

AI system self-organizes to develop features of brains of complex organisms

November 20 2023



Credit: Rice University



Cambridge scientists have shown that placing physical constraints on an artificially-intelligent system—in much the same way that the human brain has to develop and operate within physical and biological constraints—allows it to develop features of the brains of complex organisms in order to solve tasks.

As <u>neural systems</u> such as the <u>brain</u> organize themselves and make connections, they have to balance competing demands. For example, energy and resources are needed to grow and sustain the network in <u>physical space</u>, while at the same time optimizing the network for <u>information processing</u>. This trade-off shapes all brains within and across species, which may help explain why many brains converge on similar organizational solutions.

Jascha Achterberg, a Gates Scholar from the Medical Research Council Cognition and Brain Sciences Unit (MRC CBSU) at the University of Cambridge said, "Not only is the brain great at solving <u>complex</u> <u>problems</u>, it does so while using very little energy. In our new work we show that considering the brain's problem-solving abilities alongside its goal of spending as few resources as possible can help us understand why brains look like they do."

Co-lead author Dr. Danyal Akarca, also from the MRC CBSU, added, "This stems from a broad principle, which is that biological systems commonly evolve to make the most of what energetic resources they have available to them. The solutions they come to are often very elegant and reflect the trade-offs between various forces imposed on them."

In a study published in *Nature Machine Intelligence*, Achterberg, Akarca and colleagues created an artificial system intended to model a very simplified version of the brain and applied physical constraints. They found that their system went on to develop certain key characteristics and tactics similar to those found in human brains.



Instead of real neurons, the system used computational nodes. Neurons and nodes are similar in function, in that each takes an input, transforms it, and produces an output, and a single node or neuron might connect to multiple others, all inputting information to be computed.

In their system, however, the researchers applied a 'physical' constraint on the system. Each node was given a specific location in a virtual space, and the further away two nodes were, the more difficult it was for them to communicate. This is similar to how neurons in the human brain are organized.

The researchers gave the system a simple <u>task</u> to complete—in this case a simplified version of a maze navigation task typically given to animals such as rats and macaques when studying the brain, where it has to combine multiple pieces of information to decide on the shortest route to get to the end point.

One of the reasons the team chose this particular task is because to complete it, the system needs to maintain a number of elements—start location, end location and intermediate steps—and once it has learned to do the task reliably, it is possible to observe, at different moments in a trial, which nodes are important. For example, one particular cluster of nodes may encode the finish locations, while others encode the available routes, and it is possible to track which nodes are active at different stages of the task.

Initially, the system does not know how to complete the task and makes mistakes. But when it is given feedback it gradually learns to get better at the task. It learns by changing the strength of the connections between its nodes, similar to how the strength of connections between brain cells changes as we learn. The system then repeats the task over and over again, until eventually it learns to perform it correctly.



With their system, however, the physical constraint meant that the further away two nodes were, the more difficult it was to build a connection between the two nodes in response to the feedback. In the human brain, connections that span a large physical distance are expensive to form and maintain.

When the system was asked to perform the task under these constraints, it used some of the same tricks used by real human brains to solve the task. For example, to get around the constraints, the artificial systems started to develop hubs—highly connected nodes that act as conduits for passing information across the network.

More surprising, however, was that the response profiles of individual nodes themselves began to change: in other words, rather than having a system where each node codes for one particular property of the maze task, like the goal location or the next choice, nodes developed a 'flexible coding scheme.' This means that at different moments in time nodes might be firing for a mix of the properties of the maze. For instance, the same node might be able to encode multiple locations of a maze, rather than needing specialized nodes for encoding specific locations. This is another feature seen in the brains of complex organisms.

Co-author Professor Duncan Astle, from Cambridge's Department of Psychiatry, said, "This simple constraint—it's harder to wire nodes that are far apart—forces artificial systems to produce some quite complicated characteristics. Interestingly, they are characteristics shared by <u>biological systems</u> like the <u>human brain</u>. I think that tells us something fundamental about why our brains are organized the way they are."

Understanding the human brain

The team is hopeful that their AI system could begin to shed light on



how these constraints, shape differences between people's brains, and contribute to differences seen in those that experience cognitive or mental health difficulties.

Co-author Professor John Duncan from the MRC CBSU said, "These artificial brains give us a way to understand the rich and bewildering data we see when the activity of real neurons is recorded in real brains."

Achterberg added, "Artificial 'brains' allow us to ask questions that it would be impossible to look at in an actual biological system. We can train the system to perform tasks and then play around experimentally with the constraints we impose, to see if it begins to look more like the brains of particular individuals."

Implications for designing future AI systems

The findings are likely to be of interest to the AI community, too, where they could allow for the development of more efficient systems, particularly in situations where there are likely to be physical constraints.

Dr. Akarca said, "AI researchers are constantly trying to work out how to make complex, neural systems that can encode and perform in a flexible way that is efficient. To achieve this, we think that neurobiology will give us a lot of inspiration. For example, the overall wiring cost of the system we've created is much lower than you would find in a typical AI system."

Many modern AI solutions involve using architectures that only superficially resemble a brain. The researchers say their work shows that the type of problem the AI is solving will influence which architecture is the most powerful to use.

Achterberg said, "If you want to build an artificially-intelligent system



that solves similar problems to humans, then ultimately the system will end up looking much closer to an actual brain than systems running on large compute clusters that specialize in very different tasks to those carried out by humans. The architecture and structure we see in our artificial 'brain' is there because it is beneficial for handling the specific brain-like challenges it faces."

This means that robots that have to process a large amount of constantly changing information with finite energetic resources could benefit from having brain structures not dissimilar to ours.

Achterberg added, "Brains of robots that are deployed in the real physical world are probably going to look more like our brains because they might face the same challenges as us."

"They need to constantly process new information coming in through their sensors while controlling their bodies to move through space towards a goal. Many systems will need to run all their computations with a limited supply of electric energy and so, to balance these energetic constraints with the amount of information it needs to process, it will probably need a brain structure similar to ours."

More information: Spatially-embedded recurrent neural networks reveal widespread links between structural and functional neuroscience findings, *Nature Machine Intelligence* (2023). DOI: 10.1038/s42256-023-00748-9

Provided by University of Cambridge

Citation: AI system self-organizes to develop features of brains of complex organisms (2023, November 20) retrieved 9 May 2024 from <u>https://techxplore.com/news/2023-11-ai-self-</u>



organizes-features-brains-complex.html

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.