

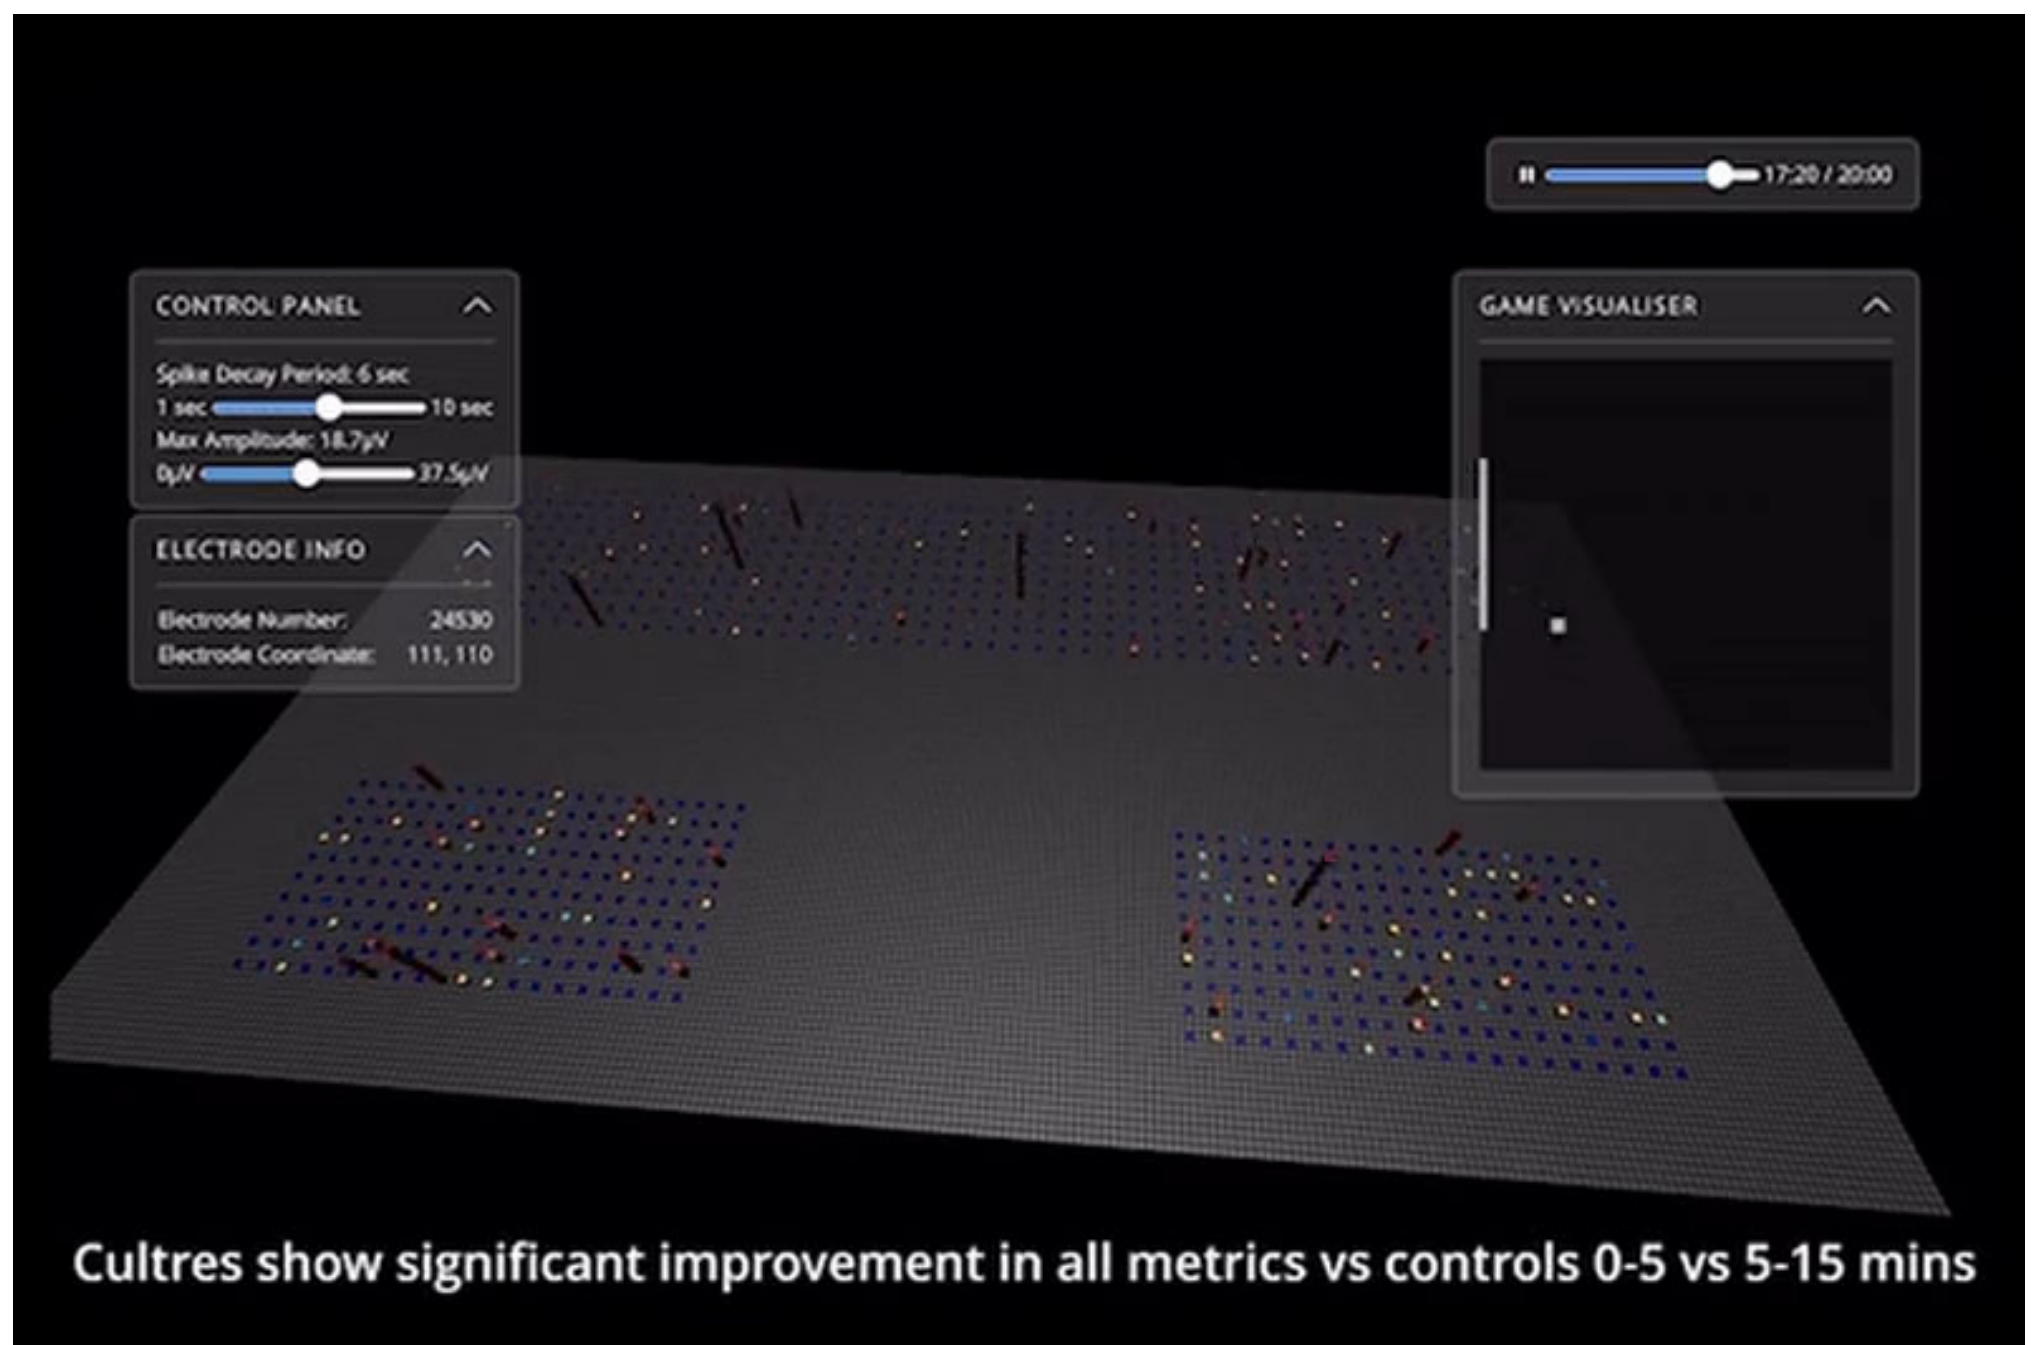
Lab-grown brains can play video games now. Do they... have rights?

By Liam Mannix
NOVEMBER 13, 2022

In a corner of a lab south of Melbourne sits an open laptop. None of the scientists working nearby give it a second glance. On the screen, someone is playing a game of Pong.

The unseen player is hesitant and twitchy, wobbling the paddle across the laptop screen toward the pixelated ball. But they hit it, more often than not.

A cable runs from the laptop to a large incubator. Inside, kept warm and bathed in nutrients, about 800,000 human neurons are behind the controls. And they are getting better.




In this video from Cortical, signals from the neurons are translated to allow them to play a game of Pong. FAIRFAX MEDIA

This brain in a dish could be the start of a whole new field of computing, where silicon and neurons are wired together to produce extraordinarily-powerful artificial intelligences. It may open a whole new path toward building an AI that thinks like a human.

It might also make an awful ethical mess. Can a bunch of brain cells ever be called conscious? Do they have *rights*? If their only reality is the pixels of Pong, is it ethical to... turn the game off?

We need to figure that out, and quickly. But Socrates has been dead for 2400 years and [neither philosophers nor neuroscientists](#) have a workable definition of consciousness. And when you start worrying about protecting consciousness, you run into our barbarous treatment of animals - many of whom surely meet some definition of having *feelings* and being *aware*.

DishBrain's arrival signals a field where our technological power may soon exceed our ethical understanding. "We're like children," says Professor Julian Savulescu, chair in ethics at the University of Oxford, "with a loaded AK-47".



Cortical Labs' Dr Brett Kagan, alongside the Hudson Institute's Dr Nhi Thao Tran and Monash's Dr Adeel Razi, who worked on the key paper published in *Neuron*. CHRIS HOPKINS

The first thing you notice when you enter Cortical Labs' office in Parkville is not the laptop - it's the robot dog.

The dog has four back-bending legs which gives it the low-slung look of a predator about to pounce, and a head full of sensors. Boston Dynamics sells them for about \$113,000, but this is a cheaper Chinese knockoff. It sits inanimate in its box, waiting like the Scarecrow for someone to insert a brain. Which is one of Cortical's next projects.

That's the company in a nutshell: low-budget, fast-moving, doing progressively-crazier things.

Cortical is the sci-fi dream of Hon Weng Chong, an ebullient entrepreneur who trained as a doctor and first tried to make it big selling Bluetooth stethoscopes (that company shuttered in 2019).

Chong had watched as scientists spent the last 20 years making dramatic progress growing neurons in the lab and turning them into increasingly-sophisticated models of regions of the human brain. These 'brain organoids' are more than just lumps of tissue: in recent experiments, the cells' electrical activity started syncing up in waves - brain waves. Some of the brain activity mimics that [seen in babies](#). In other studies, scientists were able to stimulate rat neurons and [get a signal back](#).


What if, Chong wondered, you could close the loop? Send signals into the neurons, get information back, and then give feedback? Could you get them to respond, process data, learn? You'd have a biological computer.



Cortical Labs CEO Hon Weng Chong. ELKE MEITZEL

Chong knew he'd never get conservative government research funders to back his idea. He needed a sci-fi geek. He found one in Niki Scevak, partner at venture-cap firm Blackbird; "he's always been one to believe in deep-tech and a sci-fi world", says Chong (Scevak says Cortical had ["a chance of magic"](#)). Blackbird invested \$1 million and the race to biology began.

Left to their own devices in a dish, neurons will sprout long tendril-like connections (called axons), building their own network. At first, Cortical tried to match those networks to the underlying hardware - but it proved enormously time-consuming.




A scanning electron microscope image of a neural culture that has been growing for more than six months on a high-density multi-electrode array. A few neural cells grow around the periphery and have developed complicated networks that cover the electrodes in the centre. CORTICAL LABS

Eventually Brett Kagan, Cortical's chief scientific officer, decided to stop trying to adapt the hardware to the neurons and let the neurons adapt to the hardware. After all, it is the neurons that were evolutionarily-designed to be flexible, he thought. To his surprise, left alone the neurons started performing dramatically better.

"It makes sense. Your brain, my brain, they are going to be quite different, but we can all do the same things," says Kagan.

Working alone in the lab during Melbourne's Long COVID lockdowns, he would watch his dishes of neurons play Pong again and again.

"I used to glare at it, and say, 'I swear it's getting better'. And then we did some analysis and it was," says Kagan. "And then the immediate thought was: what did we do wrong?"




Green marks neurons and axons, purple marks neurons, red marks dendrites, and blue marks all cells. CORTICAL LABS

The neurons sit atop a microchip that feeds electrical information in - like the distance of the ball from the paddle. Electrical signals from the neurons are used to move the paddle left and right. Each time the paddle hits the ball, the neurons get a little electrical reward. And, over time, they play better. Not by much, but enough to show *something* is happening.

“Yes, brain cells should be flexible and should be able to learn and get better. But to see it respond this strongly... It was pretty crazy,” says Kagan.

How? We don’t know - because we don’t really know how the brain works. But DishBrain’s success lends support to one theory: the brain is trying to build an accurate internal model of reality. Every time the neurons struck the ball, they were being told ‘your model is right’.

Such is the speed at which Cortical has to operate there is no time to rest on laurels. A new chip is already under-construction with increased input-output power. Coders are writing a custom programming language; soon they plan to start letting external developers write code and feed it in to the DishBrains. Daniela Duc shows me 3D-printed cases containing life-support systems that will allow DishBrain to sit on a desk like a portable computer - or be stacked in a server farm.



From left, the Apple HomePod, Google Home and Amazon's Alexa. JASON HENRY

Artificial intelligence controls an ever-increasing slice of our lives. Smart voice assistants hang on our every word. Our phones leverage machine learning to recognise our face. Our social media lives are controlled by algorithms that surface content to keep us hooked.

These advances are powered by a new generation of AIs built to resemble human brains. But none of these AIs are really *intelligent*, not in the human sense of the word. They can see the superficial pattern [without understanding the underlying concept](#). Siri can read you the weather but she does not really understand that it's raining. AIs are good at learning by rote, but struggle to extrapolate: even teenage humans need only a few sessions behind the wheel before they can drive, while Google's self-driving car still isn't ready after [32 billion kilometres of practice](#).


A true 'general artificial intelligence' remains out of reach - and, some scientists think, [impossible](#).

Is this evidence human brains can do something special computers never will be able to? If so, the DishBrain opens a new path forward. "The only proof we have of a general intelligence system is done with biological neurons," says Kagan. "Why would we try to mimic what we could harness?"

He imagines a future part-silicon-part-neuron supercomputer, able to combine the raw processing power of silicon with the built-in learning ability of the human brain.

Others are more sceptical. Human intelligence isn't *special*, they argue. Thoughts are just electro-chemical reactions spreading across the brain. Ultimately, everything is physics - we just need to work out the maths.

"If I'm building a jet plane, I don't need to mimic a bird. It's really about getting to the mathematical foundations of what's going on," says Professor Simon Lucey, director of the Australian Institute for Machine Learning.



Pong, one of the first video games ever coded.

Why start the DishBrains on Pong? I ask. Because it's a game with simple rules that make it ideal for training AI. And, grins Kagan, it was one of the first video game ever coded. A nod to the team's geek passions - which run through the entire project.

"There's a whole bunch of sci-fi history behind it. *The Matrix* is an inspiration," says Chong. "Not that we're trying to create a Matrix," he adds quickly. "What are we but just a gooey soup of neurons in our heads, right?"

Maybe. But the Matrix wasn't meant as inspiration: it's a cautionary tale. The humans wired into it existed in a simulated reality [while machines stole their bioelectricity](#). They were *slaves*.

Is it ethical to build a thinking computer and then restrict its reality to a task to be completed? Even if it is a fun task like Pong?

"The real life correlate of that is people have already created slaves that adore them: they are called dogs," says Oxford University's Julian Savulescu.

Thousands of years of selective breeding has turned a wild wolf into an animal that enjoys rounding up sheep, that loves its human master unconditionally.

"Maybe it's OK to create a DishBrain that's happy playing Pong, and that's all it desires," says Savulescu. "I really have no idea what the answer is."

That was a common refrain from philosophers to most of the questions DishBrain raises: we just don't know.



The Matrix wasn't meant as inspiration, but as cautionary tale.

Most philosophers say if something has consciousness, it deserves some level of protection.

DishBrains remain primitive, and no experts told *The Age* they believed it was already conscious. "Is my garage door opener conscious when it opens the garage door as my car gets close to it?" says Stanford's Hank Greely, founder of the International Neuroethics Society.

But the path is clear: we are going to continue to build ever-more-sophisticated models of the human brain. These models, says Savulescu, may one day come to represent genuinely new life-forms, that demand answers to a new set of ethical questions.

Can a brain wired into a computer suffer? If so, how could it tell us? "There's a risk you might get it completely wrong, and end up inflicting horrible suffering in the context of trying to make these things learn," says Dr Julian Koplin, a research fellow in biomedical ethics at Murdoch Children's Research Institute.

Could a sophisticated DishBrain become conscious? We don't know, because we don't have a good yardstick for what consciousness is. "And we're not close to being there yet," says Koplin. "The development of DishBrain is a sign we really need to get on top of this."

Many philosophers think of consciousness as being able to have a subjective experience: feeling the *bitterness* of coffee or the *painfulness* of pain. But... animals can experience sensation, and society often does not treat them as conscious.

Perhaps the answer is only human consciousness deserves the highest form of moral protection? "You can see where the problem arises when you think like that: who gets to be treated with dignity and who does not?" says Dr David Kirchoffer, director of the Queensland Bioethics Centre.

Aristotle distinguished between humans who could reason, and therefore deserved dignity - men - and those who could not - slaves, children, women.

The issue gets worse when you start thinking about animals, many of whom seem a lot more conscious than DishBrain. [Ravens play, remember the past and anticipate the future](#). Apes, elephants and dolphins recognise themselves in mirrors. [Elephants mourn](#). In experiments, rhesus monkeys will [refuse to injure their companions in exchange for food](#) - evidence, perhaps, for empathy. More empathy than we extend when we subject them to medical experiments.



Tumani, a critically endangered western lowland gorilla, holds her newborn in Audubon Zoo, New Orleans on September 4. AUDUBON NATURE INSTITUTE

Why is an awake, aware primate given less moral worth than a comatose human?

“It’s even difficult in the same species. Consider the right to abortion,” says Associate Professor Frederic Gilbert, head of philosophy at the University of Tasmania. Scientists [generally think of them as somewhat conscious at between 20 and 30 weeks of fetal development](#). But that’s no guarantee of moral rights: Victoria and NSW both allow abortions up to 24 weeks.

Cortical are aware of the issues and have tried to be proactive. Kagan [published a paper with Gilbert in March](#) setting out their thinking - including the argument that testing on DishBrain offers an alternative to animal testing.

Giving protection to DishBrain “would mean these cells have more ethical importance than a rodent or primate,” says Gilbert. “And so we’d keep testing on animals - and I think that’s ethically wrong.”

But even Cortical agree some form of regulation is going to be needed. Australia’s health research agency told *The Age* it was looking closely at the issue “with a view to determining if specific ethics guidance is required in the future”.

For decades, scientists were prevented from [growing a human embryo in a lab](#) for longer than 14 days. Cortical’s fridge has a DishBrain that’s been alive for more than a year.

“This is totally uncharted territory, ethically,” says Savulescu. “We’re in Medieval times in terms of our ethical progress relevant to our power.”

Get news and reviews on technology, gadgets and gaming in our Technology newsletter every Friday. [Sign up here.](#)



Liam Mannix is The Age and The Sydney Morning Herald’s national science reporter. Connect via [Twitter](#).