intuitive playing in board games. In T. Huang, Z. Zeng, C. Li and C.S. Leung (Eds.), *International Conference on Neural Information Processing*. Lecture Notes in Computer Science (Vol. 7664, pp. 282–289). doi:10.1007/978-3-642-34481-7_35.

McGrath, T., Kapishnikov, A., Tomašev, N., Pearce, A., Wattenberg, M., Hassabis, D., Kim, B., Paquet, U. & Kramnik, V. (2022). Acquisition of chess knowledge in AlphaZero. *Proceedings of the National Academy of Sciences of the United States of America*, *119*(47).

Milani, S., Juliani, A., Momennejad, I., Georgescu, R., Rzpecki, J., Shaw, A., Costello, G., Fang, F., Devlin, S. & Hofmann, K. (2023). Navigates like me: Understanding how people evaluate human-like AI in video games. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. ACM.

Milec, D., Kubíček, O. & Lisý, V. (2024). Continual depth-limited responses for computing counter-strategies in sequential games. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems. AAMAS '24* (pp. 2393–2395). International Foundation for Autonomous Agents and Multiagent Systems.

Moravčík, M., Schmid, M., Burch, N., Lisỳ, V., Morrill, D., Bard, N., Davis, T., Waugh, K., Johanson, M. & Bowling, M. (2017). DeepStack: Expert-level artificial intelligence in heads-up no-limit poker. *Science*, *356*(6337), 508–513. doi:10.1126/science.aam6960.

Morenville, A. & Piette, É. (2024). Vers une Approche Polyvalente pour les Jeux à Information Imparfaite sans Connaissance de Domaine. In *Rencontres des Jeunes Chercheurs en Intelligence Artificielle (RJCIA)* (pp. 44–46). In French.

Pálsson, A. & Björnsson, Y. (2023). Unveiling concepts learned by a world-class chess-playing agent. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence. IJCAI '23* (pp. 4864–4872).

Perolat, J., Vylder, B.D., Hennes, D., Tarassov, E., Strub, F., de Boer, V., Muller, P., Connor, J.T., Burch, N., Anthony, T., McAleer, S., Elie, R., Cen, S.H., Wang, Z., Gruslys, A., Malysheva, A., Khan, M., Ozair, S., Timbers, F., Pohlen, T., Eccles, T., Rowland, M., Lanctot, M., Lespiau, J.-B., Piot, B., Omidshafiei, S., Lockhart, E., Sifre, L., Beauguerlange, N., Munos, R., Silver, D., Singh, S., Hassabis, D. & Tuyls, K. (2022). *Science*, *378*(6623), 990–996. doi:10.1126/science.add4679.

Piette, É., Crist, W., Soemers, D.J.N.J., Rougetet, L., Courts, S., Penn, T. & Morenville, A. (2024). GameTable COST action: Kickoff report. *ICGA Journal*, *46*(1). doi:10.3233/ICG-230230.

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).

Russell, S. & Norvig, P. (2020). *Artificial Intelligence: A Modern Approach* (4th ed.). Pearson.

Schmid, M., Moravčík, M., Burch, N., Kadlec, R., Davidson, J., Waugh, K., Bard, N., Timbers, F., Lanctot, M., Holland, G.Z., Davoodi, E., Christianson, A. & Bowling, M. (2023). Student of games: A unified learning algorithm for both perfect and imperfect information games. *Science Advances*, *9*(46).

Schofield, M., Cerexhe, T. & Thielscher, M. (2012). HyperPlay: A solution to general game playing with imperfect information. In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 26, pp. 1606–1612).

Schofield, M. & Thielscher, M. (2019). General game playing with imperfect information. *Journal of Artificial Intelligence Research*, *66*, 901–935. doi:10.1613/jair.1.11844.

Schut, L., Tomasev, N., McGrath, T., Hassabis, D., Paquet, U. & Kim, B. (2023). Bridging the Human-AI Knowledge Gap: Concept Discovery and Transfer in AlphaZero. arXiv:2310.16410.

Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., Lanctot, M., Sifre, L., Kumaran, D., Graepel, T., Lillicrap, T., Simonyan, K. & Hassabis, D. (2018). A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science*, *362*(6419), 1140–1144. doi:10.1126/science.aar6404.

Sironi, C.F., Wilbik, A. & Winands, M.H.M. (2023). Explainable search: An exploratory study in SameGame. In *Proceedings of the 2023 IEEE Conference on Games* (pp. 1–4). IEEE.

Soemers, D.J.N.J., Crist, W. & Browne, C. (2019). Report on the digital ludeme project. *ICGA Journal*, *41*(3), 138–142. doi:10.3233/ICG-190118.

Soemers, D.J.N.J., Piette, É., Stephenson, M. & Browne, C. (2024). The ludii game description language is universal. In *Proceedings of the 2024 IEEE Conference on Games. Accepted*.

Sutton, R.S. & Barto, A.G. (2018). *Reinforcement Learning: An Introduction* (2nd ed.). Cambridge, MA: MIT Press.

Thielscher, M. (2010). A general game description language for incomplete information games. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence* (pp. 994–999). AAAI.

Togelius, J. (2016). How to run a successful game-based AI competition. *IEEE Transactions on Computational Intelligence and AI in Games*, *8*(1), 95–100. doi:10.1109/TCIAIG.2014.2365470.