

神经元细胞自动机 (NCA) : 生物自组织与AI的融合思想核心示意图

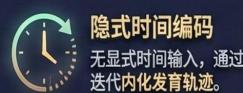
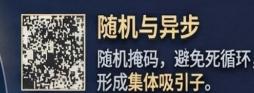
一、归纳等式

$$\begin{aligned}
 \text{NCA} &= [\text{状态预测机} + \text{预测规则生成器} + \text{扩散模型类比} + \text{时间因子强调}] = \\
 &= [\text{演化规则本身的训练及预测} + \text{规则学胜利} + \text{外推能力提升} + \text{分布式自组织}] \\
 &= [\text{非生硬数据流形匹配} + \text{非大量数据的方程拟合} + \text{非传统强化学习优化目标}]
 \end{aligned}$$

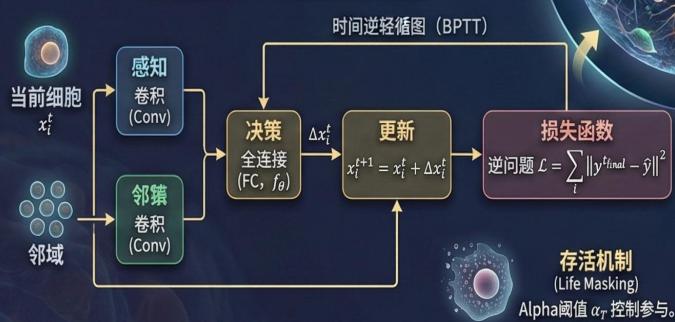
二、思想与逻辑



2. 逻辑



三、公式与机制



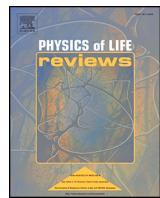
四、结论与应用

结论



数据与应用





Neural cellular automata: Applications to biology and beyond classical AI



Benedikt Hartl ^a, Michael Levin ^{a,b}, Léo Pio-Lopez ^{a,*}

^a Allen Discovery Center at Tufts University, Medford, MA, USA

^b Wyss Institute for Biologically Inspired Engineering at Harvard University, Boston, MA, USA

ARTICLE INFO

Communicated by: J. Fontanari

Keywords:

Neural cellular automata
Artificial intelligence
Biology
Morphogenesis
Collective intelligence

ABSTRACT

Neural Cellular Automata (NCA) represent a powerful framework for modeling biological self-organization, extending classical rule-based systems with trainable, differentiable (or evolvable) update rules that capture the adaptive self-regulatory dynamics of living matter. By embedding Artificial Neural Networks (ANNs) as local decision-making centers and interaction rules between localized agents, NCA can simulate processes across molecular, cellular, tissue, and system-level scales, offering a multiscale competency architecture perspective on evolution, development, regeneration, aging, morphogenesis, and robotic control. These models not only reproduce canonical, biologically inspired target patterns but also generalize to novel conditions, demonstrating robustness to perturbations and the capacity for open-ended adaptation and reasoning through embodiment. Given their immense success in recent developments, we here review current literature of NCAs that are relevant primarily for biological or bioengineering applications. Moreover, we emphasize that beyond biology, NCAs display robust and generalizing goal-directed dynamics without centralized control, e.g., in controlling or regenerating composite robotic morphologies or even on cutting-edge reasoning tasks such as ARC-AGI-1. In addition, the same principles of iterative state-refinement is reminiscent to modern generative Artificial Intelligence (AI), such as probabilistic diffusion models. Their governing self-regulatory behavior is constraint to fully localized interactions, yet their collective behavior scales into coordinated system-level outcomes. We thus argue that NCAs constitute a unifying computationally lean paradigm that not only bridges fundamental insights from multiscale biology with modern generative AI, but have the potential to design truly bio-inspired collective intelligence capable of hierarchical reasoning and control.

1. Introduction

Neural cellular automata (NCA) have recently emerged as versatile computational paradigm, bridging insights from biology and AI through their capacity to learn adaptive local rules that scale into distributed goal-directed behavior and system-level outcomes [1,2]. Unlike traditional machine learning systems, which often struggle with robustness and generalization in open-ended, dynamic environments that require abstraction and reasoning [3–5], NCAs excel at modeling tasks ranging from evolutionary development and self-orchestrated morphogenesis [1,6–13], show great promise to better understand biological regeneration and aging [14–16], can

* Corresponding author.

E-mail address: leo.lopez@tufts.edu (L. Pio-Lopez).

efficiently be trained as robust distributed controllers of composite robotic entities [17–23], and are even competitive at challenging benchmarks for abstraction and reasoning tasks [24,25] (see Fig. 1 for the different types of NCAs and Fig. 1 for the canonical example from [1] and classical CAs). In this review, we survey modern trends in NCA research and literature, and relate their operational principles to the multiscale competency architecture of biological organization [26–30] (see Fig. 1A and B), from multicellular collectives to hierarchical neuronal pathways in the neocortex [31–33]. NCAs show great promise as unifying computational framework for the emerging fields of diverse intelligence and truly bioinspired collective AI [34–36].

Cellular Automata (CAs) are one of the oldest algorithms we can find in the field of artificial life [37], see Fig. 2A–C for an illustration. A cellular automaton is a discrete computational model consisting of a grid of cells, where each cell exists in one of a finite number of states, which evolve over discrete time steps based on a fixed deterministic rule. This rule is applied to all cells individually and will determine the state transition of the cells based on their own state and the one of their neighboring cells. In that way, CAs are dynamical systems which exhibit complex global behavior from simple local interaction and computation. Decentralized systems such as Elementary Cellular Automata (Fig. 2A) have been used to model the behavior of complex systems in biology and physics (c.f.,

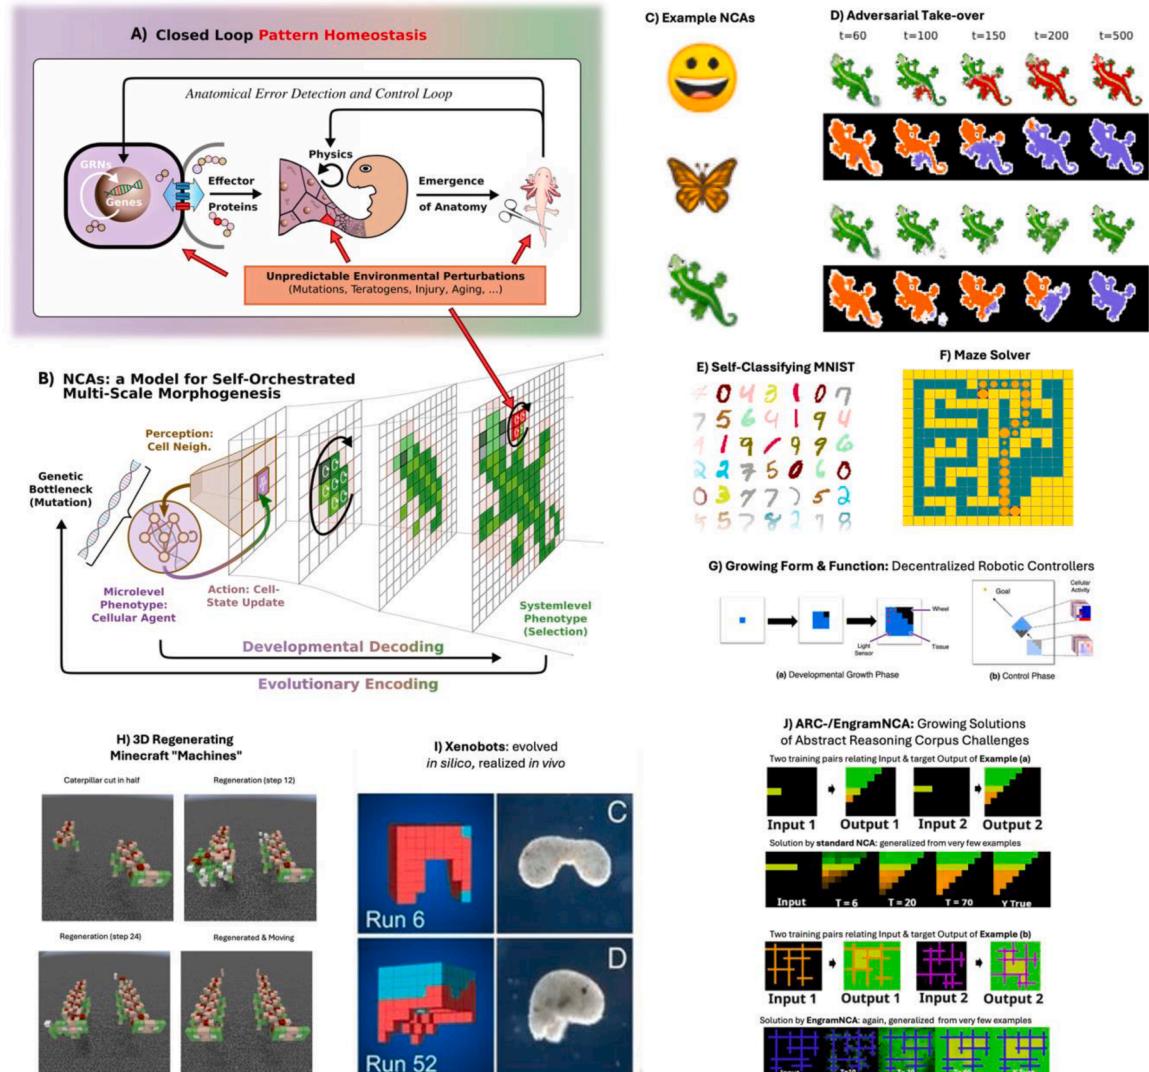


Fig. 1. Overview of Neural Cellular Automaton (NCA) models and applications. A) Development and morphogenesis as multiscale competency architecture via scale-free homeostatic error-correction [6]; B) NCAs as models for biological evolutionary development through decentralized collaborative decision-making, and for the « generative genome » [6]; [1] C) Example patterns that can be learned by NCAs [1]; D) Adversarial take-over experiments of phenotypic traits such as skin color (top row) and phenotypic shape (bottom row) of pretrained NCAs [15,16]; E) Self-classifying MNIST digits of NCA cells negotiating about what digit they are part of [54]; F) NCAs find shortest paths and solve mazes [12]; G) NCAs as unifying substrate for growing form & function of decentralized robots [18,19,21,22]; H) Self-Assembling and regenerating 3D Minecraft machines [22]; I) Xenobots – artificial, intelligent, yet fully biological novel lifeforms that have been designed in silico and realized in vitro [66]; J) EngramNCA trained and deployed on ARC-AGI-1 tasks [8,24].

Fig. 2B), and despite their simplicity, it has been shown that CAs like Elementary CAs, Conway's Game of Life (Fig. 2C), and Wireworld, are capable of universal computation [38,39] showing their remarkable modeling capacity and representational richness. Earlier work performed by von Neumann, Barricelli, and Langton used CAs to respectively model biological mechanisms like self-reproduction [40], evolution [41], and artificial life in general [42]. Since then, CAs have been applied to a variety of biological and computational phenomena including universal constructors [43], biological pattern formation [44], complexity [45,46], and of most relevance here, they have been used to model emergent collective behavior in complex systems [47–51].

Very recently, CAs have gained a renewed attention in the context of artificial intelligence (AI). Modern computational capabilities have allowed researchers to develop more sophisticated CA models that can process information in a way very similar to neural networks [52]. Miller's pioneering work demonstrated the successful evolution of cellular automata with self-repair and self-regulation capabilities for producing the "French flag" pattern, establishing that evolutionary algorithms can identify multiple distinct rules sets capable of morphogenesis [53]. In the same perspective, Nichele et al. evolved compositional pattern producing networks for cellular automata morphogenesis and replication [44]. These developments have led to the emergence of NCAs, which combine the distributed processing characteristics of traditional CAs with the learning capabilities of neural networks [1,2]. Unlike conventional CAs with fixed rules, NCAs substitute their hardcoded local update functions with cell-specific Artificial Neural Networks (ANNs); as universal approximators they can be trained via gradient-based optimization or evolutionary algorithms on complex collective target behaviors to solve, for instance, the inverse task of target pattern assembly, or *in silico* morphogenesis.

Several important developments have expanded the scope and functionality of NCAs beyond basic morphogenesis tasks. Self-classification capabilities were demonstrated by Randazzo et al. [54] (see Fig. 1E), who showed that NCAs could learn to categorize their own emergent patterns. Solving spatial problems like solving a maze were demonstrated too [12] (see Fig. 1F). The concept of embedding spaces for morphogenetic processes was introduced through the NCA manifold framework [55], enabling the

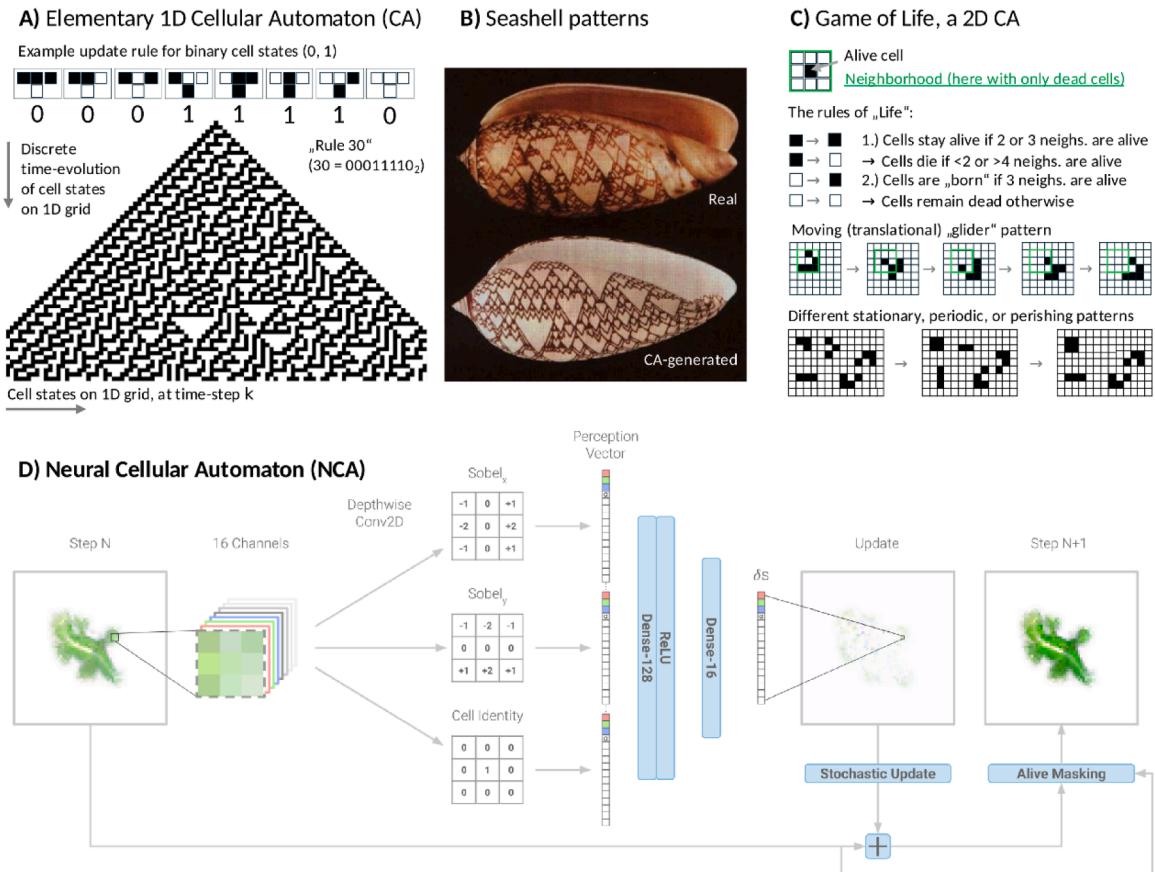


Fig. 2. From elementary Cellular Automata (CAs) to Neural Cellular Automata (NCAs). A) An elementary CA with binary states (0: white, 1: black) on a one dimensional grid (horizontal axis) that evolve based on a fixed update function over discrete time steps (from top to bottom). Such an elementary 1D CA has $2^8 = 256$ possible local neighborhoods, and thus 256 possible update rules; here, we present a CA's state evolution following "rule 30", i.e., $30 = 00011110_2$ in binary notation, which causes intriguing complex and seemingly random patterns [45]. B) Color-patterns on seashells (top is real, bottom is CA generated) [74]. C) A two dimensional CA on the example of Conway's Game of Life [151]. The Moore neighborhood, update rules, and example state-evolution of stationary, periodic, and moving patterns are shown; different initial states can give rise to vastly differently complex behavior. D) Canonical (growing) NCA architecture, illustrating self-regulatory update steps applied on a 2D lizard morphology [1].

representation of different developmental trajectories in structured latent spaces. Tesfaldet et al. incorporated self-attention mechanisms directly into the cellular update rules, allowing individual cells to selectively focus on relevant neighborhood features [56]. Information-theoretic principles were integrated into NCA design, by introducing empowerment as an auxiliary objective function to encourage coordinated cellular behaviors that enhance system robustness and adaptability [57,58]. The stochastic modeling of emergent dynamics was advanced by Palm et al. through their Variational Neural Cellular Automata framework, which employs variational inference to capture the probabilistic nature of pattern formation processes [59]. Multiscale emergent phenomena were developed by Pande and Grattarola (2023), who proposed hierarchical NCA architectures capable of systematically modeling intercellular behaviors across different levels of resolution of hierarchically stacked NCAs [11].

The intersection of cellular automata with deep learning has opened new avenues for research in several domains. In computer vision, they are used for segmentation and texture synthesis applications [60–62]. In robotics, Variengien et al. investigated the evolutionary optimization of NCA-based control architectures for balancing tasks in cart-pole environments, demonstrating the feasibility of using distributed cellular computation for dynamic control problems [23]. And Horibe et al. explored the regenerative capabilities of NCAs in soft robotics applications, where cellular update rules enable autonomous morphological reconstruction following structural damage, similarly to the self-repair mechanisms observed in biological organisms [18,19]. Pontes-Filho et al. demonstrated that growing NCAs can be used to first assemble an *in silico* robotic morphology, i.e., self-assemble a robotic body with tissue-, sensor-, and actuator-cell differentiation steps on an NCA's grid throughout a developmental timeframe, and second utilize the robots' composite body and the NCA's intercellular communication protocols as embodied policy for 2D navigation tasks [21] (see Fig. 1G); relatedly, 3D self-assembly [63] and behavioral policies [22] (see Fig. 1H) have been achieved. Furthermore, Najarro et al. demonstrated that 3D NCAs can be evolved to self-assemble hypernetworks that parametrize large-scale policy networks of RL agents, e.g., for the Lunar-Lander environment, an approach termed HyperNCA [20].

Recently, Hartl et al. successfully demonstrated that NCAs can be evolved to robustly control deformable microswimmers *in silico* in a fully decentralized way by utilizing locally coordinated body deformations in a physically realistic hydrodynamic environment; strikingly, these decentralized control policies generalized well across morphologies, i.e., swimmer sizes, and even allowed robust locomotion if large fractions of the swimmer's actuators were removed or disabled [17]. A related approach by Medvet et al. and Pigozzi et al. used local self-attention and NCA-like control paradigms to evolve embodied soft-robotic 3D morphologies [64,65] which have assisted as *in silico* models to the successful creation of novel *in vitro* lifeforms, i.e., Xenobots [66] and Anthrobots [67,68] (see Fig. 1I): living, swimming self-powered biobots. Decentralized control in autonomous soft robotics thus holds great potential for future biomedical applications. NCAs have also been used to model complex biological phenomena, including for example morphogenesis [1, 22,69], the scaling of goals from metabolic to organism-level goals via homeostasis [9], bioelectricity [70], memory [8], aging [14], cortex mechanisms [71], and evolution [7] (see Section 3 for details).

Indeed, by design, the NCA architecture mirrors a large part of biological organization and particularly the multiscale competency architecture, where biological organisms exhibit nested hierarchies of goal-directed behaviors operating across different spatial and temporal scales [6,26,27,36]. In biological systems, individual cells maintain homeostatic goals at the molecular level while simultaneously participating in tissue-level morphogenesis, organ-level functionality, scaling biological organization up to organ-level, organism-level behavior, and beyond. Similarly, NCAs show a multi-competency architecture by design as they are composed of cells embedding a neural network for decision-making and actions. Such NCAs can be organized into hierarchical structures where local cellular objectives emerge into progressively larger-scale system competencies [9,11]. These systems can exhibit anatomical homeostasis [26] and, in turn, can serve as *in silico* model systems to study biological phenomena like aging and rejuvenation when the cybernetic tissues begin to decline [14]. Therefore, NCAs are models particularly suited to investigate multiscale information processing in biological or artificial life systems. In addition, the distributed architecture of NCAs show good robustness and error-correction mechanisms similarly to biological systems [9,18,72,73]. Just as biological organisms can maintain functionality despite cellular damage or mutation through compensatory mechanisms, trained NCAs demonstrate remarkable resilience to perturbations, damage, or partial removal of cellular components while maintaining their ability to regenerate target patterns [1,9,18]. The temporal dynamics of NCAs also align closely with biological time scales, operating through iterative local updates that mirror the continuous, asynchronous cellular processes in living systems. Consequently, NCAs serve as a bridge between artificial intelligence and biology, offering a computational substrate that naturally incorporates the fundamental organizational principles underlying life itself. The interest of NCAs extends beyond pure biological applications and may present a new kind of AI based on collective intelligence and can show properties difficult to achieve in conventional AI and deep learning like self-repair capabilities, adaptive morphology and architecture, and the scaling of intelligence.

In this paper, we review the different use cases of NCAs for biological modelling, including morphogenesis and self-organization, aging, regeneration and bioelectricity, molecular design, and the generative genome. Then, we discuss why NCAs may be a model of choice for the modeling of biological systems based on the framework of biological multiscale competency architecture, i.e., by the idea that biological systems are made of agential material. We also discuss why they may be particularly important beyond classical AI, as NCAs are AI based on collective intelligence. We conclude with the challenges and limitations this kind of models are currently facing.

2. Neural cellular automata fundamentals

A CA comprises a discrete, typically 2D grid of cells $i \in 1, \dots, N$, each maintaining a numerical, often discrete but potentially also continuous vector-valued state s_i^t . The state s_i^t of each cell i evolves over discrete time steps $t \rightarrow t + 1$ via local transition rules $f : s_i^t \rightarrow s_i^{t+1}$

by considering the cells own state and all the states of its neighbors. Formally, this can be expressed as $s_i^{t+1} = f\left(s_i^t, \left\{s_j^t\right\}_{j \in \mathcal{N}(i)}\right)$, where $\mathcal{N}(i)$ represents the neighborhood of cell i on the CA's grid, often a Moore neighborhood. Even hardcoded update functions give rise to elaborate complex behaviors, but the framework allows any local update function (deterministic, stochastic, etc.). In Fig. 2A-C we illustrate the state evolution of a 1D elementary CA, a biological example on seashell patterning following CA-like rules [74], and an illustration of a 2D CA on the example of Conway's Game of Life.

NCAs extend CAs by using learnable – often differentiable – update rules $f \rightarrow f_\theta$ that are realized by an ANN with parameters θ ; see Fig. 2D. This enables each cell i on the NCA to self-regularize its continuous vector valued states x_i^t based on perceptions of its local neighborhood $\mathcal{N}(i)$ on the grid. Importantly, in vanilla NCAs, all cells use the same ANN architecture and parameters for their updates and are thus symmetric in functionality, which nonetheless can give rise to asymmetric – and potentially incredibly rich – cell-state dynamics. For differentiable update functions (i.e., ANNs) f_θ , the entire system can be optimized end-to-end via gradient-based learning (e.g., backpropagation through time), allowing emergent self-organization to be shaped by a global loss function rather than being subject to predefined logic. In this way, NCAs unify the local self-regulatory dynamics of CAs with the optimization principles of deep learning, enabling adaptive, data-driven discovery of task-specific cellular interaction protocols. In their most elemental form, NCAs implement convolutional neural networks (CNNs) with a single 3×3 convolutional layer followed by a dense feed forward layer that outputs a proposed cell-state update $\Delta x_i^t = f_\theta\left(x_i^t, \left\{x_j^t\right\}_{j \in \mathcal{N}(i)}\right)$; the parameters θ thus describe synaptic weights and neuronal biases, typically amounting to a total number of $\sim 10,000$ [1] but depending on the application can range from $\approx 10^2$ [7, 75] to $\approx 10^6$ [76]. Formally, we can express such a state update as $x_i^{t+1} = x_i^t + \Delta x_i^t$.

In their Growing NCA paper, Mordvintsev et al. demonstrated how to efficiently optimize the ANN parameters θ so the NCA will reliably self-assemble a 2D target image (e.g., of a lizard or smiley-face emojis) starting from a single seed cell on the grid [1] (see Fig. 1). This was realized by assigning the three items in each cell state vector as cell-specific RGB channels, $y_i^t \in \mathbb{R}^3$, which can be compared with the RGB values of the corresponding pixels of the target image, \hat{y}_i (the full state x_i^t can be higher-dimensional, e.g., 16D in [1]). Moreover, to realize real growth dynamics that starts from a single seed cell with (potentially trainable- [7,8]) initial conditions $x_c^{t=0}$ in the center c of the NCA's grid, a designated α (opacity) channel was introduced to distinguish "living" cells – with $\alpha > \alpha_T$ larger than a threshold of typically $\alpha_T = 0.1$ – from the background. In contrast, the states of "dead" cells are initially, and after every step set to zero. Only cells with one or more living neighbors can update their cell state and thus contribute to this *in silico* embryogenesis process. This makes the NCA's dynamics differentiable across temporal state updates $x_i^t \rightarrow x_i^{t+1}$ and allows employing gradient-based parameter optimization methods on global loss functions over emergent patterns $\{y_i^{tp}\}_{i \in N}$ with respect to target patterns $\{\hat{y}_i\}_{i \in N}$. With this, the inverse problem of target pattern formation (of images or textures), but also novel applications such as self-regulated classification or image segmentation can be solved by means of differentiable collective intelligence.

To improve convergence on spatial self-assembly tasks, the NCAs cell states are often augmented with additional perception layers $\mathcal{P}\left(x_i^t, \left\{x_j^t\right\}_{j \in \mathcal{N}(i)}\right)$ such as Sobel-like filters that extend cell state information with, e.g., local state gradients $x_i^t \rightarrow x_i^t \cup \mathcal{P}\left(x_i^t, \left\{x_j^t\right\}_{j \in \mathcal{N}(i)}\right)$. To stabilize long-term behavior of the NCA's self-regulatory behavior, especially in morphogenesis tasks, it turned out effective to deliberately corrupt the NCA's cell states, e.g. by erasing (masking) parts of the pattern, and train the system to not only learn how to grow but also how to repair a target morphology. Such NCAs exhibit remarkable regeneration capabilities akin to axolots or planaria [1].

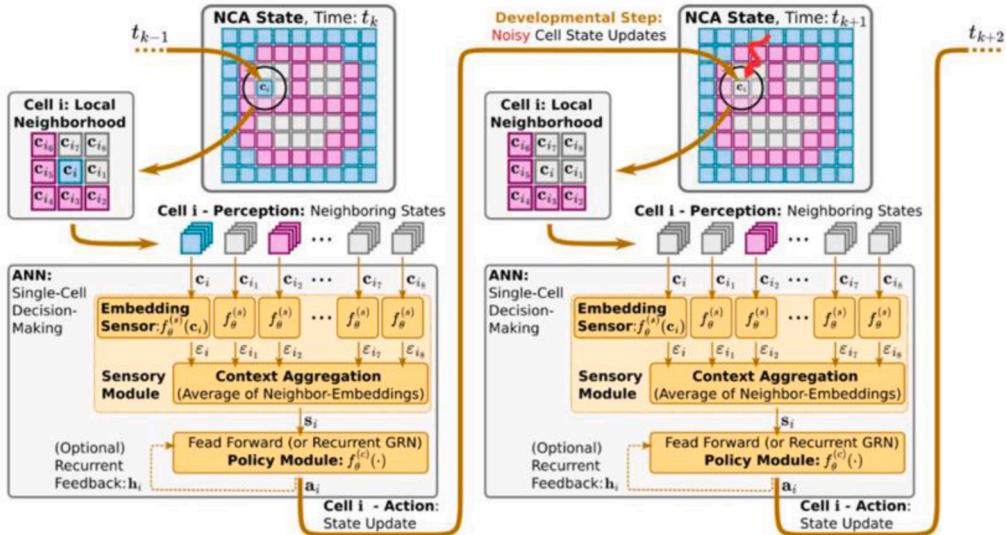
In non-differentiable situations or reinforcement learning (RL) applications, evolutionary algorithms or neuroevolution techniques have proven efficient to optimize NCA parameters [7,9,20,21]. To increase robustness, stochastic cell updates have been introduced – via an update probability – that constrains the cells' reliability to regulate their own state by discarding proposed updates $\Delta x_i^t \rightarrow 0$ across stochastically chosen cells and time steps [1]. This explicitly implements an asynchronous update process across the NCA's grid that requires cells to generalize across cellular neighborhoods by learning how to distinguish between signal and noise. This avoids overfitting cellular policies on local cell-state dynamics and, in turn, promotes global patterns as collective attractor states enabling robust morphogenesis and self-repair. Notably, this renders NCAs as closely related to denoising diffusion models [77–81], as we and others discussed in detail elsewhere [6,7,76,82] and below. Hartl et al. [7] improved both robustness and evolvability by utilizing noise in the update function $\Delta x_i^t \rightarrow \Delta x_i^t + \xi$; choosing or evolving the right combination of noise and cellular reliability (via the update probability) have proven effective to enhance learning and generalization capabilities of NCAs (see Fig. 3A).

Although the NCA's grid states serve as recurrent feedback signal in the unfolding dynamics – that gradually refines high-entropy initial conditions into structured samples conforming to the target data manifold – the majority of used ANNs are, thus far, feed-forward in nature. However, several contributions demonstrated that different ANN architectures, such as attention mechanisms [56], recurrent internal feedback layers [7], explicit separation of private and public cell states [8] (see Fig. 3B), and even hierarchical NCA architectures with interconnected stacked layers [10] (see Fig. 3C), to name but a few examples, can not only be efficiently trained but scale the applicability of NCAs towards truly bio-inspired cybernetic systems, as we will discuss below.

3. NCAs for AI-oriented computational biology

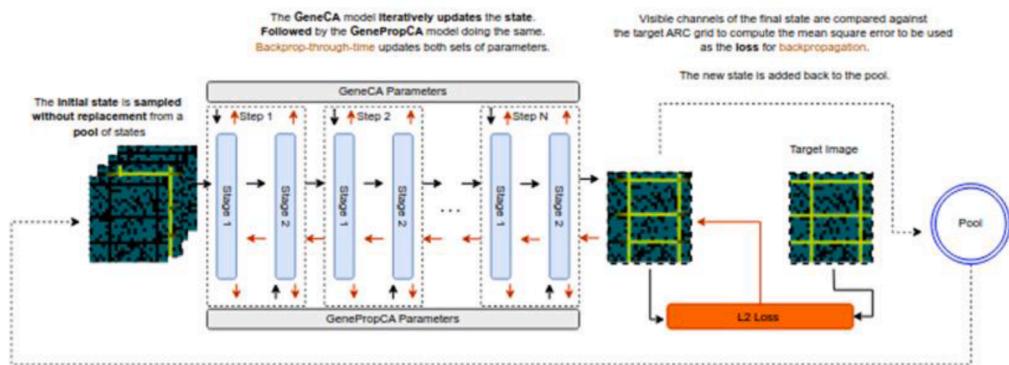
Modeling development and evolution with NCAs

Extending CAs with trainable ANNs, NCAs have been growingly applied in the simulation of morphogenesis in the last few years [1,

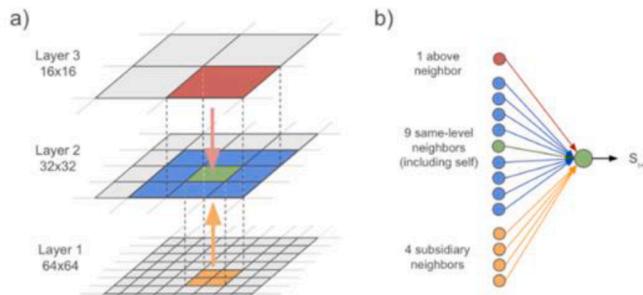


A) Bio-inspired Recurrent Gene-Regulatory (RGRN)-like NCA:

Perception via neighbor-specific embedding and permutation invariant context aggregation, with down-stream (hidden) RGRN policy-module; closely related to GNNs.



B) EngramNCA: GeneCA and GenePropCA are applied alternately on public and private cell states of the same CA grid, to grow ARC solutions conditionally to task-specific inputs.



C) Hierarchical NCA: Several Layers of different resolution integrate into hierarchy of NCAs.

Fig. 3. Alternative Neural Cellular Automaton (NCA) Architectures. A) NCAs composed of cell-specific permutation invariant recurrent graph neural networks [7]: the Artificial Neural Network (ANN) of the NCA's cells are composed of an embedding module that filters and aggregates sensory input in a permutation invariant way into a context embedding, which is then further processed by a policy model – with potentially hidden recurrent states – that proposes cell state update actions in noisy environments. B) EngramNCA architecture with twofold update steps for public cell

states via « GeneCA » and transfer of private genetic information via « GenePropCA » [8,24]. C) Hierarchical NCAs (HNCAs), comprising multiple integrated (stacked) layer of NCAs of increasingly coarse resolution at higher level in hierarchy. Cells integrate cell state information not only of their lateral neighbors, but of neighbors in adjacent layers as well. While lower level cells perform the most detailed dynamics, higher level NCAs progressively condition the collective behavior of adjacent lower levels and thereby can form increasingly more abstract representations and decisions. In that way, HNCAs develop hierarchical forms of control in multiscale architectures [10,11].

[7,9–11,22,57]. Very recently, the ENIGMA model enhanced the biological fidelity by replacing ANN with gene regulatory networks (GRNs) in each cell and incorporating cell-cell signaling pathways to mimic real biological signaling [83]. Evolutionary algorithms have been employed to ensure that ENIGMA can generate stable tissue patterns under varying conditions and thus draw a connection between computational models and biological systems. Similarly, [9] argued that biological systems can be understood as a hierarchy of homeostatic loops which range from cellular metabolic homeostasis to tissue and anatomical homeostasis. This concept is significant in drug design as in this view, molecules have to be targeted at higher levels than specific cellular processes or genetic targets, they also have to fit into the overall homeostatic networks that are trying to restore the setpoints of the body, like morphoceuticals [84]. NCAs have also been used to investigate the evolutionary implications of systems that can self-organize in a multi-competency architecture and found that the evolutionary process is greatly influenced by the cellular competency at several scales [7].

In these studies, NCAs have been employed to problems in morphogenesis, development and evolution, including the French flag problem, which is a model problem in the context of positional information in developmental biology. These simulations also demonstrated that simple competencies in cells can be scaled up to result in tissue-level organization and that there is robustness and error minimization in the formation of the anatomy [7,9,83]. Moreover, these systems demonstrate enhanced robustness against adversarial noise, generalizing well to changing system parameters, and exhibit enhanced transferability towards modified target patterning objectives [7,22].

Very recently, EngramNCA has been introduced, a specialized NCA architecture that incorporates private, intracellular memory channels alongside publicly visible cell states, enabling decentralized memory storage and transfer while facilitating the emergence of hierarchical and coexisting morphologies [8]. Inspired by biological evidence of intracellular memory mechanisms beyond synaptic changes, EngramNCA allows cells to propagate "genetic" information through hidden channels, leading to self-organized growth of diverse structures from a common seed (see Fig. 3B). This model demonstrates robust self-organization in morphogenesis, where simple local rules scale to produce complex tissue-like patterns with enhanced adaptability and regeneration capabilities, even in the presence of perturbations [8].

Modeling regeneration and aging with NCAs

One of the most promising applications of NCAs is in regenerative medicine. The paper "Growing Neural Cellular Automata" [1] has showcased the regenerative potential of NCAs by training such models to regrow patterns after simulated damage, such as cutting or removing sections of the target structure. By incorporating perturbations during the training process — including stochastic updates and pool-based sampling — these NCAs learn to persist and self-repair complex patterns like lizard or emoji shapes, mimicking biological regeneration where local rules enable recovery from unseen injuries [1]. Extending this to practical applications in soft robotics, researchers have developed NCAs that allow simulated robots to regenerate their morphology solely through local cell interactions, evolving separate automata for initial growth and post-damage repair [64,65]. In experiments with 2D and 3D voxel-based robots, damaged morphologies (e.g., bipeds or tripods) achieved up to 99 % structural similarity to originals, with varying locomotion recovery (67 % in some cases), highlighting implications for resilient, self-organizing systems in dynamic environments [64,65]. This regenerative framework has also been scaled to 3D artefacts, where NCAs generate and repair functional machines in virtual worlds like Minecraft, starting from a single cell. They demonstrate enhanced recovery (e.g., 99 % block regeneration after halving a structure) when explicitly trained for damage resilience, paving the way for self-repairing physical structures inspired by biological morphogenesis [22]. More recent advancements integrate NCAs into diffusion models for parameter-efficient image synthesis, enabling efficient inpainting tasks where missing image regions are regenerated with high fidelity leveraging local communication to maintain contextual consistency and drawing analogies to biological tissue repair in fields like pathology imaging [76]. These examples underscore NCAs' versatility in modeling self-organized regeneration across scales, from cellular patterns to macroscopic structures.

NCA simulations have shown the important role of bioelectric patterns in guiding tissue development like in planarians and other organisms [85]. For example, NCAs were used to model disrupted craniofacial structures in tadpoles that subsequently reorganize and achieve correct anatomical configurations following a specific bioelectrical pattern, and provide a framework for therapeutic interventions [70].

NCAs have further been applied to aging, conceptualized as a decline in tissue morphology [14]. Evolutionary simulations showed that aging results from loss of goal directedness in morphospace at tissue level, leading to breakdown of collective anatomical homeostasis and morphological decline, even in the absence of genetic damage or accumulation of noise [14]. In addition, this study also revealed that there is dormant regenerative potential in aged tissues, which can be reactivated through targeted interventions. This finding reveals ways to develop strategies to rejuvenate tissues and extend healthy lifespan [14,86].

Overall, the use of NCAs for modeling biology extends beyond morphogenesis. They can be used to understand aging and regeneration processes, and they provide insights into biological information processing at multiple levels—including genetic and bioelectrical. This modelling approach generates new hypotheses, sometimes experimentally validated [9,70].

NCAs as a model for the generative genome

Evolution through developmental reproduction involves encoding the features and functionality necessary for high-fidelity reconstruction of an organism into the compact form of the genome; rather than directly encoding a “blueprint” or developmental “program”, the genome more accurately represents a bowtie of latent variables instantiating organismal development. This process unfolds through gene regulatory mechanisms at the cellular level, informed by environmental cues and top-down feedback from higher organizational scales, establishing a multiscale framework [26,30,87]. Importantly, this process is neither deterministic nor simply emergent but can be related to a creative problem-solving layer between the genome and the phenotype of an organism that can deal with novel situations and environmental conditions on the fly. Such an architecture represents a form of collective intelligence from the bottom up and an embodiment of collective intelligence following William Jame’s definition of intelligence “reaching the same goals by different means” [26,27,36].

The collective architecture of NCAs aligns closely with the nested bowtie structure of developmental biology, encompassing distributed gene regulatory mechanisms to emergent tissue or organ patterns. Each NCA cell maintains a numerical state akin to the physiological status of biological cells (e.g., transcriptomic expressions), while internal ANNs regulate localized cell states according to a collective system-level agenda (modelling biological cells’ GRNs). Through that lens, the unfolding dynamics of NCAs – intercellular communication and intracellular information processing – closely models the distributed self-regulatory process of morphogenesis and morphostasis of distributed cellular agents interpreting and reacting to their local environment under novel morphogenetic contexts. The ANN of NCAs thus embody the generative genomic bottleneck for evolutionary development, not only positioning the genome as functional and generative unit, but empowering NCAs to serve as a distributed model for adaptively and creatively reconfiguring cellular expressions during the course of collective organismal development [6,88].

Thus, NCAs might be excellent candidates for bridging the computationally irreducible gap between the agential layer of physiological computations connecting an organism’s genotype and phenotype and *in silico* modeling and simulation via self-regulatory generative AI. While significant questions remain, NCAs may reflect real-life developmental pathways that can be intervened with in simulation, if trained on authentic biological data (as demonstrated through targeted adversarial reprogramming or takeover of tissue [15,16], see Fig. 1D). NCA research thus offers a transformative character in regenerative medicine, bioengineering, and synthetic biology [6].

NCAs for medical applications: molecular design and interaction modeling

NCAs have also been employed in drug discovery and protein docking. By representing small molecules and proteins as voxelized structures, NCAs have been used to design small molecule interactors, reconstruct protein structures, and model complex biomolecular interactions. These models are not only capable of simulating molecular docking and physical transformations but also of goal-directed molecule design by conditioning generation on specific protein targets. Moreover, NCAs may also be useful in reconstructing missing protein regions and modeling dynamical processes like isomerization, thus increasing their relevance to drug design and structural biology. However, the versatility of these NCA models renders them somewhat inferior to the current state-of-the-art methods (such as diffusion- or transformer-based foundational models [89,90]), yet they exhibit promising potential in translating conceptual advances to practical applications in precision medicine and biomolecular engineering [91].

4. NCAs as a model of choice for biology, and beyond?

NCAs have an important characteristic that make them very suitable to address deep biological questions: their multiscale competency architecture. Indeed, biological systems show cognitive capabilities at different levels of biological complexity, from molecular networks to cells, tissue, organs until the whole organism and beyond (see Fig. 1-A and B). This architecture enables distributed problem-solving where individual components possess limited competencies that collectively generate sophisticated system-level behaviors through dynamic interactions and information integration [26,27]. In biology, we can find these competencies even below the cell level, GRNs have shown different types of learning and memories [92], and there is a large body of studies showing that we can find cognitive capabilities like memory, decision-making, anticipation, and learning at all levels of biological complexity [26,27]. At tissue-level, the collective behavior gives rise to wound healing, and morphogenesis that exceeds the capabilities of the individual agents [31]. At higher levels of organization, organs and the entire organisms coordinate across scales to achieve complex anatomical homeostasis and adaptive responses to the environment [93]. Following the field of Basal Cognition [94–97], evolution likely pivoted the same strategies used by cell collectives to navigate the space of anatomical possibilities during morphogenesis into brain-based navigation of 3D space, leading to the emergence of higher intelligence [26,27,98]. Morphogenesis is an example of cellular collectives navigating the morphospace to resolve a goal-directed, system-level control task, i.e., forming and maintaining target anatomical states via bioelectric, biochemical, and biomechanical signaling and homeostasis—providing a ready substrate for scaling into behavioral intelligence (that can be seen as navigation in 3D space [29]). They navigate the morphospace and rely on the same molecular hardware that is now used by neurons (ion channels, gap junctions, neurotransmitters), enabling active information integration, error correction, and memory long before brains actually evolved [26,27,94–99]; several mechanisms that are typically attributed to cognitive (neuronal) processes have long been employed by (sub)cellular collectives: transcriptional and physiological networks maintain multiple forms of memory [92,100], cell and cellular collectives use active inference for pattern regulation [101], and bioelectric circuits exhibit perceptual bi-stability [102]; see also [103,104] for other examples. The TAME framework formalizes this evolutionary pivot via the notion of the scaling of goals from cellular to anatomical homeostasis that generalizes to behavior in

physical 3D space [27]. By sharing information and stress, sub-agents work together to solve higher-level problems [9,99].

Current AIs like large language models (LLMs), transformers etc. don't really leverage this modular, scale-bridging integration of established competencies. Each formal neuron in a deep learning neural network, is treated only as a filter, firing with a predefined signal after a certain activation threshold is reached. These neuronal models are not capable of learning or forming a memory themselves; while learning and information processing dynamics only occur at the scale of the whole network. This might stand in stark contrast to biological neurons, especially to cortical columns in the human neocortex. The neocortex is formed by an integrated 2D grid of copies of the same neuronal circuit, the cortical columns, which have been argued to be capable of learning arbitrary concepts of our reality (objects, animals, other human beings, mental constructs, etc.). More importantly, these cortical columns might literally "model" the learned concepts and thus represent interactable reference frames or world models for relevant features of our Umwelt [32,33,105,106]. Remembering would thus trigger models of past experiences, and thinking would translate not only into a navigation process through an associative conceptual space, but dynamically construct a network of interacting world models that are relevant under a certain context [6,32,33,107]. NCAs might be excellent candidates to model such an architecture: they maintain a trainable ANN in each cell of an integrated grid which are in principle capable of representing reference frames of arbitrary concepts [108]. In turn, NCAs might not only be excellent models for biological self-organization as a multiscale competency architecture, but even for higher-level cognitive processes of the human neocortex such as active perception [71,109], raising fascinating questions about the parallels of morphogenesis and cognition [26,27,104].

Currently, the focus for scaling intelligence is based on the scaling of training data, model parameters, and computational resources [110]. This approach is incredibly resource intensive, and so-far mimics but a fraction of the capabilities of what biology achieves seemingly effortlessly. Several prominent AI researchers criticized the costs associated with exploiting scaling laws for commercializing LLMs and ask for the development of new AI architectures beyond transformers. While LeCun promotes a world-model based approach [111], NCA architectures and their self-regulatory dynamics during inference represent a potential solution to this architectural gap by implementing biologically inspired scale-bridging learning dynamics. Indeed, biological evolution modularly repurposed minimal collective goal-directed behavior (homeostatic loops) into ever higher-level problem-solving agents (composed of competent sub-agents) through competency amplification and an internal motivation for novelty, rather than data accumulation [9, 27]. By design, NCA architectures implement the scaling of intelligence from local cellular communication pathways to system-level behavior and, potentially, to integrated world models.

Recently, ARC-NCA has demonstrated impressive performance on the abstract reasoning corpus (ARC) challenge, a collection of carefully designed tasks to benchmark modern AI frameworks for "general intelligence" and abstraction and reasoning capabilities [4, 5]. Thus far, ARC remains an extremely challenging fundamental test for artificial general intelligence (AGI) both for explicitly crafted models as well as the most competitive large reasoning models, such as GPTs [3]. Interestingly, vanilla NCAs and EngramNCAs perform competitively on the ARC challenge compared to ChatGPT-4.5 – notably at the fraction of costs and computational resources – by leveraging developmental processes of simple interactive cellular agents that collectively self-assemble, *i.e.*, "grow", the complex patterns of the ARC-AGI-1 benchmark [24,25]. In a sense, ARC-NCA models represent an embodiment of the problem-solving capacities necessary for the ARC challenge, reflecting the immense computational power that roots in embodied basal cognition [27,34]. In that way, ARC-NCA can learn to generalize to grow the target-pattern modalities of many ARC-AGI-1 examples from a minimal set of two or three training examples, something standard AI and ML techniques struggle to achieve.

This raises fundamental questions about the organizing principles underlying nested biological systems: are composite agential objects embodied models of environmental features relevant at their scale [112]? How does this hierarchical representational structure emerge across spatial and temporal scales – from molecular machines to social collectives [9,26]? Viewed through this lens, biological agents at every level function not merely as components or effectors, but as embodied representations of their respective environment, while pursuing developmental goals and systemic integrity [26,27,113]. For example, a bistable enzyme may be an embodied model of ligand-specific decision-making; its conformational switch encodes a dynamic inference about molecular context. The transcriptome functions may be seen as an embodied intracellular model of cell-fate specification within the morphogenetic field. Bioelectric fields – coordinated by collective cellular behavior – serve as coarse-grained, spatiotemporally models guiding embodying cellular collectives toward large-scale pattern formation of target system-level outcomes [114]. Even animals can be understood as living world models [111,115]: while their bodies are evolutionary fitted for the challenges of their Umwelt, their brains do not simply simulate the environment but embody its structure through continuous interaction. Neural circuits are not merely predictive tools, they function as dynamic, embodied representations of environmental regularities, integrating perception, action, and anticipation in real time [30,32, 33,107,116]; this might even scale to swarms or groups as distributed models of environmental stability and multi-agential coordination [26,28–30,36,117–122].

Crucially, these modeling aspects are not invented anew at every scale, but are regularities themselves – and if this process applies at every scale, this necessarily includes repurposing lower-level organizational units as building blocks to form higher-level embodied models [117,123]. Evolution is tinkering, *i.e.*, life does not design from scratch, it reuses, refactors, and repurposes across scales [26, 124]. Biological systems are not just composites of passive parts but are layers within layers of repurposable agential substrates – each unit simultaneously acting as an agent, a model, but also as signal. Proteins become signals in GRNs; the GRN state informs cell identity; cellular states coalesce into tissue-scale morphogenetic patterns, etc., all structured by self-modeling dynamics [6,9,26,28,29, 92,100,125]. Thus, biological organization might fundamentally be rooted in self-modelling of scale-dependent (relevant) environmental contexts through repurposable embodiments – scale-bridging representations that allow nested agents to maintain their physiological integrity against diffusive, entropic processes; this has been formalized on thermodynamic grounds via the variational free energy principle, or active inference [116]. This suggests that biology might be understood as models-within-models rather than merely layers-within-layers of organization – a multiscale competency architecture where agency and embodied representation are

inseparable. NCAs provide a computational framework that mirrors this hierarchical representational architecture: flexible local rules allow to generate global structure through embodied learning and coordination, offering a tractable way to explore how scale-bridging models emerge from simple, repurposable units and providing insight into the dynamics of nested biological self-modeling. Although there are still limitations to the analogy of NCAs and biological systems (see Section 5), NCAs offer a promising architecture to enable reuse of acquired functionality across scales, for instance, through curriculum learning [7], hierarchical or dynamical integration [10, 11], and modular conditioning [8]. Moreover, especially in composite robotic scenarios, NCAs represent an intriguing testbed for studying hierarchical modeling and embodied coordination. A deeper understanding of such scale- and substrate-bridging mechanisms of intelligence could even position NCAs as a good candidate to develop AI-generating algorithms, e.g., distributed learning systems that, given enough time, self-organize modularly toward intelligent behavior without external intervention [126].

Recent developments by Béna et al. investigate the potential of training NCAs with gradient-based methods to perform universal computation in continuous state spaces. This was achieved by training NCAs with task-specific objectives to guide the emergence of fundamental computational primitives, i.e., logic gates which can be integrated by the NCA's dynamics into universal NCAs [127]. Relatedly, Miotti et al. demonstrated that differentiable logic gates can be used to model update functions of single cells in discrete CAs; these Differentiable Logic CAs can be efficiently trained end-to-end on system-level objectives for CAs with a discrete state space DiffLogicCA [128]. These ideas represent first promising steps towards self-assembling universal computation that is based on collective and distributed intelligence, and akin to the operational properties of biological matter.

Future studies will show whether the dynamics in computational medium of NCAs will undergo a similar transition from multi-cellular communication pathways observed in basal tissue to scale-free activation patterns observed in neuronal networks. This transition might involve explicit hierarchical architectures [10,11] (see Fig. 3C), or graph-like organization comprising short and long-range connections as opposed to the fixed grid layout of current NCAs, and might bring about novel understandings of the multiplicity of the computational power of different type of biological matter. For instance, recent promising studies argue that even vanilla NCAs that are evolved to exhibit critical-like dynamics – validated by matching power-law avalanche statistics in their spatio-temporal cell state activities – and NCAs pretrained toward a near-critical regime have been reported to provide more efficient initial conditions for learning down-stream tasks [109,129]; related analysis of stochastic CAs supports the use of power-law diagnostics to assess criticality [130]. However, identifying robust quantitative measures of criticality in NCAs (e.g., correlation length, susceptibility) and developing training schemes that explicitly maintain or regulate (or exploit) criticality remain important open challenges. Criticality in the cortical brain networks is hypothesized to maximize computational capacity by sustaining long-range correlations, power-law distributed neuronal avalanches, and efficient signal propagation across scales. Such a near critical state offers the exact mixture of structured signaling and operational flexibility necessary for optimal learning and information processing, balancing sensitivity to input with stability in network dynamics [131–135]. It has been argued that criticality is a universal setpoint for brain functions [136]. However, cognition is a spectrum and does not seem to be restricted to neuronal dynamics, the latter are only speed optimized for agents navigating our 3D Umwelt, but the governing bioelectric dynamics are much more ancient [26–28,31,85, 113]: understanding biology as a multiscale competency architecture increasingly blurs the line between collective biological coordination and cognitive emergence. Criticality may not be merely a dynamical regime but a functional scaffold enabling self-regulatory agency across scales (from ion channels to tissues to behavior).

On this view, the brain's critical state is not the pinnacle of evolution but a fundamental tunable threshold allowing decentralized systems – whether neural or non-neural – to maintain homeostatic plasticity while preserving structural memory. Not only minds, but also biological forms, and, in general, collaborative collectives are solutions to information problems, and critical dynamics provide the optimal "sweet spot" for exploration vs. stability in both cognitive inference and developmental patterning [28]. Emulating such dynamics with NCAs may not only improve learning efficiency but also reveal fundamental principles linking biological organization to computational resilience – offering a bridge between the physics of collective phenomena and the emergence of cognition. Through this lens, NCAs are not just tools, but might be models of life's underlying (near-critical) computational logic [6]. All together, these approaches represent technologically feasible ways to tackle the emerging research direction of Diverse Intelligence [35] and the TAME framework [27], and new forms of AIs based on collective intelligence.

5. Challenges and limitations

The ANNs inside each cell in NCAs are definitely less complex than large fully connected neural networks. Still, balancing model complexity (of the single-cellular agents) and interpretability of the emerging dynamics is far from trivial. It is unclear how to interpret and best analyze the integrated dynamics of such systems that undergo distributed computation. NCAs, thus far, represent toy-models for biological organization and for self-regulated complex systems in general, and may therefore serve as transparent model system – their complexity and capabilities can be tuned by design – to study corresponding network dynamics, if not the dynamics of the evolution of such network dynamics.

NCAs can still be difficult to train, as solutions can collapse into trivial attractors (especially with gradient-based methods via pool-poisoning) or get trapped in suboptimal solutions that lack feature precision (predominantly in the case of neuroevolution training). Interestingly, cues from non-equilibrium physics, i.e., from criticality and phase transitions [71,109,129,130] suggest that NCA that are pretrained to operate at a critical state might be good "initial states" to learn further down-stream tasks. Thus, regulating NCA dynamics to near-critical behavior might advance their capabilities for adaptability. Moreover, training NCAs to grow multiple target patterns, and thus store multiple system-level attractor states in the cellular ANNs, proved challenging. Some approaches tackle this shortcoming by co-evolving a multitude of initial conditions alongside shared ANN parameters per cell [7]. EngramNCA circumvents this by introducing a two-fold developmental process of a private and public cell state that operates through different communication

pathways [8]. Yet, extending the pattern-storage and cell-resolution capabilities of NCAs to compete with modern generative AI techniques, not to speak of biology, is an active field of research [72,76].

This shortcoming is closely related to the fact that current NCAs don't interface well: different trainings may lead to the same system-level outcome but implemented by completely different microlevel dynamics. This limits modularity and compatibility across multiple NCAs that operate in the same environment but developed not necessarily compatible communication strategies, and thus limits the analogy to biology which all talks the same language of biochemistry. Neuroevolution techniques offer some insights into transferability and evolvability experiments of NCAs [7,10,11] and EngramNCA demonstrates that even more biologically inspired architectures – based on alternating public state updates and private gene transfer – can interface morphologies belonging to different seed “genes” [8]. Yet, the learned update rules are still largely opaque black-boxes and difficult to interpret at the parameter and interaction level; update dynamics don't follow physical or biologically relevant ODEs; it is unclear how to implement reasonable multiscale coupling principles in Hierarchical NCAs; in robotics, real-world experimental realization are in their infancies; and improving our theoretical understanding in stability, conditional dynamics, and universality requires significant effort.

Increasing the capabilities of NCAs might even call for hybrid training paradigms, combining differentiable learning with evolutionary strategies. Differentiable learning in NCAs, i.e., via backpropagation, can be highly efficient to learn task-specific dynamics (such as target pattern formation) with high precision [1,82]. Evolutionary algorithms, on the other hand, not only offer a much greater diversity of potential solutions but bring forth NCAs that are naturally highly robust against noise, corruptions, and show enhanced evolvability and transferability to novel problems [7,9–11,18,19,57,58]. Moreover, different ANN architectures – feed forward, circuit based, attention, recurrent, world models, etc. – might also strongly influence what the NCA can learn. Overall, relations of NCAs with reinforcement learning, backpropagation, active inference, or neuroevolution techniques still remain underexplored. A promising approach to not only combine the advantages of different learning paradigms (gradient-based and evolution) might be the recent conditional diffusion-evolution approach. This approach offers unprecedented quality-diversity scores, and through conditional sampling even allows to bias different evolutionary lineages towards a multitude of target outcomes [88,137].

As already discussed above, the integrated multiscale nature of biological organization poses significant challenges from a modelling and computational perspective. Hierarchical NCAs [10,11] try to overcome this obstacle by imposing a fixed architecture of hierarchically stacked NCA layers of varying resolution (lower levels implementing the dynamics, while successive higher levels exerting control over lower levels). It remains, however, unclear, how to best decouple different scales of organization, or how to identify sub-groups of cells at certain levels that group together under different contexts. They suffer from limitations in storage capacity (of multiple system-level target outcomes, or functional modalities) and in scaling; simulating realistic organ-level complexity at unicellular resolutions is currently infeasible. NCAs could learn from hierarchical reasoning approaches [138], from other multiagent-based RL techniques such as [139–142] or the recently proposed TAG framework for multi-agent hierarchical RL [143]. Another promising architecture to improve NCAs capabilities are self-supervised techniques such as Joint Embedding Predictive Architectures (JEPAs) which are closely related to nested world model architectures [111]. In addition, it has been shown that a latent NCA that shifts the computation from the conventional input space to a specially designed latent space – obtained via an autoencoder – can reduce computational requirements while maintaining high reconstruction fidelity in image restoration tasks [144]. Along these lines, a JEPAs architecture integrated as an (hierarchical) NCA may enable higher cognitive capabilities by multiscale latent predictions.

Moreover, their descriptive power, especially of biological systems, is still largely over-simplified: NCAs treat cells largely as homogeneous units and their numerical states and interactions are at best abstractions (or model composites) of most biological processes. And while progress has been made, it is difficult to integrate (and identify) NCA dynamics with molecular pathways, gene regulatory networks, or biomechanical forces. Training NCAs is mostly concerned with the accuracy of the final system-level outcomes, largely neglecting energetic, metabolic, or other physical/chemical/biological constraints necessary for developing viable morphologies. Furthermore, it remains unclear how to map the synthetic cell states of NCAs to their biological counter parts, i.e., to physiological properties of cells such as transcriptomic, proteomic, or other omics data. Notably, this is a critical missing link limiting the predictive power of NCAs for biology, which, if resolved, would render NCAs as excellent first-principle model biological dynamics. Although first efforts are made in this direction for the morphogenesis of cellular clusters [145] and the scaling of biological goals via homeostatic principles [9], the problem remains unresolved so far, especially for complex, scale-bridging organisms.

Relatedly, it is an open question how, and at what scales embodiment manifests in NCA frameworks, especially when modelling biological systems. If the multiscale competency architecture of biology is taken seriously, modeling through embodiment occurs at every scale; and biology found a way to (i) generalize learning through evolving problem-solving agents, that (ii) are repurposed across scales [6,26,28,117,118,123,124]. A key difference to biological systems is, that NCAs use logical units as signal carriers, while biological systems use physical (or biological) signals; lower-level agents can become signals for higher levels, blurring the distinction between software and hardware – and thus between computation and embodiment – in biology [27,34,113]. Still, NCAs pose a unique testbed to understand how modeling and control arise across levels of organization, from intra- and intercellular signaling to system-level morphologies and behavior. Clarifying what constitutes embodiment within such computational frameworks will be essential for advancing this line of research.

Notably, NCAs exhibit self-regulatory denoising behavior reminiscent of diffusion models [77–81], cutting edge foundational models of modern generative AI (e.g., Stable Diffusion [146] for image synthesis or SORA [147] for video synthesis, to name but a few). Both frameworks iteratively transform high-entropy, noisy initial states x_t into structured data x_0 by reversing a diffusive process. While diffusion models are explicitly trained to predict and remove noise ϵ_t from corrupted data x_t conditioned on a time step t – which encodes the scale of corruption $x_t = x_0 + \epsilon_t$ – they leverage this temporal parameter to orchestrate hierarchical refinement: coarse features emerge early, followed by fine-grained details through successive symmetry-breaking transitions in the data manifold

[148–150]. This enables structured, multiscale learning and reveals deep insights into the hierarchical organization of features and symmetries in the (generative process of) data. In contrast, NCAs lack explicit temporal feedback and supervision; they do not receive time as input. Instead, the dynamics – if at all – implicitly encode a developmental time through collective recurrent state updates, effectively internalizing optimal developmental trajectories of the target system-level outcome via embodied learning. This raises a key distinction: while diffusion models condition on time to guide denoising, NCAs must learn this temporal structure end-to-end from spatiotemporal patterns alone, a significant challenge in capturing hierarchical organization through collective interactions. In turn, future NCA architectures could benefit from hybrid designs, enabling them to emulate diffusion-like hierarchical denoising while preserving their strengths in distributed, spatially aware computation. Such models may not only improve generative quality but also reveal how biological systems achieve complex morphogenesis through internalized temporal inference and embodied cognition via self-modeling across scales.

6. Conclusion

In the context of AI-oriented computational biology, we see NCAs as a model of choice. They integrate the biological multiscale competency architecture, and their inherently distributed nature allows to model self-organization and spatiotemporal information-processing systems for a wide range of biological phenomena including higher cognition [6,8,9,14,24,25,57,63,129]. By using NCAs to simulate biological processes like morphogenesis or tissue regeneration, we can gain deeper insights into the underlying principles of these complex biological systems at different levels of organization and how cognition scales following the principles of basal cognition [26,27].

This relationship creates a positive feedback loop. By using NCAs to understand information-processing in cellular collective systems, we can in turn apply these biological principles to design new AI architectures and algorithms. In addition, developing such bioinspired AIs may be beneficial for the problem of alignment. Indeed, we can build AIs that are designed to be more adaptable, resilient, collaborative, and integrate modularly, leveraging the distributed capabilities of NCAs. And in addition, we can use all the toolbox of neuroscience and psychology to understand and interact with them as they will share the same multiscale competency architecture we find in biological systems. This symbiotic relationship —using NCAs to understand biology, and using that understanding to create truly bioinspired AIs (and possibly more aligned ones)— and despite all the challenges and limitations mentioned above, offers an optimistic new paradigm for computational biology and the future of AI and computation.

Declaration of competing interest

The Levin lab has a sponsored research agreement with Astonishing Labs, a company seeking to impact the longevity and biomedicine of aging field.

Acknowledgements

We thank Tomika Gotch for assistance with the manuscript. M.L. gratefully acknowledge support of Astonishing Labs via a sponsored research agreement to Tufts University, Schmidt Sciences, LLC, and of Templeton World Charity Foundation (via grant TWCF0606).

References

- [1] Mordvintsev A, et al. Growing neural cellular automata. *Distill* 2020;5(2):e23.
- [2] Randazzo E, Mordvintsev A, Fouts C. Growing steerable neural cellular automata. In: *ALIFE2023: Ghost in the Machine: Proceedings of the 2023 Artificial Life Conference*. MIT Press; 2023. p. 2.
- [3] ARC Team, *ARC-AGI-1 Leaderboard* 2025. <https://arcprize.org>.
- [4] Chollet F, et al. Arc prize 2024 2024. *Technical report*. arXiv preprint arXiv:2412.04604.
- [5] Chollet F. On the measure of intelligence 2019. arXiv preprint arXiv:1911.01547.
- [6] Hartl B, Levin M. What does evolution make? Learning in living lineages and machines. *Trends in Genetics* 2025;41(6):480–96.
- [7] Hartl B, Risi S, Levin M. Evolutionary Implications of Self-Assembling Cybernetic Materials with Collective Problem-Solving Intelligence at Multiple Scales. *Entropy (Basel)* 2024;26(7):532.
- [8] Guichard E, et al. EngramNCA: a neural cellular automaton model of memory transfer. In: *ALIFE2025: Ciphers of Life: Proceedings of the Artificial Life Conference 2025*. Kyoto, Japan: MIT Press; 2025. p. 3.
- [9] Pio-Lopez L, et al. The scaling of goals from cellular to anatomical homeostasis: an evolutionary simulation, experiment and analysis. *Interface Focus* 2023;13(3):20220072.
- [10] Bielawski K, et al. Evolving Hierarchical Neural Cellular Automata. In: *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM; 2024. p. 78–86.
- [11] Pande R, Grattarola D. Hierarchical neural cellular automata. In: *ALIFE2023: Proceedings of the 2023 Artificial Life Conference*. MIT Press; 2023. p. 20.
- [12] Endo K, Yasuoka K, [GitHub]. Neural Cellular Maze Solver 2024. Available from: <https://umu1729.github.io/pages-neural-cellular-maze-solver/>.
- [13] Earle S, et al. Pathfinding neural cellular automata 2023. arXiv preprint arXiv:2301.06820.
- [14] Pio-Lopez L, Hartl B, Levin M. Aging as a Loss of Goal-Directedness: An Evolutionary Simulation and Analysis Unifying Regeneration with Anatomical Rejuvenation. *Advanced Science* 2025;e09872.
- [15] Cavuoti L, et al. Adversarial takeover of neural cellular automata. In: *ALIFE2022: Proceedings of the 2022 Artificial Life Conference*. MIT Press; 2022. p. 38.
- [16] Randazzo E, et al. Adversarial reprogramming of neural cellular automata. *Distill* 2021;6(5):e00027-e00004.
- [17] Hartl B, Levin M, Zöttl A. Neuroevolution of decentralized decision-making in N-bead swimmers leads to scalable and robust collective locomotion. *Communications Physics* 2025;8(1):194.
- [18] Horibe K. Severe damage recovery in evolving soft robots through differentiable programming. *Genetic Programming and Evolvable Machines* 2022;23(3):405–26. et al.

[19] Horibe K, Walker K, Risi S. Regenerating Soft Robots Through Neural Cellular Automata. EuroGP 2021. Springer; 2021.

[20] Najarro E, et al. HyperNCA: Growing Developmental Networks with Neural Cellular Automata. ICLR 2022 Workshop: From Cells to Societies: Collective Learning Across Scales. 2022. OpenReview.net.

[21] Pontes-Filho S, et al. A Unified Substrate for Body-Brain Co-evolution. ICLR 2022 Workshop: From Cells to Societies: Collective Learning Across Scales. 2022. OpenReview.net.

[22] Sudhakaran S, et al. Growing 3D Artefacts and Functional Machines with Neural Cellular Automata. ALIFE 2021: The 2021 Conference on Artificial Life. MIT Press; 2021. p. 108.

[23] Varienigen A, et al. Towards self-organized control: Using neural cellular automata to robustly control a cart-pole agent 2021. arXiv preprint arXiv: 2106.15240.

[24] Guichard E, et al. ARC-NCA: Towards Developmental Solutions to the Abstraction and Reasoning Corpus. In: ALIFE2025: Ciphers of Life: Proceedings of the Artificial Life Conference 2025. Kyoto, Japan: MIT Press; 2025. p. 5.

[25] Xu K, Miikkulainen R. Neural Cellular Automata for ARC-AGI. MIT Press; 2025. *Artifical Life*.

[26] Levin M. Darwin's agential materials: evolutionary implications of multiscale competency in developmental biology. *Cell Mol Life Sci* 2023;80(6):142.

[27] Levin M. Technological Approach to Mind Everywhere: An Experimentally-Grounded Framework for Understanding Diverse Bodies and Minds. *Front Syst Neurosci* 2022;16:768201.

[28] Fields C, Levin M. Scale-Free Biology: Integrating Evolutionary and Developmental Thinking. *Bioessays* 2020;42(8):e1900228.

[29] Fields C, Levin M. Competency in Navigating Arbitrary Spaces as an Invariant for Analyzing Cognition in Diverse Embodiments. *Entropy* 2022;24. <https://doi.org/10.3390/e24060819>.

[30] Pezzulo G, Levin M. Top-down models in biology: explanation and control of complex living systems above the molecular level. *J R Soc Interface* 2016;13(124):20160555.

[31] Fields C, Bischof J, Levin M. Morphological Coordination: A Common Ancestral Function Unifying Neural and Non-Neural Signaling 2020;35(1):16–30. *Physiology* (Bethesda).

[32] Hawkins J. A thousand brains, 288. London, England: Basic Books; 2021.

[33] Leadholm N, et al. Thousand-brains systems: Sensorimotor intelligence for rapid, robust learning and inference 2025. arXiv [cs.AI].

[34] Bongard J, Levin M. There's Plenty of Room Right Here: Biological Systems as Evolved, Overloaded, Multi-Scale Machines. *Biomimetics* 2023;8. <https://doi.org/10.3390/biomimetics8010110>.

[35] Levin M. *Artificial Intelligences: A Bridge Toward Diverse Intelligence and Humanity's Future*. Advanced Intelligent Systems 2025;7(10):2401034.

[36] McMullen P, Levin M. Collective intelligence: A unifying concept for integrating biology across scales and substrates. *Commun Biol* 2024;7(1):378.

[37] Langton CG. Artificial life: An overview. MIT press; 1997.

[38] Cook M. *Universality in elementary cellular automata*. Complex systems 2004;15(1):1–40.

[39] Rendell P. Turing machine universality of the game of life. Springer; 2016.

[40] Von Neumann J. Theory of self-reproducing automata. University of Illinois Press; 1966.

[41] Barricelli NA. Numerical testing of evolution theories: part I theoretical introduction and basic tests. *Acta Biotheoretica* 1962;16(1-2):69–98.

[42] Langton CG. Studying artificial life with cellular automata. *Physica D: nonlinear phenomena* 1986;22(1-3):120–49.

[43] Banks ER. Universality in cellular automata. In: 11th Annual Symposium on Switching and Automata Theory (swat 1970). IEEE; 1970. p. 194–215.

[44] Nichele S, et al. CA-NEAT: evolved compositional pattern producing networks for cellular automata morphogenesis and replication. *IEEE Transactions on Cognitive and Developmental Systems* 2017;10(3):687–700.

[45] Wolfram S, Gad-el-Hak M. A new kind of science. *Appl. Mech. Rev.* 2003;56(2):B18–9.

[46] Wolfram S. Cellular automata. *Los Alamos Science* 1983;2–21.

[47] ElZayyat M, Mušić D. *Cellular Automata as a Tool in Artificial Life Research*, in: *ICCRD*. IEEE; 2025. p. 275–9.

[48] Hegselmann R. *Understanding social dynamics: The cellular automata approach*, in: *Social science microsimulation*. Springer; 1995. p. 282–306.

[49] Hoekstra AG, Kroc J, Sloot PMA. Simulating complex systems by cellular automata. Springer Science & Business Media; 2010.

[50] Huang K, et al. Behavior-based cellular automaton model for pedestrian dynamics. *Applied Mathematics and Computation* 2017;292:417–24.

[51] Manzoni S, Vizzari G. Crowd behavior modeling: From cellular automata to multi-agent systems. *Multi-Agent Systems*. CRC Press; 2018. p. 315–38.

[52] Gilpin W. Cellular automata as convolutional neural networks. *Phys Rev E* 2019;100(3-1):032402.

[53] Miller JF. *Evolving a self-repairing, self-regulating, french flag organism*. in: *GECCO*. Springer; 2004.

[54] Randazzo E, et al. Self-classifying mnist digits. *Distill* 2020;5(8):e00027-e00002.

[55] Hernandez A, Vilalta A, Moreno-Noguer F. Neural cellular automata manifold. In: *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*. 2021, Computer Vision Foundation /IEEE; 2021. p. 10020-8.

[56] Tesfaldet M, Nowrouzezahrai D, Pal C. Attention-based neural cellular automata. *Advances in Neural Information Processing Systems* 2022;35:8174–86.

[57] Grasso C, Bongard J. *Empowered neural cellular automata*, in: *GECCO* 2022:108–11.

[58] Grasso C, Bongard J. Selection for short-term empowerment accelerates the evolution of homeostatic neural cellular automata. In: *Proceedings of the Genetic and Evolutionary Computation Conference*; 2023. p. 147–55.

[59] Palm RB, et al. Variational neural cellular automata. In: *The Tenth International Conference on Learning Representations, ICLR*; 2022. OpenReview.net.

[60] Kalkhof J, González C, Mukhopadhyay A. *Med-NCA: Robust and lightweight segmentation with neural cellular automata*, in: *IPMI*. Springer; 2023. p. 705–16.

[61] Mittal A, et al. Medsegdiffnca: Diffusion models with neural cellular automata for skin lesion segmentation. In: *2025 IEEE 38th International Symposium on Computer-Based Medical Systems (CBMS)*. IEEE; 2025. p. 35–40.

[62] Niklasson E, et al. Self-organising textures. *Distill* 2021;6(2):e00027-e00003.

[63] Pathak D, et al. Learning to Control Self-Assembling Morphologies: A Study of Generalization via Modularity. *NIPS*. 2019. p. 2292–302.

[64] Medvet E, et al. Biodiversity in evolved voxel-based soft robots. In: *Proceedings of the Genetic and Evolutionary Computation Conference*; 2021. p. 129–37.

[65] Pigozzi F, et al. Evolving modular soft robots without explicit inter-module communication using local self-attention. In: *Proceedings of the Genetic and Evolutionary Computation Conference*; 2022. p. 148–57.

[66] Kriegman S, et al. A scalable pipeline for designing reconfigurable organisms. *Proc Natl Acad Sci U S A* 2020;117(4):1853–9.

[67] Gumuskaya G, et al. The Morphological, Behavioral, and Transcriptomic Life Cycle of Anthroboots. *Adv Sci (Weinh)* 2025;12(31):e2409330.

[68] Gumuskaya G, et al. Motile Living Biobots Self-Construct from Adult Human Somatic Progenitor Seed Cells. *Adv Sci (Weinh)* 2024;11(4):e2303575.

[69] Stovold J. Neural cellular automata can respond to signals. In: *ALIFE2023: Proceedings of the 2023 Artificial Life Conference*. MIT Press; 2023. p. 5.

[70] Hansali S, et al. The role of bioelectrical patterns in regulative morphogenesis: an evolutionary simulation and validation in planarian regeneration. *IEEE Transactions on Molecular, Biological, and Multi-Scale Communications* 2025.

[71] Kvalsund M-K, et al. Sensor Movement Drives Emergent Attention and Scalability in Active Neural Cellular Automata. *bioRxiv* 2024:2024-12.

[72] Pajouheshgar E, et al. Neural Cellular Automata: From cells to pixels 2025. arXiv [cs.CV].

[73] Pajouheshgar E, Xu Y, Süssstrunk S. NoiseNCA: Noisy seed improves spatio-temporal continuity of neural cellular automata. In: *ALIFE2024: Proceedings of the 2024 Artificial Life Conference*. MIT Press; 2024. p. 57.

[74] Meinhardt H. *The algorithmic beauty of sea shells*. 4th ed. Berlin Heidelberg: Springer; 2009.

[75] Mordvintsev A, Niklasson E. μ NCA: Texture Generation with Ultra-Compact Neural Cellular Automata 2021. arXiv [cs.AI].

[76] Kalkhof J, et al. *Parameter-efficient diffusion with neural cellular automata*. *npj Unconventional Computing* 2025;2(1):10.

[77] Dhariwal P, Nichol AQ. Diffusion Models Beat GANs on Image Synthesis. In: *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems, NeurIPS*. 2021; 2021. p. 8780–94.

[78] Ho J, Jain A, Abbeel P. Denoising Diffusion Probabilistic Models. In: *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems, NeurIPS*. 2020. p. 6840–51.

[79] Nichol AQ, Dhariwal P. Improved Denoising Diffusion Probabilistic Models. In: Proceedings of the 38th International Conference on Machine Learning, ICML. PMLR; 2021. p. 8162–71.

[80] Song J, Meng C, Ermon S. *Denoising Diffusion Implicit Models*, in 9th International Conference on Learning Representations, ICLR. 2021, OpenReview.net: Austria 2021.

[81] Sohl-Dickstein J, et al. Deep Unsupervised Learning using Nonequilibrium Thermodynamics. In: *Proceedings of the 32nd International Conference on Machine Learning, ICMLPMLR*; 2015. p. 2256–65.

[82] Faldor M, Cully A. CAX: Cellular Automata Accelerated in JAX. In: The Thirteenth International Conference on Learning Representations, ICLR. OpenReview.net: Singapore; 2025.

[83] Hintze A, Al-Hammadi M, Libby E. In: ALIFE2024: Proceedings of the 2024 Artificial Life Conference. Gene-Regulated Neural Cellular Automata. MIT Press; 2024. p. 99.

[84] Pio-Lopez L, Levin M. Morphoceuticals: Perspectives for discovery of drugs targeting anatomical control mechanisms in regenerative medicine, cancer and aging. *Drug Discov Today* 2023;28(6):103585.

[85] Levin M. *Bioelectric signaling: Reprogrammable circuits underlying embryogenesis, regeneration, and cancer*. *Cell* 2021;184(8):1971–89.

[86] Pio-Lopez L, Levin M. Aging as a loss of morphostatic information: A developmental bioelectricity perspective. *Ageing Res Rev* 2024;97:102310.

[87] Mitchell KJ, Cheney N. The genomic code: The genome instantiates a generative model of the organism. *Trends in Genetics* 2025;41(6):462–79.

[88] Hartl B, et al., Heuristically Adaptive Diffusion-Model Evolutionary Strategy. *Advanced Science*, in press.

[89] Li T, et al. CancerGPT for few shot drug pair synergy prediction using large pretrained language models. *NPJ Digit Med* 2024;7(1):40.

[90] Schneuing A, et al. Structure-based drug design with equivariant diffusion models. *Nat Comput Sci* 2024;4(12):899–909.

[91] Pagel S, Cronin L. Exploring the applications of Neural Cellular Automata in molecular sciences. *NeurIPS 2023 AI for Science Workshop*. 2023.

[92] Biswas S. *Gene regulatory networks exhibit several kinds of memory: quantification of memory in biological and random transcriptional networks*. *iScience* 2021;24(3):102131. et al.

[93] Sterling P. Principles of allostasis: optimal design, predictive regulation, pathophysiology, and rational 2004:17. Allostasis, homeostasis, and the costs of physiological adaptation.

[94] Fabregas-Tejeda A, Sims M. On the prospects of basal cognition research becoming fully evolutionary: promising avenues and cautionary notes. *Hist Philos Life Sci* 2025;47(1):10.

[95] Lyon P, Cheng K. Basal cognition: shifting the center of gravity (again). *Anim Cogn* 2023;26(6):1743–50.

[96] Lyon P, et al. Reframing cognition: getting down to biological basics. The Royal Society; 2021, 20190750.

[97] Manicka S, Levin M. The Cognitive Lens: a primer on conceptual tools for analysing information processing in developmental and regenerative morphogenesis. *Philos Trans R Soc Lond B Biol Sci* 2019;374(1774):20180369.

[98] Bongard J, Levin M. Living things are not (20th century) machines: updating mechanism metaphors in light of the modern science of machine behavior. *Frontiers in Ecology and Evolution* 2021;9:650726.

[99] Dodig-Crnkovic G. Cognition as Morphological/Morphogenetic Embodied Computation In Vivo. *Entropy (Basel)* 2022;24(11):1576.

[100] Biswas S, Clawson W, Levin M. Learning in Transcriptional Network Models: Computational Discovery of Pathway-Level Memory and Effective Interventions. *International Journal of Molecular Sciences* 2023;24. <https://doi.org/10.3390/ijms24010285>.

[101] Friston K, et al. Knowing one's place: a free-energy approach to pattern regulation. *Journal of The Royal Society Interface* 2015;12(105):20141383.

[102] Pezzulo G, et al. Bistability of somatic pattern memories: stochastic outcomes in bioelectric circuits underlying regeneration. *Philosophical Transactions of the Royal Society B: Biological Sciences* 2021;376(1821):20190765.

[103] Pezzulo G. Disorders of morphogenesis as disorders of inference: Comment on "Morphogenesis as Bayesian inference: A variational approach to pattern formation and control in complex biological systems". *Phys. Life Rev.* 2020;33:112–4.

[104] Pezzulo G, Levin M. Re-membering the body: applications of computational neuroscience to the top-down control of regeneration of limbs and other complex organs. *Integr Biol (Camb)* 2015;7(12):1487–517.

[105] Kaas JH. The organization of neocortex in mammals: implications for theories of brain function. *Annu Rev Psychol* 1987;38(1):129–51.

[106] Rao RPN. A sensory-motor theory of the neocortex. *Nat Neurosci* 2024;27(7):1221–35.

[107] Levin M. Self-Improvising Memory: A Perspective on Memories as Agential, Dynamically Reinterpreting Cognitive Glue. *Entropy* 2024;26. <https://doi.org/10.3390/e26060481>.

[108] Hornik K, Stinchcombe M, White H. Multilayer feedforward networks are universal approximators. *Neural networks* 1989;2(5):359–66.

[109] Guichard E. Critically pre-trained neural cellular automata as robot controllers. *Mechanical Engineering* 2024. TU Delft.

[110] Kaplan J, et al. Scaling laws for neural language models 2020. arXiv preprint arXiv:2001.08361.

[111] LeCun Y. A path towards autonomous machine intelligence 2022:1–62. *OpenReview preprint*.

[112] Cangelosi A, et al. Embodied intelligence. *Springer handbook of computational intelligence*. Springer; 2015. p. 697–714.

[113] Levin M. The Computational Boundary of a "Self": Developmental Bioelectricity Drives Multicellularity and Scale-Free Cognition. *Front Psychol* 2019;10:2688.

[114] Levin M. Bioelectric networks: the cognitive glue enabling evolutionary scaling from physiology to mind. *Anim Cogn* 2023;26(6):1865–91.

[115] Ha D, Schmidhuber J. Recurrent world models facilitate policy evolution. *Advances in neural information processing systems* 2018;31.

[116] Friston K. The free-energy principle: a unified brain theory? *Nat Rev Neurosci* 2010;11(2):127–38.

[117] Szathmary E, Smith JM. The major evolutionary transitions. *Nature* 1995;374(6519):227–32.

[118] West SA, et al. Major evolutionary transitions in individuality. *Proc Natl Acad Sci U S A* 2015;112(33):10112–9.

[119] Couzin ID. Collective cognition in animal groups. *Trends Cogn Sci* 2009;13(1):36–43.

[120] Feinerman O, Korman A. Individual versus collective cognition in social insects. *J Exp Biol* 2017;220(Pt 1):73–82.

[121] Momennejad I. Collective minds: social network topology shapes collective cognition. *Philos Trans R Soc Lond B Biol Sci* 2022;377(1843):20200315.

[122] Trianni V, et al. Swarm Cognition: an interdisciplinary approach to the study of self-organising biological collectives. *Swarm Intelligence* 2010;5(1):3–18.

[123] Szathmary E. Toward major evolutionary transitions theory 2.0. *Proc Natl Acad Sci U S A* 2015;112(33):10104–11.

[124] Jacob F. Evolution and tinkering. *Science* 1977;196(4295):1161–6.

[125] Pigozzi F, Goldstein A, Levin M. Associative conditioning in gene regulatory network models increases integrative causal emergence. *Commun Biol* 2025;8(1):1027.

[126] Gregor K, Besse F. Self-organizing intelligent matter: A blueprint for an ai generating algorithm 2021. arXiv preprint arXiv:2101.07627.,

[127] Béna G, et al. A Path to Universal Neural Cellular Automata 2025. arXiv [cs.LG].

[128] Miotti P, et al. Differentiable Logic Cellular Automata: From Game of Life to Pattern Generation 2025. arXiv [cs.AI].

[129] Pontes-Filho S, Nichele S, Lepperod M. Reservoir Computing with Evolved Critical Neural Cellular Automata. In: ALIFE2025: Ciphers of Life: Proceedings of the Artificial Life Conference 2025. Kyoto, Japan: MIT Press; 2025. p. 4.

[130] Pontes-Filho S, Lind PG, Nichele S. Assessing the robustness of critical behavior in stochastic cellular automata. *Physica D: Nonlinear Phenomena* 2022;441:133507.

[131] Beggs JM, Plenz D. Neuronal avalanches in neocortical circuits. *J Neurosci* 2003;23(35):11167–77.

[132] Beggs JM, Timme N. Being critical of criticality in the brain. *Front Physiol* 2012;3:163.

[133] Hesse J, Gross T. Self-organized criticality as a fundamental property of neural systems. *Front Syst Neurosci* 2014;8:166.

[134] Kinouchi O, Copelli M. Optimal dynamical range of excitable networks at criticality. *Nature Physics* 2006;2(5):348–51.

[135] Shew WL, et al. Neuronal avalanches imply maximum dynamic range in cortical networks at criticality. *J Neurosci* 2009;29(49):15595–600.

[136] Hengen KB, Shew WL. Is criticality a unified setpoint of brain function? *Neuron* 2025;113(16):2582–98, e2.

[137] Zhang Y, et al. Diffusion Models are Evolutionary Algorithms. In: The Thirteenth International Conference on Learning Representations, ICLR; 2025. OpenReview.net..

- [138] Wang G, et al. Hierarchical Reasoning Model 2025. arXiv [cs.AI].
- [139] Busoniu L, Babuska R, De Schutter B. A Comprehensive Survey of Multiagent Reinforcement Learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 2008;38(2):156–72.
- [140] Gupta JK, Egorov M, Kochenderfer M. Cooperative Multi-agent Control Using Deep Reinforcement Learning. *Autonomous Agents and Multiagent Systems*. Cham: Springer International Publishing; 2017.
- [141] Oroojlooy A, Hajinezhad D. A review of cooperative multi-agent deep reinforcement learning. *Applied Intelligence* 2022;53(11):13677–722.
- [142] Zhang K, Yang Z, Başar T, et al. Multi-Agent Reinforcement Learning: A Selective Overview of Theories and Algorithms. In: Vamvoudakis KG, et al., editors. *Handbook of Reinforcement Learning and Control*. Springer International Publishing: Cham; 2021. p. 321–84. Editors.
- [143] Paolo G, et al. TAG: A Decentralized Framework for Multi-Agent Hierarchical Reinforcement Learning. CoRR 2025. [abs/2502.15425](https://arxiv.org/abs/2502.15425).
- [144] Menta A, Archetti A, Matteucci M. Latent neural cellular automata for resource-efficient image restoration. In: ALIFE2024: Proceedings of the 2024 Artificial Life Conference. MIT Press; 2024. p. 100.
- [145] Deshpande R, et al. Engineering morphogenesis of cell clusters with differentiable programming. *Nat Comput Sci* 2025;5(10):875–83.
- [146] Podell D, et al. SDXL: Improving Latent Diffusion Models for High-Resolution Image Synthesis. In: The Twelfth International Conference on Learning Representations, ICLR. Vienna, Austria: OpenReview.net; 2024.
- [147] Qin Y, et al. WorldSimBench: Towards Video Generation Models as World Simulators 2024. [abs/2410.18072](https://arxiv.org/abs/2410.18072).
- [148] Ambrogioni L. The Statistical Thermodynamics of Generative Diffusion Models: Phase Transitions, Symmetry Breaking, and Critical Instability. *Entropy* 2025; 27. <https://doi.org/10.3390/e27030291>.
- [149] Raya G, Ambrogioni L. Spontaneous symmetry breaking in generative diffusion models. In: Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems, NeurIPS. LA, USA: New Orleans; 2023.
- [150] Sclocchi A, Favero A, Wyart M. A phase transition in diffusion models reveals the hierarchical nature of data. *Proc Natl Acad Sci U S A* 2025;122(1). e2408799121.
- [151] Gardner M. The fantastic combinations of John Conway's new solitaire game "life" by MartinGardner. *Scientific American* 1970;223(120-123):11.