

Unlocking Nature’s Design through Neural Cellular Automata

Frantisek Koutensky, Petr Simanek, Alexander Kovalenko

Czech Technical University in Prague, Prague, Czech Republic
koutefra@fit.cvut.cz

Abstract

This study presents Dynamics Identification via Neural Cellular Automata (DINCA), an enhancement of Neural Cellular Automata (NCA) for modeling reaction-diffusion systems. The main advantage of DINCA is its ability to estimate the parameters of the reaction-diffusion equations that govern the examined system, using minimal data. We demonstrate the method’s application potential by showing its ability to model leopard pattern formation, by learning on only three images, while revealing the governing reaction-diffusion equations. This positions NCA-based methodologies as a viable tool for inferring partial differential equations. The code is available at <https://github.com/koutefra/dinca>.

Introduction

The primary objective of studying various dynamical systems is to understand their spatial-temporal governing equations, often expressed as partial differential equations (PDE) (Leung (2013)). This task is usually non-trivial, especially when striving to adhere to Occam’s Razor, which advocates for the inclusion of only essential terms in these equations (Domingos (1999)). Although Neural Cellular Automata (NCA) (Li and Yeh (2002); Mordvintsev et al. (2020)) have shown remarkable capability in capturing the dynamics of various kinds of dynamical processes (Richardson et al. (2023); Sinapayen (2023); Niklasson et al. (2021)), they do not inherently provide a method for directly deriving the governing PDE.

In this study, we advance the capabilities of NCA by introducing a novel approach termed Dynamics Identification via Neural Cellular Automata (DINCA), which employs a data-driven approach to estimate the underlying reaction-diffusion equations of a system. DINCA refines the standard NCA framework by incorporating interactions among the cells’ channels, which address the reaction component of reaction-diffusion models. For modeling the diffusion aspect, DINCA employs three specific filters: Sobel X, Sobel Y, and Laplacian, which serve to approximate the differential operators responsible for diffusion. Unlike traditional NCA, DINCA omits the hidden layer entirely. Instead, it opts for a series of per-channel linear transformations, which

enhances interpretability by making each cell’s update value linearly dependent on the inputs.

Our study applies DINCA to the domain of pattern formation, namely modeling the development of leopard skin patterns through 2D imaging. The phenomenon of leopard pattern formation serves as a case study for this approach.




Figure 1: The synthesized leopard patterns, as generated in the study by Liu et al. (2006).

Building upon Turing’s foundational theory (Turing (1952)), Liu et al. (2006) proposed a two-stage Turing model designed to replicate leopard patterns through the manipulation of variables such as diffusion rates. This adaptation allows for the simulation of pattern evolution in leopards: from cub spots to juvenile rings, and to adult rosettes. The images synthesized from their model are depicted in Figure 1. By utilizing the capabilities of DINCA, we have successfully deduced the governing equations whose solution models the dynamics observed in the pattern development process, by training the model on the abovementioned three images.

Architecture

DINCA is a dynamical system based on a 2D grid of cells, where each cell is defined by a vector with C channels, representing both visible properties and hidden dynamic processes. A single step in the evolution of this dynamical system, illustrated in Figure 2, encompasses several key components:

- *Perception*: Each cell perceives its environment through a 3x3 convolution using three kernels: Sobel X, Sobel Y, and Laplacian. This step produces gradients which are

then merged into a perception vector with dimensions of $3 \times C$ for every cell.

- **Reaction:** The model’s ability to represent the “reaction” component of reaction-diffusion equations is enhanced in the following way: for a cell’s state vector $\mathbf{x} = (x_1, \dots, x_C)$ at a given step, DINCA considers all combinations of the vector’s components up to the third degree, each assigned a learnable weight. This is formally represented as a set S of possible products:

$$S = \{x_1^{a_1} x_2^{a_2} \dots x_C^{a_C} \mid a_i \in \{0, 1, 2, 3\}\} \quad (1)$$

- **Update Rule:** DINCA updates a cell’s value by linearly transforming its perception vector and elements from the set S into a residual vector δx of dimension C . This residual vector can then be directly added to the cell’s current value. The linear nature of the transformation ensures that each input is assigned a weight proportional to its significance in determining the output (residual vector). This process is carried out for each cell.
- **Stochastic Cell Update:** The model updates cells individually at random intervals, guided by a “fire rate” parameter. A fire rate of 0.5, used in this research, signifies a 50% chance of a cell being updated at each interval.

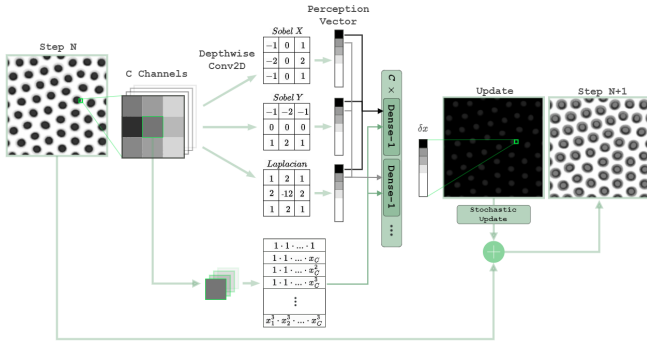


Figure 2: A single iteration step of DINCA.

By setting the fire rate to zero and assuming $C = 2$, we delineate two sets of weights; denote the weight set for the first channel as $w_{1,\cdot}$, and for the second channel as $w_{2,\cdot}$. Let s_i represent the i -th element from the set S . The dynamics of the DINCA model are thus governed by the following system of equations:

$$\begin{aligned} \frac{\partial u}{\partial t} &= w_{1,1} \frac{\partial u}{\partial x} + w_{1,2} \frac{\partial u}{\partial y} + w_{1,3} \nabla^2 u + \sum_{i=4}^{3+4^C} w_{1,i} s_{i-3} \\ \frac{\partial v}{\partial t} &= w_{2,1} \frac{\partial v}{\partial x} + w_{2,2} \frac{\partial v}{\partial y} + w_{2,3} \nabla^2 v + \sum_{i=4}^{3+4^C} w_{2,i} s_{i-3} \end{aligned} \quad (2)$$

Experiment

In our experiment, the dataset consists of three images capturing the first, second, and third developmental stages of leopard patterns, as illustrated in Figure 1. From the first channel, the second channel was created by inverting its values, and the grid was placed in a toroidal configuration. We set $C = 2$, with the aim of deriving a reaction-diffusion equation involving two variables.

The objective is to initialize DINCA with the initial “spots” pattern and then let it evolve. The number of steps, N_1 and N_2 , is randomly chosen at the start of each iteration from a range of 35 to 45 to transition the model first to a “rings” pattern and then to a “rosettes” pattern. At steps N_1 and N_2 , the model calculates the pixel-wise L1 loss for backpropagation compared to the target patterns. To promote sparsity, we apply weight decay with L1 loss and use a pruning strategy to remove weights that fall to zero during training.

The results show that the model successfully identifies the reaction-diffusion equations that govern a process that closely mirrors the original process of leopard pattern formation. The equations specifically include a diffusion term only for the variable v . This choice is not unreasonable, considering that the diffusion of v indirectly affects u through channels interactions. The equations read as follows:

$$\begin{aligned} \frac{\partial u}{\partial t} &= 0.044v + 0.002u^3 + 0.205u^3v - 0.028 \\ \frac{\partial v}{\partial t} &= -0.005\nabla^2 v - 0.034v + 0.079v^3 + 0.001u - \\ &\quad -0.036uv + 0.022u^2 \end{aligned} \quad (3)$$

The structure of the resulting reaction-diffusion equation significantly diverges from traditional models. The reason behind this discrepancy can be attributed to the fact that nearly every theory, including reaction-diffusion, is an approximation to some extent. The process of developing these approximations typically relies on established principles, which inherently limits how the theory can evolve based on its foundational underpinnings. On the other hand, in our case, the model operates in a higher level of flexibility as the main constraint it has is an approximation capacity given by the model’s size.

Conclusion

This study presents DINCA, a novel approach that advances the modeling of dynamical systems through Neural Cellular Automata. By deducing partial differential equations governing leopard pattern formation from limited data, DINCA showcases the potential of NCA-based methodologies for capturing complex dynamics. Our findings highlight DINCA’s versatility in simulating systems and offering insights into underlying physical laws, suggesting a promising avenue for future research in dynamical system modeling.

References

- Domingos, P. (1999). The role of occam's razor in knowledge discovery. *Data mining and knowledge discovery*, 3:409–425.
- Leung, A. W. (2013). *Systems of nonlinear partial differential equations: applications to biology and engineering*, volume 49. Springer Science & Business Media.
- Li, X. and Yeh, A. G.-O. (2002). Neural-network-based cellular automata for simulating multiple land use changes using gis. *International Journal of Geographical Information Science*, 16(4):323–343.
- Liu, R. T., Liaw, S. S., and Maini, P. K. (2006). Two-stage turing model for generating pigment patterns on the leopard and the jaguar. *Phys. Rev. E*, 74:011914.
- Mordvintsev, A., Randazzo, E., Niklasson, E., and Levin, M. (2020). Growing neural cellular automata. *Distill*. <https://distill.pub/2020/growing-ca>.
- Niklasson, E., Mordvintsev, A., Randazzo, E., and Levin, M. (2021). Self-organising textures. *Distill*. <https://distill.pub/selforg/2021/textures>.
- Richardson, A. D., Antal, T., Blythe, R. A., and Schumacher, L. J. (2023). Learning spatio-temporal patterns with neural cellular automata. *arXiv preprint arXiv:2310.14809*.
- Sinapayen, L. (2023). Self-Replication, Spontaneous Mutations, and Exponential Genetic Drift in Neural Cellular Automata. In *ALIFE 2023: Ghost in the Machine: Proceedings of the 2023 Artificial Life Conference*, ALIFE 2023: Ghost in the Machine, page 14.
- Turing, A. M. (1952). The chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 237(641):37–72.