



IDEO-LAB

An Ideo-Lab publication

HUMAN BRAIN × ARTIFICIAL INTELLIGENCE

Visual Atlas — Synapses, Parameters, Attention, Context and Inference



AUTHOR

Guillaume O'Neill




VERSION

Version 1.0



DATE

15 August 2026



HUMAN BRAIN

x ARTIFICIAL INTELLIGENCE

Synapses, parameters, attention, context, and inference: a comparative anatomy of two radically different forms of computation.

VISUAL ATLAS



38 PAGES

2026 EDITION

OPENING

The wrong number?

Why counting connections is not enough to explain intelligence.

**BRAIN**

A living biological network: electrochemistry, plasticity, time, metabolism, body, and environment.

AI

A digital system: learned parameters, temporary activations, context, attention, inference, and tools.

KEY IDEA

The most fruitful comparison is not to convert synapses into parameters, but to study how each system dynamically organizes its resources.

SOURCE / REFERENCE Guiding principle of this atlas: functional analogy, never biological identity.

NAVIGATION

Four acts, one question

What turns an immense number of connections into useful intelligence?



THE BRAIN

numbers, synapses, energy, connectome



AI

Transformer, attention, context, inference



COMPARE

memory, learning, limits, scaling



ORCHESTRATE

systems, tools, adaptive compute, future

KEY IDEA


Almost any page should work on its own and immediately tell one idea, mechanism, or paradox.

ACT I - THE BRAIN

86 billion: a useful figure, not an odometer

The total number of neurons is an estimate. What matters even more is how they are distributed across regions and cell types.

Illustrative indices - units are not comparable



The chart shows why comparing raw numbers can be misleading.

The widely cited figure of about 86 billion comes from robust quantitative methods, but it should not be confused with an exact cell-by-cell inventory. More recent estimates remind us that a margin of uncertainty remains.

The brain is not a bag of identical neurons: the cerebellum concentrates an enormous number of cells, while the cortex deploys extraordinary connectivity and dendritic arborization.

KEY IDEA


The total tells us about scale. The distribution tells us far more about architecture.

SOURCE / REFERENCE Azevedo et al. (2009), J. Comp. Neurol.; later methodological debates.

ACT I - THE BRAIN

100 to 500 trillion synapses?

The best answer is a range, not a single number.



Estimates of the total number of synapses vary with brain region, age, local density, and the experimental definition of what counts as a synapse. A value on the order of 10^{14} is common; values several times higher remain plausible under different assumptions.

This uncertainty is not embarrassing: it reveals the nature of the system. Connections appear, disappear, strengthen, and weaken. The graph itself is alive.

KEY IDEA

Asking “how many synapses does the human brain have?” is asking for a snapshot of a network that is constantly changing.

SOURCE / REFERENCE Orders of magnitude from quantitative neuroanatomy and connectomics.

ACT I - THE BRAIN

An extremely heterogeneous synaptic density

Two regions of the same volume can have radically different architectures.

DIMENSION	BRAIN	AI
Substrate	cells + biochemistry	numerical matrices
State	continuous, living	temporary activations
Memory	integrated	weights + context + external
Adaptation	plasticity	training + orchestration

The cortex, cerebellum, hippocampus, and deep nuclei use neither the same neuronal types nor the same connection motifs. Purkinje cells illustrate this diversity with spectacular dendritic arborizations.

Function is therefore never determined by quantity alone. Wiring patterns, distances, projection direction, and local microcircuits can turn the same number of cells into very different computations.

KEY IDEA

Network topology is information in its own right.

SOURCE / REFERENCEClassical neuroanatomy; principles of regional connectivity.



ACT I - THE BRAIN

1 mm³ of cortex: 1.4 petabytes of data

Connectomics reveals the staggering density of brain tissue.



BIOLOGICAL



DIGITAL

A nanoscale reconstruction published in 2024 of roughly one cubic millimeter of human temporal cortex contained about 57,000 cells, 230 mm of blood vessels, and 150 million synapses. The dataset reached 1.4 petabytes.

This result gives a powerful intuition: a connectome is not merely a list of edges. Fine geometry, membranes, organelles, and contacts multiply the dimensions required to describe the tissue.

KEY IDEA


A fragment smaller than a grain of rice can already become a massive data problem.

SOURCE / REFERENCE Shapson-Coe et al., Science (2024), "A petavoxel fragment..."

ACT I - THE BRAIN

A synapse is a tiny machine

The word “connection” hides a dynamic biochemical system.



A synapse can modulate its release probability, the amount of neurotransmitter, the number and type of receptors, and its response to recent activity. It operates across multiple time scales.

That alone invalidates the simplistic equation “1 synapse = 1 parameter.” A parameter in an artificial network is a numerical coefficient; a synapse is a biological device whose state and behavior can evolve.

KEY IDEA

A synapse is closer to an adaptive micro-system than to a simple weighted wire.

SOURCE / REFERENCE Cellular neuroscience and synaptic plasticity.

ACT I - THE BRAIN

A neuron is not a transistor

Dendritic geometry participates in computation.



A neuron integrates signals in space and time. The location of a synapse on a dendrite, local ionic properties, and soma dynamics all change the response. A neuron is therefore not an abstract adder.

This richness explains why counting neurons and synapses does not directly yield a “computational capacity.” Part of the computation already happens inside the cell itself.

KEY IDEA


A biological neuron is an architecture, not a point.

SOURCE / REFERENCE Dendritic neurophysiology and neuronal integration.

ACT I - THE BRAIN

About 20 watts: the energy feat

The brain performs general cognition under a severe energy constraint.



The metabolic power of the human brain is commonly placed around 20 watts at rest. That budget supports cellular maintenance, membrane potentials, spikes, synaptic transmission, and recycling mechanisms.

A raw comparison with AI is difficult: computing infrastructures perform different operations and their consumption depends on hardware, model size, and workload. Yet the biological efficiency gap remains a major scientific question.

KEY IDEA

The brain cannot “burn” unlimited computation: energy shapes its architecture.


SOURCE / REFERENCE Classical references on cerebral metabolism; order of magnitude ~20 W.

ACT II - AI

The Transformer changes the comparison

Modern AI no longer tries to literally copy the brain.

Illustrative indices - units are not comparable



The chart shows why comparing raw numbers can be misleading.

The Transformer proposed in 2017 replaces dependence on recurrence with attention mechanisms that process relationships among elements of a sequence in a highly parallelizable way.

The historical importance of the Transformer is less that it resembles a brain than that it provides an efficient mechanism for building contextual relationships on demand.

KEY IDEA


The useful innovation was not better imitation of biology, but better organization of digital computation.

SOURCE / REFERENCE / aswani et al. (2017), "Attention Is All You Need".

ACT II - AI

Attention: creating relationships on demand

Each element can dynamically weight the relevance of other elements.



In attention, representations play the roles of queries, keys, and values. Similarities determine how information is aggregated. The same sequence can therefore produce different interaction patterns depending on its content.

This is where the analogy with “functional activation” becomes interesting: the parameters remain the same, but the path taken by information changes with context.

KEY IDEA

Learned weights provide structure; attention helps decide how to use that structure now.

SOURCE / REFERENCE / aswani et al. (2017) and Transformer literature.

ACT II - AI

Parameters are not explicit memories

Model knowledge is distributed across a vast numerical transformation.

DIMENSION	BRAIN	AI
Substrate	cells + biochemistry	numerical matrices
State	continuous, living	temporary activations
Memory	integrated	weights + context + external
Adaptation	plasticity	training + orchestration

A large model is not a database containing one record per concept. Learned information is distributed across many parameters and becomes expressed through the activations generated by an input.

This distribution explains both the power of generalization and the difficulty of interpretation. Asking exactly “where” a piece of knowledge is stored often has no simple answer.

KEY IDEA

A parameter is closer to a local constraint in a vast transformation space than to an individual memory.


SOURCE / REFERENCEGeneral concept from deep neural networks.




ACT II - AI

Context becomes working memory

Providing code or a document does not add new parameters.



BIOLOGICAL



DIGITAL

When a user provides files, code, or a conversation, that information becomes temporary context. The model can then compute relationships among elements that were not present in its weights as specific objects.

This is crucial for understanding recent progress: quality depends greatly on useful context length, the ability to retrieve the right information, and the stability of reasoning across that context.

KEY IDEA


Context is not the model's long-term memory; it is its workbench.

SOURCE / REFERENCE Extended-context language model architectures.

ACT II - AI

70,000 lines of code: what is the model really doing?

It looks for a compact representation of the software.



On a large project, reasoning line by line would be inefficient. The system gradually reconstructs a map: modules, functions, objects, dependencies, data flows, conventions, invariants, and risk areas.

That representation is necessarily imperfect, but it supports global questions: “if this function changes, what breaks?”, “where is this data produced?”, “which module owns this responsibility?”.

KEY IDEA

Effective analysis moves from raw reading to structural compression.

SOURCE / REFERENCEFunctional analogy with modern code analysis.

ACT II - AI

From code to a dependency graph

A software base becomes a network of objects and relationships.



Imports form a first graph. Function calls, data contracts, routes, SQL models, templates, and events form others. The challenge is to superimpose these graphs without losing useful causality.

This is exactly the kind of task where machines gain breadth: they can maintain many textual relationships at once and rapidly revisit the relevant regions.

KEY IDEA


Understanding software means reconstructing an approximate causal graph.

SOURCE / REFERENCE Static analysis, call graphs, and software architecture.


ACT II - AI

Inference becomes a process

For difficult problems, a direct answer is no longer enough.



Illustrative attention matrix



Solving can be decomposed into subproblems: identify constraints, propose a hypothesis, test it, search for contradictions, correct it, and only then formulate the answer.

Part of the improvement in inference engines comes from allocating intermediate computation, using verifiers, and calling tools when internal reasoning is not sufficient.

KEY IDEA

Quality depends on the computation path, not only on the model's final state.


SOURCE / REFERENCE Recent literature on reasoning and inference-time computation.

ACT II - AI

Test-time compute: thinking longer

The compute budget can be adapted to problem difficulty.

Illustrative indices - units are not comparable



The chart shows why comparing raw numbers can be misleading.

Recent work shows that allocating more computation at inference time can improve performance, especially when the system chooses how to spend that budget according to prompt difficulty.

This breaks with the view that all intelligence is “frozen” at training time. The model becomes a component that can search, sample, verify, and compare multiple trajectories.

KEY IDEA


The same model can deliver different performance levels depending on how much computation it is given.

SOURCE / REFERENCE Snell et al. (2024), scaling test-time compute.

ACT II - AI

Tools change the unit of measurement

The model is no longer necessarily the whole system.



A modern AI can consult a source, execute code, query a database, calculate an expression, or call a service. Final performance then depends on the interface between model and tools.

This architecture recalls a classic engineering principle: combining specialized components often achieves more than asking one component to be perfect at everything.

KEY IDEA

The relevant unit becomes the “AI system,” not just the “model.”

SOURCE / REFERENCE Tool-using systems and software agents.

ACT III - COMPARISON

Brain vs AI: five useful dimensions

Compare mechanisms, not slogans.

DIMENSION	BRAIN	AI
Substrate	cells + biochemistry	numerical matrices
State	continuous, living	temporary activations
Memory	integrated	weights + context + external
Adaptation	plasticity	training + orchestration

The brain has continuous plasticity, metabolic constraints, a body, sensory signals, and distributed memory. AI has numerical parameters, explicit contextual memory, parallel computation, and programmable external tools.

Some functional similarities are useful - activation, routing, working memory - but none justifies a biological identity.

KEY IDEA


The analogy is productive when it illuminates; it becomes dangerous when it erases differences.



ACT III - COMPARISON

Biological memory vs artificial memory

The same word refers to very different architectures.



BIOLOGICAL



DIGITAL

Human memory combines episodes, knowledge, skills, emotions, and bodily context. It consolidates, reconstructs, and forgets. In AI, weights, context, and external stores form separable memory layers.

That artificial separation can be an advantage: we can choose what persists, what remains temporary, and what should be retrieved from an external source.

KEY IDEA


Artificial memory is more modular; biological memory is more integrated.

SOURCE / REFERENCE Memory neuroscience; RAG and memory architectures.

ACT III - COMPARISON

Humans learn from remarkably few examples

Adaptation efficiency remains an impressive biological advantage.



A human can learn a rule, concept, or procedure from a few demonstrations and integrate it with prior experience. Training large models, by contrast, requires enormous amounts of data and computation.

Models compensate through breadth: once trained, they can reuse learned regularities across many tasks and exploit context without changing their weights.

KEY IDEA

The brain excels at frugal plasticity; AI excels at massive reuse of learned representations.

SOURCE / REFERENCE Conceptual comparison of human and machine learning.

ACT III - COMPARISON

Superhuman on some axes, not all

Machines already exceed individuals along several measurable dimensions.



Reading speed, documentary coverage, pattern search, generation of alternatives, exact calculation through tools, and simultaneous exploration of many paths are areas where an AI system can greatly exceed an isolated human.

But that breadth coexists with fragilities: misreading ambiguous context, excessive confidence, factual errors, or limited physical understanding of the world.

KEY IDEA


“Superhuman” should always be followed by the question: on which dimension, and under what conditions?

SOURCE / REFERENCEComparative synthesis.


ACT III - COMPARISON

Coherence is not truth

A plausible answer can still be wrong.



Illustrative attention matrix



Language models optimize the production of useful, coherent continuations according to their training and orchestration. That does not guarantee that every proposition corresponds to the real world.

Search tools, automated tests, calculators, and databases therefore play an essential role in turning a good intuition into a verified answer.

KEY IDEA

Even the most convincing reasoning remains subject to the test of facts.


SOURCE / REFERENCE General verification principle for generative systems.

ACT III - COMPARISON

Bigger is not always better

Raw size is only one term in the equation.

Illustrative indices - units are not comparable



The chart shows why comparing raw numbers can be misleading.

Scaling laws showed that performance depends on parameter count, data volume, and computation. Compute-optimal work then showed that an oversized but undertrained model wastes part of its potential.

The conclusion extends beyond AI: a resource has value only if the architecture knows how to feed, train, and use it effectively.

KEY IDEA


System optimization matters more than fascination with the largest number.

SOURCE / REFERENCE Kaplan et al. (2020); Hoffmann et al. (2022).

ACT III - COMPARISON

Chinchilla: balancing model and data

A smaller model can beat a much larger one.



Hoffmann et al. trained Chinchilla, a 70-billion-parameter model, on far more data than Gopher 280B under a comparable compute budget. Chinchilla outperformed much larger models on many evaluations.

That experiment changed the conversation: “parameter count” can no longer serve by itself as a proxy for capability.

KEY IDEA

Useful intelligence depends on compute allocation, not only on the size of the container.

SOURCE / REFERENCE Hoffmann et al. (2022), Training Compute-Optimal Large Language Models.

ACT IV - ORCHESTRATION

Intelligence as orchestration

The right computation, in the right place, at the right time.

DIMENSION	BRAIN	AI
Substrate	cells + biochemistry	numerical matrices
State	continuous, living	temporary activations
Memory	integrated	weights + context + external
Adaptation	plasticity	training + orchestration

In the brain, different regions participate in networks that reconfigure with the task. In modern AI, the model can likewise select contextual information, trigger tools, and devote more computation to particular stages.

The analogy is abstract but fruitful: in both cases, not all resources are mobilized in the same way at every moment.

KEY IDEA

Observable performance comes from the dynamic coordination of a set of resources.


SOURCE / REFERENCEConceptual synthesis: brain / AI systems.



ACT IV - ORCHESTRATION

Clusters rather than a monolith

The future looks like an ecosystem of components.



BIOLOGICAL

A central model can be surrounded by memory, document retrieval, code execution, databases, verifiers, simulators, and control policies. Each component contributes specialization.



DIGITAL

This shifts the frontier of AI: the question is no longer only “which model?” but “which system architecture?”.

KEY IDEA


The model becomes the cortex of a broader software system - a useful metaphor, but only a metaphor.

SOURCE / REFERENCE Architecture of tool-augmented AI systems.

ACT IV - ORCHESTRATION

Adaptive compute changes the feeling of power

A simple task and a difficult task should not cost the same effort.



When a system learns to recognize problem difficulty, it can accelerate easy cases while reserving more search, verification, or computation for hard ones. This improves overall efficiency.

This is also what the user feels: not merely a “more knowledgeable answer,” but an ability to keep track of a complex problem, revisit a hypothesis, and consolidate the solution.

KEY IDEA

The feeling of reasoning comes partly from intelligent management of the compute budget.

SOURCE / REFERENCE Literature on inference-time scaling and verification systems.

ACT IV - ORCHESTRATION

Emergence: real discontinuity or measurement threshold?

Some abilities seem to appear suddenly as scale increases.



The word “emergence” is appealing, but it can hide two different phenomena: a genuine regime transition, or continuous improvement that suddenly crosses a benchmark threshold.

To understand progress, we therefore need to examine curves, metrics, system composition, and protocol changes - not only spectacular anecdotes.

KEY IDEA


A capability may seem to appear all at once because our measuring instrument is discrete.

SOURCE / REFERENCE Scientific debates on emergent abilities.

ACT IV - ORCHESTRATION

The deeper analogy: a temporary functional map

Brain and AI both construct, by different means, a task-relevant state.



In the brain, active networks vary with perception, memory, and goals. In a Transformer, activations and attention patterns depend on the present context.

We can therefore speak of a “temporary functional map,” provided we do not confuse it with a biological structure. That map links the elements needed for the solution.

KEY IDEA

The useful comparison concerns coordination dynamics, not physical substrate.


SOURCE / REFERENCEFunctional analogy, not biological identity.

ACT IV - ORCHESTRATION

What if we have been watching the wrong number?

The number of connections fascinates; the organization of computation explains more.

Illustrative indices - units are not comparable



The chart shows why comparing raw numbers can be misleading.

One hundred or five hundred trillion synapses: the scale is staggering. But a more connected brain is not automatically more intelligent, just as a model with more parameters is not automatically better.

The most fruitful questions concern topology, plasticity, energy, memory, routing, adaptive computation, and the ability to integrate new tools or new experiences.

KEY IDEA


Size is a condition of possibility; orchestration turns that possibility into observable intelligence.

SOURCE / REFERENCE Synthesis conclusion.

ACT IV - ORCHESTRATION

The question that remains open

How far can a system go if it learns to mobilize its own resources more effectively?



Future breakthroughs may come as much from system architecture as from model size: more reliable memories, verifiers, planning, simulation, specialized tools, better perception of the world, and adaptive computation.

The brain comparison will remain valuable if it stays disciplined: seek principles of organization without forcing an equivalence between biology and computing.

KEY IDEA

The fascinating question is not when AI will “have as many synapses as we do.” That statement probably has no scientific meaning.

SOURCE / REFERENCE Research perspective.

SYNTHESIS

Brain vs AI: dashboard

Ten dimensions for comparison without confusion.

DIMENSION	BRAIN	AI
Substrate	cells, membranes, chemistry	transistors, memory, matrix computation
Connection	dynamic synapse	numerical parameter
Short-term state	neural activity	activations + context
Memory	episodic, semantic, procedural	weights, context, external storage
Learning	continuous plasticity	training, adaptation, memory
Energy	~20 W order of magnitude	varies with hardware and load
Time	multi-scale	digital compute cycles
Body	embodied	often external to the model
Tools	hands, language, culture	APIs, code, search, databases
Error	bias, forgetting, illusion	hallucination, poor generalization

KEY IDEA

The table shows one thing above all: identical words often describe different mechanisms. The analogy must remain functional.

SOURCE / REFERENCE Synthesis of concepts from neuroscience and AI system architecture.



MYTHS

Seven shortcuts to avoid

Seductive phrases that make the brain-AI comparison false or useless.

1

1 synapse = 1 parameter

False: a synapse is a dynamic biochemical system.

2

More parameters = more intelligent

False: data, compute, and architecture matter.

3

Context = human memory

No: it is temporary digital working memory.

4

Attention = consciousness

No: Transformer attention is a mathematical operation.

5

86 billion = exact value

No: it is a reference estimate.

6

AI only recites

Too simplistic: it recombines representations and can reason through tools.

7

Superhuman = superior everywhere

No: superiority is dimensional and conditional.

SOURCE / REFERENCE These myths summarize the conceptual precautions developed in the previous chapters.

GLOSSARY

Ten words, two worlds

The same vocabulary can hide radically different mechanisms.

Synapse

plastic cellular junction

Parameter

learned coefficient

Activation

temporary state

Context

information available now

Attention

weighted digital routing

Plasticity

lasting biological change

Inference

computation of an output

Connectome

map of connections

Memory

family of mechanisms

Orchestration

coordination of resources

SOURCE / REFERENCE Conceptual glossary, intentionally concise.

REFERENCES

A few scientific anchors

A short selection for further reading, not a systematic review.

Azevedo et al. (2009)

Equal numbers of neuronal and nonneuronal cells make the human brain an isometrically scaled-up primate brain. *Journal of Comparative Neurology*.

Shapson-Coe et al. (2024)

A petavoxel fragment of human cerebral cortex reconstructed at nanoscale resolution. *Science*.

Vaswani et al. (2017)

Attention Is All You Need. *arXiv:1706.03762*.

Kaplan et al. (2020)

Scaling Laws for Neural Language Models. *arXiv:2001.08361*.

Hoffmann et al. (2022)

Training Compute-Optimal Large Language Models. *arXiv:2203.15556*.


Snell et al. (2024)

Scaling LLM Test-Time Compute Optimally can be More Effective than Scaling Model Parameters. *arXiv:2408.03314*.

Connectomics and neurophysiology

Reviews of quantitative neuroanatomy, synaptic plasticity, cerebral metabolism, and functional connectivity.

SOURCE / REFERENCE The references were selected to anchor the key figures and structural concepts of the atlas.



CONCLUSION

Perhaps we have been watching the wrong number for a long time.

The brain is not extraordinary because it has exactly one hundred, three hundred, or five hundred trillion synapses. It is extraordinary because it turns a relevant fraction of that immense network into perception, memory, hypothesis, decision, and action.

And the most fascinating question for AI is probably not: “when will it have as many synapses as we do?” The real question is: how far can a system go when it learns to mobilize its resources intelligently at the moment they are needed?