Learning Graph Cellular Automata

Daniele Grattarola, Lorenzo Livi, Cesare Alippi

Neural Information Processing Systems 2021

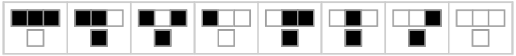

(a) Transition rule

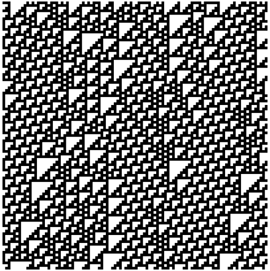

(b) Evolution of CA

Cellular automata

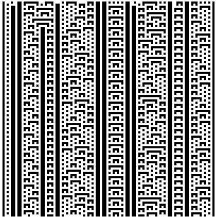

(a) Rule 73

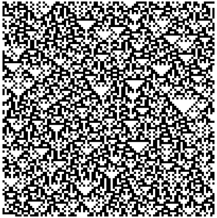

(b) Rule 90

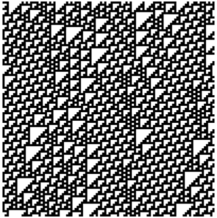

(c) Rule 110

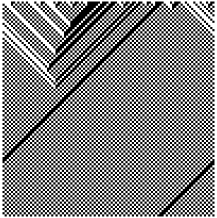

(d) Rule 184

Cellular automata

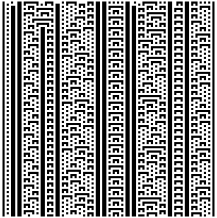

(a) Rule 73

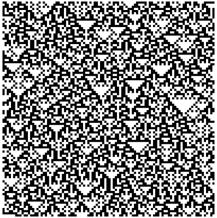

(b) Rule 90

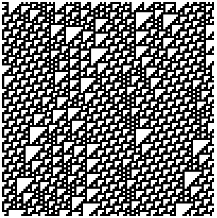

(c) Rule 110

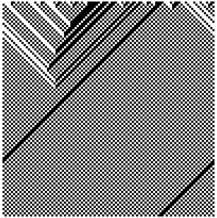

(d) Rule 184

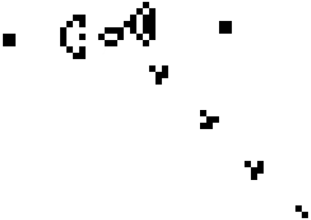

(e) Conway's Game of life

Cellular automata

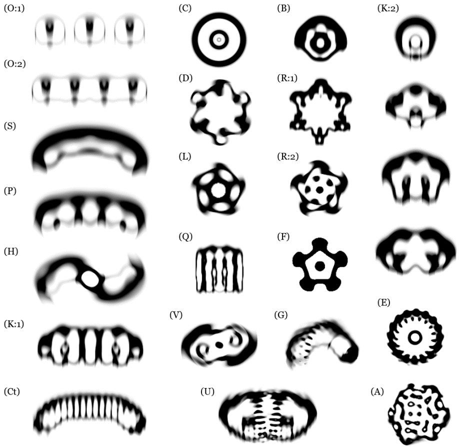

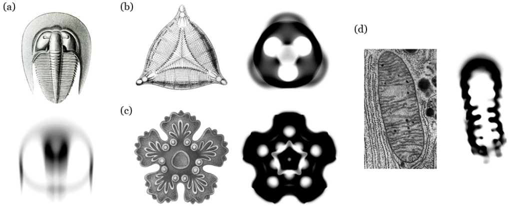

(a) Life-like creatures in Lenia

[1] B. W.-C. Chan, “Lenia-biology of artificial life,” 2018.

Graph cellular automata

A Graph CA is the generalisation of typical CA:

- Cells arranged in a graph;
- State space is any vector space;
- Transition rule τ as a function of neighbours $\mathcal{N}(i)$;

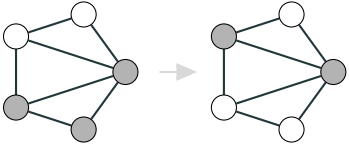

Graph cellular automata

The transition rule has the form:

$$\tau(s_i) : \{s_i\} \cup \{s_j \mid j \in \mathcal{N}(i)\} \mapsto s_i'.$$

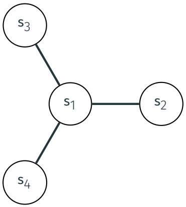

Graph cellular automata

The transition rule has the form:

$$\tau(\mathsf{s}_i) : \{\mathsf{s}_i\} \cup \{\mathsf{s}_j \mid j \in \mathcal{N}(i)\} \mapsto \mathsf{s}'_i.$$

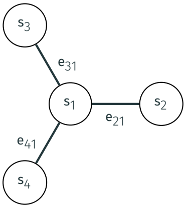

However, we also allow **anisotropic** rules:

$$\tau(\mathsf{s}_i) : \{\mathsf{s}_i\} \cup \{\mathsf{s}_j, \mathsf{e}_{ji} \mid j \in \mathcal{N}(i)\} \mapsto \mathsf{s}'_i,$$

where e_{ji} encodes type, distance, direction, or unique ID of neighbour.

Learning GCA

Problem: how to design a useful rule?

Learning GCA

- **Wulff and Hertz [2]**: learning 1D and 2D CA using "$\sum$-$\prod$ networks with weight sharing";

[2] N Wulff *et al.*, "Learning cellular automaton dynamics with neural networks," 1992.

Learning GCA

- **Wulff and Hertz [2]**: learning 1D and 2D CA using "$\sum$-$\prod$ networks with weight sharing";
- **Elmenreich and Fehérvári [3], Nichele, Ose, Risi, *et al.* [4]**: neuroevolution to learn rules with target behaviour;

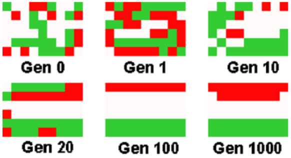

[3] W. Elmenreich *et al.*, "Evolving self-organizing cellular automata based on neural network genotypes," 2011.

[4] S. Nichele *et al.*, "CA-NEAT: evolved compositional pattern producing networks for cellular automata morphogenesis and replication," 2017.

Learning GCA

- **Wulff and Hertz [2]:** learning 1D and 2D CA using "$\sum$-$\prod$ networks with weight sharing";
- **Elmenreich and Fehérvári [3], Nichele, Ose, Risi, *et al.* [4]:** neuroevolution to learn rules with target behaviour;
- **Gilpin [5]:** universal CNN architecture for M-state 2D CA;

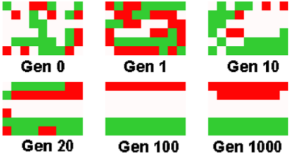

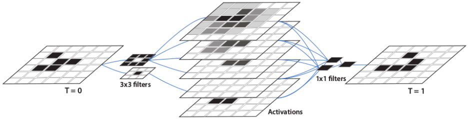

[5] W. Gilpin, "Cellular automata as convolutional neural networks," 2019.

Learning GCA

- **Wulff and Hertz [2]**: learning 1D and 2D CA using "$\sum$-$\prod$ networks with weight sharing";
- **Elmenreich and Fehérvári [3], Nichele, Ose, Risi, *et al.* [4]**: neuroevolution to learn rules with target behaviour;
- **Gilpin [5]**: universal CNN architecture for M-state 2D CA;
- **Mordvintsev, Randazzo, Niklasson, *et al.* [6]**: learning to *grow* a given configuration (inspired by flatworms).

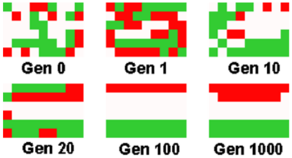

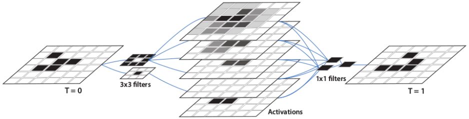

[6] A. Mordvintsev *et al.*, "Growing neural cellular automata," 2020.

Learning GCA

GCA transition rules are message-passing functions:

$$\mathbf{s}_i' = \gamma\left(\mathbf{s}_i, \sum_{j\in\mathcal{N}(i)} \phi\left(\mathbf{s}_i, \mathbf{s}_j, \mathbf{e}_{ji}\right)\right).$$

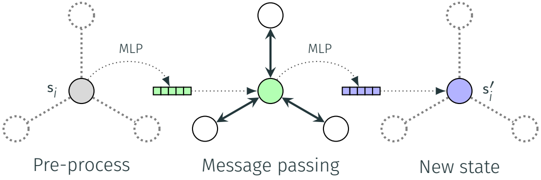

Graph Neural Cellular Automata: GCA with GNN transition rule.

Learning GCA

Extend the results of Gilpin [5] to implement any M-state GCA:

- MLP for one-hot encoding states;
- Message-passing for pattern matching;

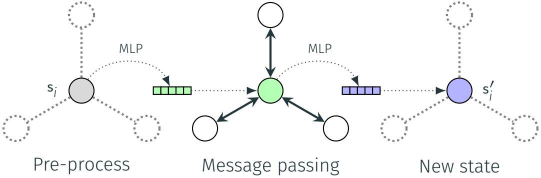

[5] W. Gilpin, "Cellular automata as convolutional neural networks," 2019.

GNCA on Voronoi tessellation

Binary GCA on Voronoi tessellation (equiv. Delaunay triangulation).

Simplest extension to GCA.

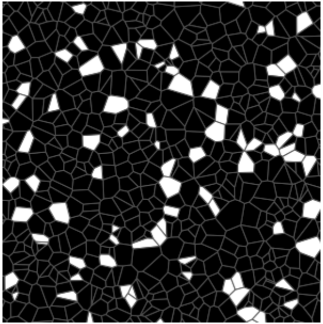

GNCA on Voronoi tessellation

Outer-totalistic rule only depends on the density ρ_i of alive neighbours:

$$\tau(\mathrm{s}_i) = \begin{cases} \mathrm{s}_i, & \text{if } \rho_i \leq \kappa \\ 1 - \mathrm{s}_i, & \text{if } \rho_i > \kappa. \end{cases}$$

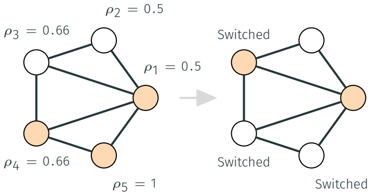

(a) Example transition

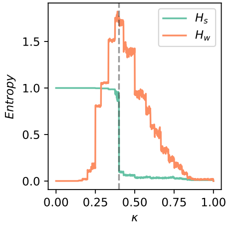

(b) Edge of chaos at $\kappa = 0.4$

GNCA on Voronoi tessellation

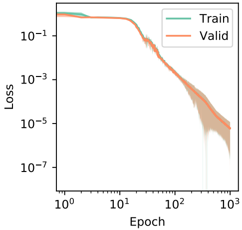

(a) Loss.

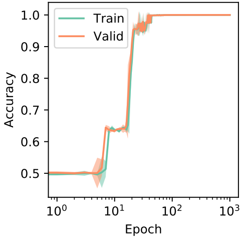

(b) Accuracy.

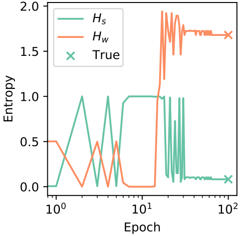

(c) Entropy.

Continuous-state GCA with dynamic graph based on the Boids algorithm [7]:

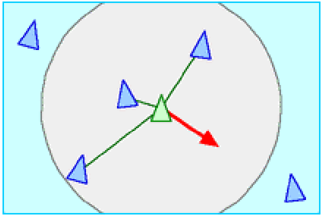

(a) Separation

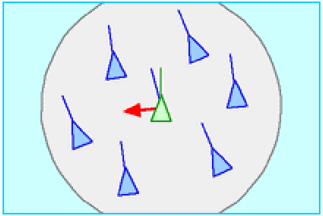

(b) Alignment

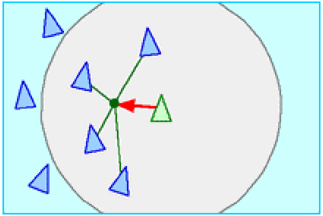

(c) Cohesion

[7] C. W. Reynolds, "Flocks, herds and schools: A distributed behavioral model," 1987.

GNCA for agent-based modelling

Loss goes to 10^{-6} almost immediately, but approximation is not good:

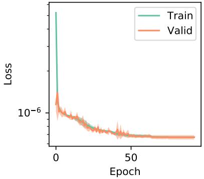

(a) Loss.

Use sample entropy and correlation dimension to evaluate how good the learned rule is.

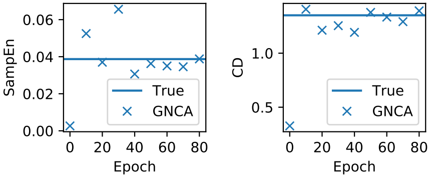

(a) SampEn and CD during training.

GNCA for agent-based modelling

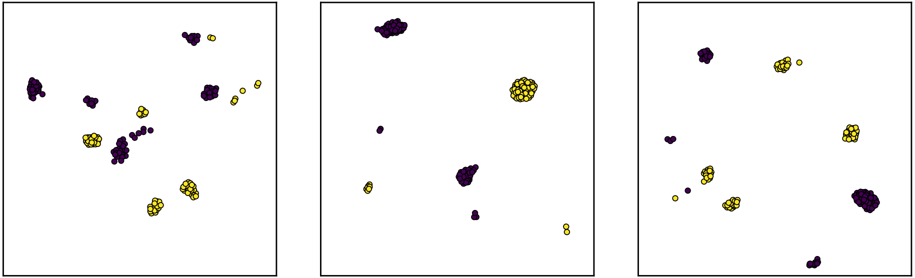

(a) Examples of flocks from the true system and the GNCA.

GNCA that converge to a fixed target

Goal: design a rule with a desired behaviour (converging to target).

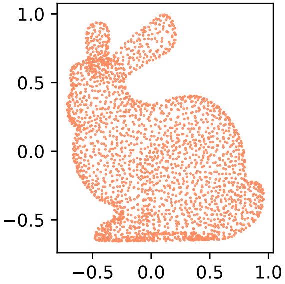

GNCA that converge to a fixed target

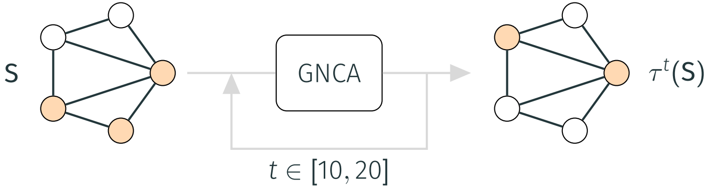

Apply rule for t steps starting from S.

GNCA that converge to a fixed target

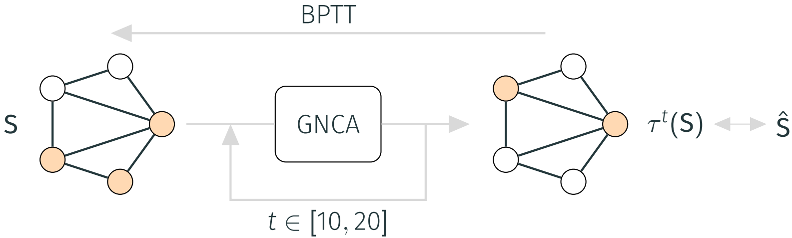

Compute loss w.r.t. target state $\hat{\mathsf{S}}$.

GNCA that converge to a fixed target

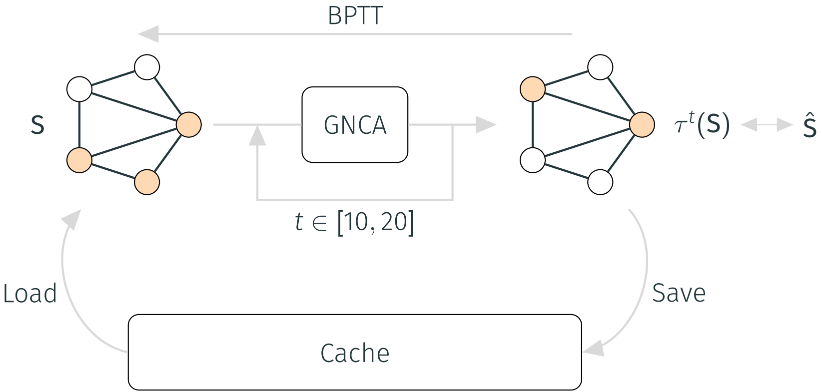

Use a cache to ensure stable attractor and adequate state space exploration [6].

[6] A. Mordvintsev *et al.*, "Growing neural cellular automata," 2020.

GNCA that converge to a fixed target

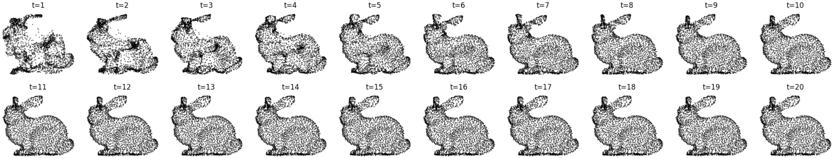

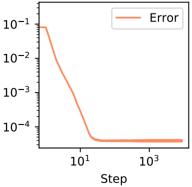

(a) Logo, GNCA trained with $t = 20$, example of convergence.

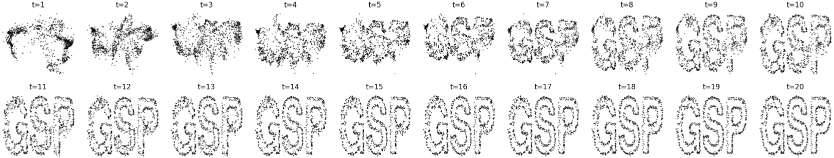

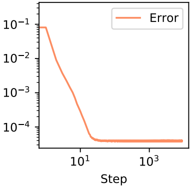

(a) Logo, GNCA trained with $t = 20$, example of convergence.

GNCA that converge to a fixed target

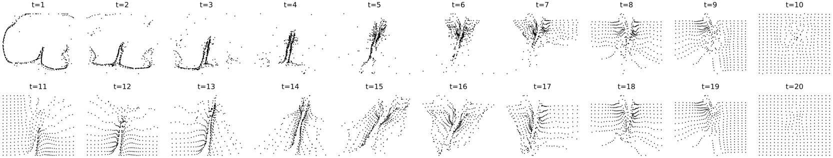

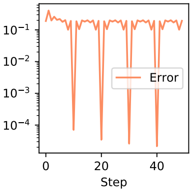

(a) Grid, GNCA trained with $t = 10$, example of periodic behaviour.

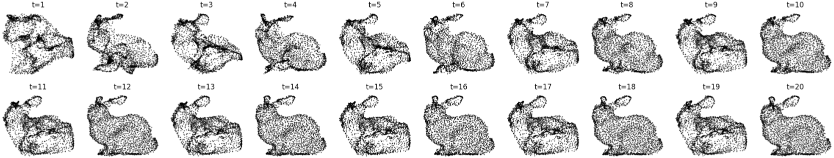

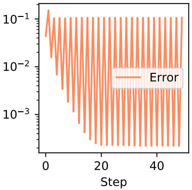

(a) Bunny, GNCA trained with $t = 20$, example of periodic behaviour.

Future research

- Predict global properties (e.g. graph classification);
    - Something similar was done by Gori, Monfardini, and Scarselli [8];

[8] M. Gori *et al.*, “A new model for learning in graph domains,” 2005.

Future research

- Predict global properties (e.g. graph classification);
  - Something similar was done by Gori, Monfardini, and Scarselli [8];
- Decentralised control (e.g. IoT, autonomous vehicles);

[8] M. Gori *et al.*, “A new model for learning in graph domains,” 2005.

Future research

- Predict global properties (e.g. graph classification);
    - Something similar was done by Gori, Monfardini, and Scarselli [8];
- Decentralised control (e.g. IoT, autonomous vehicles);
- Epidemiological networks;
    - Work in progress...

[8] M. Gori *et al.*, “A new model for learning in graph domains,” 2005.

Future research

- Predict global properties (e.g. graph classification);
  - Something similar was done by Gori, Monfardini, and Scarselli [8];
- Decentralised control (e.g. IoT, autonomous vehicles);
- Epidemiological networks;
  - Work in progress...
- Modelling biological systems (e.g. neurons).

[8] M. Gori *et al.*, "A new model for learning in graph domains," 2005.

Conclusion

Learning Graph Cellular Automata

Daniele Grattarola, Lorenzo Livi, Cesare Alippi

Published at NeurIPS 2021

`github.com/danielegrattarola/GNCA`

Summary:

- Learn GCA rules with GNNs;
- Universal architecture for M-state GCA;
- Enable design of GCA from high-level specification;
- Solve tasks through emergent computation.

Get in touch: daniele.grattarola@gmail.com or @riceasphait

References i

[1] B. W.-C. Chan, "Lenia-biology of artificial life," *arXiv preprint arXiv:1812.05433*, 2018.

[2] N Wulff and J. A. Hertz, "Learning cellular automaton dynamics with neural networks," *Advances in Neural Information Processing Systems*, vol. 5, pp. 631–638, 1992.

[3] W. Elmenreich and I. Fehérvári, "Evolving self-organizing cellular automata based on neural network genotypes," in *International Workshop on Self-Organizing Systems*, Springer, 2011, pp. 16–25.

[4] S. Nichele, M. B. Ose, S. Risi, and G. Tufte, "Ca-neat: Evolved compositional pattern producing networks for cellular automata morphogenesis and replication," *IEEE Transactions on Cognitive and Developmental Systems*, vol. 10, no. 3, pp. 687–700, 2017.

[5] W. Gilpin, "Cellular automata as convolutional neural networks," *Physical Review E*, vol. 100, no. 3, p. 032 402, 2019.

[6] A. Mordvintsev, E. Randazzo, E. Niklasson, and M. Levin, "Growing neural cellular automata," *Distill*, vol. 5, no. 2, e23, 2020.

[7] C. W. Reynolds, "Flocks, herds and schools: A distributed behavioral model," in *Proceedings of the 14th annual conference on Computer graphics and interactive techniques*, 1987, pp. 25–34.

[8] M. Gori, G. Monfardini, and F. Scarselli, "A new model for learning in graph domains," in *Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005.*, IEEE, vol. 2, 2005, pp. 729–734.