




THE SUBSTRATE ISSUE · NEUROMORPHIC COMPUTE

The Neuron Was Never a *Switch*

A new PNAS study measured how hard it is for a deep network to imitate one human brain cell — and found the cell wins. The result is being read as a story about intelligence. It is really a story about wiring, and it lands in the middle of a \$403 billion quarter for silicon.

AGAINST THE GRAIN

Why the “brains are complicated” takeaway is the wrong one.

SIGNALS

Chips clear \$403.3B in a quarter; a black hole wearing a star.

BY THE NUMBERS

20 watts, eight layers, and one quintillion.

03	Feature · The Computing Neuron	<i>What an artificial twin revealed about a real cell</i>
04	Feature · Inside the Tree	<i>Membrane area, NMDA nonlinearity, and depth</i>
05	Feature · So What	<i>From simulation to silicon · Field Notes · Glossary</i>
06	Against the Grain	<i>The wire is the computer — and why most disagree</i>
07	Signals	<i>Semiconductors · Space · Energy</i>
08	By the Numbers	<i>Seven figures worth carrying into Monday</i>
09	The Long View	<i>The container was a participant all along</i>

DISPATCH · FROM THE EDITOR

The medium stopped being furniture.

Engineering is mostly the art of drawing a box around the thing you care about and calling everything outside it "environment." The transistor is the computer; the wire is plumbing. The plasma is the reactor; the wall is a container. The star is the engine; the gas is scenery. This decomposition is not laziness — it is the single most productive simplification of the last century, and it built everything on your desk.

Three papers surfacing this week all say the same inconvenient thing from three unrelated directions: the environment is not furniture. It is a participant, and often the dominant one.

In a Hebrew University study published in PNAS, the branching structure surrounding a neuron's cell body turns out to carry so much of the computation that an artificial network needs several internal layers just to impersonate a single cell. In *Nature Communications*, the metal lattice surrounding two deuterium nuclei changes their fusion probability by eighteen orders of magnitude. In *Nature*, a hydrogen cocoon surrounding an infant black hole makes it look, from twelve billion light-years away, like an enormous star.

Substrate, lattice, cocoon. Three fields, one lesson, and a useful question for anyone who runs an engineering organization: which part of your stack are you still treating as a container when the data says it is doing the work?

None of these results overturn a law. What they overturn is a habit — the reflex of assigning causation to the object in the middle of the diagram and treating the rest as boundary conditions. That reflex is cheap, which is why it is universal, and it fails in exactly one situation: when the surroundings are dense, structured, and coupled. Which is to say, when the system is interesting.

Our deep dive this week is the neuron paper, because it is the one with a bill attached. The world is currently spending several hundred billion dollars a quarter on machines built from the opposite premise — that the switch computes and the wire merely delivers. It is worth knowing what the other architecture looks like, even if nobody can fabricate it yet.

The contrarian column on page six argues the finding is not really about intelligence at all. Read it as a floorplan.

The Computing Neuron

To find out how much arithmetic one brain cell performs, a team in Jerusalem built an artificial twin and made it try to keep up. The number of layers the twin required is the measurement — and the human cell demanded more of them than any rat's.

BY THE INFLECTION DESK · REPORTED FROM THE LITERATURE

For seventy years the working cartoon of a neuron has been a doorbell. Charge accumulates, a threshold is crossed, a spike goes out. Everything interesting was supposed to live in the pattern of connections between the doorbells — which is exactly the assumption that made artificial neural networks tractable, and exactly the assumption a paper published this month in the *Proceedings of the National Academy of Sciences* spends thirty pages dismantling.

The team — Ido Aizenbud, Daniela Yoeli, David Beniaguev, Michael London and Idan Segev at the Hebrew University's Edmond and Lily Safra Center for Brain Sciences, with Christiaan de Kock of the Free University of Amsterdam — set out to answer a question that sounds unanswerable. How much computation does one cell do?

You cannot count it directly. A neuron is not a circuit diagram you can trace; it is a wet analog device with thousands of inputs arriving on a branching structure whose every fork changes how a signal decays. So the team did something oblique. They took detailed biophysical models of cortical pyramidal neurons, fed them realistic synaptic input, recorded the output, and then trained a deep artificial network to reproduce that input-output mapping at millisecond resolution.

Then they asked how big the imitator had to be.

The trick: complexity as difficulty-of-imitation.

That ratio is the paper's contribution. They call it the Functional Complexity Index, or FCI: a general, deep-learning-based framework for assessing the input-output complexity of a neuron. The harder the artificial twin has to work — the more layers, the more capacity — the more computation the biological cell was doing all along.

It is a lovely move, because it converts a question about biology into a question about machine learning, where we already have an honest ruler. And it has a precedent. In 2021, Beniaguev, Segev and London showed in *Neuron* that faithfully mimicking a single layer-5 cortical pyramidal cell at spiking resolution required a temporally convolutional network five to eight layers deep. Not five to eight neurons. Five to eight layers.

The new work turns that one-off demonstration into a measuring instrument, then points it across species — and the gap it finds is the story.

The methodology is unusual in a specific way. Most attempts to quantify biological computation smuggle in an assumption about what counts — bits per spike, information rate, channel capacity — each requiring you to decide in advance what the neuron is for. FCI declines to decide. It asks only how much machinery is needed to be indistinguishable from the cell.

The cost of that elegance is a strange unit. "Five layers" means five layers of the architecture this team happened to train. What survives is the comparison — human against rat, intact tree against ablated tree — and comparisons are where the evidence lives.

Inside the tree

Applying the FCI to reconstructed human and rat cortical pyramidal neurons, the team reports that human cells are significantly more functionally complex than their rat counterparts. That result is not surprising in direction. What matters is the attribution: *which* physical properties buy the complexity.

01 · Surface, not switches

The dominant factor is dendritic membrane area and branching pattern. Human cortical pyramidal neurons carry larger, longer, more elaborately arborised dendritic trees than rodent cells. In the FCI framework, that extra surface is not passive real estate for more synapses to land on. It is the compute.

A dendritic branch is a lossy transmission line. A signal arriving at a distal tip attenuates on its way to the cell body, and how much it attenuates depends on where its neighbours landed and when. Every fork in the tree is therefore an operator — a place where two inputs interact nonlinearly before anything reaches the soma. Add branches, add operators. Depth in a neural network and branching in a dendrite are, functionally, the same currency.

02 · The NMDA receptor as a multiplier

The second factor the paper isolates is the density and non-linearity of NMDA-mediated synaptic receptors. An NMDA receptor is unusual: it only conducts if the membrane is already depolarised *and* its transmitter is bound. That is a coincidence detector — an AND gate with analog gain, scattered by the thousand across the tree.

Strip the NMDA nonlinearity out of the simulated cell and the required twin shrinks. The receptors are not modulation. They are instruction set.

“People often think of a neuron as a simple switch that either turns on or off. What we show is that a single human neuron is itself an extraordinarily sophisticated computing device.”

IDAN SEGEV · HEBREW UNIVERSITY OF JERUSALEM

03 · One cell, one classifier

The most vivid demonstration in the accompanying materials is almost provocative: a single human cortical neuron, given visual input, can perform discriminations of the sort we normally assign to whole networks — telling images of cats from images of dogs. Not well, and not by itself in any behaviourally meaningful sense. But the fact that the question is even coherent is the finding.

04 · A ruler, not just a result

The framework generalises. FCI links measurable physical geometry — surface area, branch order, receptor density — to a number describing computational depth. That is the piece with a half-life longer than the headline: a way to ask of any cell, or any circuit element, how much work its shape is doing.

Note what the paper does *not* claim. It does not say human neurons are faster; they are not. It does not say the human cortex has more cells; per unit volume it has fewer than a mouse's. It says the same evolutionary budget was spent differently — on morphology rather than multiplicity. Given a fixed skull and a fixed calorie allowance, depth-per-unit was apparently the better trade.

That is a design decision, not a miracle, and design decisions are the kind of thing engineers can copy.

The uncomfortable corollary is that “number of neurons” — the quantity every brain-equivalence estimate has leaned on for forty years — was measuring the wrong axis. It counted the nodes and ignored the depth inside each one. Whether that makes the brain a harder target or a better blueprint depends entirely on which business you are in.

What it means for the people buying racks

Nothing here ships. The FCI is measured against biophysical simulations, not living tissue, and no one is fabricating a dendrite this decade. But the result arrives at a specific commercial moment, and that moment gives it teeth.

Global semiconductor sales hit \$403.3 billion in the second quarter of 2026 alone, up 35.1% on Q1. An SIA-DeLoitte analysis puts semiconductors at 95% of the value of an AI data-server rack, with annual revenue from AI-data-center chips potentially exceeding \$1.2 trillion by 2028. Essentially all of that spend is on hardware built from units that are, individually, arithmetically trivial — a multiply, an add, a rectifier. The entire capability of the system comes from having a very large number of them and moving data between them very fast.

The brain runs the other configuration. Roughly 100 billion neurons, each apparently worth several layers of network, on a metabolic budget of about 20 watts. Not twenty watts per rack. Twenty watts, total, including the parts that keep you breathing.

The near-term product of this line of research is unlikely to be a chip. It is more likely to be a *unit*: architectures whose basic element is internally deep — multi-compartment neurons, dendritic layers, per-node nonlinear cascades — rather than a scalar activation. Segev's group frames the paper explicitly as a template for brain-inspired AI built from artificial units that are themselves computationally deep, unlike the highly simplified units in today's systems.

If that direction works, parameter counts stop being comparable across architectures — a billion deep units is not a billion floats — and the industry loses its favourite scoreboard. Watch for it first in edge inference and always-on sensing, where the energy constraint is real and the workloads are small enough that an unfamiliar training story is survivable.

The practical near-term signal to track is not a chip announcement. It is a training result: the first credible demonstration that a network of internally deep units can be optimised end-to-end at scale. Until that exists, this remains a beautiful measurement in search of a gradient.

FIELD NOTES · HOW THE MEASUREMENT WORKS

1

Simulate the real cell.

Build a detailed biophysical model of a reconstructed cortical pyramidal neuron and drive it with realistic synaptic input. Record every spike it produces.

2

Train an imitator.

Fit a deep artificial network to reproduce that same input-to-output mapping at millisecond resolution. Grow it until the fit is faithful.

3

Read the size.

The capacity the twin needed is the cell's Functional Complexity Index. Then ablate features — branches, NMDA receptors — and watch the required size fall.

GLOSSARY

Dendrite

The branching input structure of a neuron. Historically modelled as a passive funnel; increasingly modelled as a computing element in its own right.

NMDA receptor

A synaptic receptor that conducts only when transmitter binding and membrane depolarisation coincide — a biological AND gate with analog gain.

Pyramidal neuron

The dominant excitatory cell type of the cortex, named for its shape. Layer-5 pyramidal cells are the workhorses of these studies.

Functional Complexity Index

The paper's new metric: the size of the deep network required to imitate a neuron's input-output behaviour.

This is not a paper about intelligence. It is a floorplan.

The consensus reading of the Jerusalem result is a story about scale. Brains are more complicated than we assumed; therefore our estimates of brain-equivalent compute were too low; therefore artificial general intelligence is further away than the optimists say. Sobering, tidy, and — I think — pointed in exactly the wrong direction.

Read the attribution again. The complexity did not come from more neurons, faster neurons, or better neurons. It came from *dendritic membrane area and branching*. It came from the wire.

Every computing architecture we have built since 1945 rests on a clean separation between the part that computes and the part that carries. The ALU thinks; the bus fetches. We have spent the AI era paying an escalating tax on that separation — HBM stacks, chip-to-chip fabrics, co-packaged optics, entire product categories that exist purely to move a number from where it is stored to where it will be multiplied. The interconnect is treated as overhead: necessary, expensive, and definitionally not where value is created.

The dendritic tree is the counterexample sitting inside your skull. It is unambiguously interconnect — it is the input cabling of the cell — and the measurement says it is doing several layers of nonlinear work in transit. Transport and arithmetic are the same operation. That is not an efficiency trick bolted onto a Von Neumann machine. It is a different answer to the question of where computation lives.

Which reframes the twenty-watt figure. The brain is not efficient because its switches are cheap. It is efficient because it almost never *moves data without computing on it*. The energy the rest of us spend on transport, biology spends once and gets arithmetic thrown in.

If you run infrastructure, the actionable inversion is this: stop modelling the fabric as a cost centre to be minimised and start asking what it could compute on the way. In-memory and in-network compute have been filed for a decade under "efficiency." This paper suggests they belong under "capability."

The test of a Move 37 is not that it sounds clever. It is that it changes what you would fund. Under the consensus reading, you fund more of the same and wait. Under this one, you fund the group that has been quietly arguing your switch fabric should have arithmetic units in it — and you stop letting them describe the project as a power-saving initiative.

WHY THE CONSENSUS DISAGREES — AND HAS A CASE

It is simulation, not tissue. The FCI is computed against biophysical models. Faithful ones, but models. No electrode measured a cat-versus-dog discrimination in a living human cell.

Hard to fit is not the same as useful. Turbulence is hard to fit. Difficulty-of-imitation is a measure of expressivity, not of computational value, and the paper cannot distinguish the two.

Depth may be untrainable. We have no credible method for training networks whose every unit is itself five to eight layers deep. Gradient flow through such a thing is an open problem, and the training bill could eat the inference saving whole.

Neuromorphic has lost before. TrueNorth and Loihi were elegant and efficient and were beaten by brute-force GPUs on economics, not physics. Substrate arguments have a poor record against supply chains.

\$403.3_B

Q2 2026 CHIP SALES

SEMICONDUCTORS

The demand curve has stopped looking like a cycle

The Semiconductor Industry Association reported global semiconductor sales of \$403.3 billion for the second quarter of 2026 — a 35.1% increase over Q1 of the same year. June alone accounted for \$134.5 billion, up 123.6% year-on-year and 9.7% month-on-month. Sequential growth of that magnitude inside a single quarter is not a recovery pattern; it is a capacity-constrained market clearing at whatever price it must. The industry's own projection has annual sales topping \$1.5 trillion globally in 2026, and a separate SIA-Deloitte study finds semiconductors now constitute 95% of the value of an AI data-server rack. The bill of materials for intelligence has become, almost entirely, silicon.

SOURCE · SEMICONDUCTOR INDUSTRY ASSOCIATION, PRESS RELEASE, 6 AUGUST 2026

100,000

SOLAR MASSES, AT COSMIC DAWN

SPACE

A black hole wearing a star as a coat

Writing in *Nature*, a team led by MIT's Rohan Naidu described MoM-BH*-1: an object the size of our solar system, radiating like a star but 100 billion times brighter than nuclear fusion can physically manage. Their reconstruction is a black hole of roughly 100,000 solar masses wrapped in a hydrogen envelope so dense it behaves like a stellar photosphere — producing the deepest Balmer break ever recorded and almost no elements heavier than helium. The team argues this "black hole star" may explain the little red dots that appear in nearly every deep JWST image and vanish by the present day. If so, the growth channel for supermassive black holes was hiding in plain sight, disguised by its own gas.

SOURCE · NAIDU ET AL., NATURE, 12 AUGUST 2026 · MIT NEWS

10¹⁸

FUSION-YIELD ENHANCEMENT

ENERGY

The lattice is not a bystander in the reaction

A UC Davis and Berkeley Lab collaboration loaded palladium and titanium foils with deuterium, then bombarded them with low-energy deuterium ions. Below 2 keV, where theory predicts fusion yield should collapse exponentially, they measured a plateau — a finite, non-vanishing floor. At the lowest energies probed, yields exceeded bare-nucleus expectations by more than a factor of 10¹⁸. The authors attribute it to electron screening within the metal hydride partially cancelling the Coulomb barrier, and describe it as a previously unrecognised regime in which "materials degrees of freedom can fundamentally renormalize tunneling probabilities." Nobody is claiming net energy. They are claiming the reactor wall has a vote.

SOURCE · KARAHADIAN ET AL., NATURE COMMUNICATIONS, 18 JULY 2026 · DOI 10.1038/S41467-026-74421-1

5–8

Layers of a temporally convolutional deep network required to reproduce the input–output behaviour of a *single* layer-5 cortical pyramidal cell at millisecond resolution.

20 W

Approximate metabolic power budget of the entire human brain — about 20% of the body's oxygen and calorie consumption, on roughly 2% of its mass.

100 B

Neurons in the human brain, the figure on which every brain-equivalent-compute estimate has been built. The new work argues the multiplier per neuron was wrong, not the count.

95%

Share of the value of an AI data-server rack accounted for by semiconductors, across the full stack of chip technologies.

\$403.3 B

Global semiconductor sales in Q2 2026 — up 35.1% on the previous quarter, with June sales up 123.6% year-on-year.

 10^{18}

Factor by which deuterium–deuterium fusion yields in metal hydrides exceeded bare-nucleus expectations below 2 keV.

100 B

Times brighter MoM-BH*-1 is than any star powered by nuclear fusion can physically be — the observation that forced a new class of astrophysical object.

\$1.2 T

Projected annual revenue from chips deployed in AI data centres by 2028 — close to a tenfold increase over four years.

The container was a participant all along

There is a particular kind of scientific result that does not add a fact so much as relocate a boundary. This week produced three of them, and they rhyme.

A neuron's dendritic tree was the packaging around the computation; it turns out to be several layers of the computation. A metal foil was the container for a fusion experiment; it turns out to renormalise the tunnelling probability by eighteen orders of magnitude. A cloud of hydrogen was the dust obscuring an early black hole; it turns out to be the photosphere that makes the object a new class of thing entirely.

In each case the error was the same, and it was not a factual error. It was a modelling convenience that hardened into an ontology. We drew the box where it was analytically tractable to draw it, got a century of good results, and forgot the line was ours.

This should be uncomfortable for anyone building large systems right now, because the AI stack is made of exactly these boundaries. Model and inference server. Compute and storage. Weights and data. Training and deployment. Each of those lines was drawn for tractability. Each is now load-bearing in a way nobody chose.

The useful discipline is not to erase the boundaries — they are still how anything gets built — but to keep a written list of them, and to revisit it when the numbers stop making sense. When your system is inexplicably slow, or inexplicably expensive, or inexplicably good, the answer is disproportionately often in the part of the diagram you labelled "environment" and stopped looking at.

The dendrite was in the picture the whole time. It was just drawn in grey.

SOURCES & FURTHER READING

Aizenbud, Yoeli, Beniaguev, de Kock, London & Segev. "Dendritic morphology and synaptic nonlinearities enhance functional complexity in human cortical neurons." *PNAS* 123(28), 2026. DOI 10.1073/pnas.2533168123 — <https://www.pnas.org/doi/10.1073/pnas.2533168123>

Hebrew University of Jerusalem / ScienceDaily. "Human brain cells are far more powerful than scientists thought," 13 August 2026 — <https://www.sciencedaily.com/releases/2026/08/260812015216.htm>

Neuroscience News. "The Branching Dendrite Secret Behind Human Cognitive Superiority," with full abstract — <https://neurosciencenews.com/branching-dendrite-human-cognition-31015/>

Beniaguev, Segev & London. "Single cortical neurons as deep artificial neural networks." *Neuron* 109(17):2727–2739, 2021 — <https://pubmed.ncbi.nlm.nih.gov/34380016/>

Semiconductor Industry Association. "Global Semiconductor Sales Increase 35.1% from Q1 2026 to Q2 2026," 6 August 2026 — <https://www.semiconductors.org/global-semiconductor-sales-increase-35-1-from-q1-2026-to-q2-2026/>

SIA & Deloitte. "Semiconductors Account for 95% of an AI Data Server Rack's Value," 1 June 2026 — <https://www.semiconductors.org/new-report-finds-semiconductors-account-for-95-of-an-ai-data-server-racks-value-encompassing-the-full-stack-of-chip-technologies/>

Naidu, Simcoe, Sun et al. "A Gas Enshrouded and Gas Reddened Black Hole at Cosmic Dawn." *Nature*, 12 August 2026 — <https://www.nature.com/articles/s41586-026-10846-4>


MIT News. "Astronomers discover a brand-new type of astrophysical object: A black hole star," 12 August 2026 — <https://news.mit.edu/2026/astrophysicists-discover-brand-new-type-astrophysical-object-black-hole-star-0812>

Karahadian, Colborne, Persaud, Schenkel & Munday. "Enhanced nuclear fusion in the sub-keV energy regime." *Nature Communications*, 18 July 2026. DOI 10.1038/s41467-026-74421-1 — <https://www.nature.com/articles/s41467-026-74421-1>

Attwell & Laughlin. "An energy budget for signaling in the grey matter of the brain" / "Appraising the brain's energy budget," *PNAS & PMC* — <https://pmc.ncbi.nlm.nih.gov/articles/PMC124895/>

INFLECTION.

THE WEEKLY MAGAZINE OF INNOVATION



THE LENS · A QUESTION TO CARRY

Which part of your stack are you still calling “environment” — while it quietly does the work?

NEXT ISSUE · FRIDAY

Issue 03 turns to a different domain. Same method: one deep dive, three signals, and a reading the room has not arrived at yet.

COLOPHON

Researched, written & designed with Claude. Typeset in Poppins & Lora on the Anthropic palette. Issue No. 02 · 21 August 2026. Every figure in this issue is drawn from a cited primary source.