Neural Cellular Automata Can Respond to Signals

James Stovold

School of Computing and Communications
Lancaster University Leipzig
Nikolaistraße 10
Leipzig 04109, Germany
Email: j.stovold@lancaster.ac.uk

Abstract

Neural Cellular Automata (NCAs) are a model of morphogenesis, capable of growing two-dimensional artificial organisms from a single seed cell. In this paper, we show that NCAs can be trained to respond to signals. Two types of signal are used: internal (genomically-coded) signals, and external (environmental) signals. Signals are presented to a single pixel for a single timestep.

Results show NCAs are able to grow into multiple distinct forms based on internal signals, and are able to change colour based on external signals. Overall these contribute to the development of NCAs as a model of artificial morphogenesis, and pave the way for future developments embedding dynamic behaviour into the NCA model.

Code and target images are available through GitHub: https://github.com/jstovold/ALIFE2023

Introduction

Signals are the central component of animal interaction, and an essential part of animal life (Maynard Smith and Harper, 1995). We can distinguish between internal signals (signals within the body itself), societal signals (external to the body but internal to the social group of the animal), and external signals (from the environment or from other animals).

Internal signals include hormonal or neural signals within the body (Widmaier et al., 2006), bioelectrical (Adams and Levin, 2013) and genomic signals during development (Jaenisch and Bird, 2003; Wang et al., 2004). External signals include the basic senses and detecting gradients (Parent and Devreotes, 1999). Societal signals include the famous waggle dance (Lindauer, 1957; Camazine et al., 1999) and its inhibitory headbutt (Seeley et al., 2012), the pheromones released by foraging ants (Deneubourg et al., 1989), identity formation signals (Stovold et al., 2014), aggregation signals of *Dictyostelium discoideum* (Gross et al., 1976), dog sneezes (Walker et al., 2017), and—of course—human language and communication. Understanding and studying signals is clearly an important topic; as researchers in artificial life, it is reasonable to ask how this can be included in our models.

Mordvintsev et al. (2020) developed the 'Neural Cellular Automata' (NCA) as a model of morphogenesis, to study the growth of an artificial organism from a single cell. By exploiting the power of artificial neural networks to derive the rules of interaction in a cellular automata (Wolfram, 1984; Adamatzky, 2010) the NCA is able to produce two-dimensional artificial organisms. What NCAs lack, however, is a mechanism for responding to signals.

This paper presents an approach to training NCAs such that they are able to respond to signals. Two types of signal are provided to the NCA: internal and external. For internal signals, we encode the signal into the seed cell, showing that the NCA can grow into different shapes and colours according to the signals provided. For external signals, we provide the signal after the NCA organism has fully grown, and show that the organism can consistently and repeatedly respond to the external signal to change its colour.

Previous approaches to including signal responses in NCAs consist of moving the entire organism up an environmental gradient (Kuriyama et al., 2022), and replacing the organism with a new organism that has a different property (say, colour) by introducing a new seed cell (Cavuoti et al., 2022). In contrast to these approaches, the work presented in this paper uses a single NCA to respond to a signal that is only presented through a single cell for a single timestep. This means that the grown organism is *responding* to a signal rather than the signal acting as a new seed cell triggering the growth of a distinct organism.

Neural Cellular Automata

A conventional Cellular Automata (CA) consists of a two-dimensional grid of cells. Each cell can be in one of two states: alive or dead. Each cell contains a single automaton which updates the state of the cell based on the state of its neighbouring cells. The rules of the automaton were originally designed by hand (Izhikevich et al., 2015), but a variety of techniques (such as genetic algorithms) have also been employed to find CAs with various interesting characteristics (Mitchell et al., 1996). Multiple variants

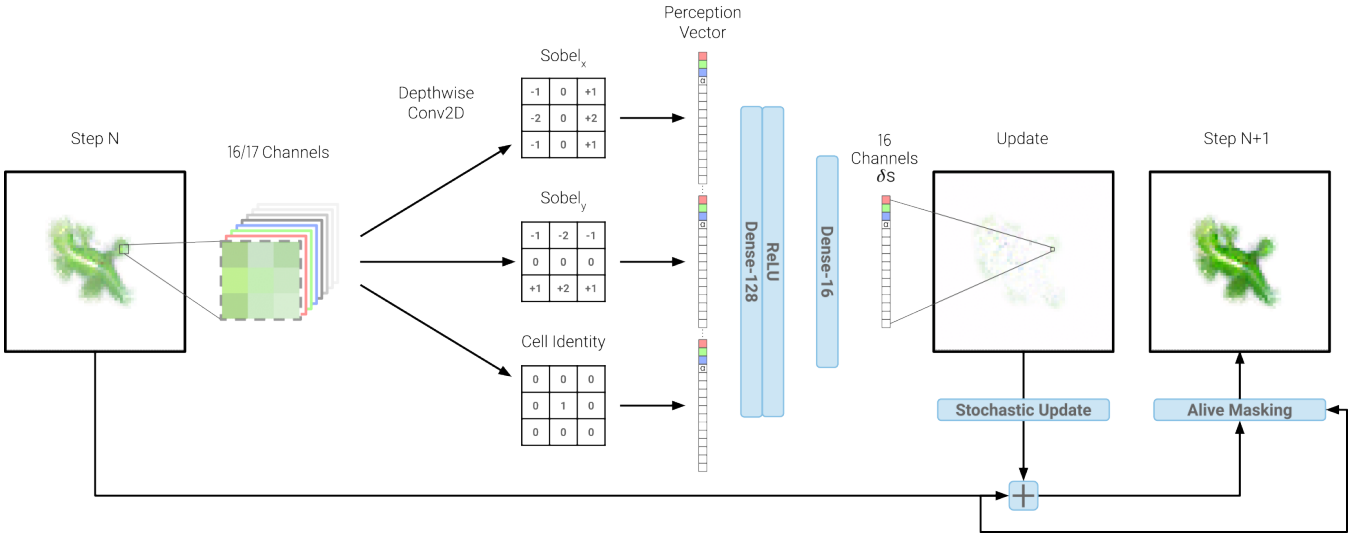

Figure 1: Diagram depicting one pass of the NCA update step. The diagram also shows the structure of the neural network. Image adapted from (Mordvintsev et al., 2020), licenced under CC BY 4.0.

of CAs have been developed, with many focussing on their application to biological phenomena (Ermentrout and Edelstein-Keshet, 1993; Alber et al., 2003; Chan, 2019),

Neural Cellular Automata (NCAs) (Mordvintsev et al., 2020) extend the conventional 2D CA so that the state of each cell is now a vector of 16 real-valued numbers rather than just alive/dead and replaces the update rule with an artificial neural network. The artificial neural network is trained to update the state of each cell in such a way that a certain macroscopic form will emerge from a single seed cell.

In previous work on NCAs, the state vector in each cell consists of the same form: three visible channels representing red, green, and blue of each pixel in the organism, an alpha channel which represents the maturity stage of the cell (see fig. 2), and 12 'hidden' channels which cells use to communicate. The seed cell for a new organism is then a single cell with all values set to 1, except for the RGB channels, which are set to 0. There are occasionally alternative approaches employed, such as multiple seed cells to orient the organism in space (Mordvintsev et al., 2022); or multiple alpha channels to combine two NCAs into one (Cavuoti et al., 2022).

Training an NCA consists of running the CA from a seed cell for a set number of steps (typically in the range 64–200), then comparing the grown organism to a target image. The comparison is used to produce a loss signal which is fed back to the neural network for training via the backpropagation algorithm (Rumelhart et al., 1986). Figure 1 depicts one forward step of the NCA, showing how the neural network acts as the update rule for the CA.

Mordvintsev et al. (2020) introduce three approaches to training the NCA, labelled 'growing', 'persistent', and

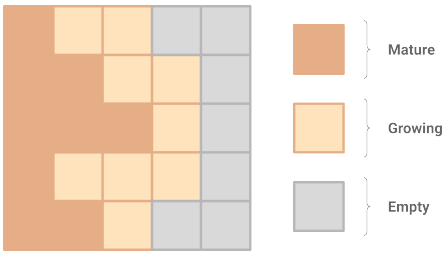

Figure 2: Diagram depicting the role of the alpha channel in the NCA for indicating which cells are alive and which are eligible to come to life in the next step. Image from Mordvintsev et al. (2020), licenced under CC BY 4.0.

'regenerating'. While all three use a batch training approach, the manner in which the batch is manipulated differs between them. In 'growing', the NCA is trained as described above, with the CA iterated for a set number of steps, then the organism compared with a target image.

In 'persistent' and 'regenerating', however, the CA is grown from a seed, but once it is grown it is placed back into a training pool. At each training step, the batch is drawn from the training pool by random sample, meaning some fully-grown organisms are included in the batch (see fig. 3). These fully-grown organisms are iterated in the same way as the original CA, training the NCA to produce organisms that are stable over extended periods of time.

For the 'regenerating' training regime, the same batch approach is used as with 'persistent', but a small portion of the grown organisms extracted from the training pool are damaged before they are iterated. This damage introduces a new signal to the neural network that it should learn how to maintain the organisms even in the presence of external

interference. The 'regenerating' approach to training NCAs forms the basis of our external signal training approach (detailed below).

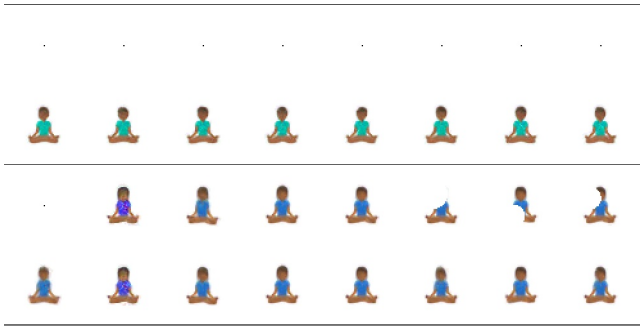

Figure 3: Example batches from the 'growing' (top), and 'regenerating' (bottom) training regimes. Each example comprises two rows, the upper row shows the initial state of the CA, the lower row shows the state of the CA after it has been iterated for between 64 and 200 steps (the actual number of steps is randomly picked each time). The damage caused to the organism is clearly seen in the three rightmost NCAs of the regenerating example. Horizontal lines added to help distinguish the two examples.

Methods

In this paper, we present two types of signal to the NCA: *internal signals*, which are provided through the seed cell (similar to genomic signals during development); and *external signals*, which are provided after the organism is fully developed.

Internal signals are presented to the NCA without requiring any structural changes to be made. For the external signals, however, we extend the state vector to include a read-only 'environment' channel. This channel allows us to provide a signal to a specific cell by setting that channel to 1 for a single timestep. As this increases the number of channels which the NCA has on its input (but not its output), we end up with a slightly different structure to the neural network.

Fig. 1 shows the overall structure of a single update step for the NCA. This is mostly the same as for the original NCA update step defined by Mordvintsev et al. (2020) except that we have 17 input channels and a 'perception vector' of size 51, rather than 16 inputs and a perception vector of 48. The number of output channels remains the same at 16, the 17th channel is read-only so the NCA is not able to make changes, only to receive signals from it.

Fig. 4 shows the structure of our updated seed cells. In both cases the first four channels are the same as other implementations of NCAs, in that c_0–c_2 code for the current RGB colour value of the cell, and the alpha channel (c_3) codes for the maturity of the cell.

When considering internal signals, we encode information into the seed cell to signal to the NCA which organism it should grow. The following n channels (c_4–c_{4+n-1}) encode this length-n binary 'genome'. The remaining channels are set to 1.0 in the seed cell and represent the space available for cells to communicate with one another. It is worth highlighting that, while the seed cell has the genome encoded into the n channels c_4–c_{4+n-1}, this is only provided to the seed cell, and it is up to the NCA to remember and use that information—it is not provided to every cell at every time step.

For external signals, we add an extra channel to the end of the cell state vector which the neural network can read from, but not write to. This allows us to introduce signals to the organism as if from the environment. For the present paper, this channel remains at 0 until we introduce a signal at which point a single cell will have the environment channel set to 1 for a single time step.

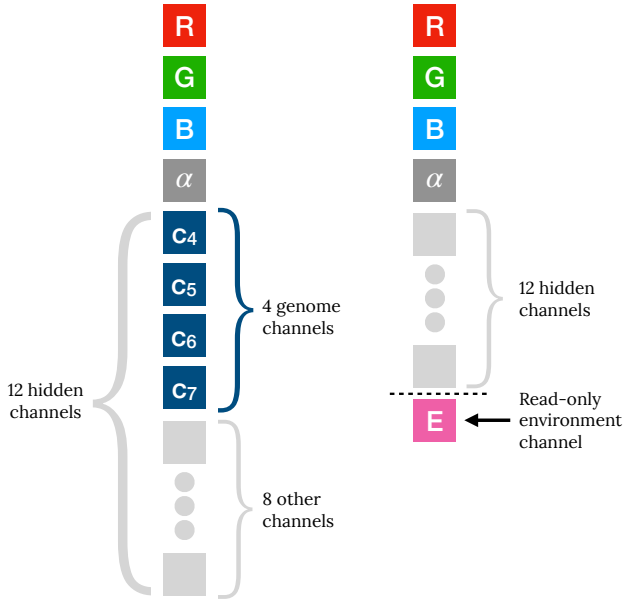

Figure 4: Structure of the NCA seed cell (a) for internal signals, and (b) for reacting to external signals. In (a), we have 4 genome channels, but this could be increased or decreased by adjusting the number of remaining hidden channels accordingly. In no experiments were more than 12 hidden channels used.

Training Approach

All the models presented in this paper were trained from a single seed cell, with a single neural network trained for each experiment (rather than training a different network for each organism).

Internal Signals

To train the NCA to respond to internal signals, we provide different target images depending on the genomic sequence in the seed cell (as depicted in fig. 4). As the signal is only provided at the start, the NCA needs to learn how to remember that signal as it grows, because it will only be evaluated against the target organism after it has finished growing (64–200 steps later). The error between the target and the grown organism is then used to train the neural network via backpropagation.

An early result is shown in fig. 5, where we used a single genomic bit (c_4) to distinguish between growing two distinct forms of meditation emoji. The signal is only provided to the NCA at the start of the run (i.e. in the seed cell).[1]

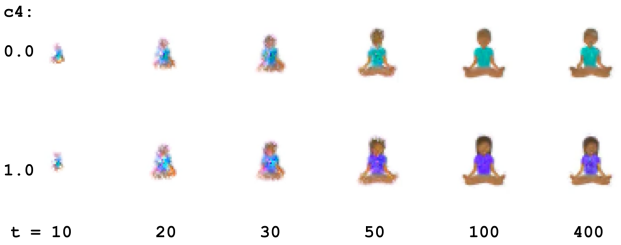

Figure 5: Growth of two meditation emoji from a single NCA trained to produce one of two organisms depending on a single change in the seed cell. The only difference between the two is $c_4 = 0$ for the top example, $c_4 = 1$ for the bottom example. The NCA is trained to produce the organism within 200 timesteps. Snapshots are shown from times $t = 10, 20, 30, 50, 100, 400$.

External Signals

When training the NCA to respond to external signals, we had to adjust the training regime so that the new signal could be introduced during training. The approach taken was to co-opt the concept of damaging the organism from the 'regenerating' regime to introduce an external signal instead. Furthermore, we kept track of the number of times we provided the signal to the organism, to push the NCA to learn how to switch back and forth between two similar organisms. This was in direct response to an early attempt which saw the NCA learn to change in one direction, but when it changed back again it looked similar but had lost the ability to respond to signals (meaning it was actually a different organism). The other benefit of co-opting the damage part of training was that it prevented the signal being inadvertently lost by providing the signal at the same time as the damage.

[1] interestingly, a separate experiment was run to try and encourage the NCA to keep a copy of the genome within each cell, similar to 'real' cells and their DNA, but while this trained more quickly, it appeared to produce less stable organisms

A batch of size 12 was used for training the signal response: the 12 organisms were sorted based on their current loss, with the 3 worst performing being removed from the batch and replaced by a new seed cell. The rest of the batch were split into (a) those which should 'persist' (i.e. they were not provided a signal, but had to continue to exist for another round, preventing the organism from degrading) and (b) those which were provided a signal, and so had to change some property (typically their colour). Organisms that receive the external signal had target images changed at the same time in order for them to learn the response.

Results

This section presents the models that were trained, and shows how the organism varies based on the signal provided. Internal signals, provided through the seed cell, are used in the first two tests to indicate which organism should be grown by the NCA; external signals, provided to the grown organism through the environment, are used in the third test to indicate to the organism that it should change colour.

A subset of the target images used for training the NCAs are provided in fig. 6 (there are 16 variants of the green gecko with different limbs removed, only 5 are shown here).

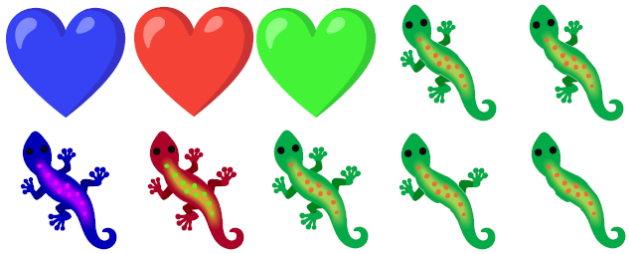

Figure 6: A subset of the target images used for training the NCAs. There are a further 11 geckos with different missing legs that are not shown.

Internal Signals I: Different Forms

This test aimed to see whether a single NCA was able to grow into different colours and shapes based on information encoded into the seed cell—our internal signals. Using the 'growing' training approach, we tested two scenarios: (i) growing the same shape (a heart) in different orientations and different sizes, and (ii) growing two shapes (heart and gecko) in three colours (red, green, and blue).

Figure 7 shows the outcome of scenario (i). For fig. 7a we trained the NCA to produce a small heart if we set c_4 to 0.0 and a large heart if we set c_4 to 1.0. Fig. 7a shows the growth of the NCA at $t = 30, 50, 150$. The top row is when $c_4 = 0.0$, the bottom is $c_4 = 1.0$, and the middle is $c_4 = 0.5$. It is clear that the NCA has learnt to associate the size of the organism with that one internal signal (it is worth emphasising that at no point was the NCA shown the middle-sized heart).

For fig. 7b we similarly trained the NCA to produce an upright heart when $c_4 = 0.0$ and a rotated heart if $c_4 = 1.0$. Unlike the size example, when we give the NCA a signal of $c_4 = 0.5$ it is unable to produce something half-way between the two. This is not unexpected, given the heart was rotated by 180° but we also tested 90° and 270° and found similar behaviour.

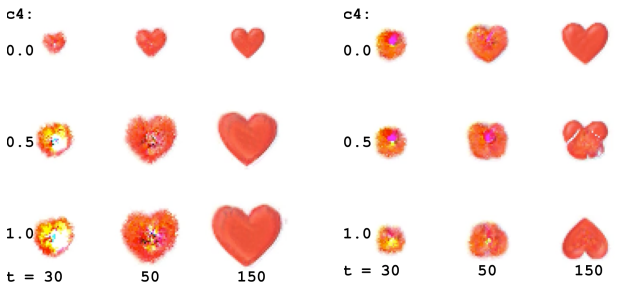

(a) Growth of heart shape with different sizes.

(b) Growth of heart shape with different orientations.

Figure 7: In both cases, from left-to-right: $t = 30, 50, 150$; from top to bottom: $c_4 = 0.0, 0.5, 1.0$. Only $c_4 = 0.0, 1.0$ were trained into the NCA, meaning $c_4 = 0.5$ is an out-of-training test in both cases.

For scenario (ii), we were interested in pushing the limits of the neural network's capacity—how many distinct organisms could be trained into a single NCA? We trained the NCA to grow six organisms: red, green, and blue for both geckos and hearts (see leftmost 6 target images in fig. 6). To encode this into the seed cell, we used four of the hidden channels: c_4–c_7. Table 1 shows the mapping of channel to phenotypic trait.

Channel	Trait
c_4	shape (gecko or heart)
c_5	red
c_6	green
c_7	blue

Table 1: Genome mapping to encode phenotypic traits into the NCA seed cell. For example, a genome of $(0, 0, 1, 0)$ would encode a green gecko and $(1, 1, 0, 0)$ would encode a red heart.

Fig. 8 shows the outcome of this test. The NCA successfully learnt to grow all six different shapes with different colours. The NCA was trained to exist for up to 200 timesteps; fig. 9 shows that past this point, the organisms have a tendency to either disintegrate, sprout other organisms that are trained into the network, or switch to another organism entirely (e.g. the blue heart switching to red after sprouting a green heart on the bottom row).

Similar to the out-of-training tests we performed on the heart scenario above, we tested the ability of the NCA to generalise to other colours by providing seeds with multiple colour genes set to 1. Fig. 10 shows the outcome of this test. While some of the organisms produced were part way between two colours—especially the purple gecko in the second column—closer inspection shows that this was not due to the information provided through the seed cell. In this particular example, the seed included green and blue, rather than red and blue as would have been expected for a purple gecko. Furthermore, it appears that the NCA also produces less-stable organisms when we include multiple colours in the seed cell. It is clear from these images that—with the current approach to training—the NCA is unable to generalise to other colours.

Figure 8: Growth of NCA trained with six distinct figures (three hearts and geckos of different colours); from left to right: $t = 20, 30, 40, 50, 100$ timesteps.

Internal Signals II: Gecko Legs

We wanted to study how the NCA remembers the signal as it grows. One possible mechanism is that it does not remember the signal itself, but that the properties of the organism are sufficiently different for each target that once it starts growing, the growing organism itself embodies the signal. For example, if we code for a blue gecko, once there are blue cells in the NCA, we know that it is either going to be a blue gecko or a blue heart.

To better understand how the NCA remembers the internal signal, we trained lots of very similar target

Figure 9: Extended stability test for multiple figures stored in NCA. The five columns represent the state of the NCA after $t = 200, 300, 400, 500, 1000$ timesteps.

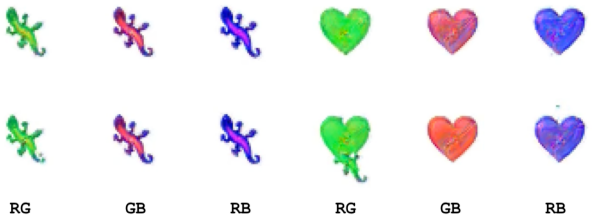

Figure 10: Out of training colour test. From left to right: geckos with: red and green (RG), green and blue (GB), red and blue (RB); similar for hearts; top is $t = 100$, bottom is $t = 400$.

organisms into the same NCA. Specifically, we designed the target organisms in such a way that one organism could easily be changed into another. The target images were all green geckos, but each with a different arrangement of legs (see rightmost four images of fig. 6).

By restricting the difference between geckos to just the legs, the core body shape, head, and tail were all present in all targets. We then use the damage process of the 'regenerating' training regime (whereby we remove a chunk of the image to see how well it can grow back) to remove one of the legs, effectively changing the image into one of the other targets. If the NCA is able to re-grow the correct legs, then we know that it is able to remember which organism is being grown based on the information provided in the seed cell rather than relying solely on the existing shape of the grown organism.

In this test, we encode 16 variants of the gecko into four channels in the seed cell, where each gene signals whether a particular leg l should be grown ($c_l = 1$) or should not ($c_l = 0$); see table 2.

Channel	Trait
c_4	front-left leg
c_5	front-right leg
c_6	back-left leg
c_7	back-right leg

Table 2: Genome mapping to encode phenotypic traits into the NCA seed cell. For example, $(1, 1, 1, 1)$ encodes for all four legs.

Fig. 11 shows the result of training the NCA. Each organism is grown from a distinct seed cell, according to the mapping given in table 2. All sixteen variants of the gecko are trained into the same NCA, demonstrating that an NCA is able to grow many similar organisms as well as the highly-distinct organisms shown before.

Figure 11: Fully-grown form ($t = 200$) of all sixteen distinct geckos from a single NCA, trained using the genomic mapping given in table 2.

To test whether the NCA relies on the current shape of the organism to encode which organism it is, we damaged the grown organisms and tested whether they could regenerate back to their original forms. Fig. 12 shows the outcome of this test. In this example, all the organisms were grown with

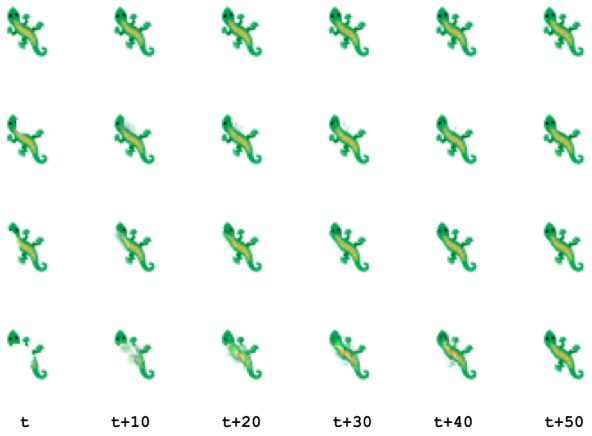

Figure 12: Effect of introducing damage to fully-formed geckos; columns are at $t = 0, 10, 20, 30, 40, 50$ after damage introduced. From top to bottom: no damage (reference case), front-right leg removed, front-left leg removed, both left legs and part of body removed. All seeds are as top row (i.e. all legs).

a seed cell that encodes for all four legs. Upon damaging the organism, the damage is repaired within 40–50 timesteps, but often results in a different organism to which it started. This shows that—in this case—the NCA is not encoding information about the organism into the cells, it relies on the current form of the organism to encode the information instead.

This section has presented a range of different tests showing how NCAs can be trained to respond to internal (genomically-coded) signals. This provides a reasonable starting point for further investigation into the role of the seed cell for providing a genotype for an NCA-grown artificial organism.

External Signals: Colour Changes

While the previous tests were focussed on the ability of the NCA to respond to internal signals encoded in the seed cell of the organism, this section focusses on the ability of the NCA to respond to external signals provided to the fully-grown organism from the environment. For this, we will make use of the new read-only channel added to the NCA state vector (see fig. 4), and will train the NCA to change colour when it receives a signal on this channel. The external signal will be provided to the fully-grown organism through one pixel for one timestep.

Fig. 13 shows the result of training an NCA with this regime. In early training attempts, the NCA would learn to change from green to red, and back to green, but was unable to react to any further signals. This suggests that while the NCA appeared to be responding to the signal, it was actually replacing the organism with a different organism that was unable to respond to subsequent signals. This is not unlike the 'adversarial takeover' NCAs of Cavuoti et al. (2022), but with a single neural network and a single pixel, rather than using multiple trained networks and multiple pixels.

To address this problem, we provide the NCA with the signal multiple times throughout the training process (e.g. if the organism is sampled from the training pool into the current training batch, we might provide another signal to it). The signal is still only presented for a single timestep, and we give sufficient time for the organism to change between signals. This resulted in much more stable behaviour, with the organism able to switch between the two colours when it receives the signal, seemingly as many times as it receives the signal.[2]

To confirm that this behaviour would work on more complicated organisms (i.e. with more than just one block colour), we also tested using gecko target images. As shown in fig. 13b, this works just as well as the heart targets.

In order to ensure the NCA was not just growing a new organism over the top of the old organism, we introduced a small amount of 'jitter' to the provided signal, i.e. we randomly moved the location of the signal within a small region. This jitter pushed the NCA to learn to *respond* to the signal. To ensure that the NCA had not just memorised all the possible locations at which the signal might be provided, we tested the response of the organism with signals outside the region used in training. While these generally took longer to respond, the NCA still responded in the expected manner. If the location was too far away from the originally-trained regions, the response did not occur, but this is not dissimilar to ensuring a signal is provided to the correct nerve to trigger a muscle contraction.

Discussion

In this paper we have demonstrated that NCAs can be trained to respond to signals. We have trained the NCA to respond to two types of signal: those encoded in the seed cell (internal) and those received from the environment (external). In both cases the signal is transient—it is only provided for a single timestep and through a single pixel, the main difference is whether the organism is fully-formed when the signal is received.

Given that NCAs were originally conceived as a model of morphogenesis, the inclusion of signal response into the model opens new avenues for research along these lines, especially focussed on programmed cell death (apoptosis) and cell division (mitosis). We have already started looking at two areas: innervation of muscles in the NCA model, and how we might be able to adapt the training mechanism presented here to develop mitotic behaviour, where a fully-grown cell receives a signal and splits into two cells.

Furthermore, this work contributes to the development

[2] informally, successfully tested up to 28 signals

(a) Fully-grown heart organism responding to two ($t = 0, 100$) external signals.

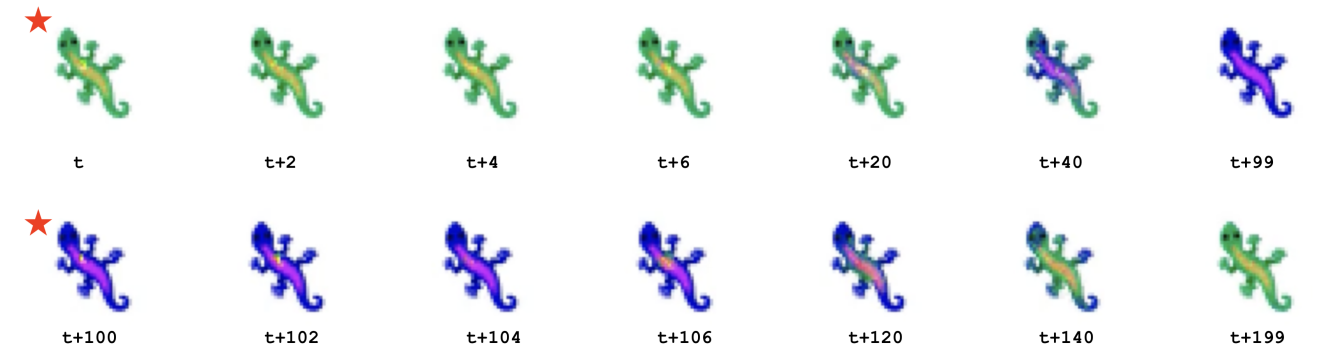

(b) Fully-grown gecko organism responding to two ($t = 0, 100$) external signals.

Figure 13: Two-way colour change as a response to an external signal. The organism used is the fully-trained form of the NCA which can change colour back and forth repeatedly. Column labels indicate number of timesteps after signal provided. Organism on bottom row is continuation of top row. Red star indicates timestep with signal provided.

of an NCA in which the neural network acts as the machinery for expressing genetic information stored in the seed cell (Dreyfuss et al., 2002). This 'generalised' NCA should then be able to produce the corresponding phenotype for any seed cell provided to it, and may allow us to study the computational mechanisms around gene expression in an artificial substrate. This could also be a reasonable point at which to consider training the NCA on real cellular data.

One of the disadvantages of using NCAs to study biological phenomena is the lack of transparency in neural network models. This is a reasonable consideration that needs further work, but in this case we would argue that the lack of transparency can be considered a tradeoff with our ability to engineer the emergent behaviour of the system—a notoriously difficult task (Polack and Stepney, 2005; Stepney et al., 2006).

To conclude, we have presented an approach for training NCAs such that they are able to respond to internal and external signals. We presented results focussed on the ability of the NCA to grow into different organisms based on genomically-coded signals, and to change colour based on environmental signals. Overall these contribute to the development of NCAs as a model of artificial morphogenesis, and pave the way for future developments embedding dynamic behaviour into the NCA model.

References

Adamatzky, A. (2010). *Game of Life Cellular Automata*. Springer Publishing Company.

Adams, D. S. and Levin, M. (2013). Endogenous voltage gradients as mediators of cell-cell communication: strategies for investigating bioelectrical signals during pattern formation. *Cell and Tissue Research*, 352(1):95–122.

Alber, M. S., Kiskowski, M. A., Glazier, J. A., and Jiang, Y. (2003). On cellular automaton approaches to modeling biological cells. In Rosenthal, J. and Gilliam, D. S., editors, *Mathematical Systems Theory in Biology, Communications, Computation, and Finance*, pages 1–39. Springer New York.

Camazine, S., Visscher, P. K., Finley, J., and Vetter, R. S. (1999). House-hunting by honey bee swarms: collective decisions and individual behaviors. *Insectes Sociaux*, 46:348–360.

Cavuoti, L., Sacco, F., Randazzo, E., and Levin, M. (2022). Adversarial Takeover of Neural Cellular Automata. In *Proc. ALIFE 2022: The 2022 Conference on Artificial Life*.

Chan, B. (2019). Lenia: Biology of Artificial Life. *Complex Systems*, 28(3):251–286.

Deneubourg, J. L., Goss, S., Franks, N., and Pasteels, J. M. (1989). The blind leading the blind: Modeling chemically mediated army ant raid patterns. *Journal of Insect Behavior*, 2(5):719–725.

Dreyfuss, G., Kim, V. N., and Kataoka, N. (2002). Messenger-rna-binding proteins and the messages they carry. *Nature Reviews Molecular Cell Biology*, 3(3):195–205.

Ermentrout, G. B. and Edelstein-Keshet, L. (1993). Cellular automata approaches to biological modeling. *Journal of Theoretical Biology*, 160(1):97–133.

Gross, J., Peacey, M., and Trevan, D. (1976). Signal emission and signal propagation during early aggregation in Dictyostelium discoideum. *Journal of Cell Science*, 22(3):645–656.

Izhikevich, E. M., Conway, J. H., and Seth, A. (2015). Game of Life. *Scholarpedia*, 10(6):1816.

Jaenisch, R. and Bird, A. (2003). Epigenetic regulation of gene expression: how the genome integrates intrinsic and environmental signals. *Nature Genetics*, 33(3):245–254.

Kuriyama, S., Noguchi, W., Iizuka, H., Suzuki, K., and Yamamoto, M. (2022). Gradient Climbing Neural Cellular Automata. In *Proc. ALIFE 2022: The 2022 Conference on Artificial Life*.

Lindauer, M. (1957). Communication in swarm-bees searching for a new home. *Nature*, 179(4550):63–66.

Maynard Smith, J. and Harper, D. (1995). Animal signals: Models and terminology. *Journal of Theoretical Biology*, 177(3):305–311.

Mitchell, M., Crutchfield, J. P., and Das, R. (1996). Evolving cellular automata with genetic algorithms: A review of recent work. In *Proceedings of the First International Conference on Evolutionary Computation and its Applications (EvCA '96)*, volume 8.

Mordvintsev, A., Randazzo, E., and Fouts, C. (2022). Growing Isotropic Neural Cellular Automata. In *Proc. ALIFE 2022: The 2022 Conference on Artificial Life*.

Mordvintsev, A., Randazzo, E., Niklasson, E., and Levin, M. (2020). Growing neural cellular automata. *Distill*.

Parent, C. A. and Devreotes, P. N. (1999). A cell's sense of direction. *Science*, 284(5415):765–770.

Polack, F. and Stepney, S. (2005). Emergent properties do not refine. *Electronic Notes in Theoretical Computer Science*, 137(2):163–181. Proceedings of the REFINE 2005 Workshop (REFINE 2005).

Rumelhart, D. E., Hinton, G. E., and Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, 323(6088):533–536.

Seeley, T. D., Visscher, P. K., Schlegel, T., Hogan, P. M., Franks, N. R., and Marshall, J. A. R. (2012). Stop signals provide cross inhibition in collective decision-making by honeybee swarms. *Science*, 335(6064):108–111.

Stepney, S., Polack, F., and Turner, H. (2006). Engineering emergence. In *11th IEEE International Conference on Engineering of Complex Computer Systems (ICECCS'06)*, pages 1–9.

Stovold, J., O'Keefe, S., and Timmis, J. (2014). Preserving swarm identity over time. In Sayama, H., Rieffel, J., Risi, S., Doursat, R., and Lipson, H., editors, *Proceedings of the 14th international conference on the synthesis and simulation of living systems (ALIFE '14)*, pages 728–735.

Walker, R. H., King, A. J., McNutt, J. W., and Jordan, N. R. (2017). Sneeze to leave: African wild dogs (*lycaon pictus*) use variable quorum thresholds facilitated by sneezes in collective decisions. *Proceedings of the Royal Society B: Biological Sciences*, 284(1862):20170347.

Wang, Q. T., Piotrowska, K., Ciemerych, M. A., Milenkovic, L., Scott, M. P., Davis, R. W., and Zernicka-Goetz, M. (2004). A genome-wide study of gene activity reveals developmental signaling pathways in the preimplantation mouse embryo. *Developmental Cell*, 6(1):133–144.

Widmaier, E. P., Raff, H., and Strang, K. T. (2006). *Vander's human physiology*, volume 5. McGraw-Hill New York, NY.

Wolfram, S. (1984). Cellular automata as models of complexity. *Nature*, 311(5985):419–424.