Investigating dynamics of Neural Cellular Automata applied to image data in diverse complex systems

Vaibhav Mahajan
Department of Computer Science and Technology
University of Cambridge, Cambridge
vm439@cam.ac.uk

Soumya Banerjee *
Department of Computer Science and Technology
University of Cambridge, Cambridge
sb2333@cam.ac.uk

Abstract

Neural cellular automata (NCA) provide a powerful computational paradigm for modelling morphogenetic processes through local interactions and self-organization. We apply NCAs to a number of prototypical complex systems ranging from morphogenesis to reaction-diffusion systems. We explore the capacity of NCA to not only replicate complex visual patterns, but also to learn the underlying update rules of dynamic systems from spatiotemporal image sequences. We reproduce the behaviour of a morphogenesis system through various training regimes and demonstrate how training strategies critically influence the ability of the NCA to grow, persist, and regenerate patterns from visual data. We find that NCAs cannot be applied "out of the box" to these diverse problems but must be adapted. We introduce a stratified multi-step training process that can be used to train NCAs to replicate diverse complex systems from image observations. Our approach demonstrates the potential for learning complex system dynamics from purely visual observations, a key capability for imageomics applications. Lastly we find that NCAs use the hidden channels to generalize to novel behaviour. We further analyse the role of hidden channels in encoding spatial memory and guiding complex pattern formation. Our experiments provide new insights into how neural CA can be adapted as general-purpose models for learning, replicating, and possibly innovating system dynamics from image-based observations. Our findings illustrate the versatility of NCA as a self-organising and rule-learning system (albeit with complex training regimes) and suggest broader applications in modelling natural and artificial systems through visual pattern analysis.

1 Introduction

Cellular automata have long served as a compelling framework for studying complex systems through simple local rules. Recent work by Mordvintsev et al. [2020] introduced a neural formulation of cellular automata (NCA) that combines the interpretability of traditional automata with the flexibility of neural networks. This approach encodes the state of each cell in multiple channels - including both visible (e.g. RGB) and hidden dimensions - and evolves the system through convolutional update rules derived from local perceptions of neighbouring cells.

*Corresponding author

39th Conference on Neural Information Processing Systems (NeurIPS 2025) Workshop: The 3rd Workshop on Imageomics: Discovering Biological Knowledge from Images Using AI.

The original NCA model demonstrated impressive capabilities in self-organizing and regenerating images from seed states, offering a novel framework for differentiable morphogenesis. However, it remains an open question whether such systems can generalize beyond handcrafted targets and be used to infer rules underlying arbitrary spatiotemporal phenomena.

In this paper, we explore the potential of NCA to learn the dynamics of a system from observation alone. We investigate how different training regimes impact the ability of NCA to capture growth, persistence, and regeneration behaviours. In particular, we study how carefully curated datasets and loss strategies can induce desired behaviours such as long-term stability and structural recovery. We further analyze the role of hidden channels in encoding spatial memory and guiding complex pattern formation. Our experiments provide new insights into how neural CA can be adapted as general-purpose models for learning, replicating, and possibly innovating system dynamics.

2 Inferring the Rules

We investigate whether neural cellular automata can be used to reproduce complex behaviour by inferring the rules of a system given only snapshots of how the system evolves, investigating the morphogenesis system of Mordvintsev et al. [2020], the BZ reaction and Conway's Game of Life Gardner [1970a].

The first training regime that was used was to simply provide the model with a snapshot given as input and its successor as the target, so that the model would learn a single step. However, training based only on a single step of the model meant that although the model was quite accurate when predicting a single step, it failed to learn the long-term dynamics of the system. Additionally, the models do not learn to control their behaviour, and running the models for a large number of iterations displays instability.

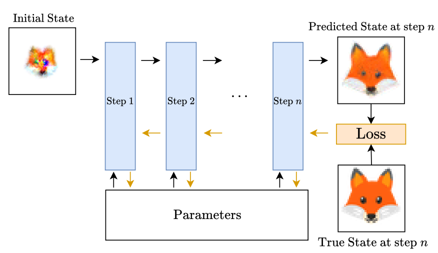

Figure 1: Schematic of the multi-step training strategy.

To allow the model to learn to exhibit long-term dynamics, the loss must be computed over multiple steps. This training regime randomly chooses a number n and computes the loss between the model prediction of the state n time-steps into the future and the true state of the frame at that time-step. This training regime is much more effective than the single-step regime, exhibiting the correct growth and stability behaviour; however, the model still suffered from the fact that the growth stage in the training data was relatively small, which meant that the model was not able to reproduce the pattern to the level of detail which the original NCA produced. Additionally, since the training data did not contain any regeneration data, the model did not learn how to regenerate damaged images.

Thus, the training strategy has to be further modified to emphasise the growing stage and to include regeneration behaviour.

We can broadly categorise the behaviour of the morphogenesis system into two parts: the growing stage, where the seed grows into the pattern over successive updates, and the stable stage, where once the target pattern has been created, successive updates do not change the pattern.

Whether the system is in the growing stage or the stable stage is controlled by the values of the hidden channels. A middle ground is when the system is regenerating a pattern: here, the system maintains the persistence behaviour for most of the image, but exhibits growth on the damaged section.

In order to ensure that the NCA we are training also exhibits this behaviour, multiple runs of the model were used as training data: one where the system is growing, one where the system is stable

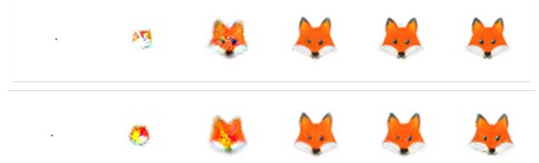

Figure 2: Top: evolution of the original morphogenesis system. Bottom: evolution of the trained NCA model using the stratified multiple step training.

and one where the system is regenerating. Stratified sampling from each of these phases was used per batch, with an emphasis on the growing phase. This ensures that the model does not just learn the simpler behaviour of persisting an already grown image, and but is also able to grow it from scratch like the original morphogenesis system.

3 Influence of Hidden Channels

The hidden channels of the model are crucial to controlling the behaviour of the morphogenesis system. Although the morphogenesis model is only directly trained by taking the L2 loss over the RGBA channels, the hidden channels govern the behaviour of the model. During the growth phase, all channels are active in the pattern. However, in the stable phase, different channels are active in different physical locations of the pattern, using the channels to encode the final pattern.

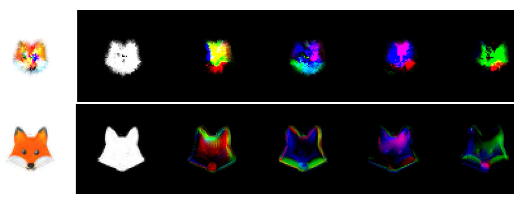

Figure 3: Snapshot of the morphogenesis model displaying all 16 channels at the growing (top) and stable (bottom) phases. The first image is the RGB channels, the second image is of the alpha channel which is used to determine whether a cell is considered alive or dead, and the four other images display three of the hidden channels each.

A key finding of this work is insight into how information about the spatial location of cells is encoded in the hidden channels, allowing the different features of the pattern to be grown at different locations despite using a common update rule for all the cells.

Figure 3 shows a visualisation of the channels for when the system is in the growing and stable phases. When the system is still growing, the expression of all the channels is still changing for each cell, and there is overlap between the locations where each channel takes a high value. However, in the stable state the cells differentiate themselves, using the different channels to encode information about role within the pattern. The channels become active in separate regions of the pattern, with less overlap than in the growing stages (visually, this corresponds to the fact that the colours displayed in Figure 3 are mostly the primary colours). For example, the cells which have large values of channel 8 (corresponding to the green channel in the third image from the right in Figure 3) are concentrated in the lower half of the pattern, but not in the nose, whereas cells which have large values of channel of channel 9 (the blue channel in the third image from the right) are concentrated in the nose and the ears of the fox pattern.

To analyse the spatial distribution of cells with similar states, clustering is performed on the cell states. The k-means clustering algorithm was applied to cluster the cell states (real valued vectors) over multiple iterations and across all spatial locations. Figure 4a visualises the result of clustering the cells into five clusters[2], showing how the cluster membership varies over time as the system grows and then persists the pattern. During the growing stage, the cells are similar, belonging mainly to the

[2]The number of clusters was varied between two and fifteen, and the value of five clusters chosen by the elbow method for both the fox and cat morphogenesis systems.

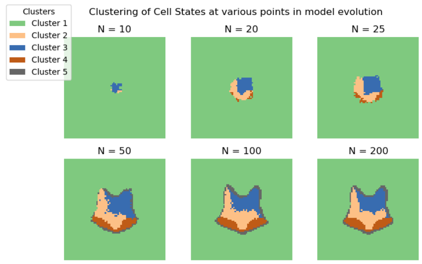

(a) Fox morphogenesis system

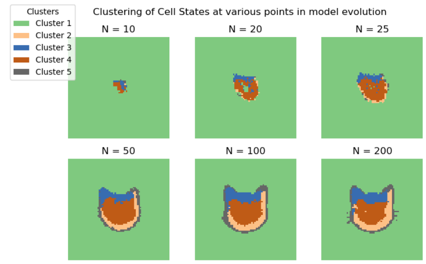

(b) Cat morphogenesis system

Figure 4: Visualisation of clustering of the cells of the stable state of the two systems into five clusters.

second and third cluster. However, as the system grows, the cells differentiate and the fourth and fifth clusters appear. The cells which form the boundary all belong to the same cluster, suggesting that the pattern is prevented from growing further due to the boundary cells encoding that information in the channels. Similar behaviour is illustrated in Figure 4b, which shows the clustering of the cell states in the cat morphogenesis system as the system evolves.

This insight into the behaviour of the morphogenesis system helps explain previous results, which show that reducing the number of channels available to an NCA model diminishes its ability to perform the task of image morphing Richardson et al. [2024]. Using fewer channels limits the model's capacity to encode information about the pattern features, thereby affecting the model's ability to produce detailed patterns. Thus, as the number of available channels is decreased, the model loses the capacity to exhibit complex behaviour.

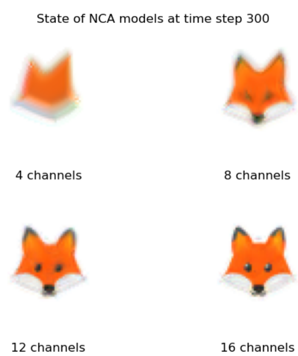

(a) State reached by models using different numbers of channels after 300 iterations

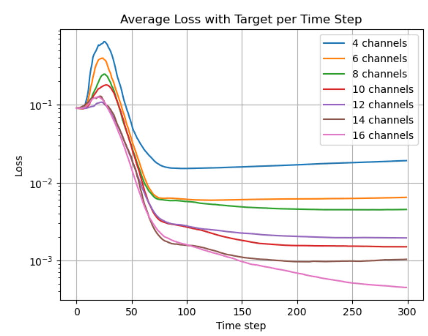

(b) Loss with the target at each iteration of the morphogenesis models

Figure 5: Comparison of morphogenesis models that use different numbers of channels to encode the cell state

4 Conway's Game of Life

Compared to the morphogenesis and BZ reaction systems, Conway's Game of Life Gardner [1970b] has much simpler rules. However, an added complexity is that the system is discrete, whereas NCA are continuous.

The training data was collected by using different well-known patterns of the Game of Life such as the glider gun and pulsar, as well as including a randomly initialised run of the system.

Training using the single step regime achieves loss of 1e-6, suggesting that in this case the single step dynamics are simple enough that they can be learnt. However, we see that when running the model for multiple steps, the model fails to replicate the behaviour of the Game of Life system. This is due to the fact that each step of the model adds noise, and so over multiple steps the noise continuously accumulates and destroys the pattern.

Thus, even for incredibly simple systems such as this one, single step training is not sufficient. To train the model to avoid this accumulation of noise, we again have to use the multiple step training regime, training the model to predict between 1 and 8 steps into the future. It is effective to first train the model using the single step training, and then perform the multi-step training, since the model has effectively learnt the update rule, and only needs to learn to avoid adding noise.

5 Conclusion

Our study reaffirms the ability of neural cellular automata to replicate, persist, and regenerate complex patterns through localised interactions. By designing training strategies that emphasise different phases of morphogenetic evolution, growth, stability, and regeneration, we demonstrate that NCAs can internalise sophisticated rule-based behaviours from data alone. The stratified multi-step training regime proves especially effective in balancing these phases and in preventing models from degenerating into trivial persistence or overfitting to short-term dynamics.

We further highlight the critical role of hidden channels in modulating emergent behaviours, suggesting that these internal representations function analogously to latent fields in biological development. Our experiments show that NCAs not only reconstruct known dynamics but can also generalize to novel behaviours such as pattern reproduction, hinting at their broader applicability in artificial life, system identification, and generative modeling.

Because all of our training and evaluation use spatiotemporal image sequences (RGB/alpha frame observations and derived perception filters), the results speak directly to imageomics workflows: NCAs can be used to learn, compress, and forecast dynamics from raw image data in domains such as microscopy, remote sensing, and experimental imaging. In particular, NCAs present a lightweight and interpretable inductive bias for (i) identifying latent spatial memory fields from images, (ii) performing robust short- and long-horizon forecasting, and (iii) enabling image-level interventions such as targeted regeneration or controlled pattern perturbation.

As neural CA continue to blur the boundaries between rule-based simulation and learned behaviour, future work should explore their use in more abstract domains, from agent-based modelling to distributed computation, and investigate the extent to which they can discover interpretable or symbolic representations of the systems they emulate.

References

Alexander Mordvintsev, Ettore Randazzo, Eyvind Niklasson, and Michael Levin. Growing neural cellular automata. *Distill*, 2020. doi: 10.23915/distill.00023. URL `https://distill.pub/2020/growing-ca`.

Martin Gardner. The fantastic combinations of john conway's new solitaire game life. *Scientific American*, 223:120–123, 1970a.

Alex D. Richardson, Tibor Antal, Richard A. Blythe, and Linus J. Schumacher. Learning spatio-temporal patterns with neural cellular automata. *PLOS Computational Biology*, 20(4):1–27, 04 2024. doi: 10.1371/journal.pcbi.1011589. URL `https://doi.org/10.1371/journal.pcbi.1011589`.

Martin Gardner. The fantastic combinations of john conway's new solitaire game "life". *Scientific American*, 223(4):120–123, 1970b.