Discovering and Controlling Diverse Self-Organised Patterns in Cellular Automata Using Autotelic Reinforcement Learning

Marko Cvjetko, Gautier Hamon, Clément Moulin-Frier, Pierre-Yves Oudeyer
Flowers AI & CogSci Lab, Inria, France
marko.cvjetko@inria.fr

Abstract

Autotelic AI algorithms, which pursue self-generated goals, have proven to be effective as automated discovery assistants in cellular automata (CAs). Previous work in this domain focused on algorithms which produce diverse behaviors by setting the automaton's initial conditions. Here, we extend these methods beyond initial-condition search and adapt them to systems that support sequences of closed-loop interventions. Using Lenia as a test environment, we train goal-conditioned reinforcement learning (RL) agents to perform targeted interventions during the system's evolution, guiding it towards desired states. The resulting agent behaviors are robust and diverse, demonstrating the potential of closed-loop interaction for discovery and control. Furthermore, we show that goal-conditioned RL agents performing interventions can discover novel self-organising patterns and generalize to previously unseen and noisy environments.

Companion website: developmentalsystems.org/rl-for-ca

Introduction

Cellular automata (CAs) have long been used as computational models in complex systems science, artificial life, and the study of diverse real-world phenomena. Despite being defined by simple local rules, they can generate remarkably rich and unexpected patterns. In artificial life, CAs have been extensively used to study the questions of self-replication, life, and open-ended evolution (Sayama and Nehaniv, 2025). Many CA systems are chaotic: they are highly sensitive to initial conditions and external perturbations introduced as the system evolves (). More recently, interest in CAs has grown with the development of new expressive models, such as Lenia and Flow-Lenia (Chan, 2018; Plantec et al., 2025), which are continuous extensions of Conway's Game of Life that display lifelike behaviors, and neural cellular automata (Mordvintsev et al., 2020), which incorporate deep learning architectures and optimization methods into the traditional CA framework.

For many years, researchers and hobbyists alike have explored the vast space of possible CA configurations in search of interesting patterns. Although this exploration traditionally relied on manual tuning and random search methods (Wolfram, 2025), recent approaches have also used AI-assisted automated discovery. Among these are autotelic diversity search methods, which generate a sequence of experiments to explore the parameters of a dynamical system by targeting a diversity of self-generated goals (Etcheverry, 2023). Lenia, in particular, has been a favored CA testbed of automated discovery methods, which have uncovered many novel, lifelike patterns (Hamon et al., 2024; Etcheverry et al., 2020; Faldor and Cully, 2024; Kumar et al., 2024).

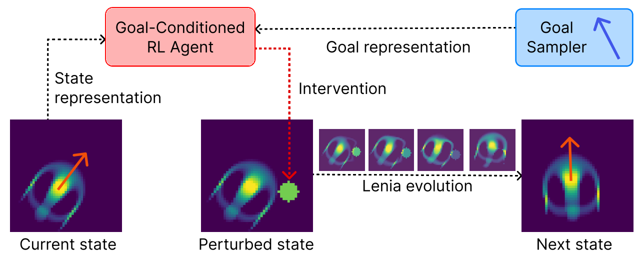

Figure 1: Visualization of how goal-conditioned RL agents can guide the CA grid towards some target state. This process is repeated to form an episode.

To the best of our knowledge, all automated discovery methods applied to CAs so far have been open-loop. In such methods, the algorithm only specifies the initial state of the automaton and observes the outcome after the full rollout. Because CAs are inherently chaotic, the discovered patterns are likely to be highly unstable. Furthermore, open-loop methods do not intervene during rollout, which restricts their usefulness for control problems, even though CAs are frequently used to model real-world systems where control is critical, such as epidemics and ecosystem dynamics. This highlights the need for closed-loop discovery methods that can adjust the automaton's state through targeted interventions during its evolution. Such methods would enhance the expressiveness of the system, potentially enable the discovery of more robust and diverse phenomena, and open a new pathway towards applications in real-world control problems.

We propose a closed-loop approach for AI-driven scien-

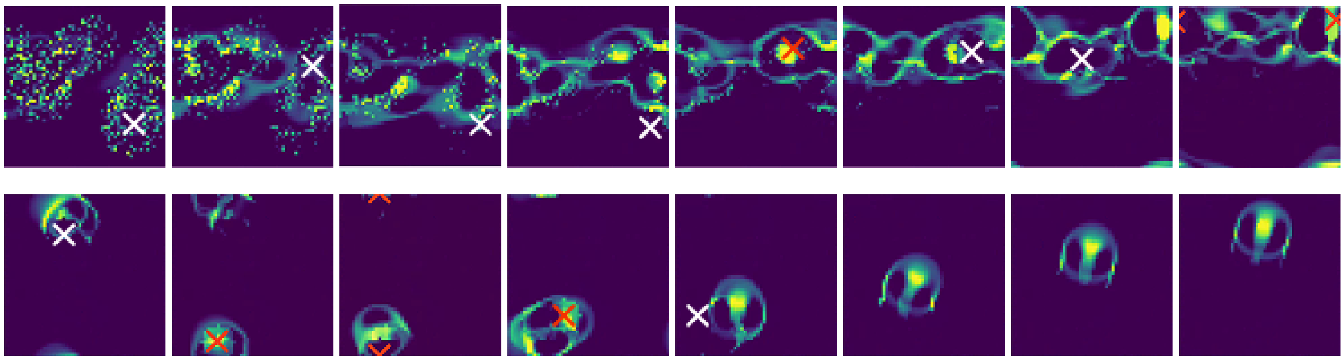

Figure 2: Example of post-training behavior. Adding and removing activations is denoted with red and white colors, respectively. No cross indicates no action is taken. Top: an agent maintains the grid population by constantly removing activations in order to prevent an exploding Turing pattern from forming. Bottom: an agent directs an Orbium towards a target direction.

tific discovery in CAs, consisting of autotelic reinforcement learning (RL) agents that learn to directly perturb CA's states to achieve diverse targets. Although recent CA models have occasionally been used as RL environments (Earle and Togelius, 2024; Sánchez-Fibla et al., 2024), and supervised machine learning methods have guided neural CA pattern growth (Sudhakaran et al., 2022), control and interactive AI-assisted discovery remain largely unexplored.

Experiments

We present first experimental steps towards the larger project of closed-loop AI-assisted discovery[1]. Specifically, we conducted two experiments to investigate whether goal-conditioned RL agents can make effective interventions to guide the CA's evolution to reach a desired state. At each timestep, the agent observes the current grid state and the goal embedding, then selects a location of the CA grid on which to do an intervention, which is either adding or subtracting an amount of activations within a fixed radius. The grid then evolves for k CA steps, after which the agent receives a reward (see Fig. 1). This process repeats for N steps, constituting one episode, with predefined goals sampled uniformly at the start of each episode. We use Lenia as our testbed and train agents using double deep Q-learning with a U-Net–like neural network architecture.

Controlling Grid Population

In the first experiment, we train an agent to maintain the total number of living cells on the grid at a target level. The initial state contains a random number of activation clusters. The reward is the mean absolute error between the current grid sum and the sampled target value. For large targets, the agent learns to add activations in a way that produces self-sustaining, growing Turing patterns. Intermediate targets are handled by continuously reshaping the pattern, while preventing a collapse or an explosion (top of Fig. 2). For smaller targets, the agent learned to construct novel self-sustaining patterns whose size is close to the target value.

[1]Experiment videos are available on the companion website.

Steering an Orbium

The second experiment investigates whether the agent can direct an Orbium towards a target direction. The initial state contains a single Orbium at a random location and orientation on a toroidal grid. The reward is based on the centroid movement along a target direction. The action space additionally includes a *do-nothing* option, and actions that perturb the grid incur a small penalty, encouraging the agent to leverage the CA's intrinsic dynamics. After training, we observe that the agent can steer an Orbium towards any target direction. During a typical episode, the agent first steers the chaotic 'non-Orbium' configuration roughly toward the target direction before discovering a state that naturally converges into a properly oriented Orbium, after which it ceases interventions and allows the pattern to move autonomously (bottom of Fig. 2). We also observe that agents successfully complete the task in noisy environments, where random perturbations are added to the grid between iterations.

Conclusion, Future Work

We have presented two experiments demonstrating that goal-conditioned RL agents can perform targeted interventions during Lenia's evolution and guide it towards desired states. While intervening, trained agents often create self-sustaining patterns, showcasing their ability to discover and exploit intrinsic dynamics of the complex system. Moreover, the agents generalize effectively to previously unseen, noisy environments. Future experiments may explore other environments, goal sampling within episodes, and automatic curriculum learning. Overall, these results represent an encouraging step towards robust closed-loop AI-assisted scientific discovery and control in complex systems, with promising potential for application to real-world systems.

References

Chan, B. W.-C. (2018). Lenia - Biology of Artificial Life. https://arxiv.org/abs/1812.05433v3.

Earle, S. and Togelius, J. (2024). Autoverse: An evolvable game language for learning robust embodied agents.

Etcheverry, M. (2023). *Curiosity-Driven AI for Science: Automated Discovery of Self-Organized Structures*. PhD thesis, Université de Bordeaux. Accessed: 2025-09-06.

Etcheverry, M., Moulin-Frier, C., and Oudeyer, P.-Y. (2020). Hierarchically organized latent modules for exploratory search in morphogenetic systems. *Advances in Neural Information Processing Systems*, 33:4846–4859.

Faldor, M. and Cully, A. (2024). Toward Artificial Open-Ended Evolution within Lenia using Quality-Diversity. In *ALIFE 2024: Proceedings of the 2024 Artificial Life Conference*. MIT Press.

Hamon, G., Etcheverry, M., Chan, B. W.-C., Moulin-Frier, C., and Oudeyer, P.-Y. (2024). Discovering Sensorimotor Agency in Cellular Automata using Diversity Search.

Kumar, A., Lu, C., Kirsch, L., Tang, Y., Stanley, K. O., Isola, P., and Ha, D. (2024). Automating the search for artificial life with foundation models.

Mordvintsev, A., Randazzo, E., Niklasson, E., and Levin, M. (2020). Growing neural cellular automata. *Distill*. https://distill.pub/2020/growing-ca.

Plantec, E., Hamon, G., Etcheverry, M., Chan, B. W.-C., Oudeyer, P.-Y., and Moulin-Frier, C. (2025). Flow-Lenia: Emergent evolutionary dynamics in mass conservative continuous cellular automata. *Artif. Life*, 31(2):228–248.

Sánchez-Fibla, M., Moulin-Frier, C., and Solé, R. (2024). Cooperative control of environmental extremes by artificial intelligent agents. *J. R. Soc. Interface*, 21(220):20240344.

Sayama, H. and Nehaniv, C. L. (2025). Self-reproduction and evolution in cellular automata: 25 years after evoloops. *Artificial Life*, 31(1):81–95.

Sudhakaran, S., Najarro, E., and Risi, S. (2022). Goal-guided neural cellular automata: Learning to control self-organising systems. In *From Cells to Societies: Collective Learning across Scales*.

Wolfram, S. (2025). What can we learn about engineering and innovation from half a century of the game of life cellular automaton? Stephen Wolfram Writings.