### 3.5.3 Further Work

Our next steps are to clean up and open source these systems, along with publishing our initial survey of the speedrunning landscape with respect to AI. Beyond that, we believe Celeste in PICO-8 represents a good domain for an AI competition. Designing the framework and rules for this competition will also help us clarify what challenges are most interesting, and begin to grow academic interest around this area.

**References**

1 Speedrun.com – About. http://www.speedrun.com/about Accessed 3 July 2024.
2 Manuel Lafond. *The complexity of speedrunning video games*. In 9th International Conference on Fun with Algorithms (FUN 2018). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018.
3 Scully-Blaker, Rainforest. *Re-curating the Accident: Speedrunning as Community and Practice*. Masters thesis, Concordia University, 2016.
4 gonengazit/Celia. https://github.com/gonengazit/Celia Accessed 9 July 2024.

## 3.6 Skill-Discovery in (Strategy) Games

*Alexander Dockhorn (Leibniz Universität Hannover, DE), Manuel Eberhardinger (Hochschule der Medien – Stuttgart, DE), Chengpeng Hu (Southern Univ. of Science and Technology – Shenzen, CN), and Matthias Müller-Brockhausen (Leiden University, NL)*



Strategy games present a unique challenge in artificial intelligence (AI) research. They can broadly be classified into two types: turn-based and real-time strategy games. Both types typically require the player or AI to manage multiple units or resources, often with incomplete information about the opponent's actions. The large branching factor and long game duration make it difficult for AI to explore all possible strategies, which is further complicated by the need to plan several moves ahead. The state-of-the-art methods in AI for strategy games include search-based algorithms and reinforcement learning (RL), but these often rely on human-defined strategies or subgoals, limiting their scalability and generalizability.

While the work on AlphaStar [2, 3] and OpenAI Five [5] have shown that it is possible to train strong AI agents for complex games such as Starcraft 2 and Dota 2, both works required massive amounts of compute resources until satisfying results have been achieved. For the purpose of speeding up the learning process, the working group on skill discovery in (strategy) games has been formed to evaluate the applicability of skill discovery methods to this special domain. We particularly emphasize works on skill discovery as part of RL algorithms. In this context, skill discovery refers to identifying and learning sub-policies or strategies that can be applied to achieve or identify specific subgoals within a game, thereby enabling more efficient and scalable AI systems.

### 3.6.1 Preliminaries of Skill Discovery

Skill discovery in AI remains an open problem, particularly when it comes to discovering skills without human intervention. The interpretation of what constitutes a "skill" in a given context is still unclear, making it challenging to develop a unified approach to skill discovery.

Particularly within the RL framework, a skill is typically defined as a policy aimed at achieving a specific subtask or goal. The options framework [7, 6] formalizes this by learning policies for subtasks, identifying the start and end points of these tasks, and using these learned skills to simplify the overall decision-making process. However, the distinction between a skill and a task is not always clear, especially when a subgoal can only be reached through a single deterministic policy.

### 3.6.2 Proposed Approaches and Ideas

Current research explores various methods for skill discovery, including hierarchical approaches, bottom-up skill learning, and the application of relational representations of game elements. Given an initial literature review, our working group has identified the following promising approaches for skill discovery in strategy games:

- **Text-based Task Decomposition:** Strategy games often have a simple goal, e.g. defeating the opponent's units or destroying its base. However, doing so involves plenty of subtasks. Those can be defined on varying ranges of granularity. Given the increasing capabilities of large language models, task descriptions such as "defend the base" could be decomposed into "train at least 3 units" and "patrol the surroundings of your base". Such enriched descriptions may directly represent sub-goals and allow for a more interpretable and scalable approach to skill discovery. Further, it allows to define more fine-grained reward functions given the descriptions [12].
- **Relational Representations:** Game state representations in strategy games can become quite complex. Units, abilities, weapons, buildings, and resources are just a few of the typical systems included in strategy games. Attempts to create general vectorized state representations have recently been studied [11], however, those create a unique representation for every game mapping all of its subsystems. While they allow the definition of state-space abstractions, transferring results from one game to the other is hindered by the granularity of this state representation. Similarly, matrix-based or image-based representations as used in AlphaStar [2] enabled the training through large-scale reinforcement learning but due to the complexity of the input at the cost of enormous computational resources.
  One possibility to overcome this problem is the use of a relational representation of game elements, such as "workers – mine – resources." Defining low-level systems for the execution of such relations allows to focus the agent's training on high-level strategic decision-making. At the same time, the high-level relation allows the transfer of knowledge in between games with different low-level controls. Using such a representation in combination with relational reinforcement learning [10, 9, 8] may therefore improve the efficiency of training agents in complex strategy games.
- **Pattern Mining and Clustering:** Given a data set of successful and unsuccessful play traces, pattern mining and clustering algorithms may be used to cluster them into groups of similar elements and extract abstract prototypes. Techniques like the KRIMP algorithm[1] or Skid Raw [19], which extract patterns from previous action sequences, could be applied to discover meaningful skills. Similarly, time sub-series mining [4], used to measure the distance between interaction sequences, may reveal underlying patterns that represent skills.
- **Bottom-Up Skill Discovery:** While most existing methods rely on top-down approaches, our group was exploring approaches for reversing this process by discovering skills from the ground up. Current methods are able to learn skills in the latent space of neural networks from collected demonstrations [21, 22, 23, 24]. However, these skills

are not interpretable, and only after their execution can one infer what the learned skills represent. Another disadvantage of these methods is that they are only suitable for small environments and toy tasks, where the agent needs to navigate to multiple goals. To overcome the limitations of simple tasks and apply these methods to more complicated domains, we propose to learn skills from sequences of actions and iteratively refine these skills to handle the large search spaces inherent in strategy games. [18].

- **Skill Discrimination:** Effective skill discovery requires a discriminator to identify whether a discovered policy qualifies as a skill and if it is any different than already known skills [13]. Quality Diversity Optimization as in the Diversity Policy Gradient algorithm [20] introduces a method for discovering a diverse set of skills by balancing the exploration of different strategies with maintaining high-quality solutions. Recently, Wang et al. [25] proposed to incorporate a regularization term into the RL objective that maximizes the negative correlation to increase the diversity of RL policies via assembling multiple sub-policies. Notably, this diversity pertains to the behavior of the derived policy rather than the parameter space, as minor variations in parameters could lead to significant differences in behavior. Although this algorithm was initially verified in the context of game content generation, it could be adapted for skill discrimination to maximize the diversity of discovered skills.

### 3.6.3 Conclusion

Skill discovery in strategy games is a critical area of research that has the potential to significantly enhance the capabilities of AI systems and speed up their training. By exploring new methods for discovering and learning skills, our working group reviewed recent works on skill discovery in other domains than game-playing and identified interesting areas for further research. From here on, we outline several future actions to advance skill discovery in strategy games:

- We plan to investigate hybrid/iterative approaches that combine a bottom-up and top-down search for skills. Such hybrid approaches may offer a more robust solution to the challenges of skill discovery in large and complex game environments.
- Leveraging existing game platforms, such as GVGAI [17], Stratega [14, 15], and Griddly [16], could facilitate the testing and validation of new skill discovery methods. These platforms provide standardized environments for benchmarking AI performance, which is crucial for comparing the effectiveness of different approaches.
- Given the limitations of current methods, particularly in terms of scalability, we propose to produce a comprehensive survey paper.

#### References

1 Vreeken, J., Leeuwen, M. & Siebes, A. Krimp: mining itemsets that compress. *Data Mining And Knowledge Discovery.* **23**, 169-214, 2010 (10). https://doi.org/10.1007/s10618-010-0202-x.

2 Mathieu, M., Ozair, S., Srinivasan, S., Gulcehre, C., Zhang, S., Jiang, R., Paine, T., Powell, R., Żołna, K., Schrittwieser, J., Choi, D., Georgiev, P., Toyama, D., Huang, A., Ring, R., Babuschkin, I., Ewalds, T., Bordbar, M., Henderson, S., Colmenarejo, S., Oord, A., Czarnecki, W., Freitas, N. & Vinyals, O. *AlphaStar Unplugged: Large-Scale Offline Reinforcement Learning*, 2023. https://arxiv.org/abs/2308.03526.

3 Arulkumaran, K., Cully, A. & Togelius, J. Alphastar: An evolutionary computation perspective. *Proceedings Of The Genetic And Evolutionary Computation Conference Companion.* pp. 314-315, 2019.

4 Mörchen, F. *Time series knowlegde mining*, Görich und Weiershäuser, 2006.

**5** OpenAI, :, Berner, C., Brockman, G., Chan, B., Cheung, V., Dębiak, P., Dennison, C., Farhi, D., Fischer, Q., Hashme, S., Hesse, C., Józefowicz, R., Gray, S., Olsson, C., Pachocki, J., Petrov, M., O. Pinto, H., Raiman, J., Salimans, T., Schlatter, J., Schneider, J., Sidor, S., Sutskever, I., Tang, J., Wolski, F. & Zhang, S. *Dota 2 with Large Scale Deep Reinforcement Learning*, 2019. https://arxiv.org/abs/1912.06680.

**6** Al-Emran, M. Hierarchical reinforcement learning: a survey. *International Journal Of Computing And Digital Systems*. **4**, 2015.

**7** Stolle, M. & Precup, D. Learning options in reinforcement learning. *Abstraction, Reformulation, And Approximation: 5th International Symposium, SARA 2002 Kananaskis, Alberta, Canada August 2–4, 2002 Proceedings 5*. pp. 212-223, 2002.

**8** Zambaldi, V., Raposo, D., Santoro, A., Bapst, V., Li, Y., Babuschkin, I., Tuyls, K., Reichert, D., Lillicrap, T., Lockhart, E. & Others. *Relational deep reinforcement learning*, 2018. *ArXiv Preprint ArXiv:1806.01830*.

**9** Morales, E., Scaling up reinforcement learning with a relational representation. *Proc. Of The Workshop On Adaptability In Multi-agent Systems*. pp. 15-26, 2003.

**10** Džeroski, S., De Raedt, L. & Blockeel, H., Relational reinforcement learning. *Inductive Logic Programming: 8th International Conference, ILP-98 Madison, Wisconsin, USA, July 22–24, 1998 Proceedings 8*. pp. 11-22, 1998.

**11** Dockhorn, A., Hurtado-Grueso, J., Jeurissen, D., Xu, L. & Perez-Liebana, D., Game state and action abstracting monte carlo tree search for general strategy game-playing. *2021 IEEE Conference On Games (CoG)*. pp. 1-8, 2021.

**12** Luketina, J., Nardelli, N., Farquhar, G., Foerster, J., Andreas, J., Grefenstette, E., Whiteson, S. & Rocktäschel, T., A survey of reinforcement learning informed by natural language. *ArXiv Preprint ArXiv:1906.03926*, 2019.

**13** Eysenbach, B., Gupta, A., Ibarz, J. & Levine, S., Diversity is all you need: Learning skills without a reward function. *ArXiv Preprint ArXiv:1802.06070*, 2018.

**14** Dockhorn, A., Grueso, J., Jeurissen, D. & Liebana, D. STRATEGA: A General Strategy Games Framework. *AIIDE Workshops*, 2020.

**15** Perez-Liebana, D., Dockhorn, A., Grueso, J. & Jeurissen, D. The design of "Stratega": A general strategy games framework. *ArXiv Preprint ArXiv:2009.05643*, 2020.

**16** Bamford, C. Griddly: A platform for AI research in games. *Software Impacts*. **8**, 2021.

**17** Perez-Liebana, D., Liu, J., Khalifa, A., Gaina, R., Togelius, J. & Lucas, S., General video game ai: A multitrack framework for evaluating agents, games, and content generation algorithms. *IEEE Transactions On Games*. **11**, 195-214, 2019.

**18** Liu, G., Hu, E., Cheng, P., Hung-Lee & Sun, S., Hierarchical Programmatic Reinforcement Learning via Learning to Compose Programs, 2023. https://arxiv.org/abs/2301.12950.

**19** Tanneberg, D., Ploeger, K., Rueckert, E. & Peters, J., Skid raw: Skill discovery from raw trajectories. *IEEE Robotics And Automation Letters*. **6**, 4696-4703, 2021.

**20** Pierrot, T., Macé, V., Chalumeau, F., Flajolet, A., Cideron, G., Beguir, K., Cully, A., Sigaud, O. & Perrin-Gilbert, N., Diversity policy gradient for sample efficient quality-diversity optimization. *Proceedings Of The Genetic And Evolutionary Computation Conference*, 2022 (7). https://doi.org/10.1145/3512290.3528845.

**21** Nieto, J. J., Castanyer, R. C., & Giro-i-Nieto, X., *Unsupervised Skill-Discovery and Skill-Learning in Minecraft*. In ICML 2021 Workshop on Unsupervised Reinforcement Learning, 2021.

**22** Kim, T., Ahn, S., & Bengio, Y., Variational temporal abstraction. Advances in Neural Information Processing Systems, 32, 2019.

**23** Kipf, T., Li, Y., Dai, H., Zambaldi, V., Sanchez-Gonzalez, A., Grefenstette, E., ... & Battaglia, P., *Compile: Compositional imitation learning and execution*. In International Conference on Machine Learning, pp. 3418-3428, 2019.

24 Jiang, Y., Liu, E., Eysenbach, B., Kolter, J. Z., & Finn, C., *Learning options via compression.* Advances in Neural Information Processing Systems, 35, 21184-21199, 2022.
25 Wang, Z., Hu, C., Liu, J., & Yao, X., *Negatively correlated ensemble reinforcement learning for online diverse game level generation.* In The Twelfth International Conference on Learning Representations, 2024.

## 3.7 Introducing AI Experience: Games UX in the Age of Generative AI

*Anders Drachen (University of Southern Denmark – Odense, DK), Paolo Burelli (IT University of Copenhagen, DK), Leonie Kallabis (TH Köln, DE), and David Melhart (University of Malta – Msida, MT)*



### 3.7.1 Introduction

The work group considered the evolving role of User Experience (UX) research in the context of digital games as they transition towards using generative AI. Traditionally, game design has been a manual process where designers meticulously craft environments, narratives, and interactions to shape a predictable and controllable user experience. However, with the integration of procedural content generation and generative AI models, such as Large Language Models (LLMs), the landscape of game development is potentially shifting towards an era where games can be dynamically created and adapted, not just during production but also in real-time as players engage with them.

This shift presents new challenges and opportunities for UX research. The essay outlines how existing UX research methods, which rely on controlled testing environments and predictable user interactions, are increasingly inadequate for understanding and evaluating experiences in games that are generated on-the-fly by AI. Traditional UX frameworks are built on the premise that game environments and player interactions can be pre-defined and tested empirically. However, in a generative game context, where AI can autonomously create complex, responsive environments and narratives tailored to individual players, the very foundations of UX research are called into question.

### 3.7.2 How Generative Games Impact UX Research

Several key dimensions of generative games that impact UX research: conversion, complexity, timing, staticness, social complexity, and personalization. These dimensions describe the extent to which game elements are generated, their complexity, when generation occurs (pre-production, at game start, or in real-time), how static or dynamic the generated content is, the number of players involved, and the level of personalization to individual players. Each of these factors adds layers of variability that challenge traditional UX evaluation methods, making it harder to predict and measure user experience outcomes.

If we consider a future where AI could create entire gaming experiences from scratch, adapting continuously to user behavior and preferences, what role is left for human designers? In such a scenario, the role of human designers and traditional UX researchers could diminish, replaced by AI systems that not only generate games but also simulate user responses to test and refine these experiences. This raises profound questions about the future of UX research. Will traditional concepts like sample sizes and controlled environments become obsolete? Will UX researchers need to transform into AI Experience (AIX) engineers who design the parameters and constraints that guide AI-generated experiences?