Self-Organizing Models of Brain Wiring: Developmental Programs for Evolving Intelligence

Jamieson Warner Risto Miikkulainen

University of Texas at Austin and Cognizant AI Lab

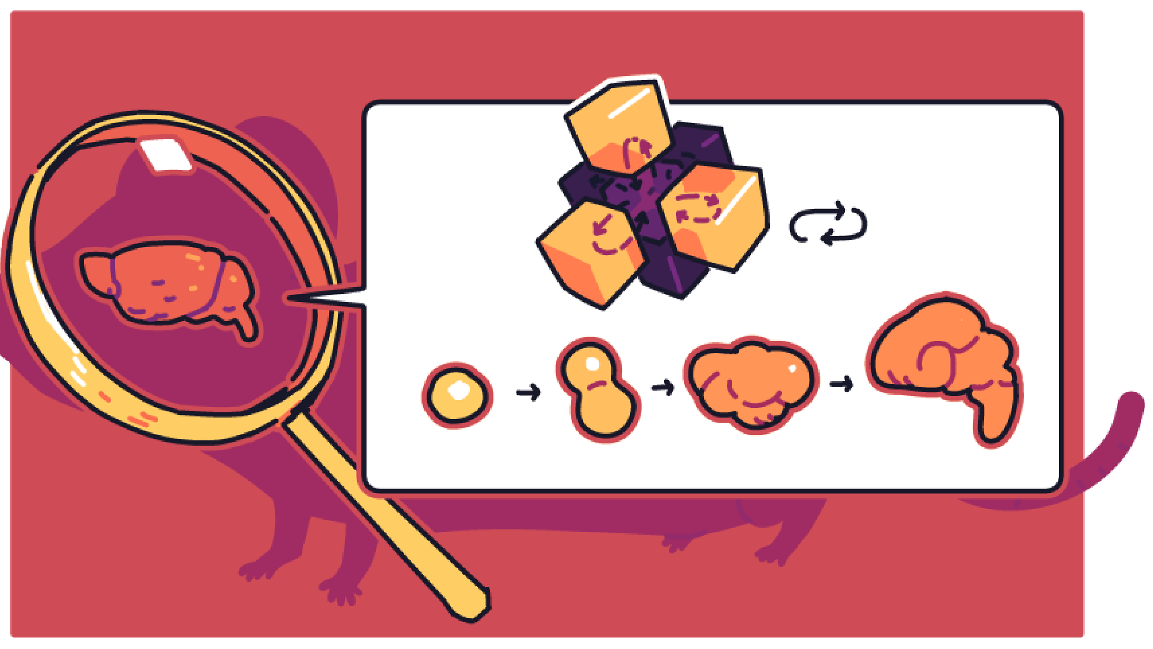

Why I Trained an NCA to Produce Mouse Brain Connectivity

Can We Derive Where Connectivity Comes From by Looking at the Connectivity Data Only?

- Previous studies identify genetic correlates of connectivity (Figure 1, Top)
- Can connectivity be explained from first principles by a developmental process model?
- **No access to gene expression** during training — only connectivity
- After training: test whether the model's state aligns with gene expression (Figure 1, Bottom)
- If so: **computational evidence** that connectivity is derived from a developmental process

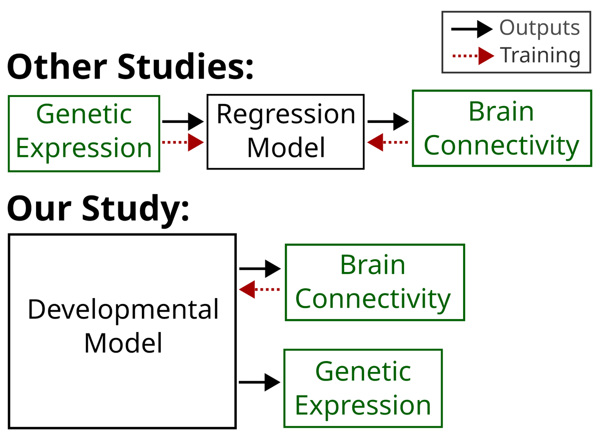

Figure 1

Explaining the Connectivity Data

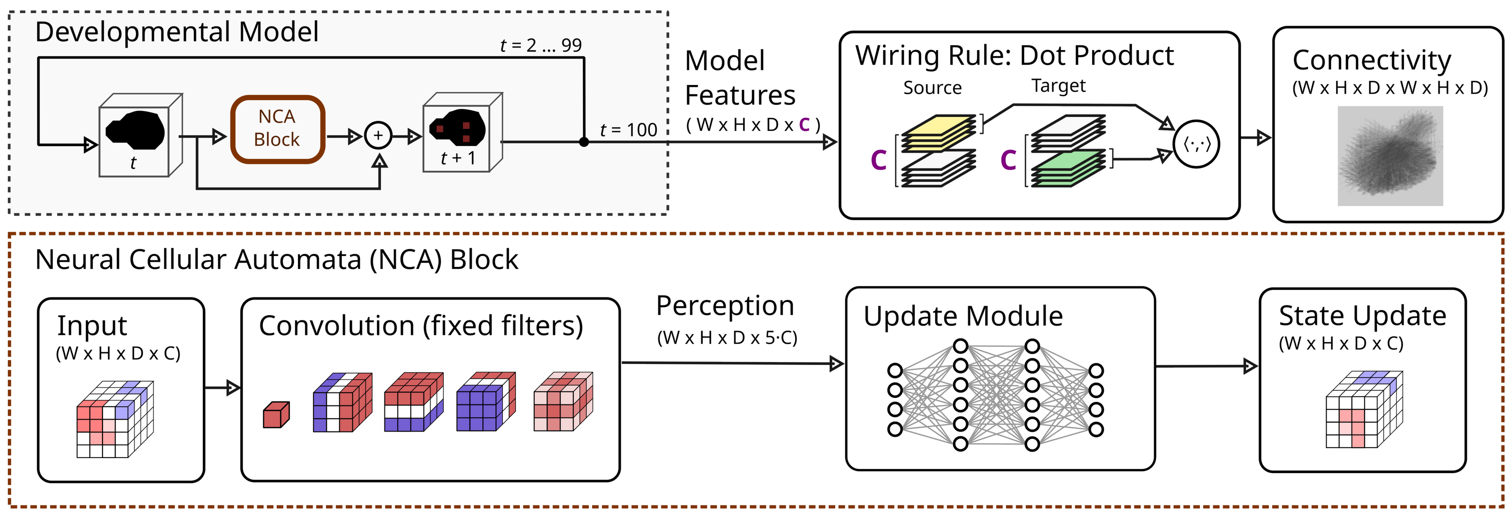

Figure 2

- NCA (Neural Cellular Automata) [4]: a differentiable model of biological pattern formation
- We adapt NCA for generative connectomics using a **homophilic wiring rule** [1] — connections form between units with similar features
- Connectivity is computed as a dot product between the **first half of the source's features** and the **second half of the target's features** (Figure 2)

What Can NCA Tell Us About Biological Self-Organization?

- NCA features align more strongly with gene expression than those from a static model (Figure 3)
- Gene score = correlation between predicted and actual gene expression across held-out brain regions
  1. Leave-one-out cross-validation across 12 major regions
  2. Linear map from model features → gene expression (20,000 genes)
  3. Scores averaged across test voxels
- Static model: a **phenomenological baseline** with no developmental dynamics
  - Directly learns model features per voxel
  - No spatial constraints; voxels treated independently
  - Equivalent to an SVD of the connectivity matrix
- **Result: NCA recovers biologically meaningful structure not captured by static models**

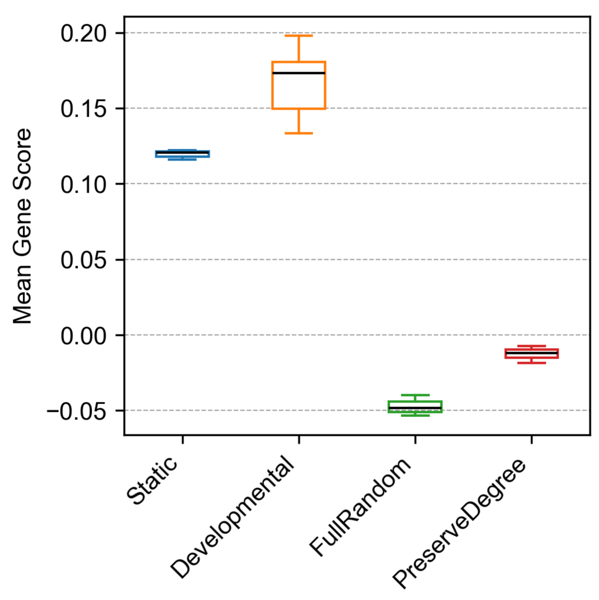

Figure 3

Surprising Finding: Simplicity is Key

- NCA shows strongest alignment with gene expression when using a **lightweight update module** (few hidden units) (Figure 4, Left; see Figure 2)
- The same pattern holds for the static model — simpler representations yield stronger biological correspondence (Figure 4, Right)
- Suggests that **simple programs** best capture the developmental logic encoded in gene expression

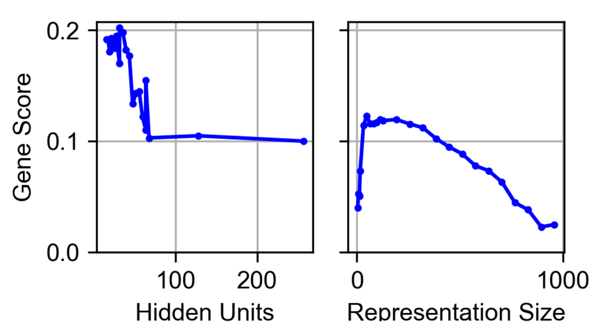

Figure 4

Abstract

In developing brains, axonal projections follow chemical gradients shaped by local interactions. This paper asks whether such a process can be inferred from its outcome: For instance, given observed mouse brain connectivity, can one recover the developmental program that produced it? If such developmental programs can be recovered, they not only explain how biological connectivity arises, but also offer a biologically grounded search space for artificial intelligence, in which architectures emerge through the evolution of genetic encodings that produce plausible wiring diagrams. A framework is proposed that uses the biological connectome itself as a beacon to guide this search, referred to as the Connectome-Generating, AI-Generating Algorithm (*CONGA*). Implemented as neural cellular automata (NCA), a model was trained to reproduce axon-tracing data in the mouse connectome, and its internal representations were compared to gene expression patterns measured in the same spatial coordinates. The result demonstrates how the brain of an intelligent organism may self-assemble through an indirect encoding of connectivity. The model outperformed a static linear baseline, but only when constrained in size, suggesting that compact developmental programs better align with biological mechanisms.

Challenge for Self-Organization Research: Reproduce Connectome Using Developmental and Ecological Priors

Nobody Knows What Brains Do or How to Make Them

- Ultimate goal is **human-like intelligence**, or AGI — i.e., to make a brain
- **Manual approach**: identify cognitive modules, and assemble them (Figure 5, Top)
- **AI-generating algorithm approach** [2]: identify an ecological reward function, and train intelligent agents (Figure 5, Bottom)

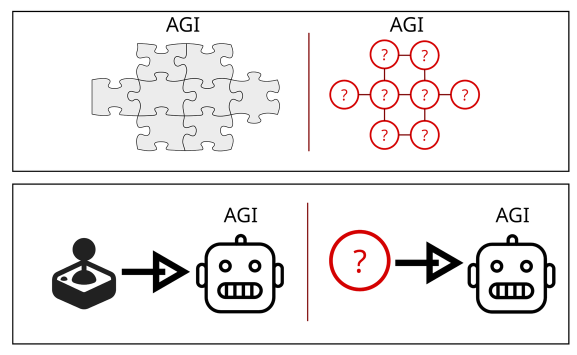

Figure 5

Connectome as a Beacon for Brain-like Intelligence

- AGI remains unrealized; we lack a complete ecological or cognitive model
- Proposal: **use the data we do have — the connectome** — to guide model development (Figure 6)

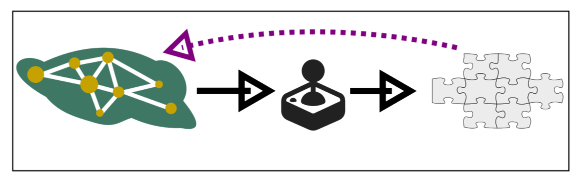

Figure 6

Proposal: Connectome-Generating, AI-Generating Algorithms

> *"What is the reward function (aka loss function or fitness function) for the environment generator? This is one of the key questions for AI-GA research."*
> *— Jeff Clune [2]*

- Agents evolve to solve environments, yielding both behavior and **circuits**
- Environments evolve to **induce** circuits that match the biological **connectome**
- **Wiring becomes the reward signal** for the environment generator (Figure 7)
  - Behavior fitness is defined by the environment
  - Environment fitness is defined by similarity to the connectome
- If an evolved model recapitulates brain-like wiring, it likely recapitulates brain-like function
- **Thus: the connectome offers a biologically grounded path toward AGI**

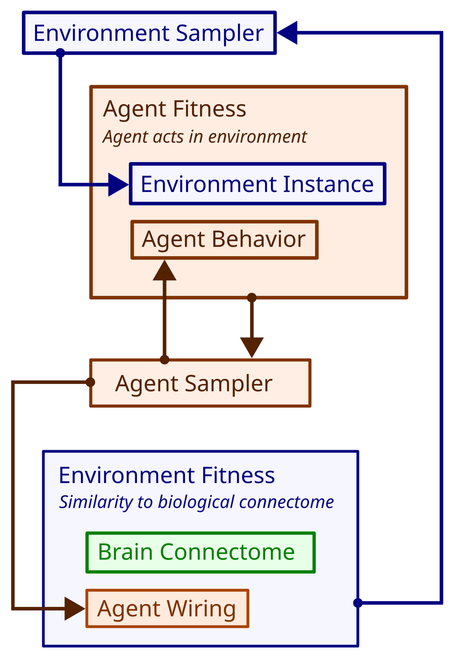

Figure 7

Development + Ecology as a Path to AGI

- A **functional account of the connectome** may be the key to AGI
- This requires searching over brain-like architectures
- **Developmental programs** provide structured priors over wiring
- These programs must generate both circuits and behavior
- We also need a design space of environments that model ecological pressure
- **Together, development and ecology define the search space for brain-like intelligence**

References

[1] Dániel L. Barabási and Albert-László Barabási.
A Genetic Model of the Connectome.
Neuron, 105(3):435–445.e5, February 2020.

[2] Jeff Clune.
AI-GAs: AI-generating algorithms, an alternate paradigm for producing general artificial intelligence, February 2020.

[3] Bernadette Meeker.
Illustration work featured in this poster. See: https://www.bernummm.com/, 2025.

[4] Alexander Mordvintsev, Ettore Randazzo, Eyvind Niklasson, and Michael Levin.
Growing Neural Cellular Automata.
Distill, 5(2):e23, February 2020.