Building brains for the future:
The neuroscience school leaders need now

Barbara Oakley, PhD, PE, FAAAS, FIEEE, FAIMBE
Coursera's Inaugural Innovation Instructor
Winner of McGraw Prize in Education (2023)
Distinguished Professor of Engineering
Oakland University, Rochester, Michigan
Oakley@oakland.edu

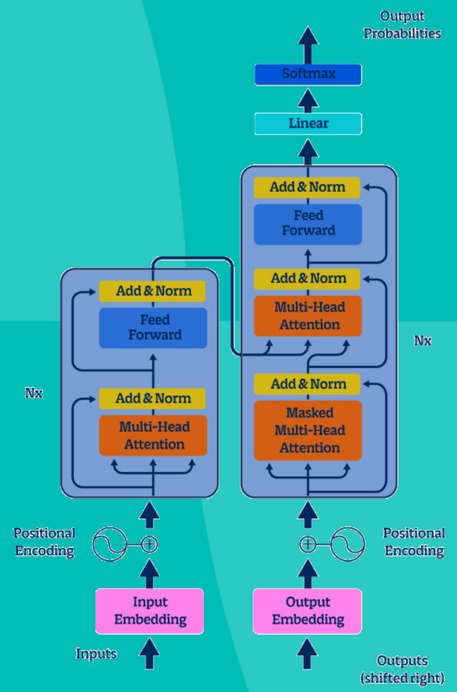


CANADA
Edmonton
Vancouver
Regina
Winnipeg
Québec
Montréal
Ottawa
Toronto
Seattle
Olympia
Washington
Salem
Oregon
Montana
Helena
Idaho
Boise
North Dakota
Bismarck
South Dakota
Pierre
Minnesota
St. Paul
Wisconsin
Madison
Michigan
Lansing
Detroit
Wyoming
Cheyenne
Nebraska
Lincoln
Iowa
Des Moines
Chicago
Illinois
Springfield
Indiana
Indianapolis
Ohio
Columbus
Sacramento
San Francisco
Carson City
Nevada
Salt Lake City
Utah
Denver
Colorado
Topeka
Kansas
St. Louis
Jefferson City
Missouri
California
Las Vegas
Los Angeles
Arizona
Phoenix
Santa Fe
New Mexico
Oklahoma
Oklahoma City
Arkansas
Little Rock
Memphis
Nashville
Tennessee
Kentucky
Frankfort
West Virginia
Charleston
Virginia
Richmond
Raleigh
North Carolina
Columbia
South Carolina
Atlanta
Georgia
Mississippi
Jackson
Montgomery
Alabama
Dallas
Texas
Austin
Houston
Louisiana
Baton Rouge
New Orleans
Tallahassee
Florida
Miami
Maine
Augusta
Montpelier
Vt.
N.H.
Concord
Albany
New York
Boston
Mass.
Providence
R.I.
Hartford
Conn.
New York
Trenton
New Jersey
Pennsylvania
Philadelphia
Harrisburg
Dover
Delaware
Md.
Annapolis
Washington, D.C.
North Pacific Ocean
North Atlantic Ocean
Gulf of Mexico
Mexicali
Hermosillo
Chihuahua
MEXICO
Monterrey
THE BAHAMAS
Nassau
National capital
U.S. State
State capital
Scale 1:13,500,000
Albers Equal-Area Projection
0
500 Kilometers
0
500 Miles
130
120
110
100
90
80
70
50
40
30




Philip Oakley

The transformer

Attention Is All You Need

Ashish Vaswani* Google Brain avaswani@google.com

Noam Shazeer* Google Brain noam@google.com

Niki Parmar* Google Research nikip@google.com

Jakob Uszkoreit* Google Research usz@google.com

Llion Jones* Google Research llion@google.com

Aidan N. Gomez* † University of Toronto aidan@cs.toronto.edu

Łukasz Kaiser* Google Brain lukaszkaiser@google.com

Illia Polosukhin* ‡ illia.polosukhin@gmail.com

Abstract

The dominant sequence transduction models are based on complex recurrent or convolutional neural networks that include an encoder and a decoder. The best performing models also connect the encoder and decoder through an attention mechanism. We propose a new simple network architecture, the Transformer, based solely on attention mechanisms, dispensing with recurrence and convolutions entirely. Experiments on two machine translation tasks show these models to be superior in quality while being more parallelizable and requiring significantly less time to train. Our model achieves 28.4 BLEU on the WMT 2014 English-to-German translation task, improving over the existing best results, including ensembles, by over 2 BLEU. On the WMT 2014 English-to-French translation task, our model establishes a new single-model state-of-the-art BLEU score of 41.0 after training for 3.5 days on eight GPUs, a small fraction of the training costs of the best models from the literature.

31st Conference on Neural Information Processing Systems (NIPS 2017), Long Beach, CA, USA.

Encoder

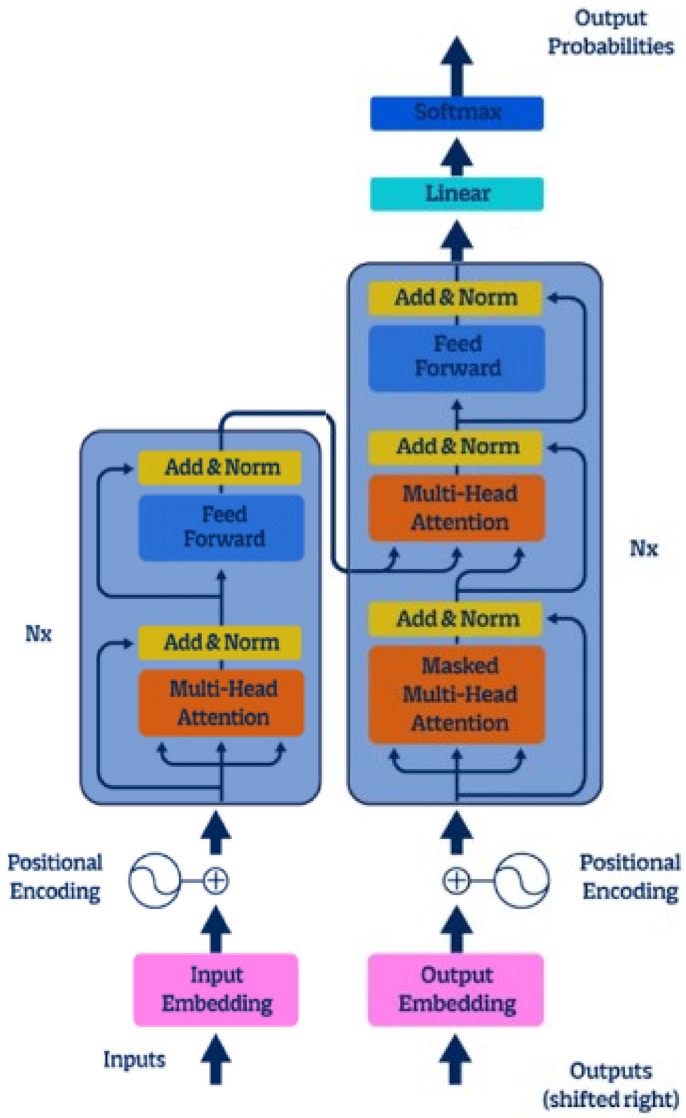

Decoder

Flow of information through transformers in AI

Flow of information in the brain

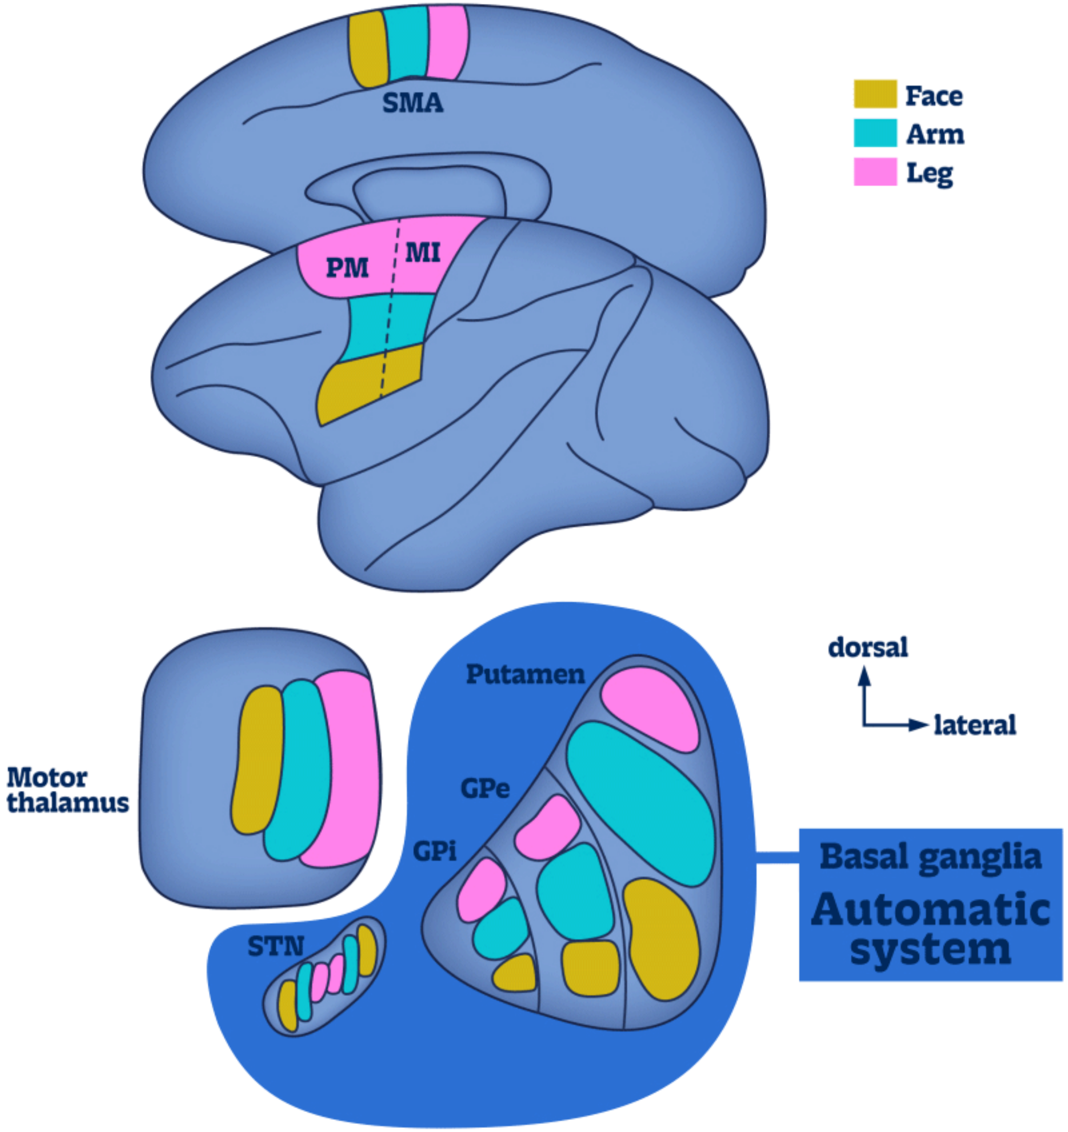

Axon
Dendritic spines

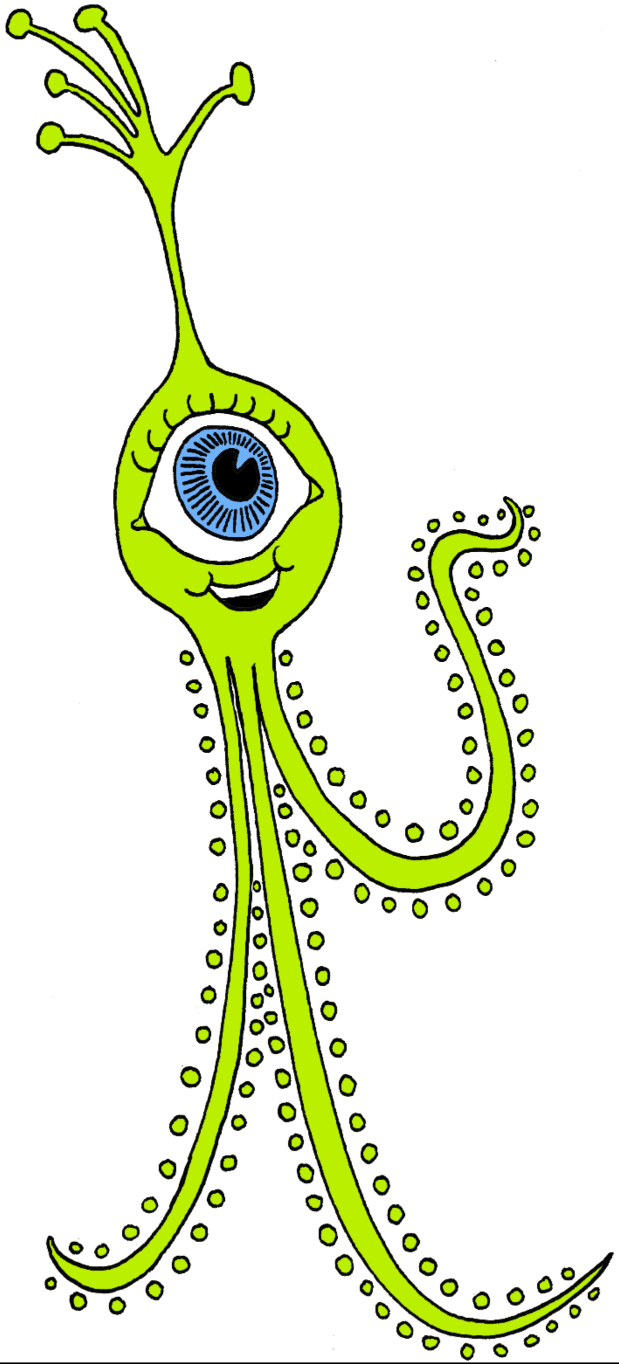

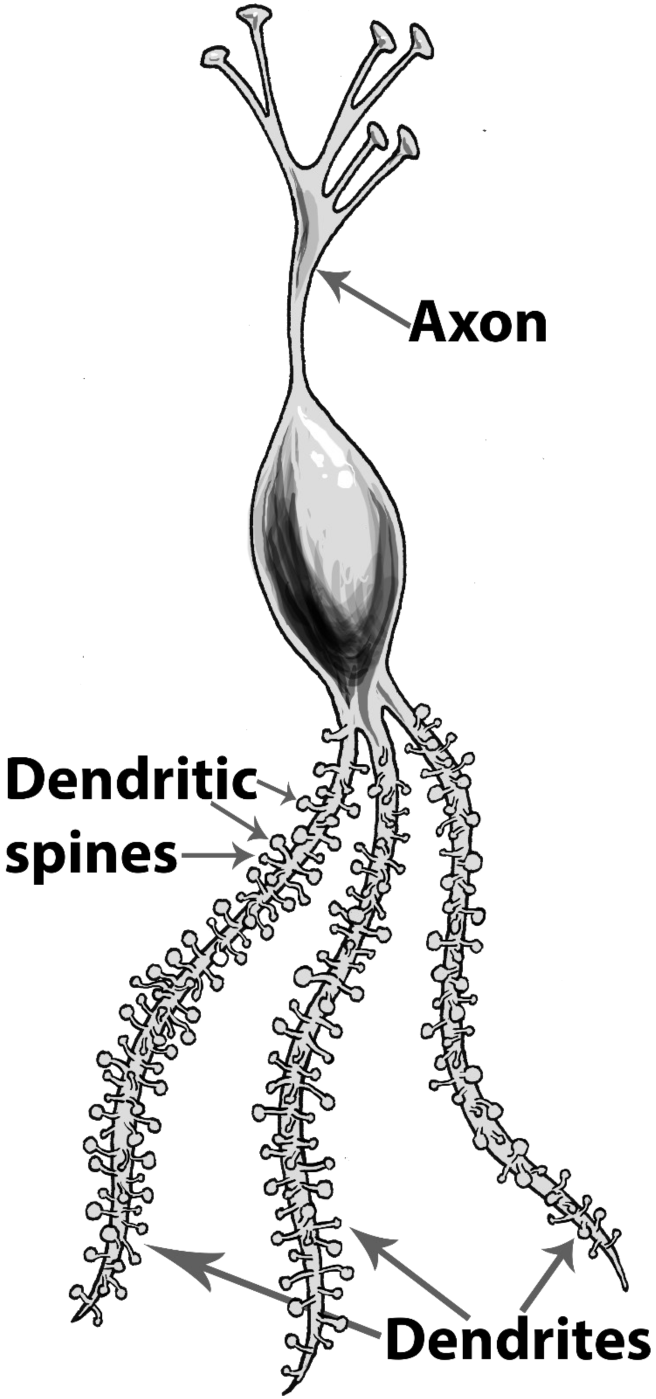
Axon
Dendritic
spines
Dendrites

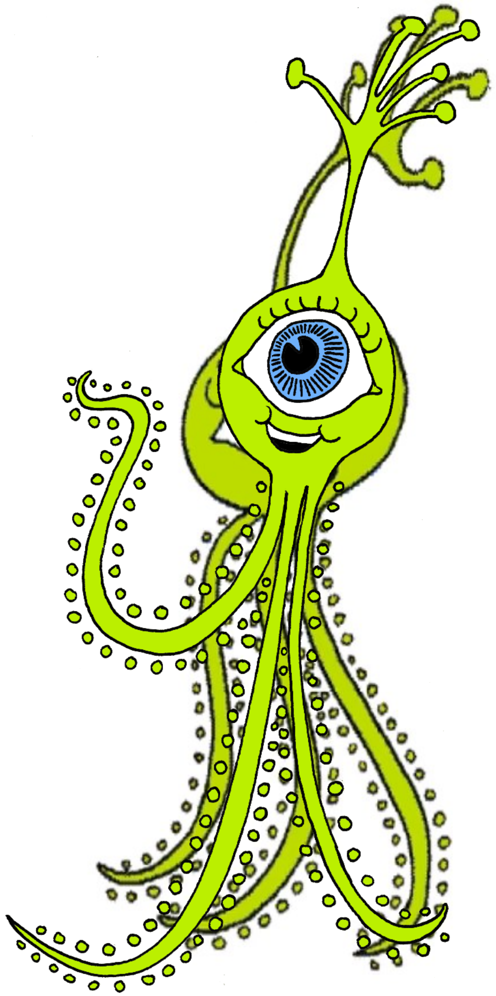

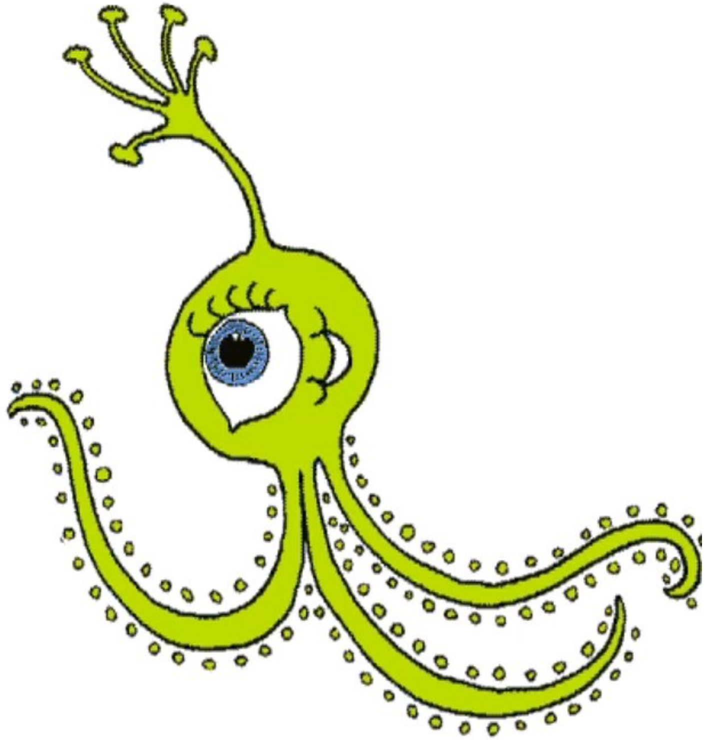

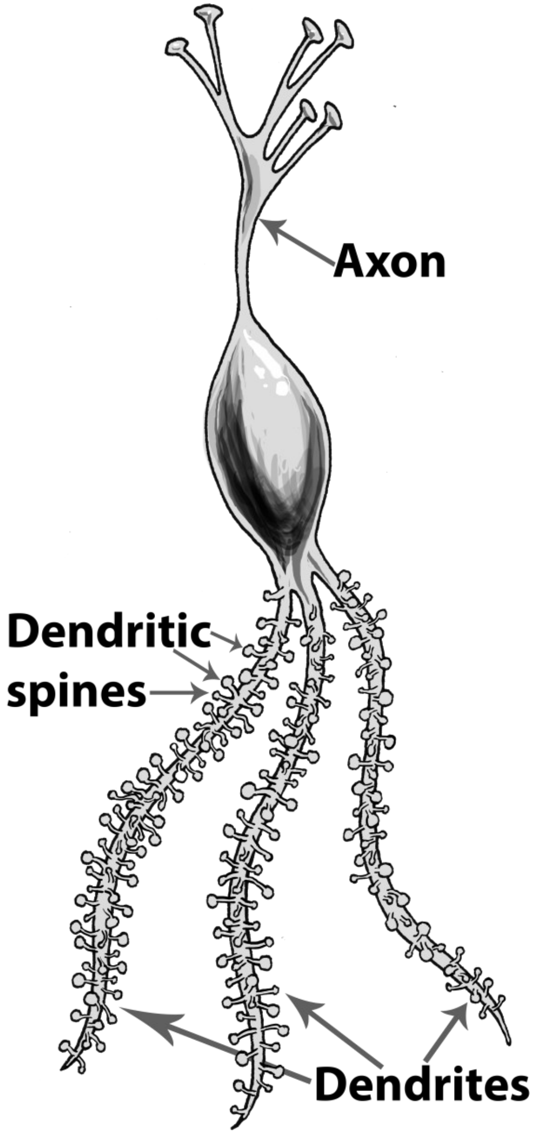
Axon
Dendritic
spines
Dendrites

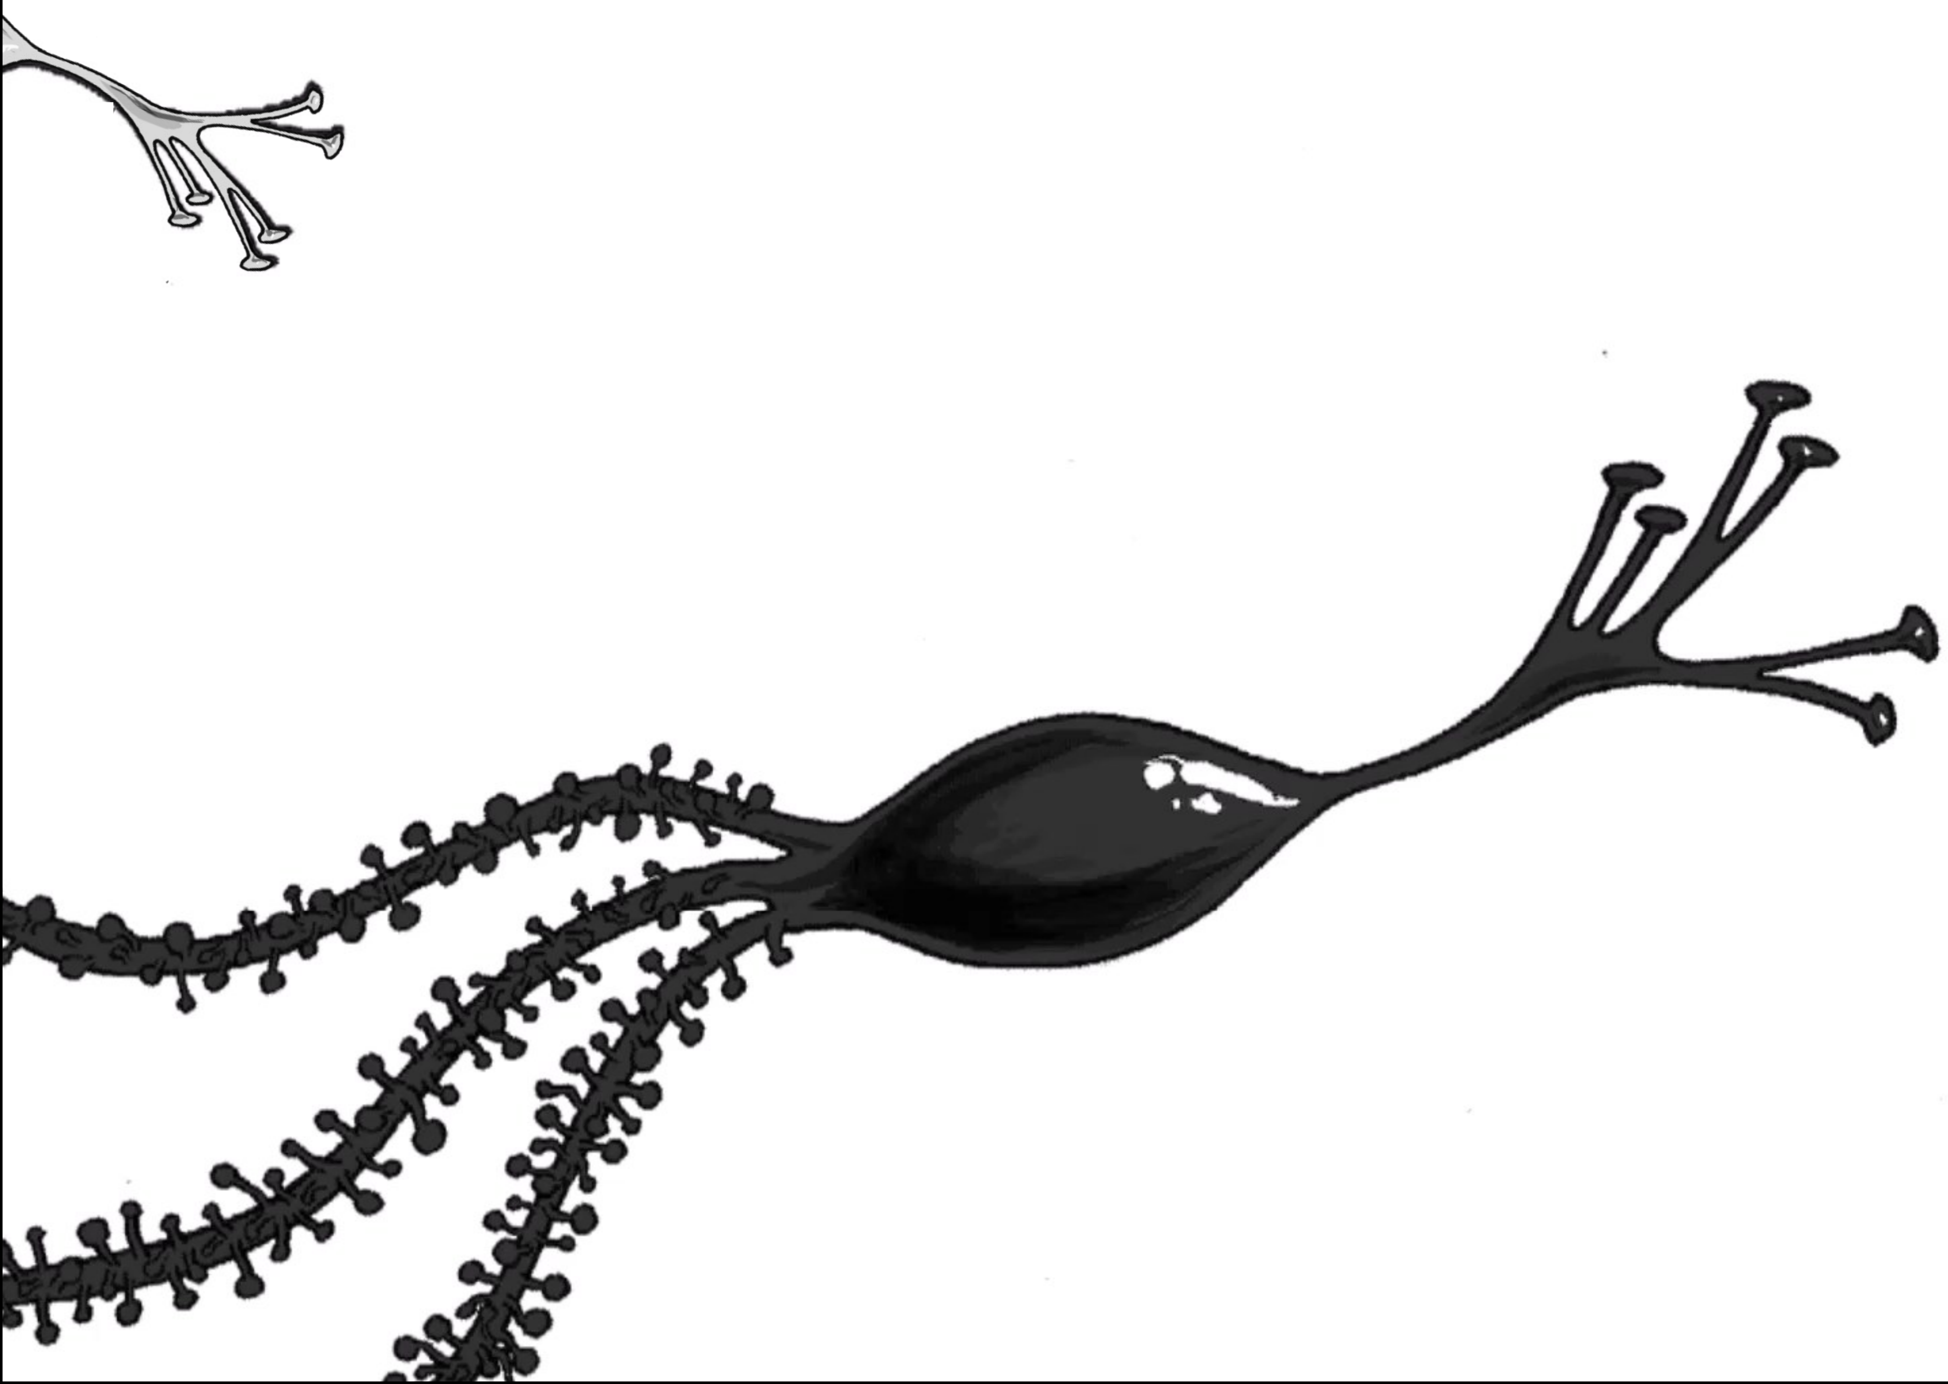

Neurons
send
signals

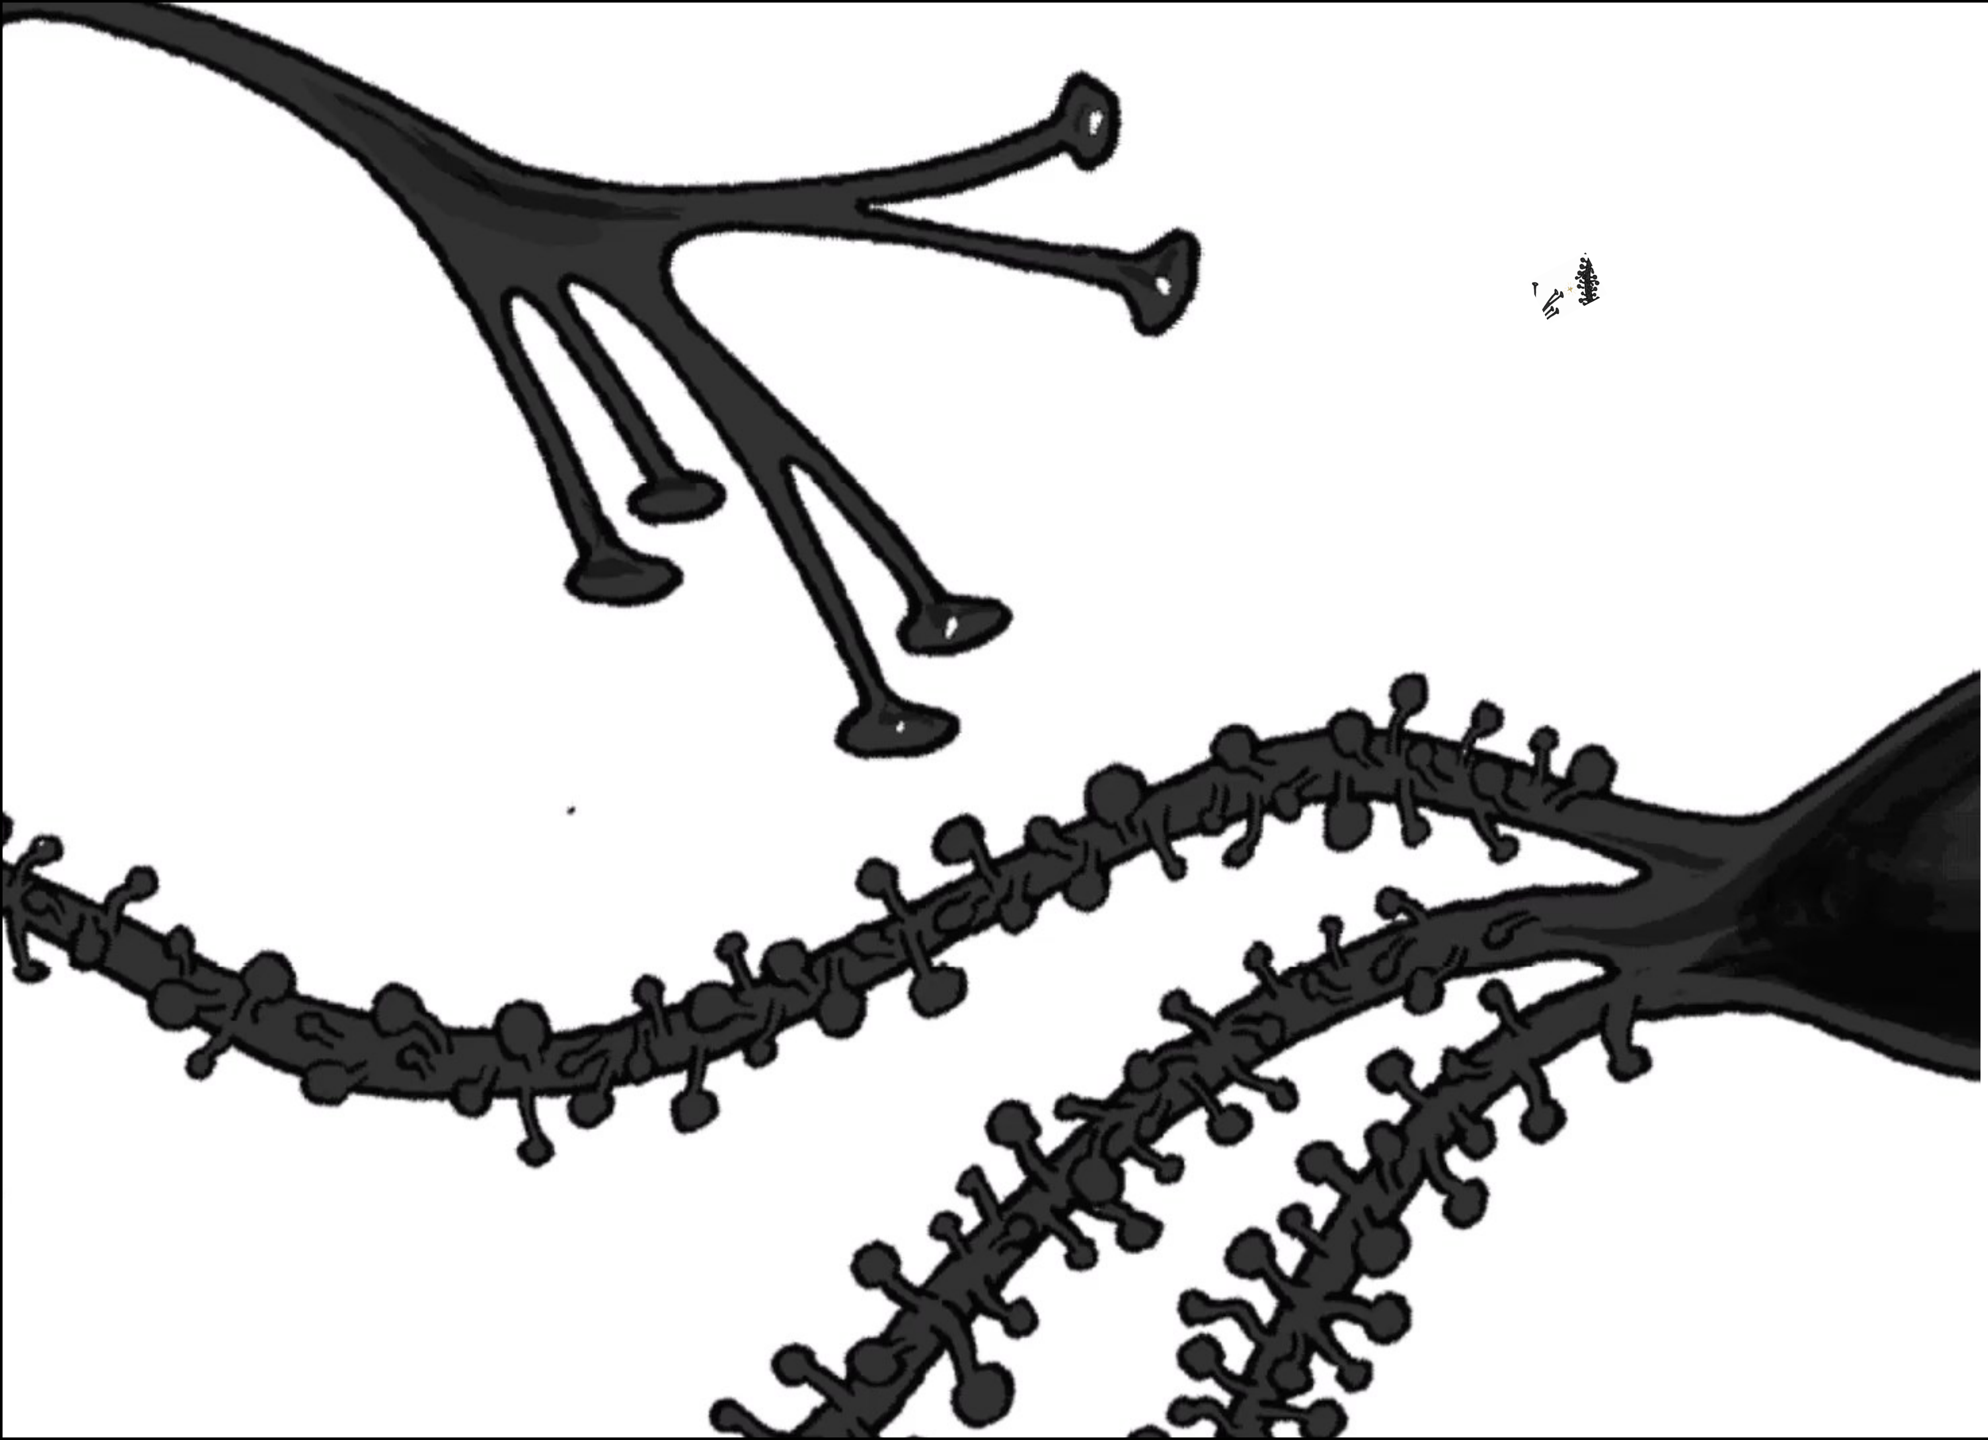

Neurons create sets of links when you learn something.

Connections strengthen with practice

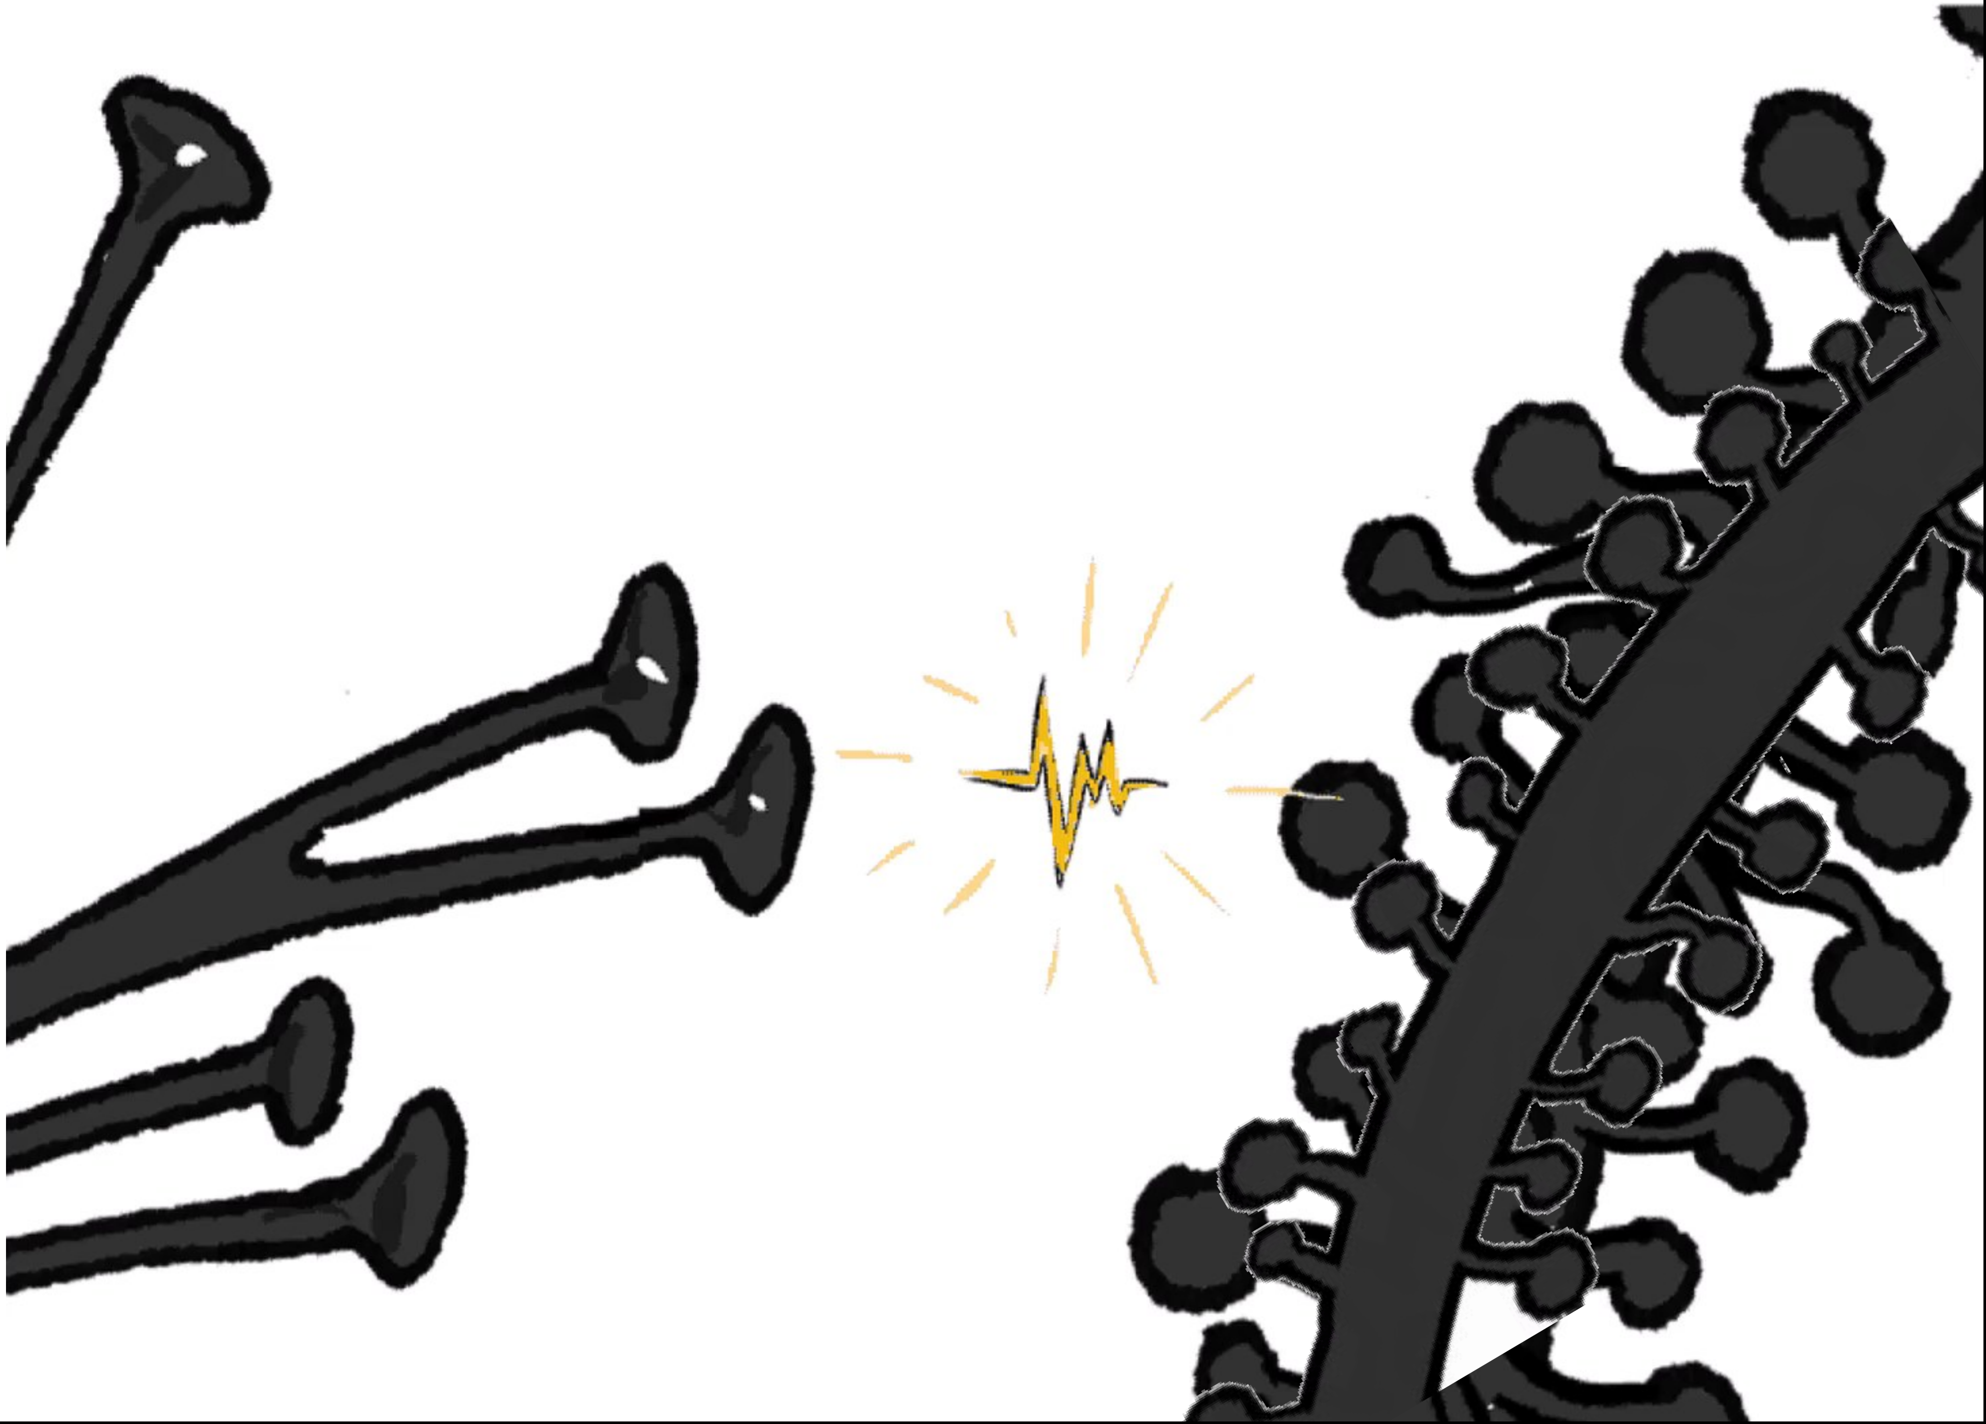

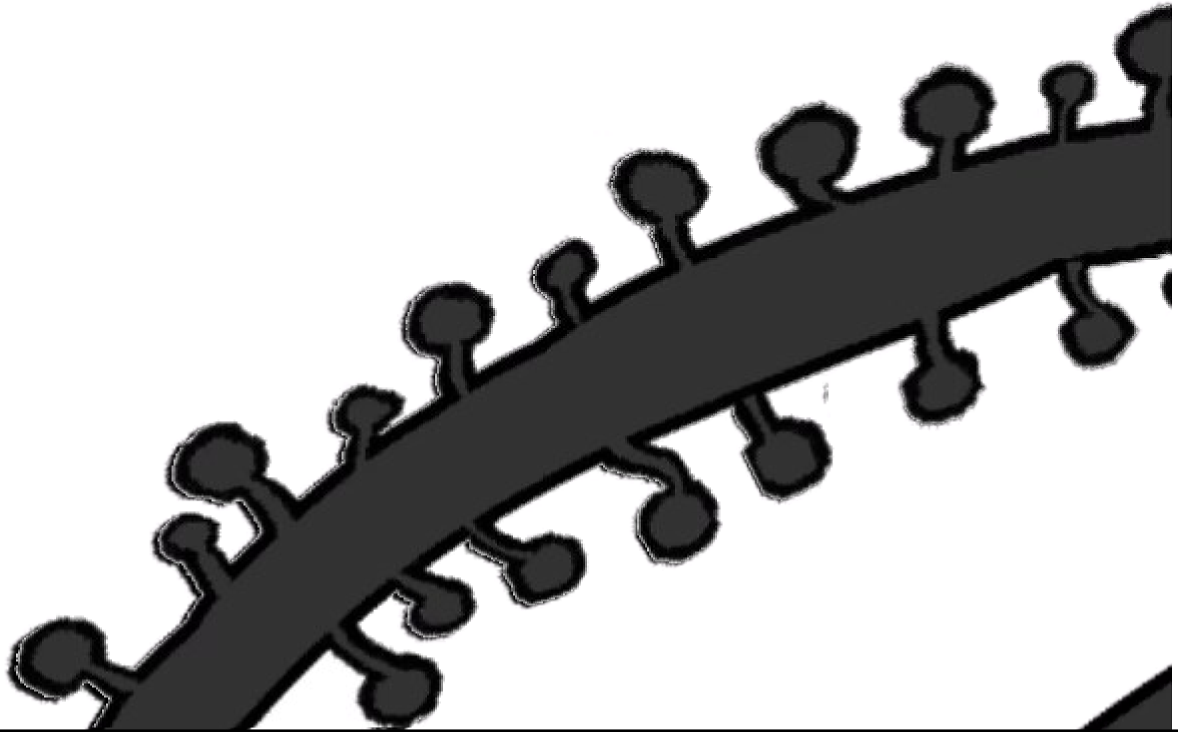

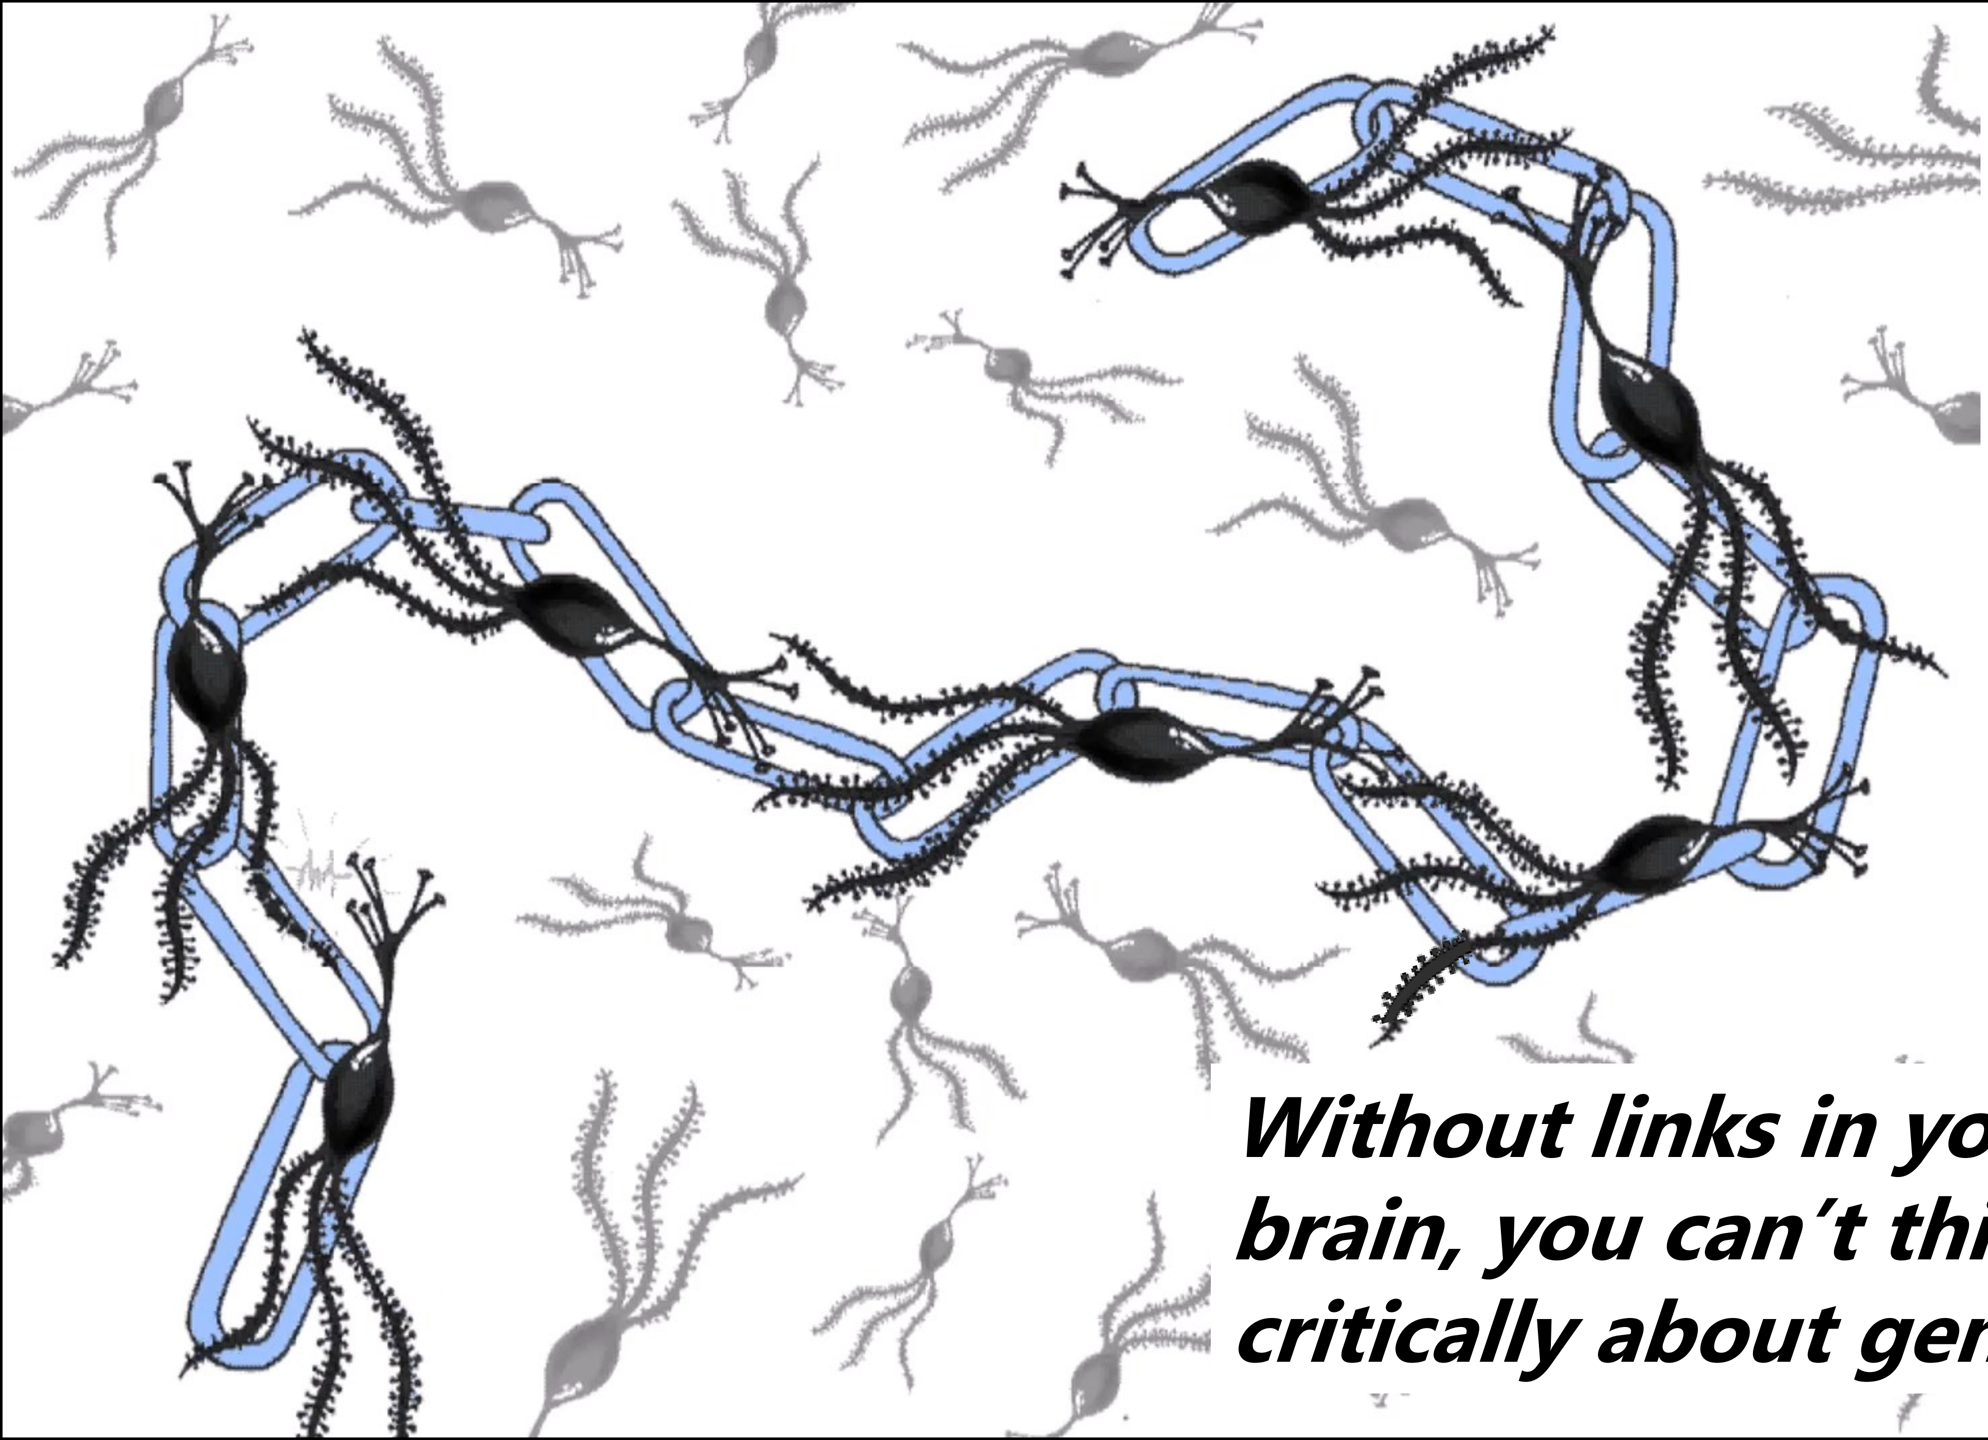

Without links in your own brain, you can't think critically about genAI!

RAY M. DOLBY, 1957 · AGE 24

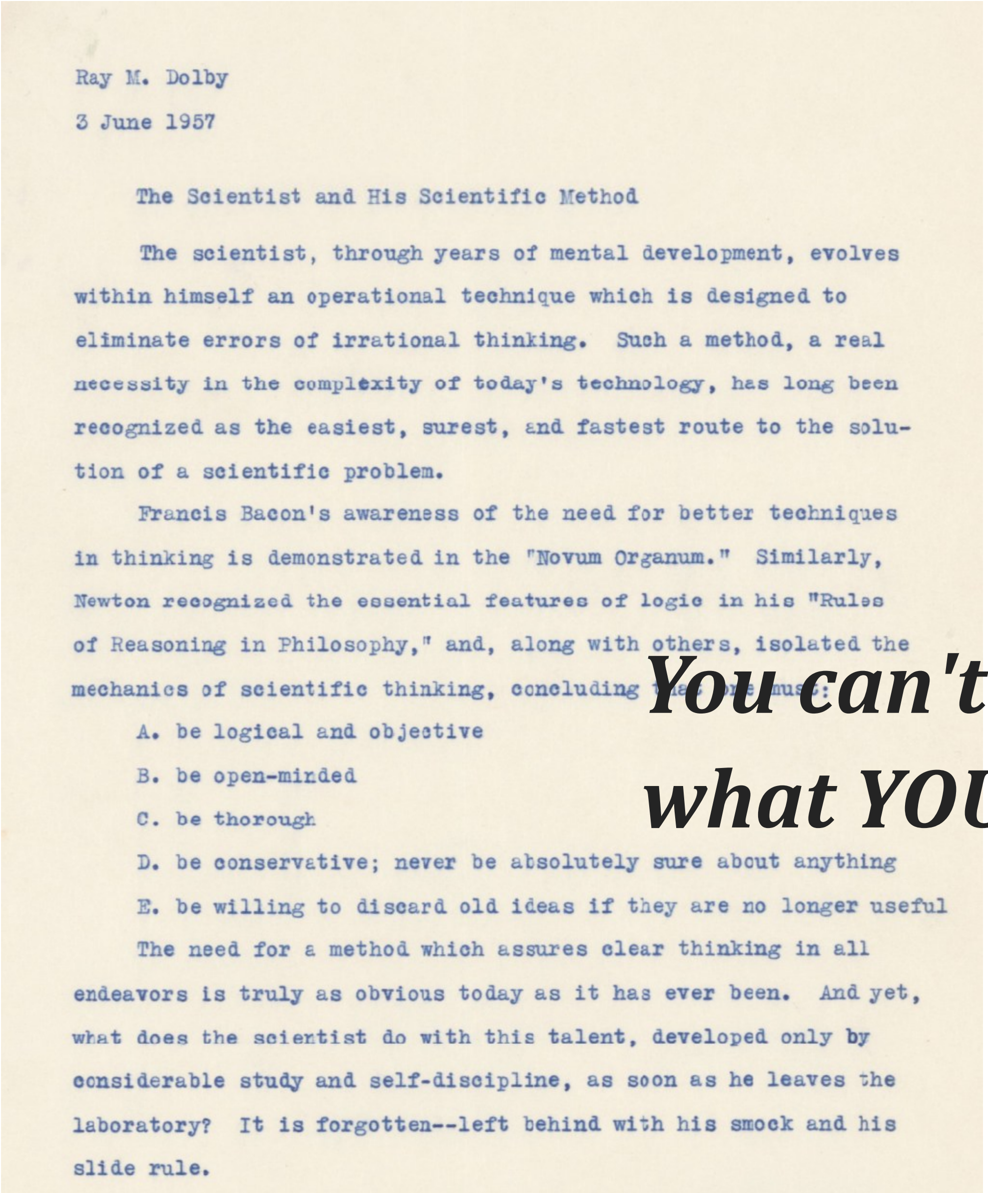

Ray M. Dolby
3 June 1957

The Scientist and His Scientific Method

The scientist, through years of mental development, evolves within himself an operational technique which is designed to eliminate errors of irrational thinking. Such a method, a real necessity in the complexity of today's technology, has long been recognized as the easiest, surest, and fastest route to the solution of a scientific problem.

Francis Bacon's awareness of the need for better techniques in thinking is demonstrated in the "Novum Organum." Similarly, Newton recognized the essential features of logic in his "Rules of Reasoning in Philosophy," and, along with others, isolated the mechanics of scientific thinking, concluding that one must:

A. be logical and objective

B. be open-minded

C. be thorough

D. be conservative; never be absolutely sure about anything

E. be willing to discard old ideas if they are no longer useful

The need for a method which assures clear thinking in all endeavors is truly as obvious today as it has ever been. And yet, what does the scientist do with this talent, developed only by considerable study and self-discipline, as soon as he leaves the laboratory? It is forgotten--left behind with his smock and his slide rule.

You can't reason well about what YOU don't yet know well.

Without links in your own brain, you can't think critically about genAI!

Stanford | University Libraries

Archival Collections at Stanford

Ray Dolby papers, [c. 1940s - 2000s]

What is the most powerful technique to help students learn most efficiently?

- Reread
- Highlight or underline
- Retrieval practice ("recall")
- Create a concept map

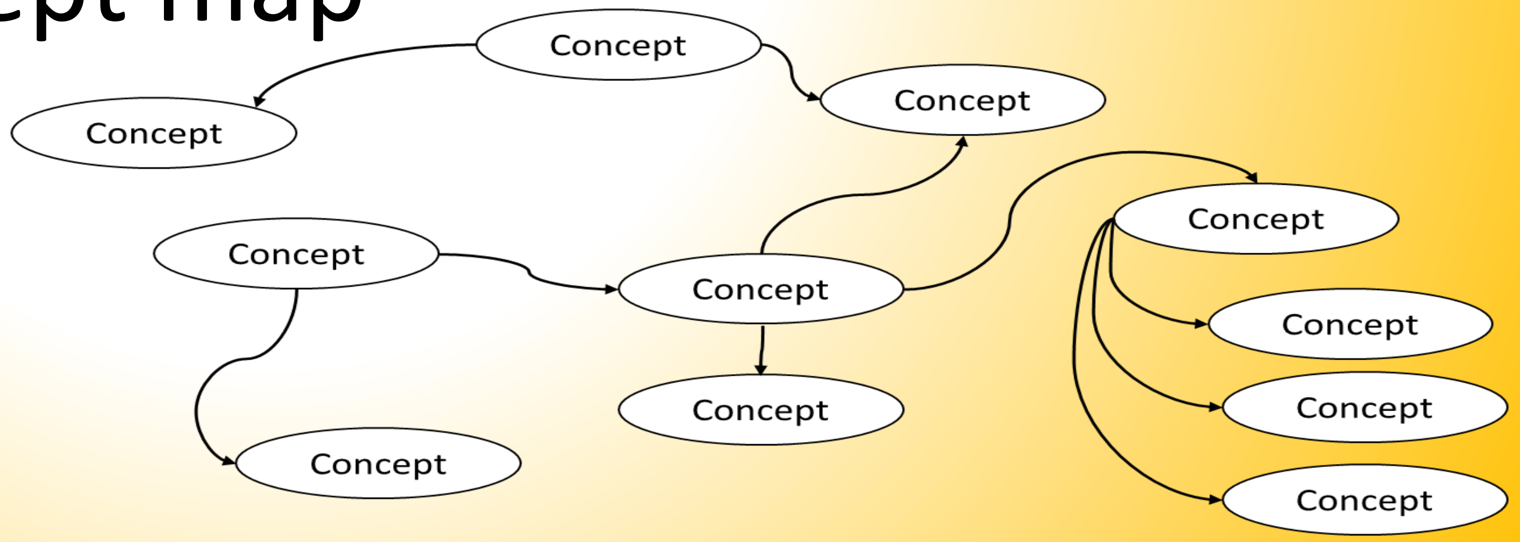

Retrieval Practice

Retrieval Practice

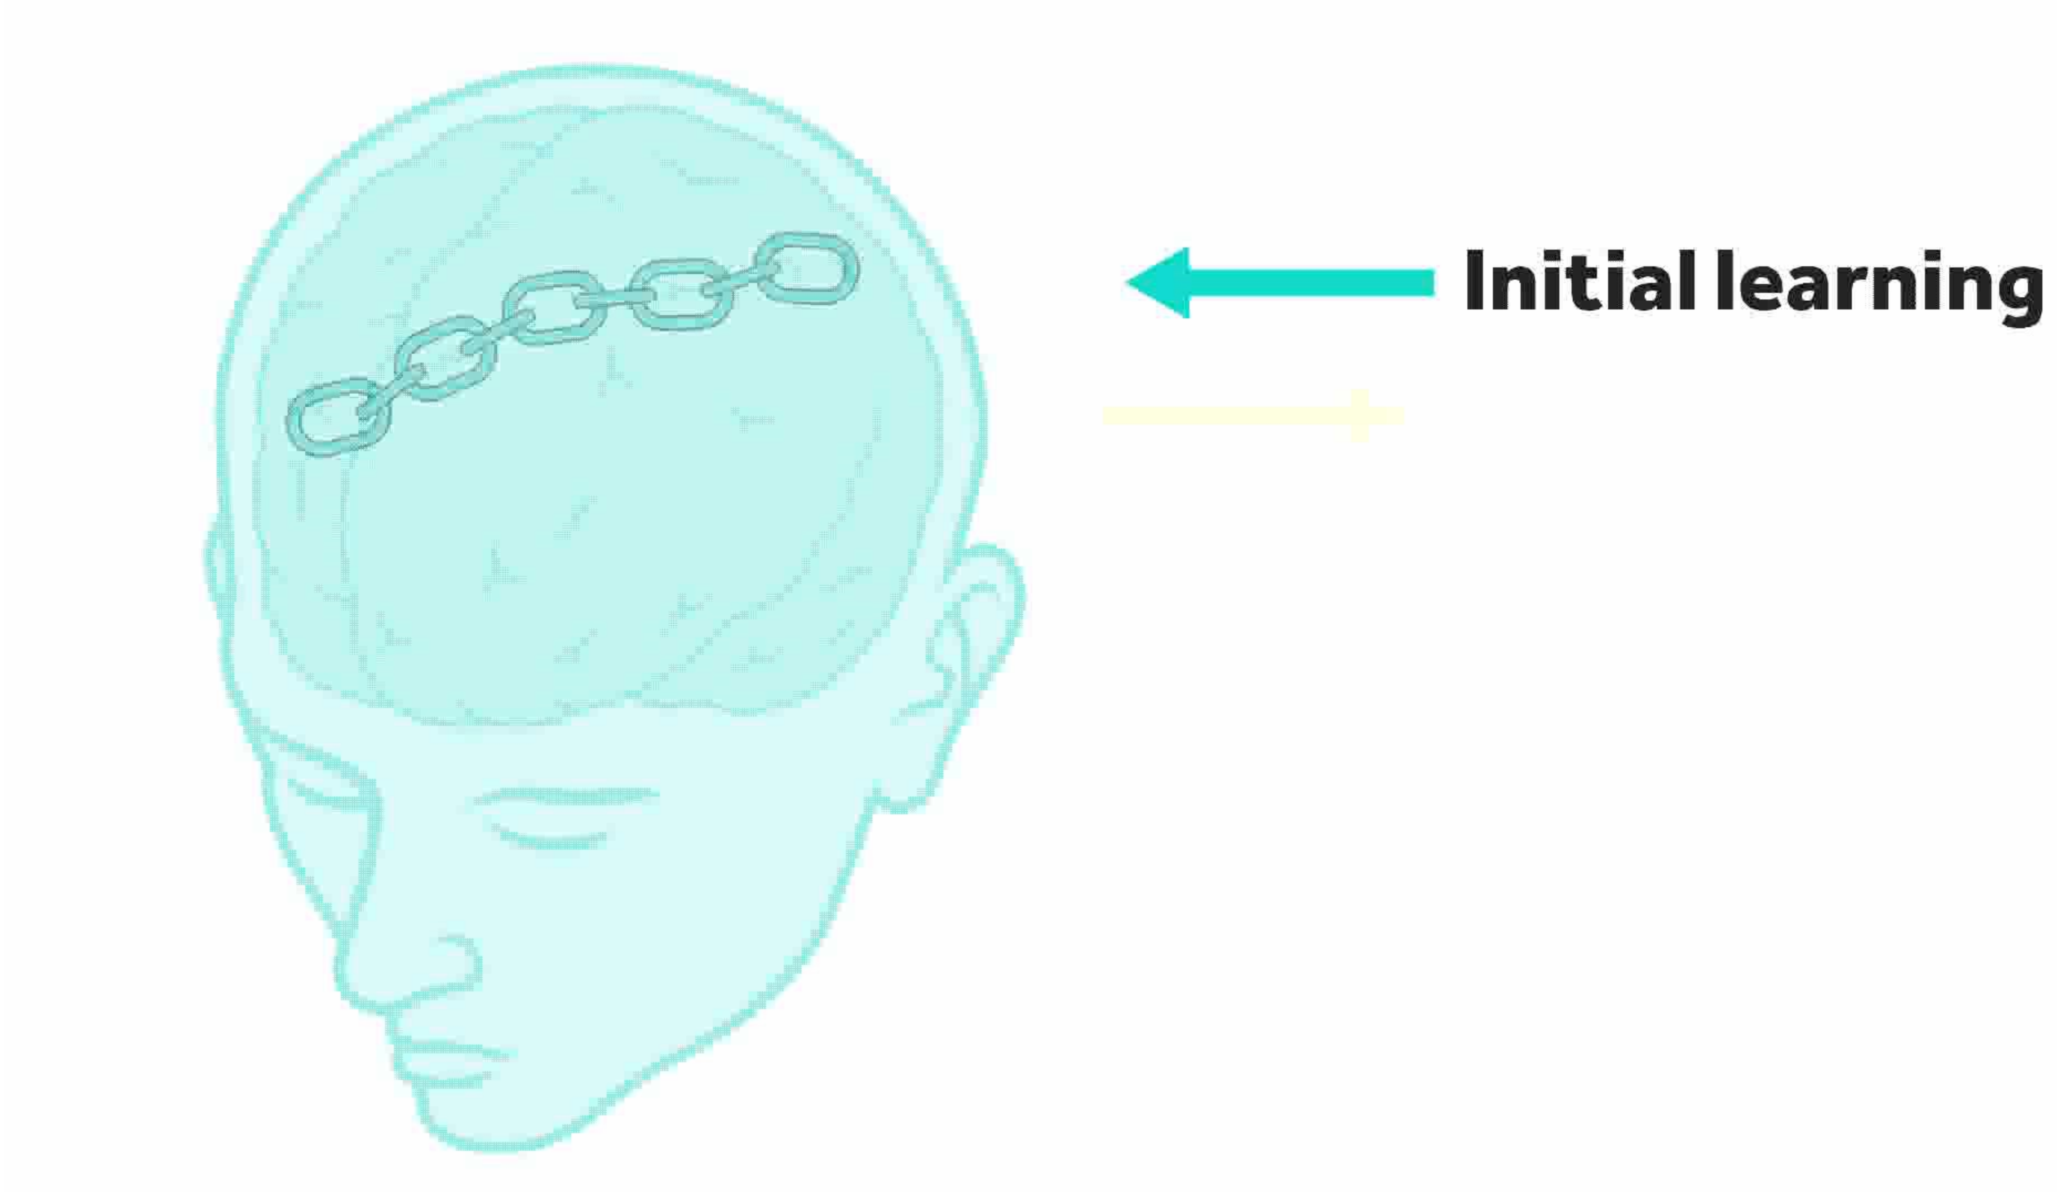

Mapping

. Blunt

activities that encourage elaborative studyin
ractice retrieving and reconstructing knowle
cticing retrieval produces greater gains in m
cept mapping. The advantage of retrieval pra

AI shows *But I’m the only one who KNOWS!*

Large Language Model

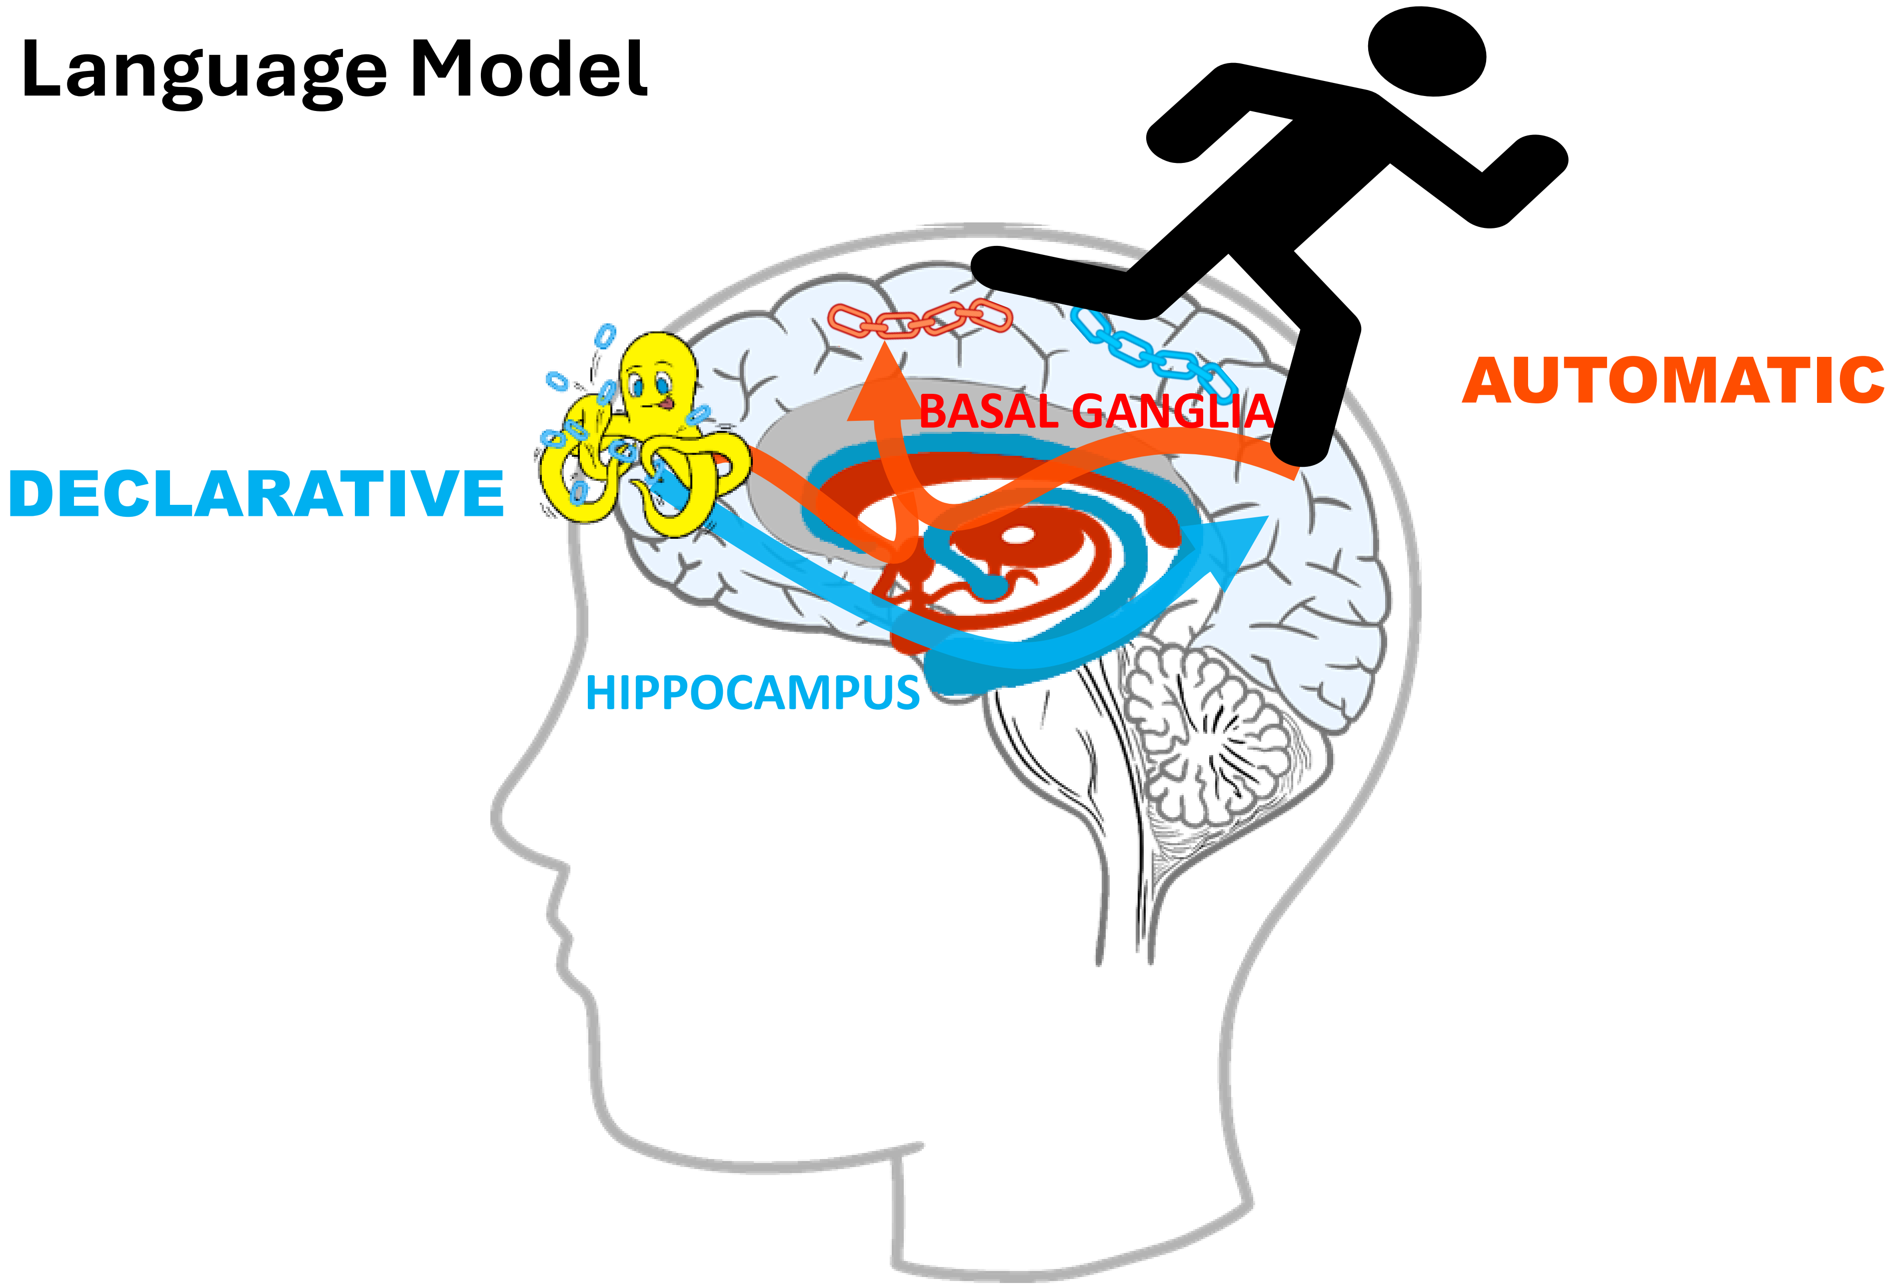

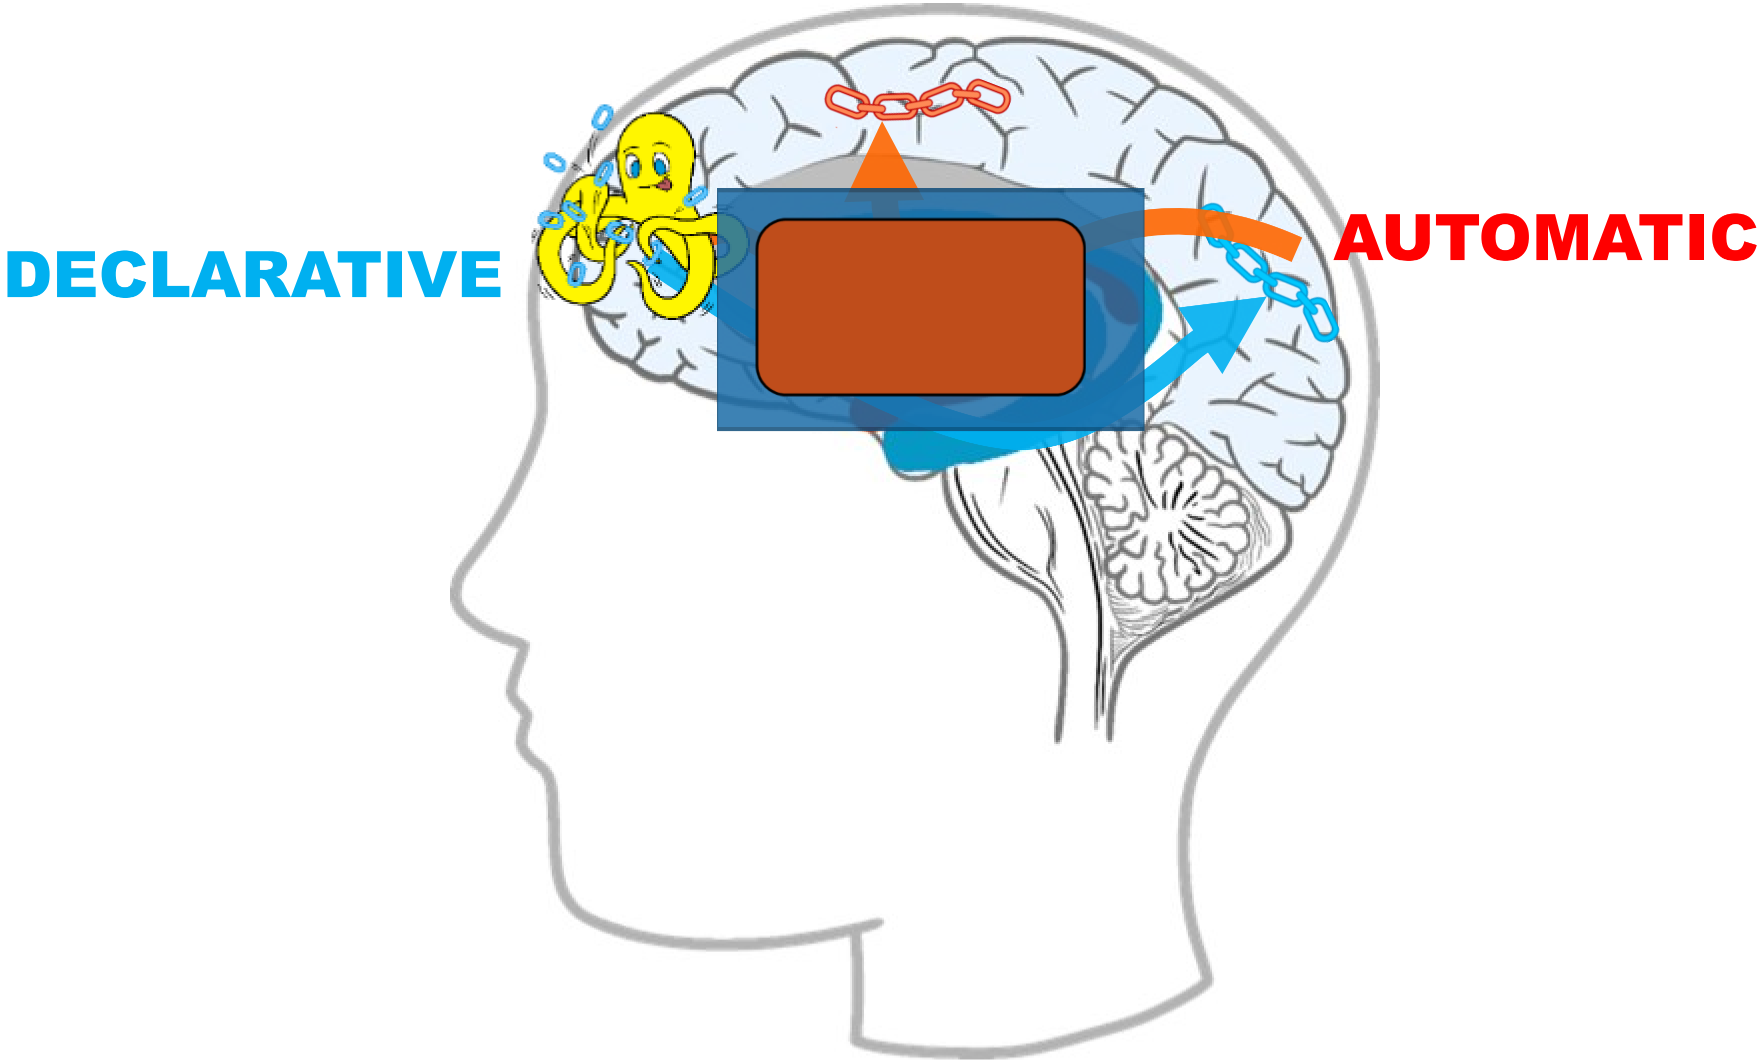
DECLARATIVE
AUTOMATIC

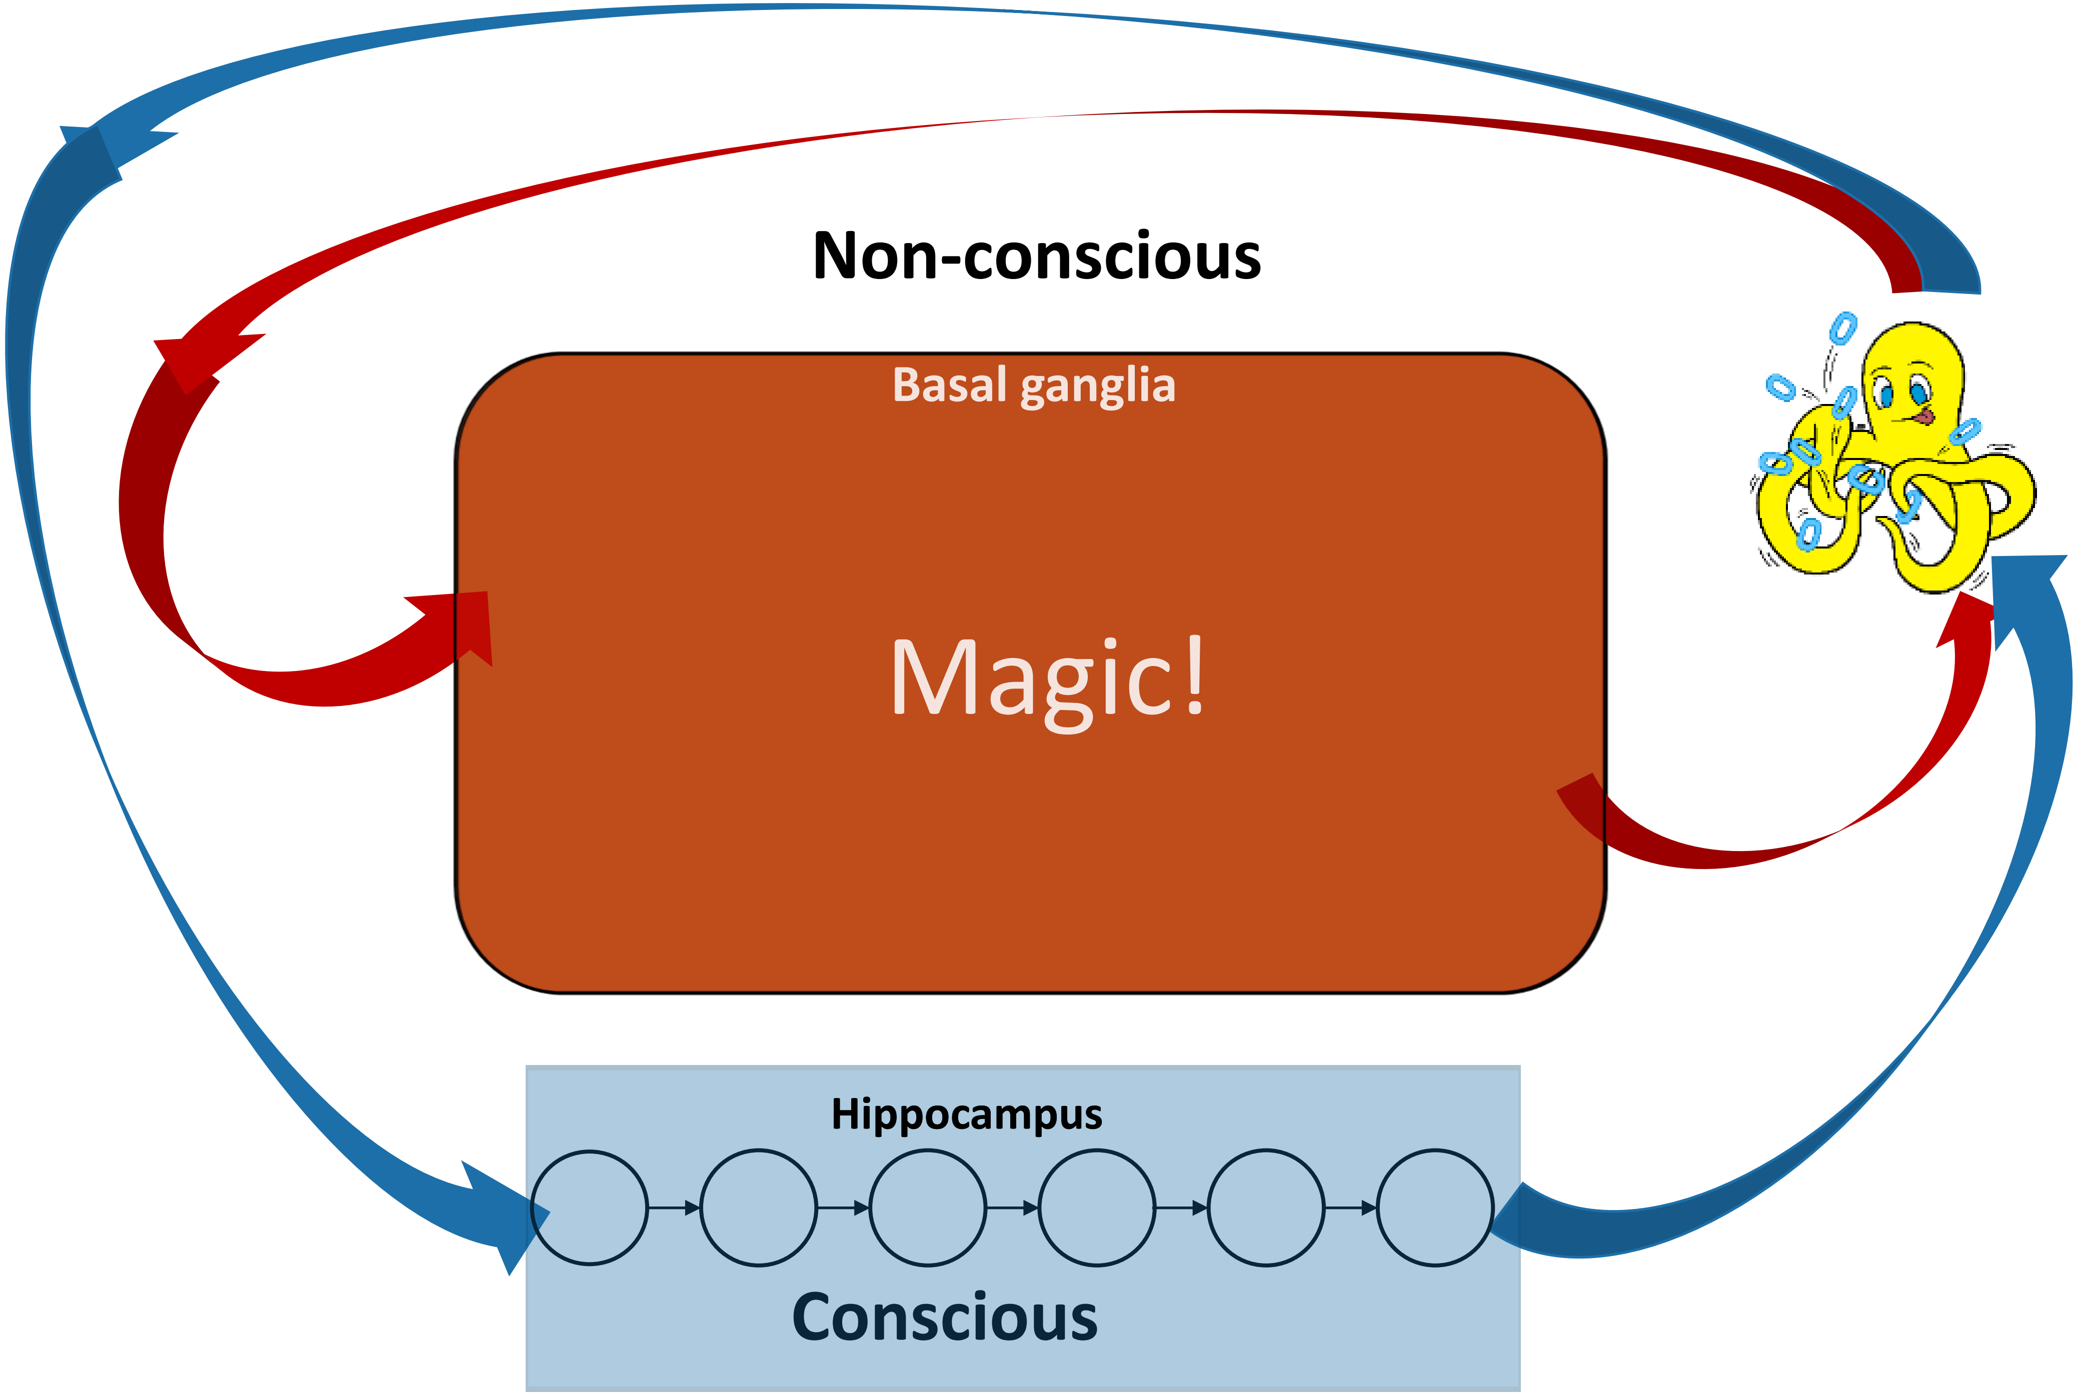
Non-conscious
Basal ganglia
Magic!
Hippocampus
Conscious

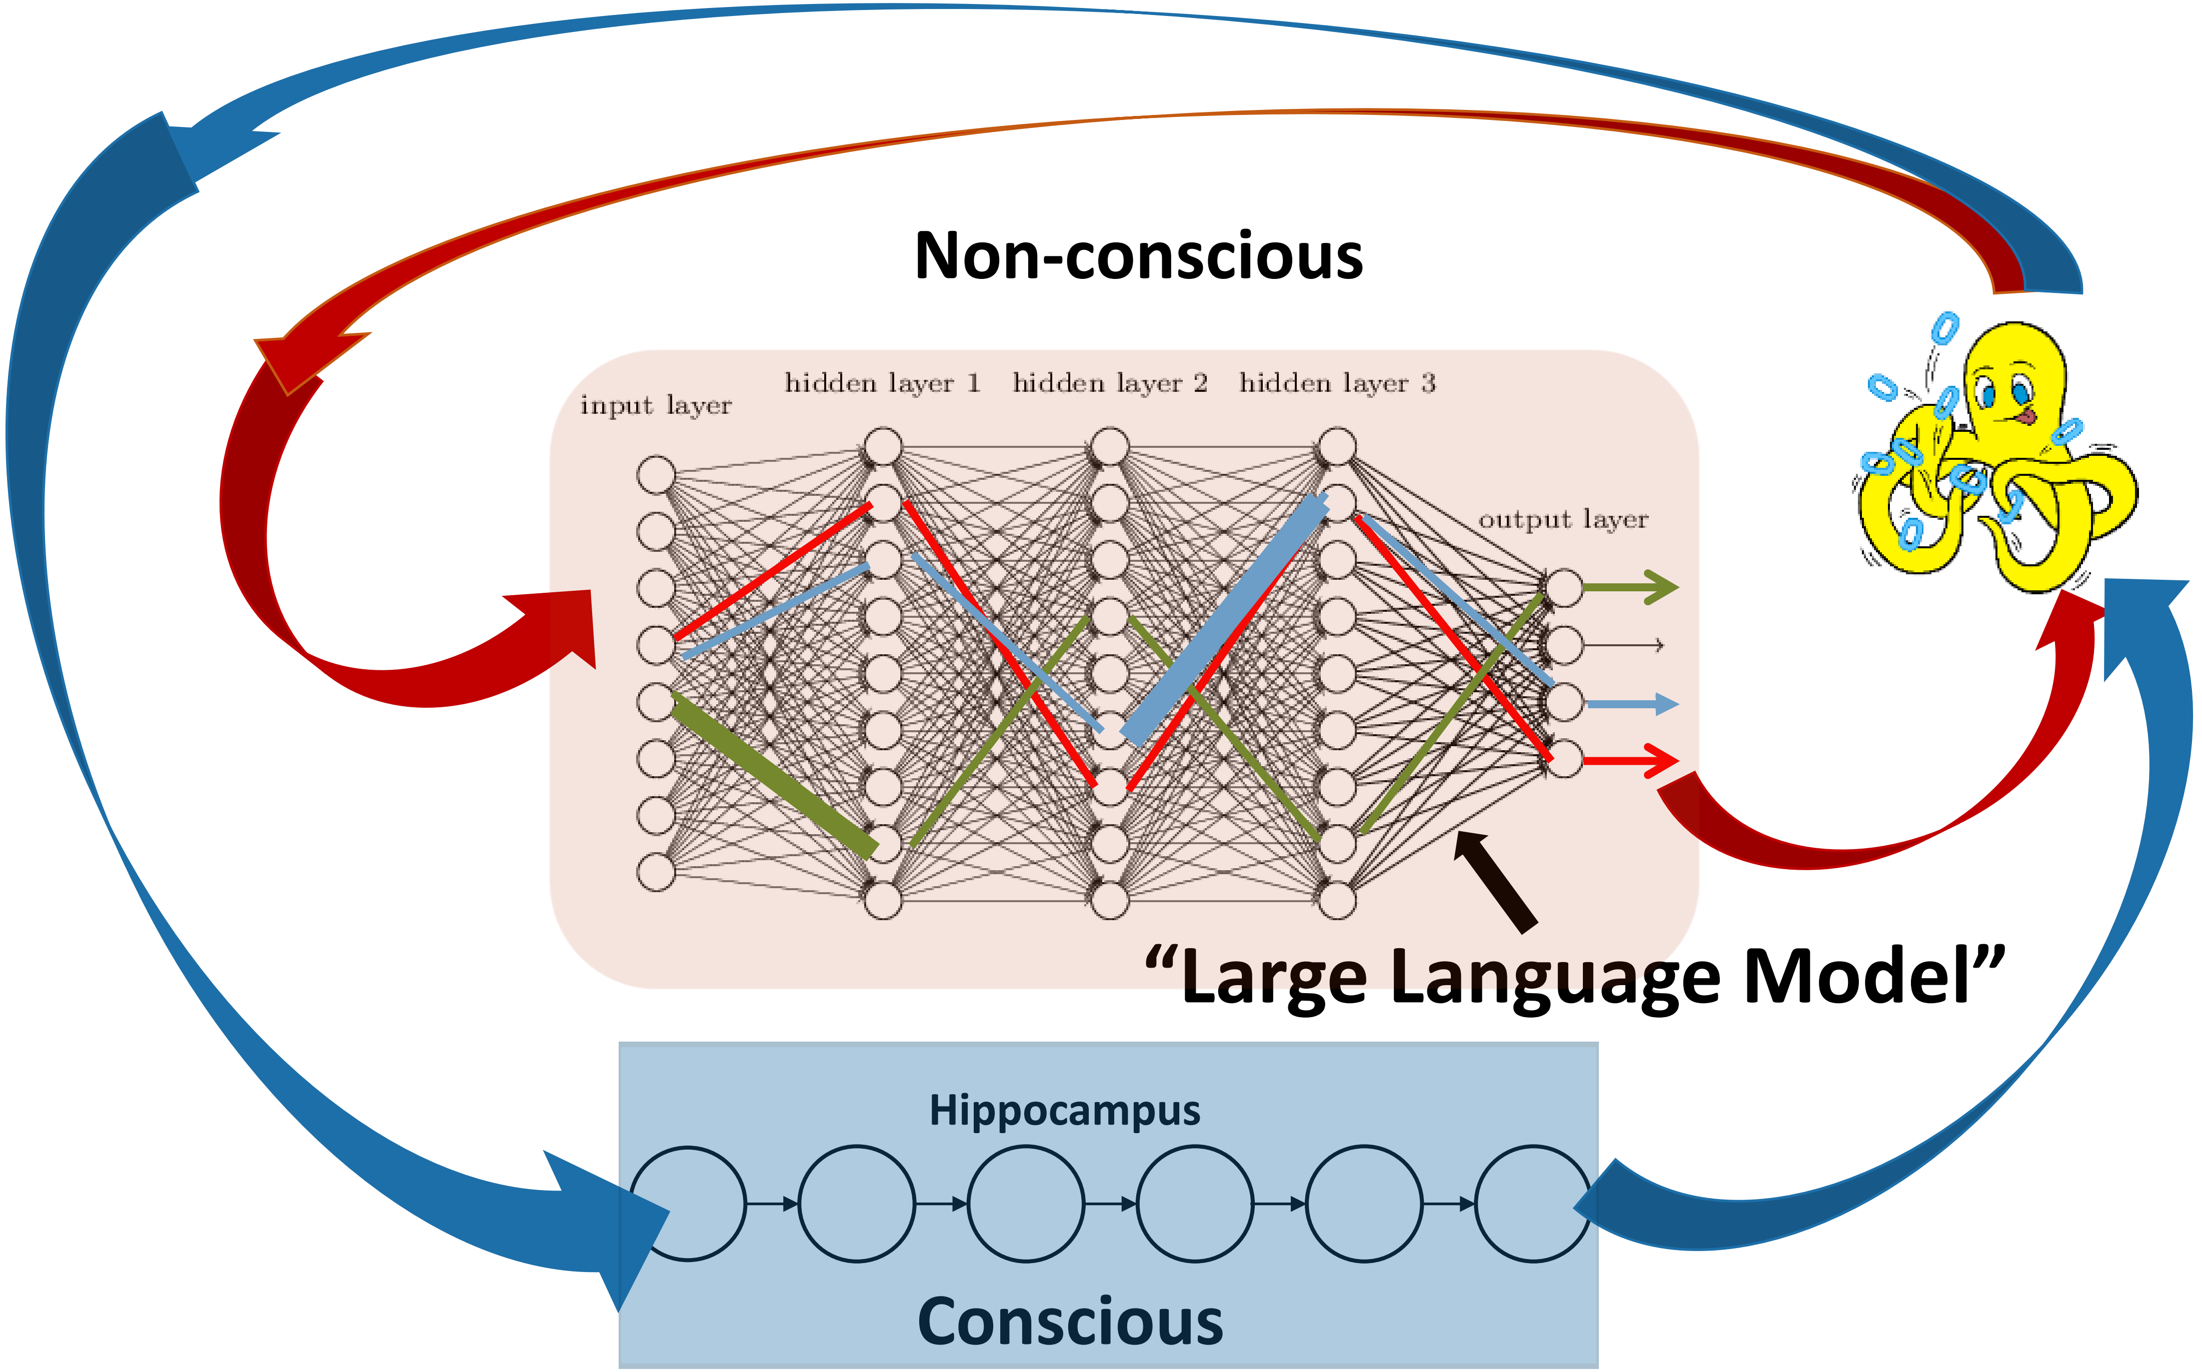
Non-conscious
input layer
hidden layer 1
hidden layer 2
hidden layer 3
output layer
"Large Language Model"
Hippocampus
Conscious

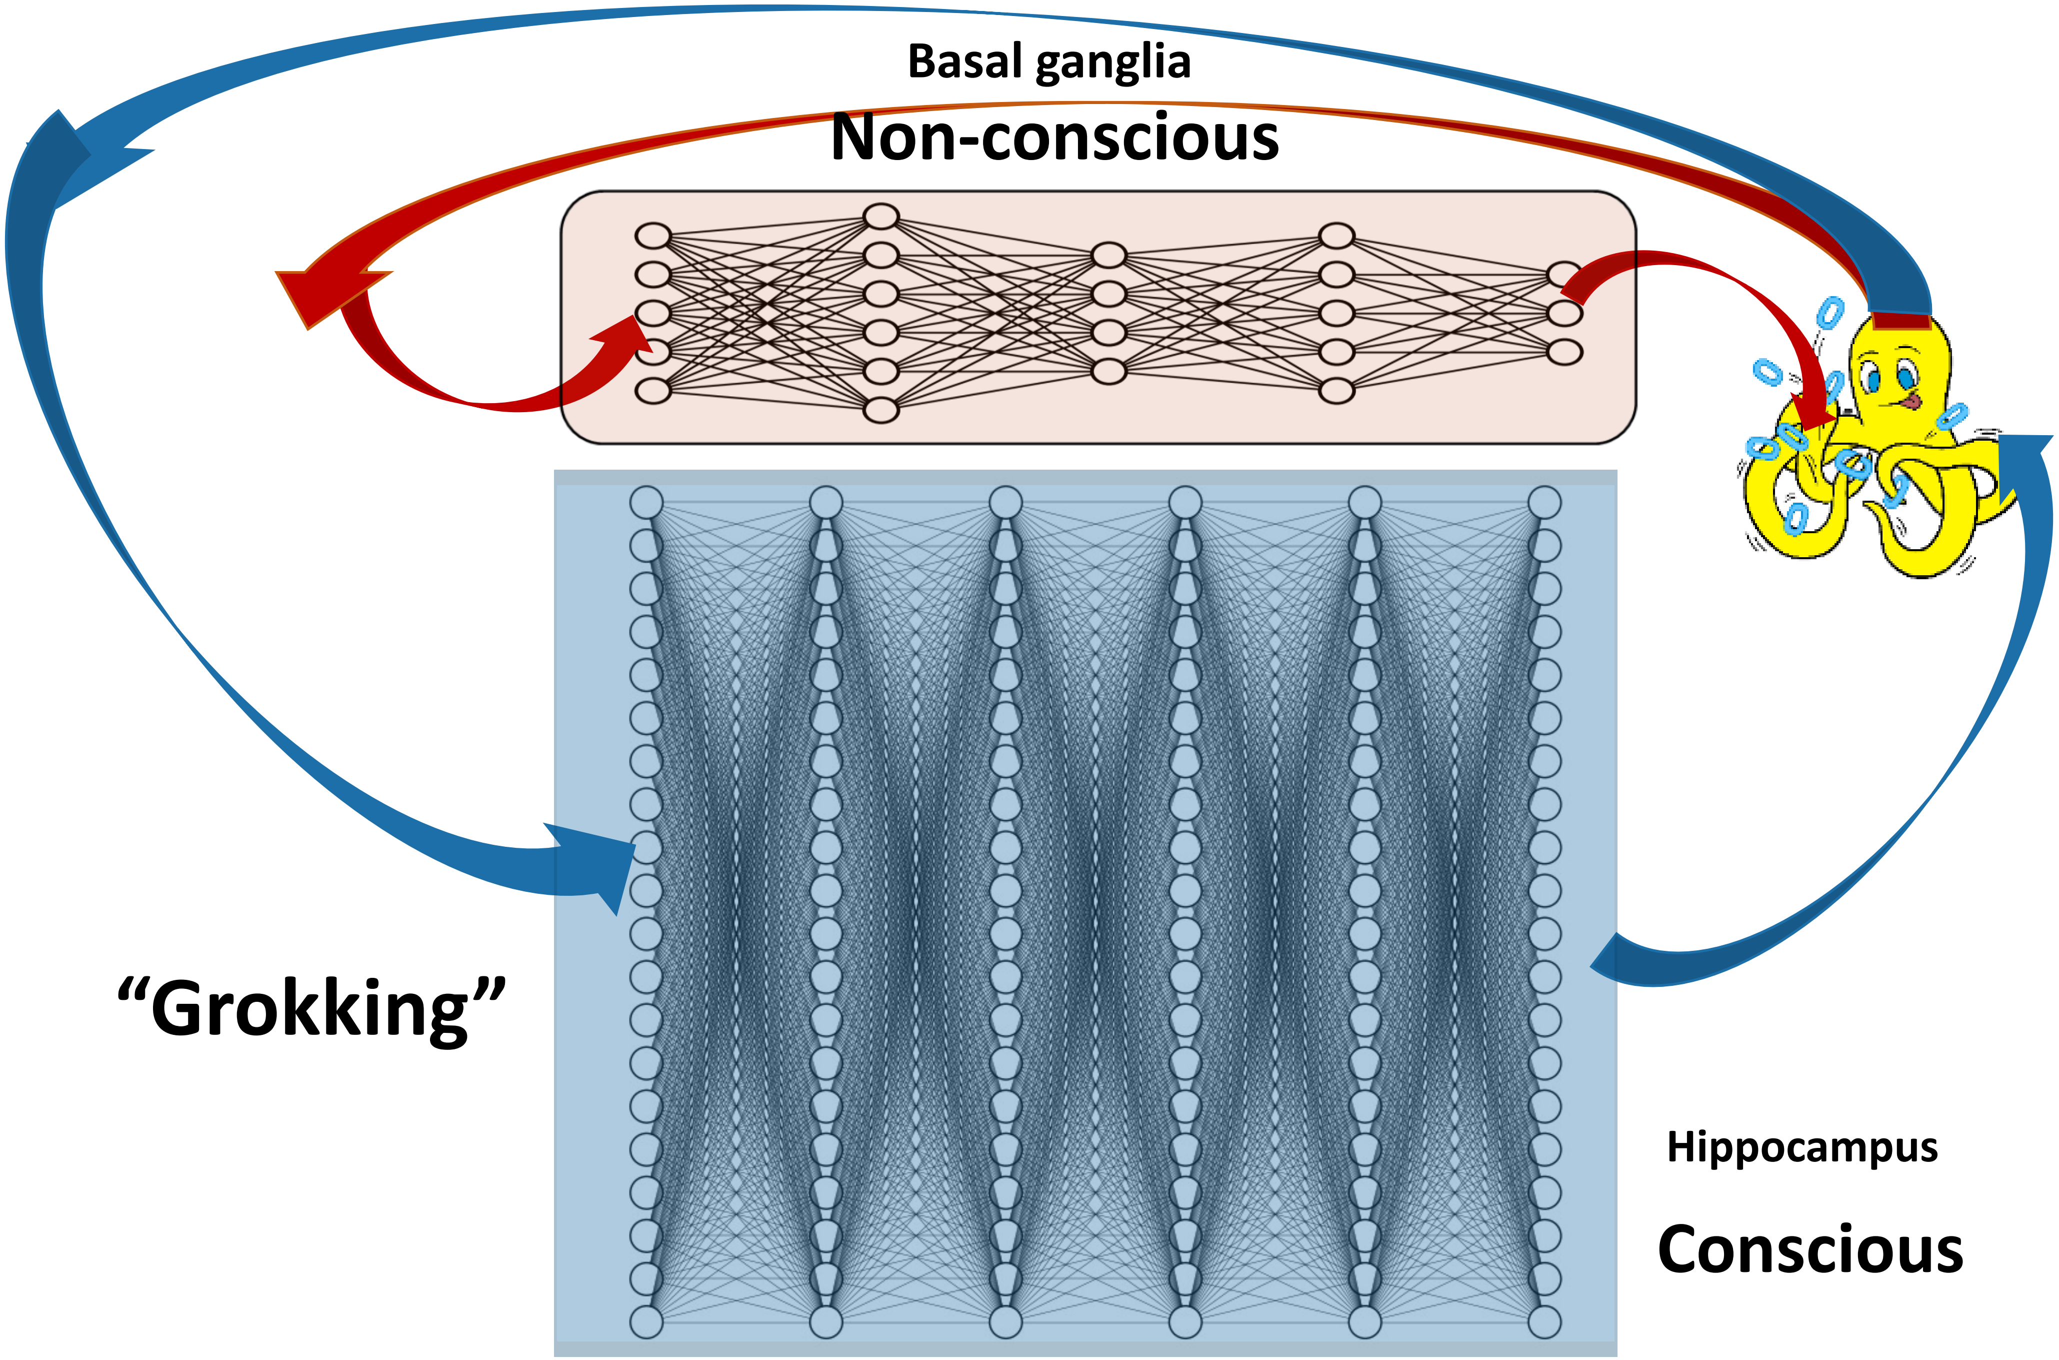
Basal ganglia
Non-conscious
"Grokking"
Hippocampus
Conscious

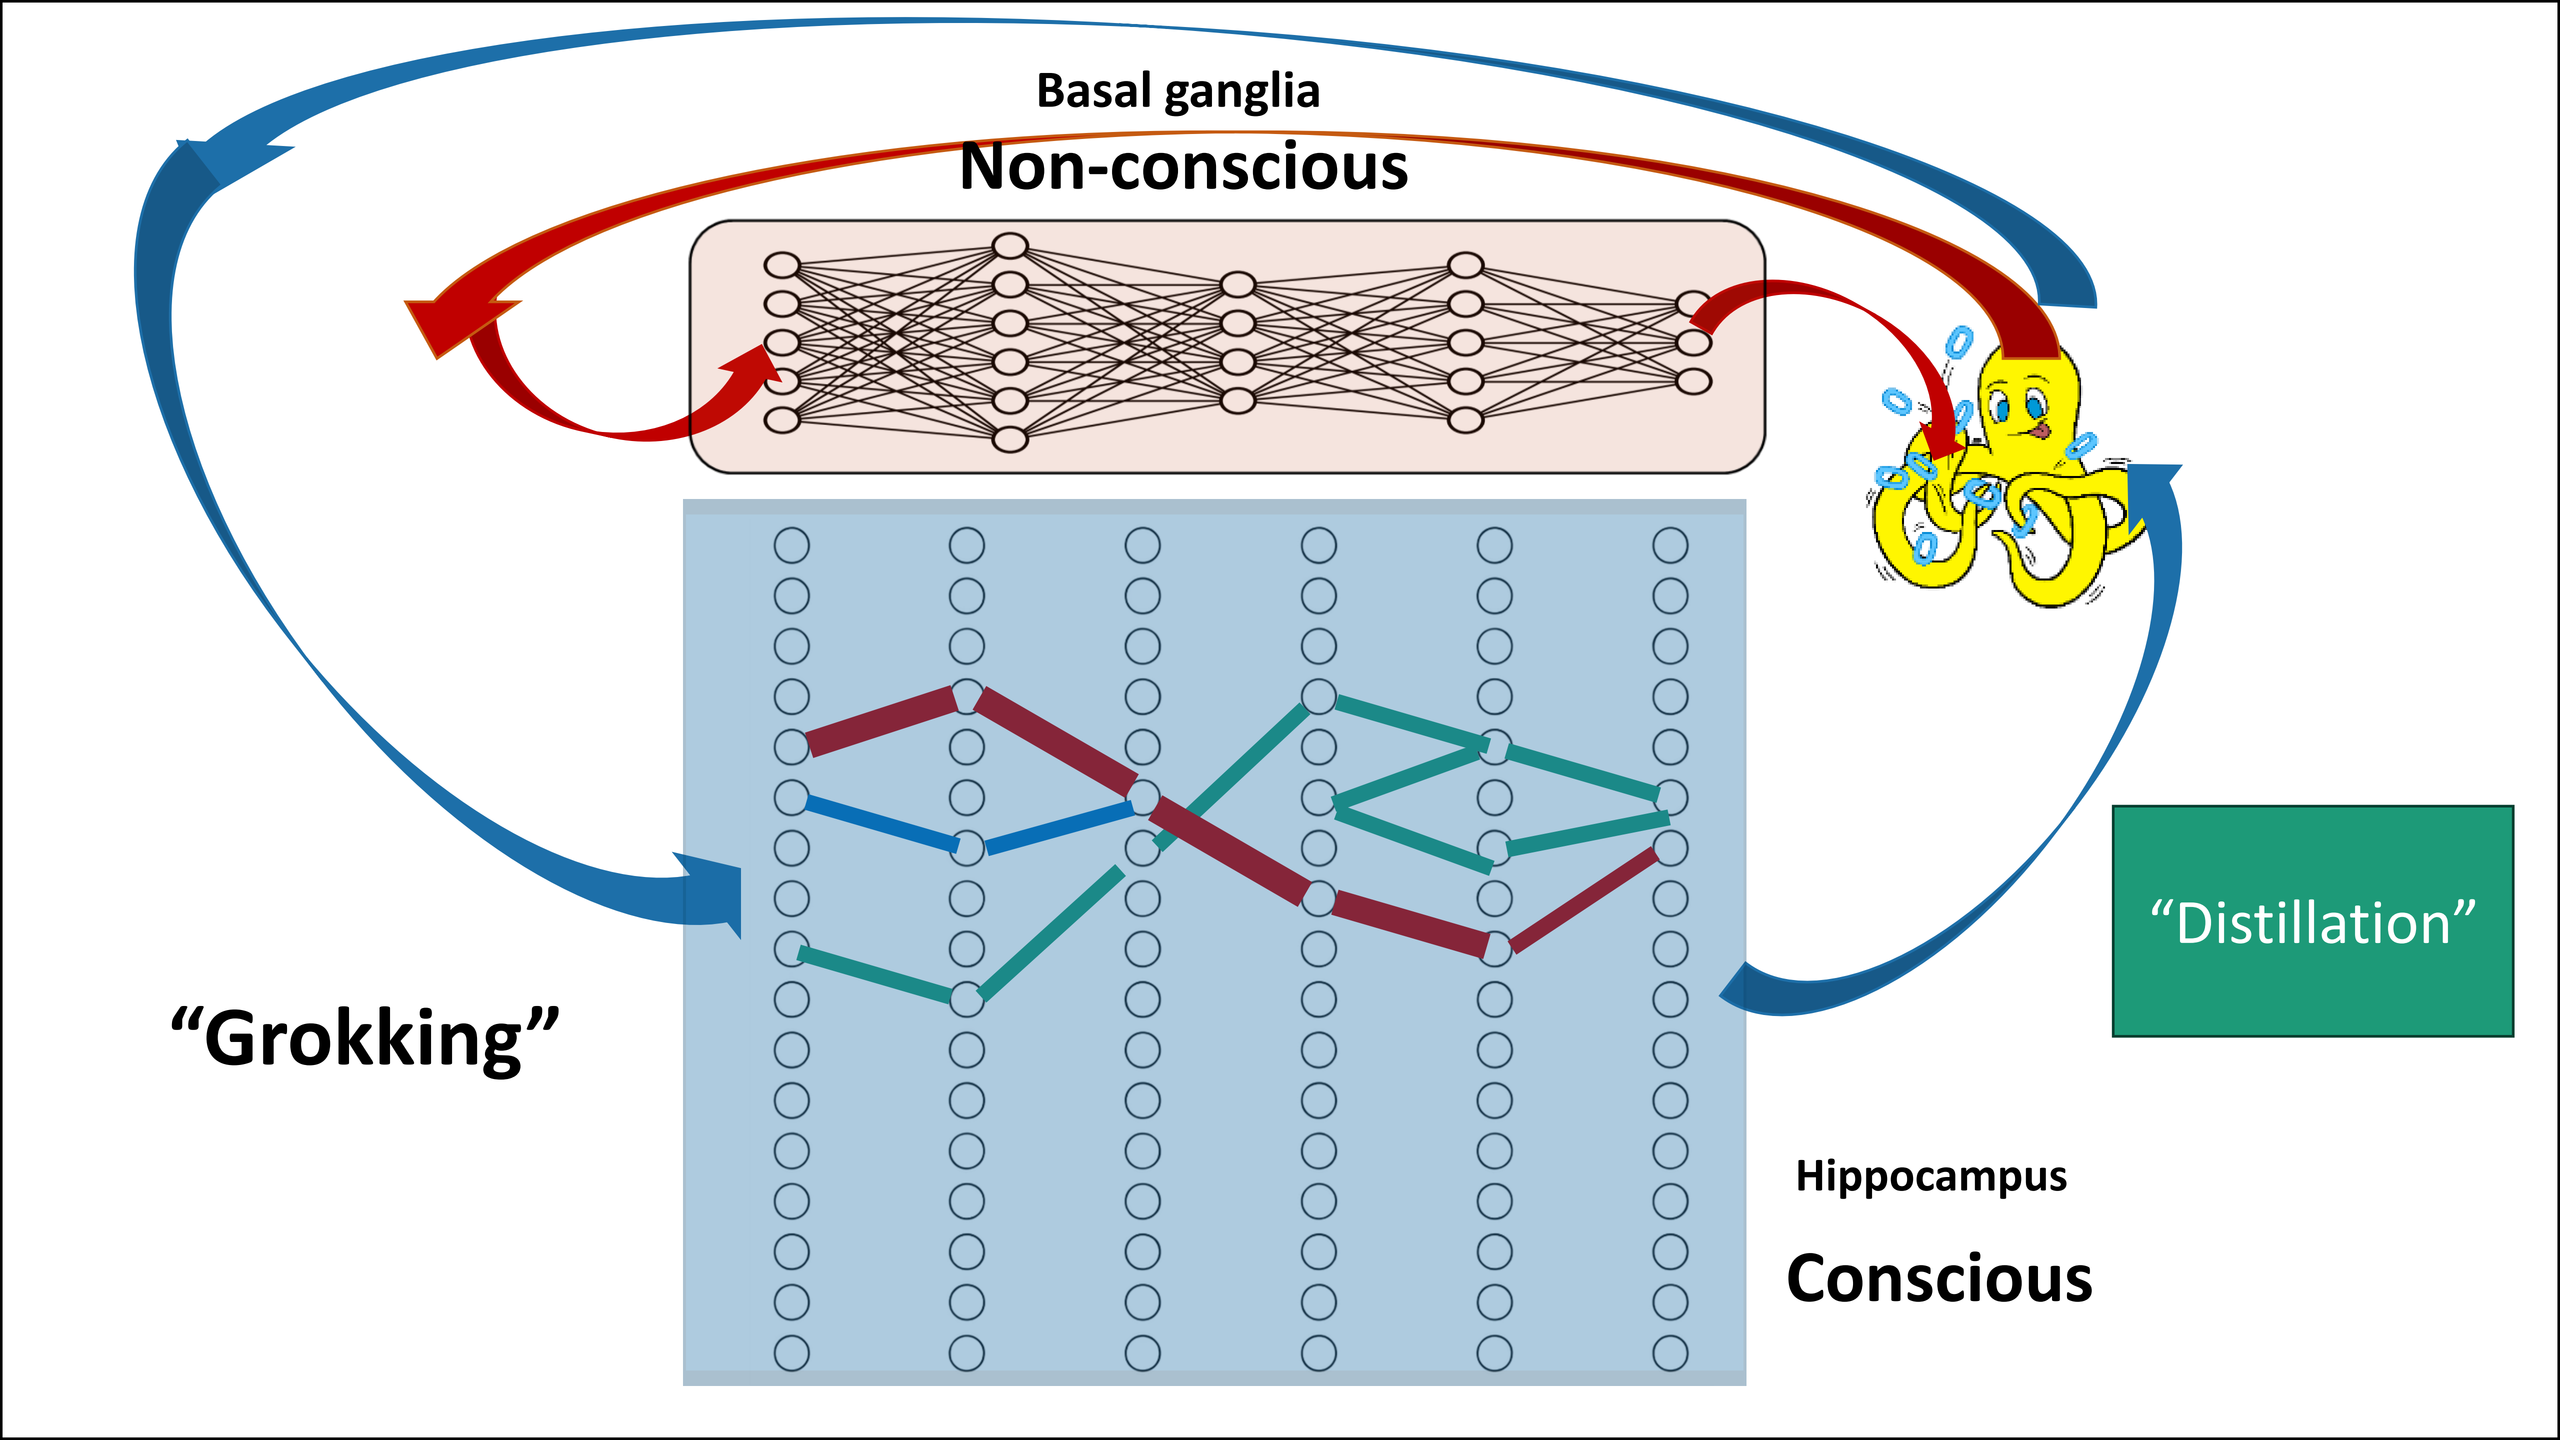
Basal ganglia
Non-conscious
"Grokking"
"Distillation"
Hippocampus
Conscious

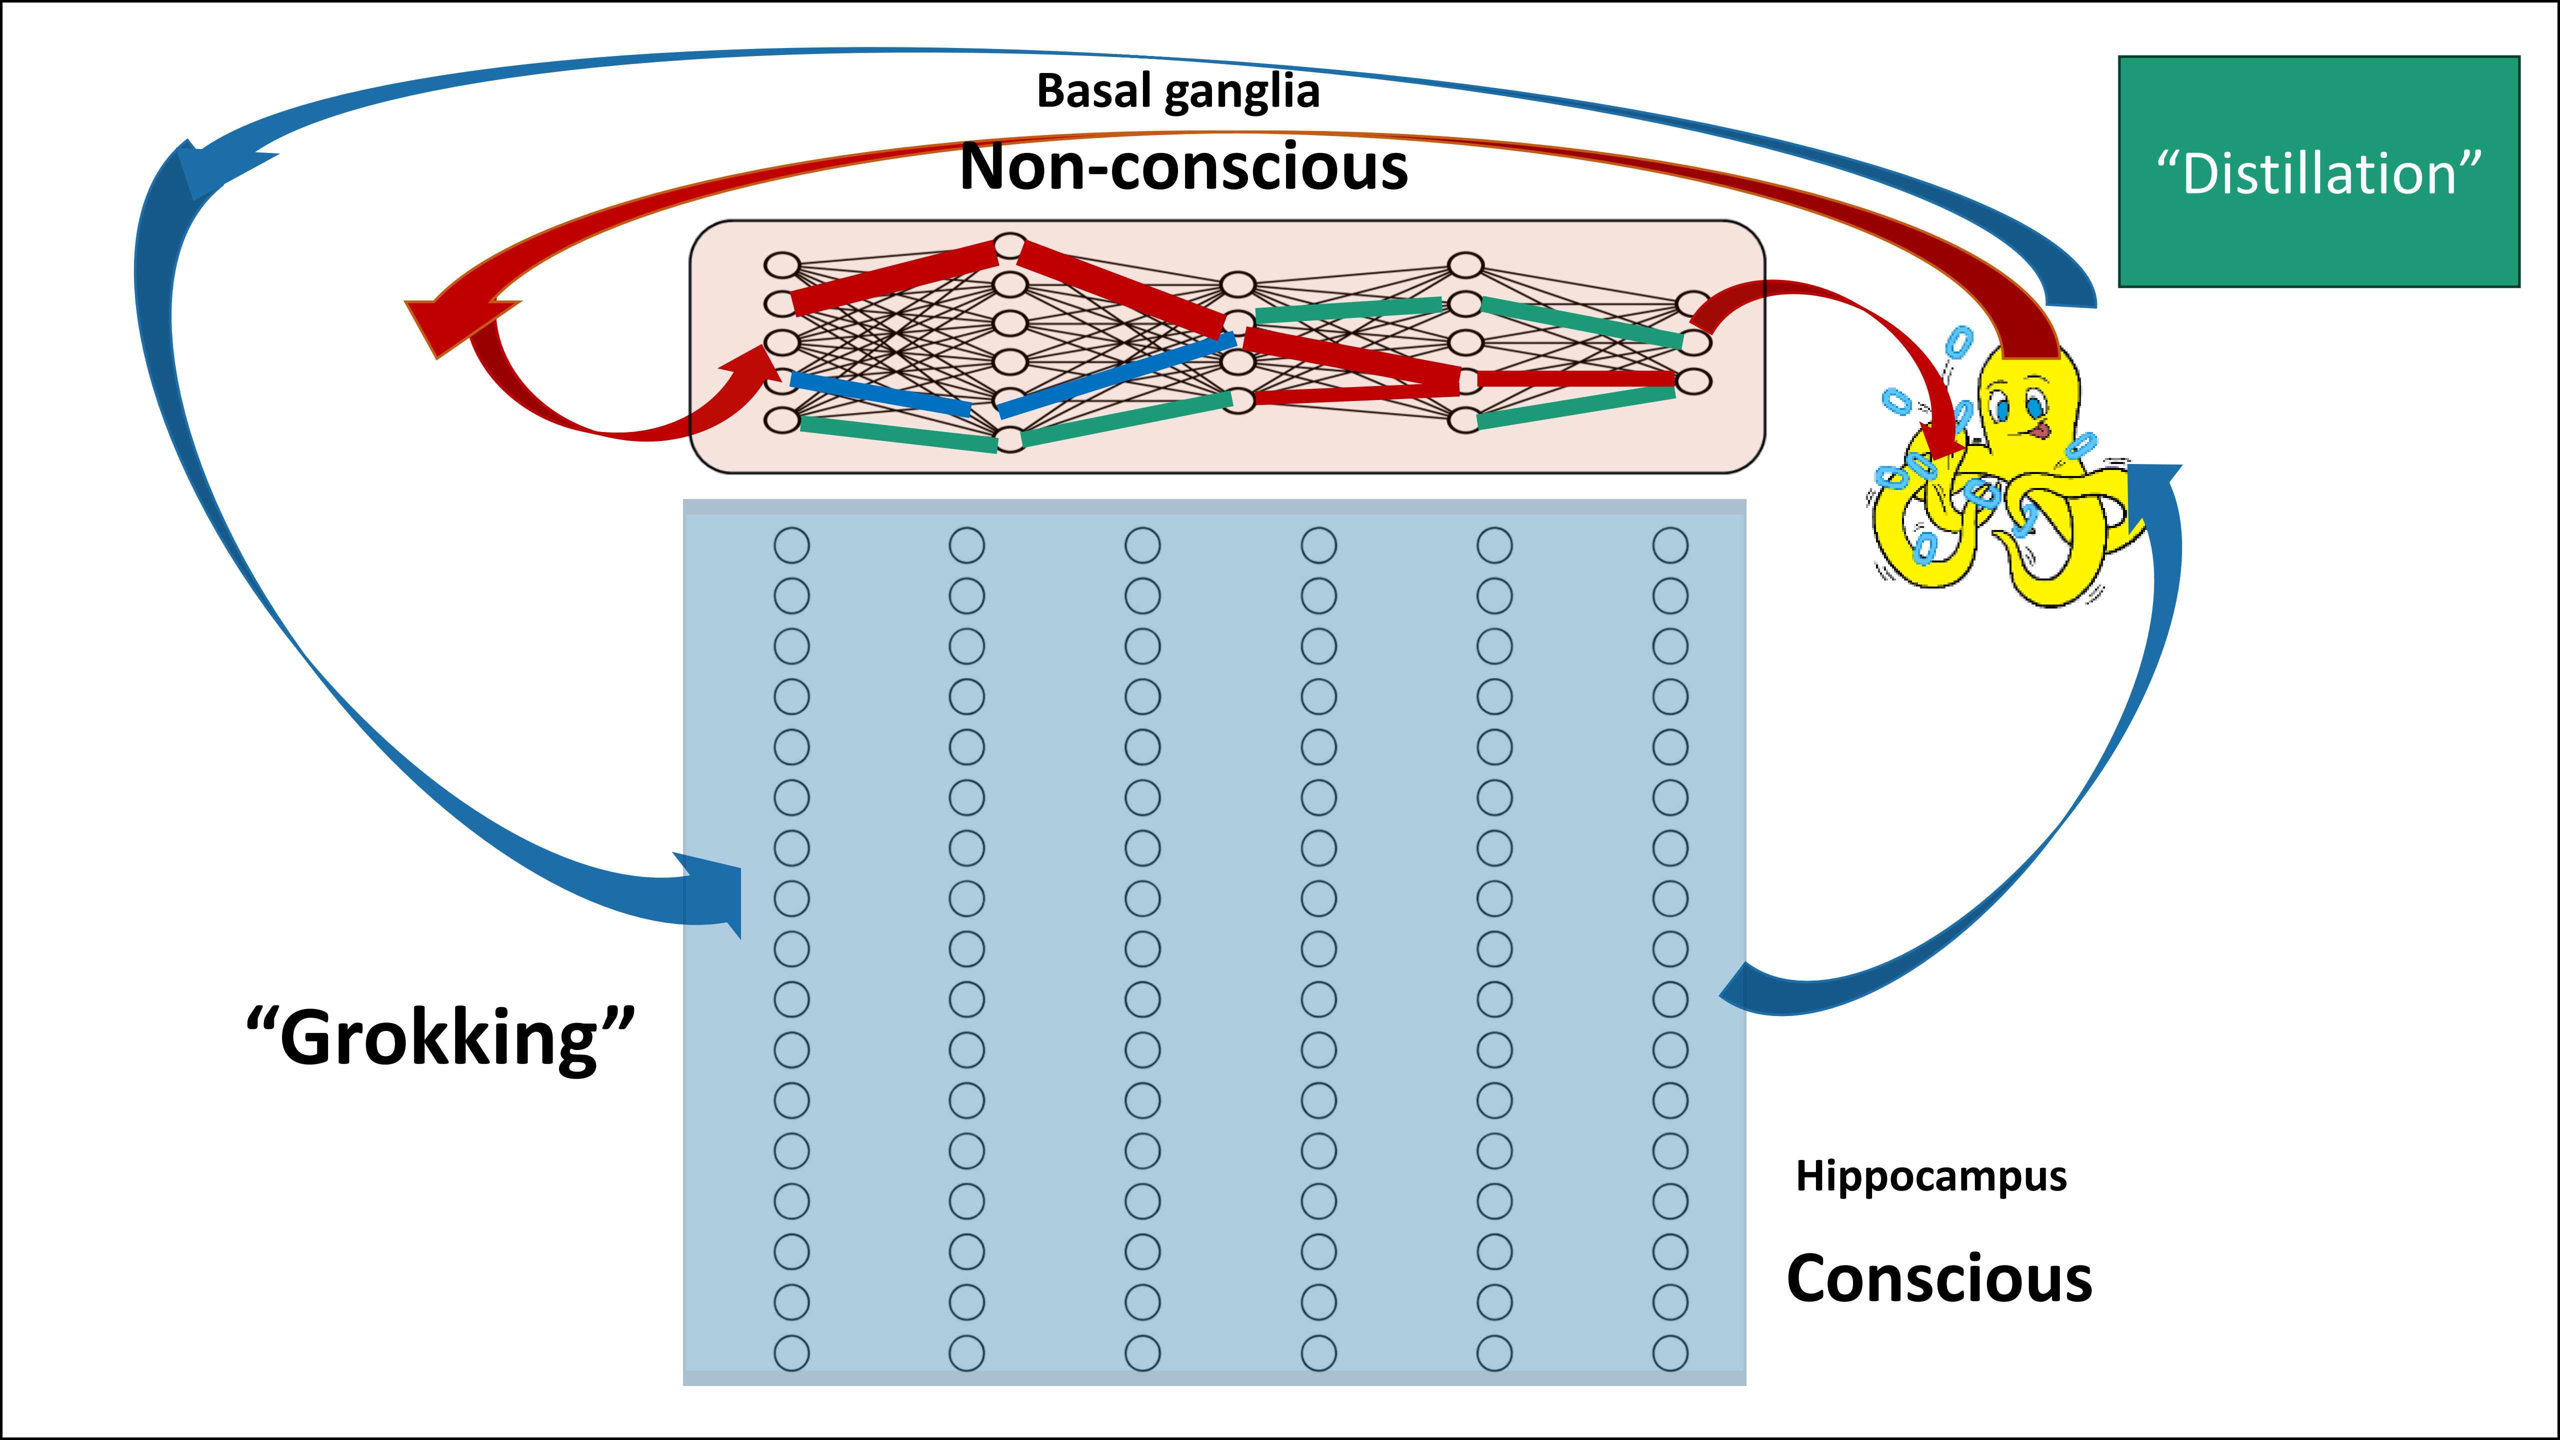
Basal ganglia
Non-conscious
"Distillation"
"Grokking"
Hippocampus
Conscious

nature COMMUNICATIONS

(2019)10:4646

https://doi.org/10.1038/s41467-019-12552-4 OPEN

Let's look deeper at grokking

The Eighty Five Percent Rule for optimal learning

Robert C. Wilson[1,2]*, Amitai Shenhav[3,4], Mark Straccia[5] & Jonathan D. Cohen[6]

Researchers and educators have long wrestled with the question of how best to teach their clients be they humans, non-human animals or machines. Here, we examine the role of a single variable, the difficulty of training, on the rate of learning. In many situations we find that there is a sweet spot in which training is neither too easy nor too hard, and where learning progresses most quickly. We derive conditions for this sweet spot for a broad class of learning algorithms in the context of binary classification tasks. For all of these stochastic gradient-descent based learning algorithms, we find that the optimal error rate for training is around 15.87% or, conversely, that the optimal training accuracy is about 85%. We demonstrate the efficacy of this 'Eighty Five Percent Rule' for artificial neural networks used in AI and biologically plausible neural networks thought to describe animal learning.

Let's look deeper at grokking

- Power, A., Burda, Y., Edwards, H., Babuschkin, I. & Misra, V. (2022). *"Grokking: Generalization beyond overfitting on small algorithmic datasets."* arXiv:2201.02177.

  — **Original discovery:** neural networks suddenly generalize long after memorizing training data.

- Davies, X., Langosco, L. & Krueger, D. (2023). *"Unifying grokking and double descent."* arXiv:2303.06173.

  — **Three competing patterns** (memorization, fast heuristic, true understanding) with different learning speeds. Math behind the visualization.

- Fedorenko, E., Piantadosi, S. T. & Gibson, E. A. F. (2024). *"Language is primarily a tool for communication rather than thought."* Nature 630, 575–586.

  — **Language network is for communication, not thinking.** Thinking happens in separate brain networks. People with severe aphasia retain math, logic, causal reasoning. Language is always the surface layer.

- Ullman, M. T. (2020). *"The Declarative/Procedural Model: A neurobiologically motivated theory of first and second language."* In B. VanPatten et al. (Eds.), Theories in Second Language Acquisition (3rd ed., pp. 128–161). Routledge.

  — **Two memory systems:** hippocampal (declarative—fast, explicit, words and facts) and basal ganglia (procedural—slow, implicit, grammar and skills). Grammar shifts from declarative to procedural with extended practice. Strikingly parallel dynamic to grokking curves.

- Sundaram, S. et al. "*Teaching Models to Teach Themselves: Reasoning at the Edge of Learnability*." arXiv preprint arXiv:2601.18778 (2026).

  — **Models plateau** when problems exceed their learnability frontier. SOAR meta-RL framework: teacher generates synthetic stepping-stone problems, rewarded by student progress on hard datasets. Self-generated curricula unlock learning where direct training fails.

Alt-tab/F11

Download the interactive visualization: https://claude.ai/public/artifacts/263b8968-7c3d-4dec-a2be-d2d8e7ee1832

Grokking: From Memorization to Understanding

Training accuracy ON

Accuracy Over Training

dashed = train, solid = test

Grokking: From Memorization to Understanding

- Training accuracy ON
- Test accuracy ON

Accuracy Over Training dashed = train, solid = test

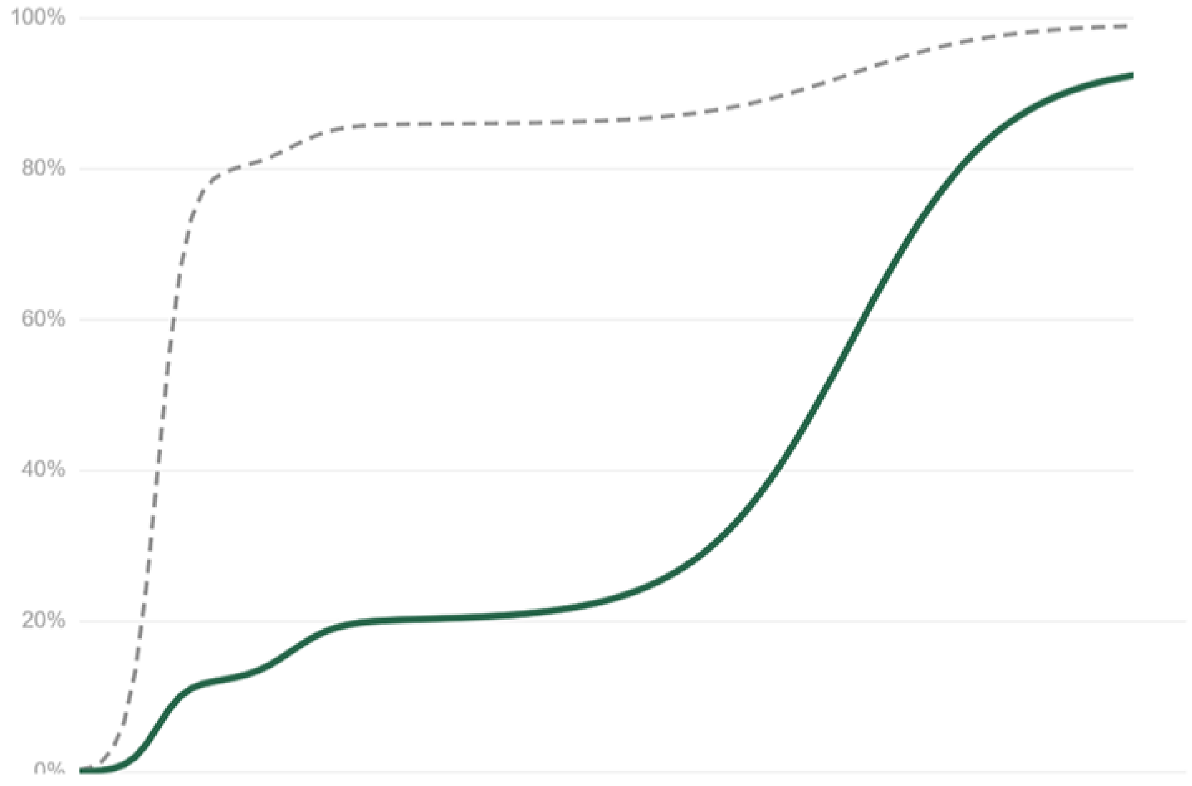

Grokking: From Memorization to Understanding

- Training accuracy ON
- Test accuracy ON
- Memorization ON
- Fast heuristic OFF
- True understanding OFF

Memorization

"How are you?"

reveal ▶

Accuracy Over Training dashed = train, solid = test

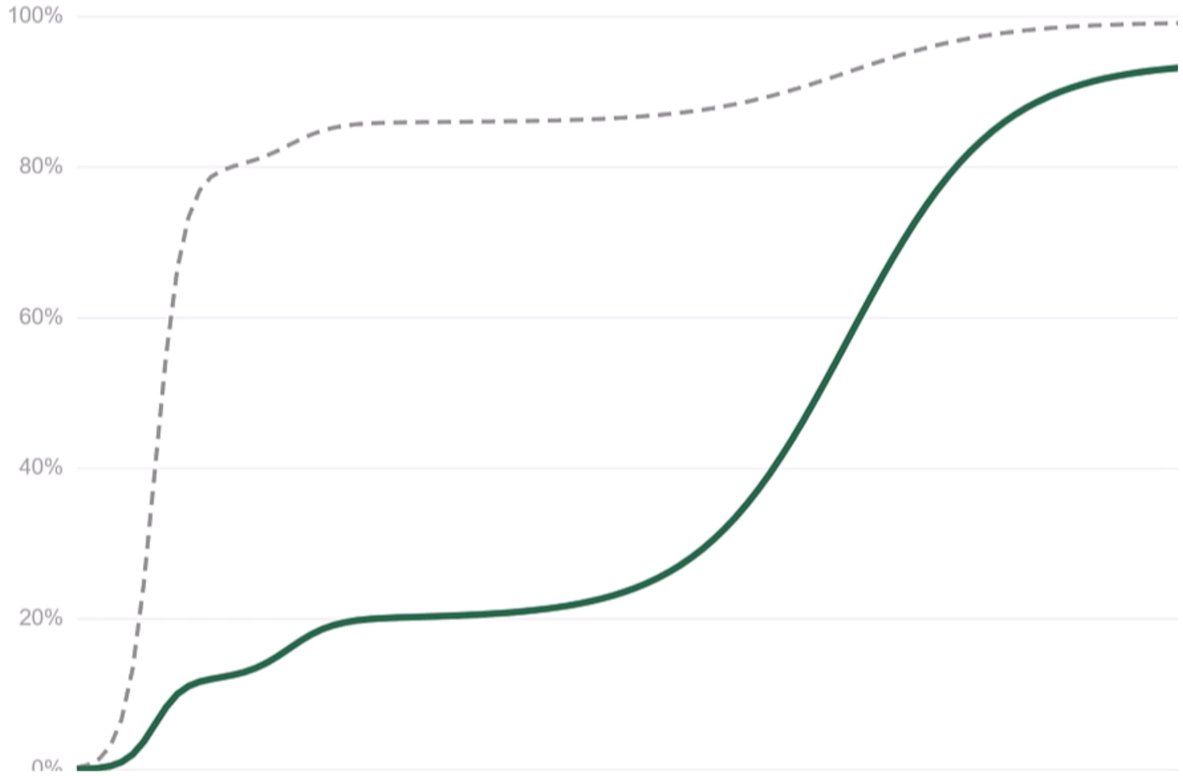

Grokking: From Memorization to Understanding

- Training accuracy ON
- Test accuracy ON
- Memorization ON
- Fast heuristic OFF
- True understanding OFF

Memorization

"How are you?" *I am fine, thank you*
"How is your mother?" *Huh?*

Accuracy Over Training dashed = train, solid = test

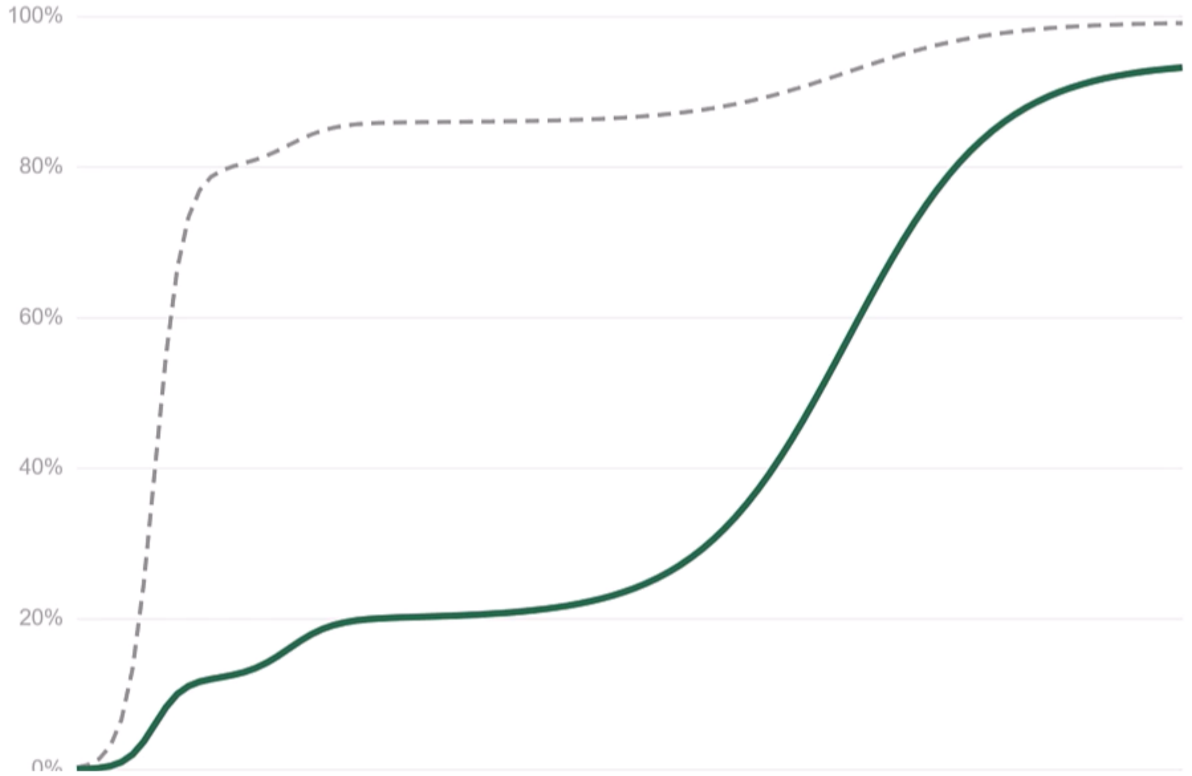

Pattern Strength Over Training what's happening inside the learner?

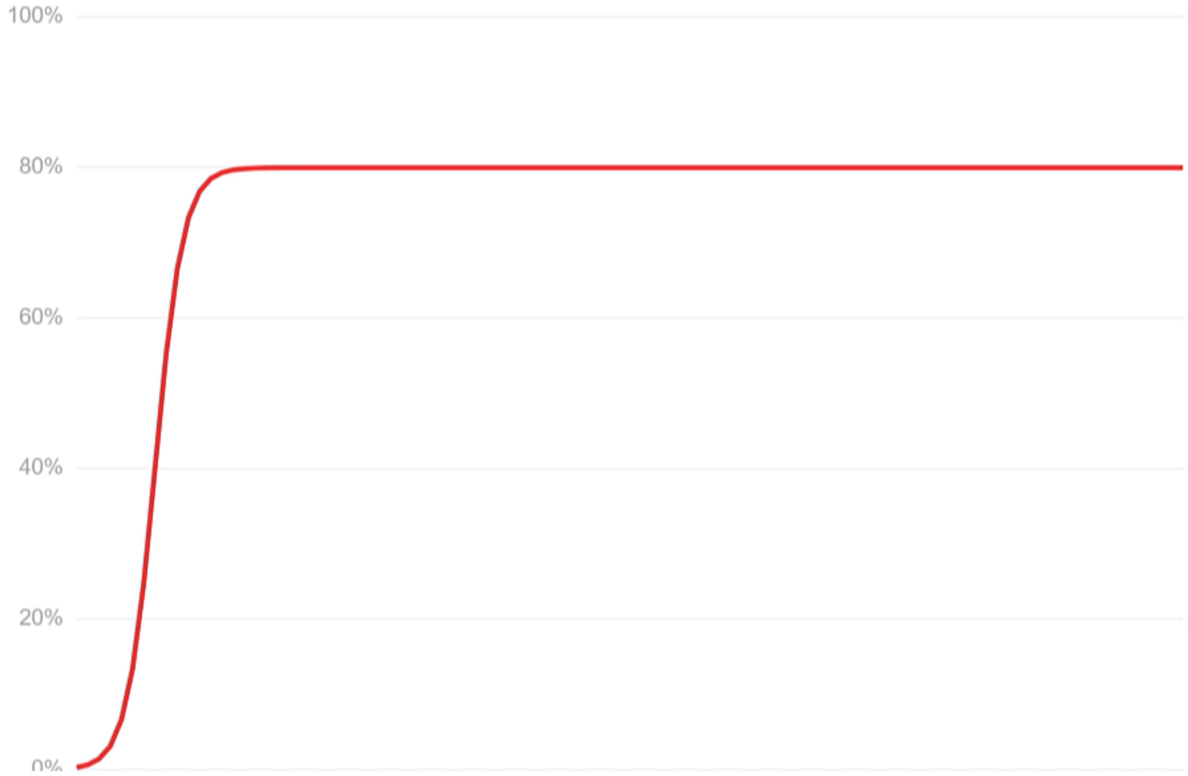

Grokking: From Memorization to Understanding

Training accuracy ON

Test accuracy ON

Memorization ON

Fast heuristic ON

True understanding OFF

Memorization

"How are you?" *I am fine, thank you*
"How is your mother?" *Huh?*

Fast Heuristic

Accuracy Over Training dashed = train, solid = test

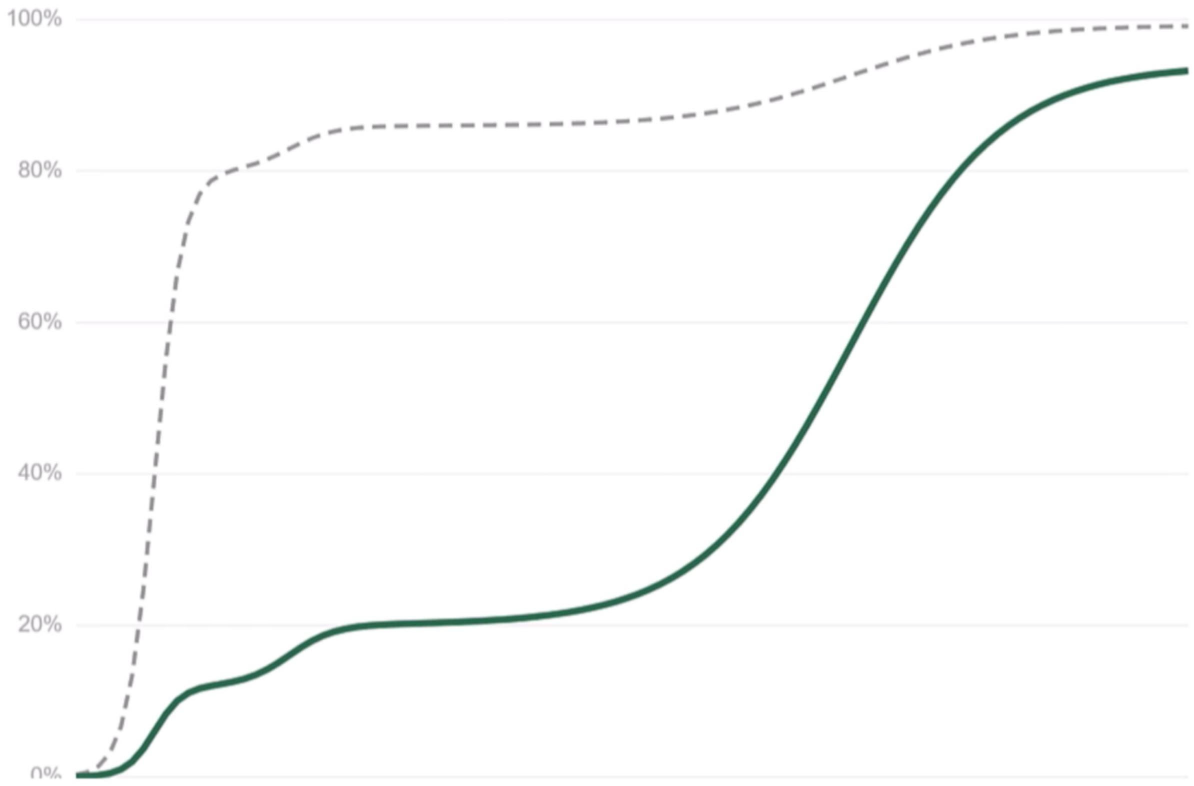

Pattern Strength Over Training what's happening inside the learner?

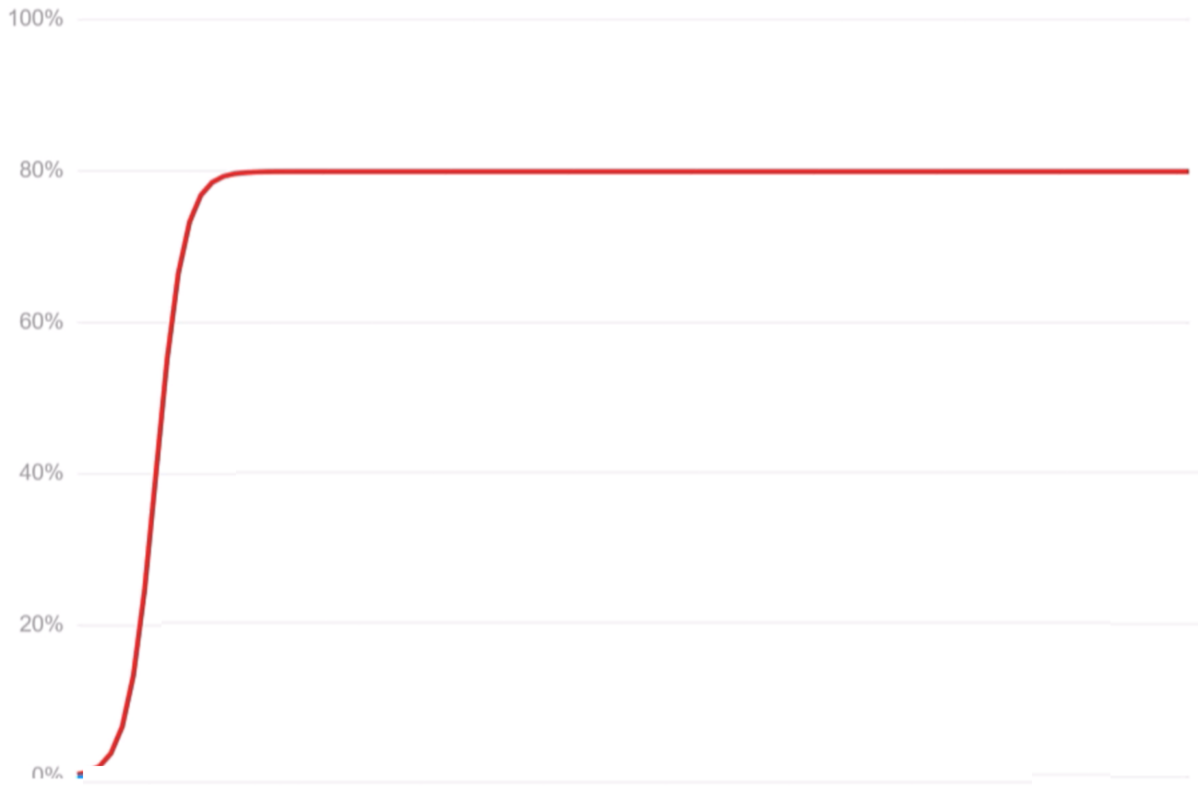

Grokking: From Memorization to Understanding

- Training accuracy ON
- Test accuracy ON
- Memorization ON
- Fast heuristic ON
- True understanding OFF

Memorization

"How are you?" *I am fine, thank you*
"How is your mother?" *Huh?*

Fast Heuristic

Rule: *add -ed for past tense walked, talked ✓*
But: *goed, runned ✗*

Accuracy Over Training dashed = train, solid = test

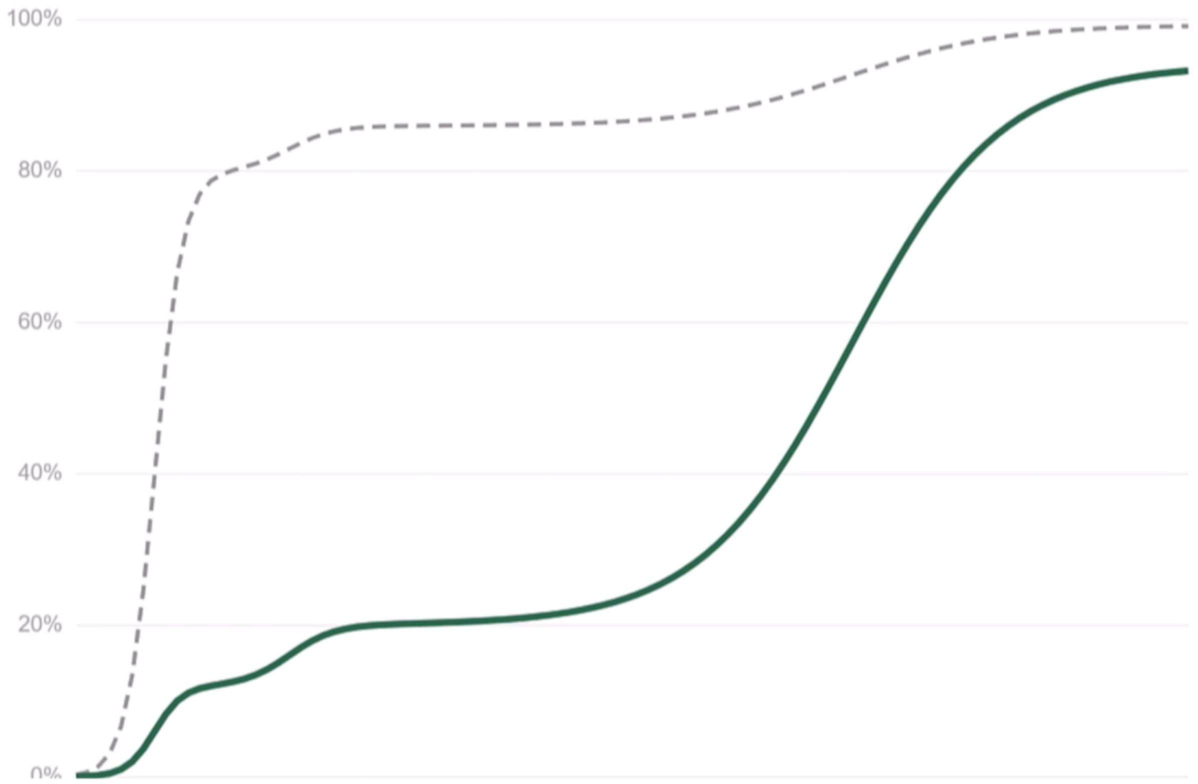

Pattern Strength Over Training what's happening inside the learner?

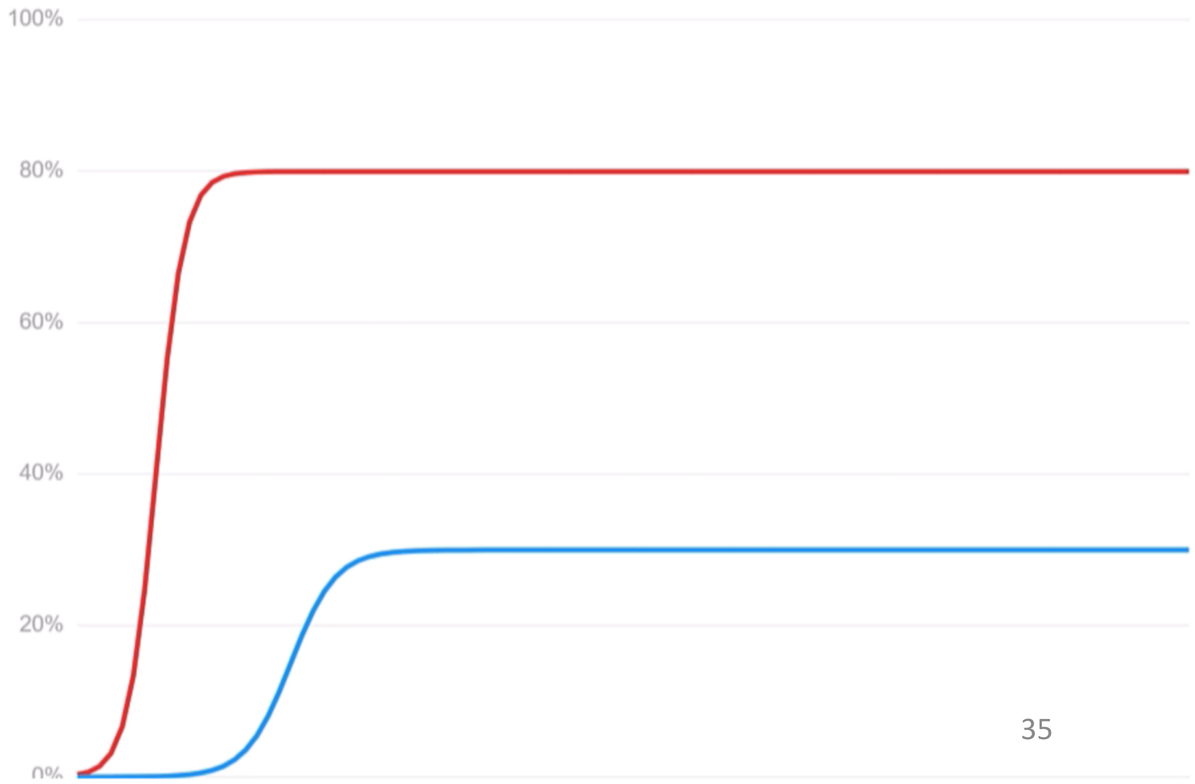

Grokking: From Memorization to Understanding

- Training accuracy ON
- Test accuracy ON
- Memorization ON
- Fast heuristic ON
- True understanding OFF

Memorization

"How are you?" *I am fine, thank you*
"How is your mother?" *Huh?*

Fast Heuristic

Rule: *add -ed for past tense walked, talked ✓*
But: *goed, runned ✗*

True Understanding

Handles anything new—even sentences never heard before

Accuracy Over Training dashed = train, solid = test

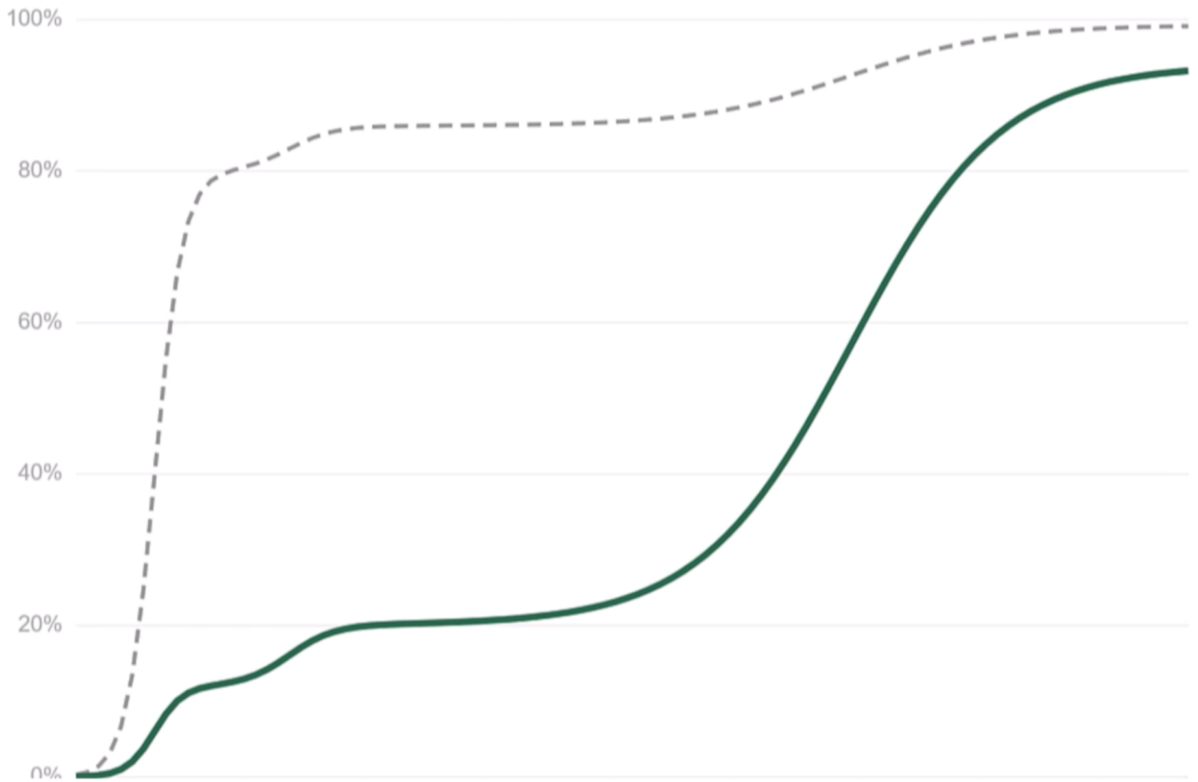

Pattern Strength Over Training what's happening i

Grokking: From Memorization to Understanding

- Training accuracy ON
- Test accuracy ON
- Memorization ON
- Fast heuristic ON
- True understanding ON

Well-scaffolded Teaching 100%

Memorization

"How are you?" *I am fine, thank you*
"How is your mother?" *Huh?*

Fast Heuristic

Rule: *add -ed for past tense walked, talked ✓*
But: *goed, runned ✗*

True Understanding

Handles anything new—even sentences never heard before

Accuracy Over Training dashed = train, solid = test

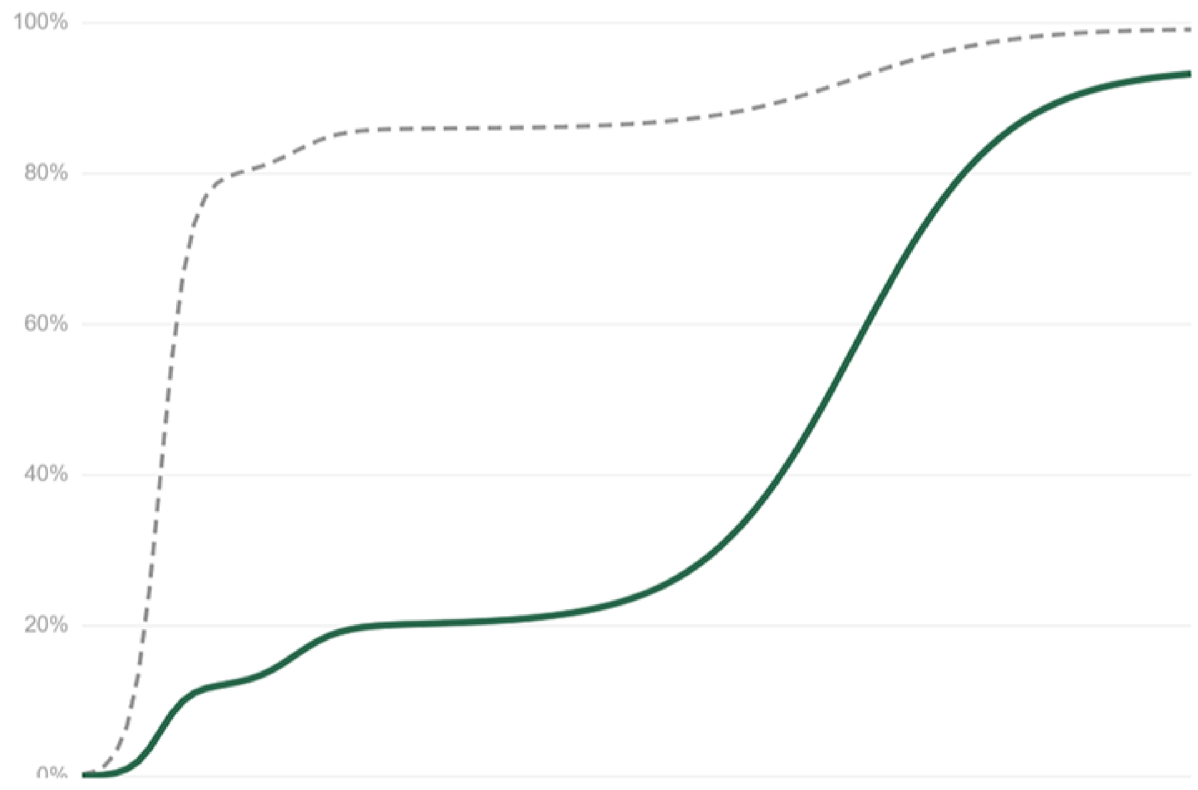

Pattern Strength Over Training what's happening i

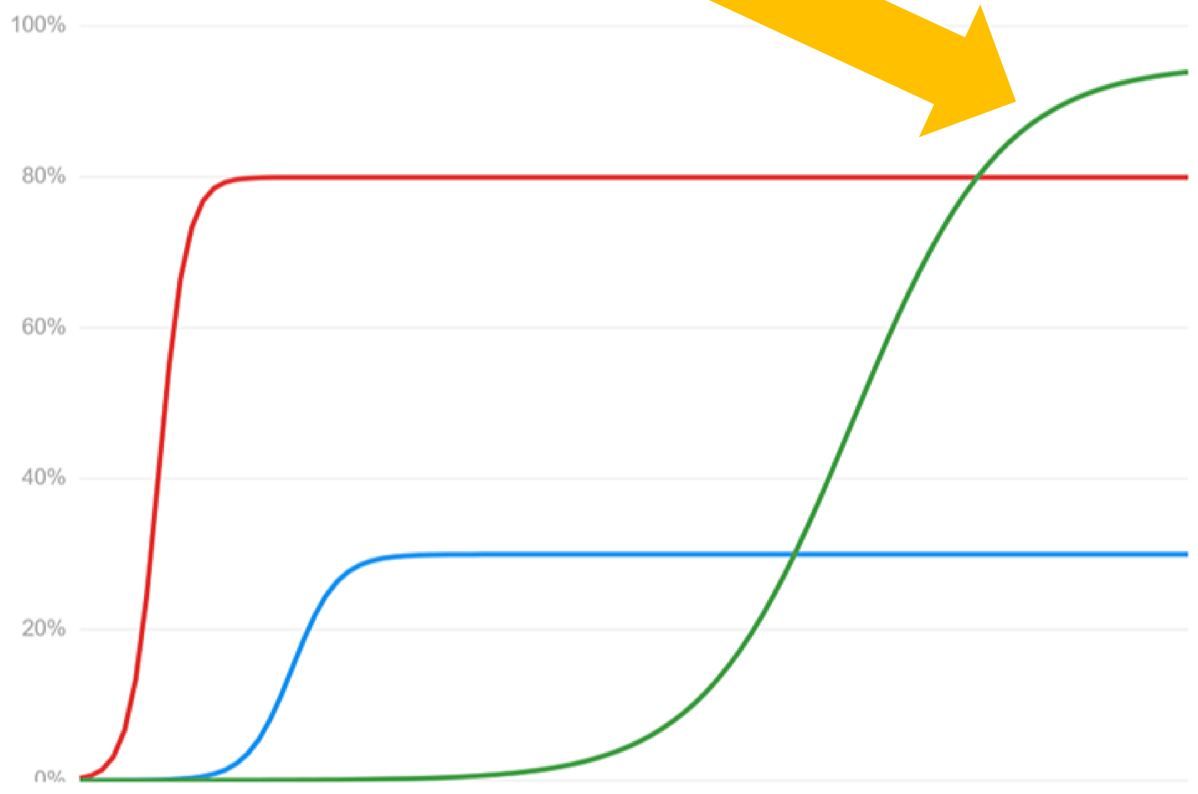

Grokking: From Memorization to Understanding

- Training accuracy ON
- Test accuracy ON
- Memorization ON
- Fast heuristic ON
- True understanding ON

Well-scaffolded Teaching 100%

Memorization

"How are you?" *I am fine, thank you*
"How is your mother?" *Huh?*

Fast Heuristic

Rule: *add -ed for past tense walked, talked ✓*
But: *goed, runned ✗*

True Understanding

Handles anything new—even sentences never heard before

Accuracy Over Training dashed = train, solid = test

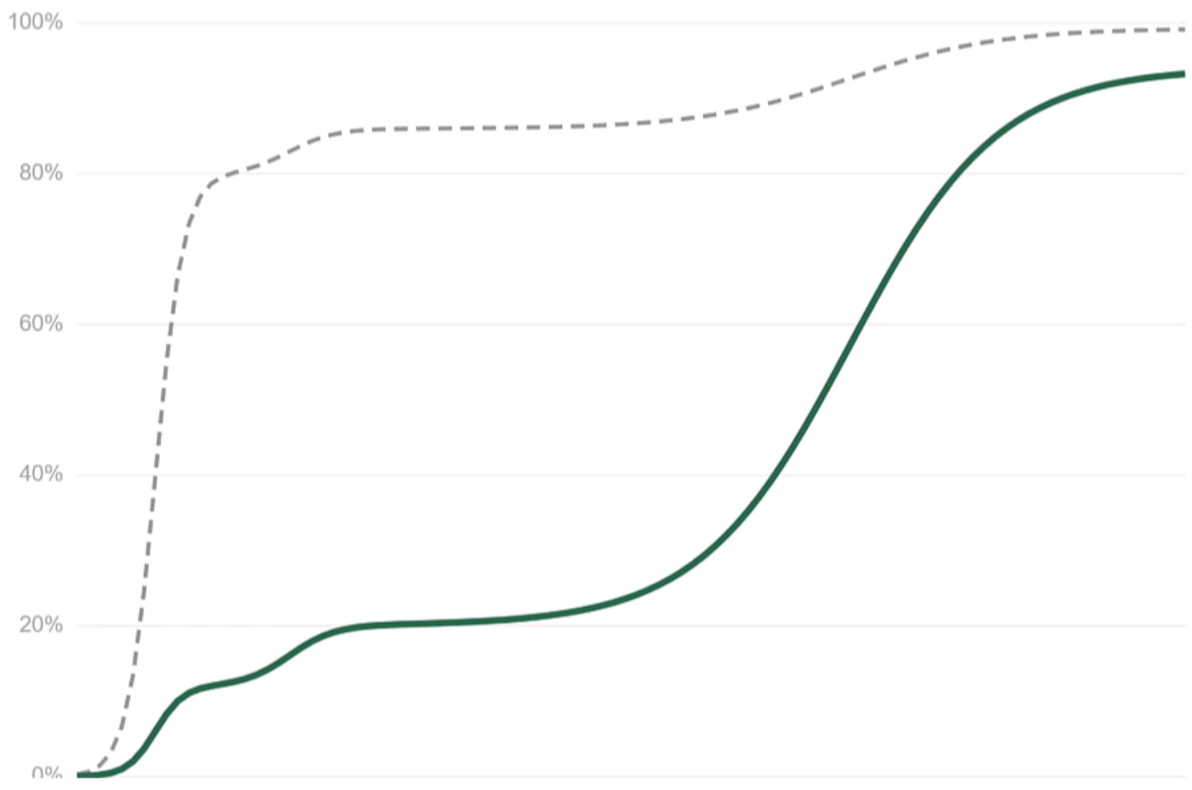

Pattern Strength Over Training what's happening inside the learner?

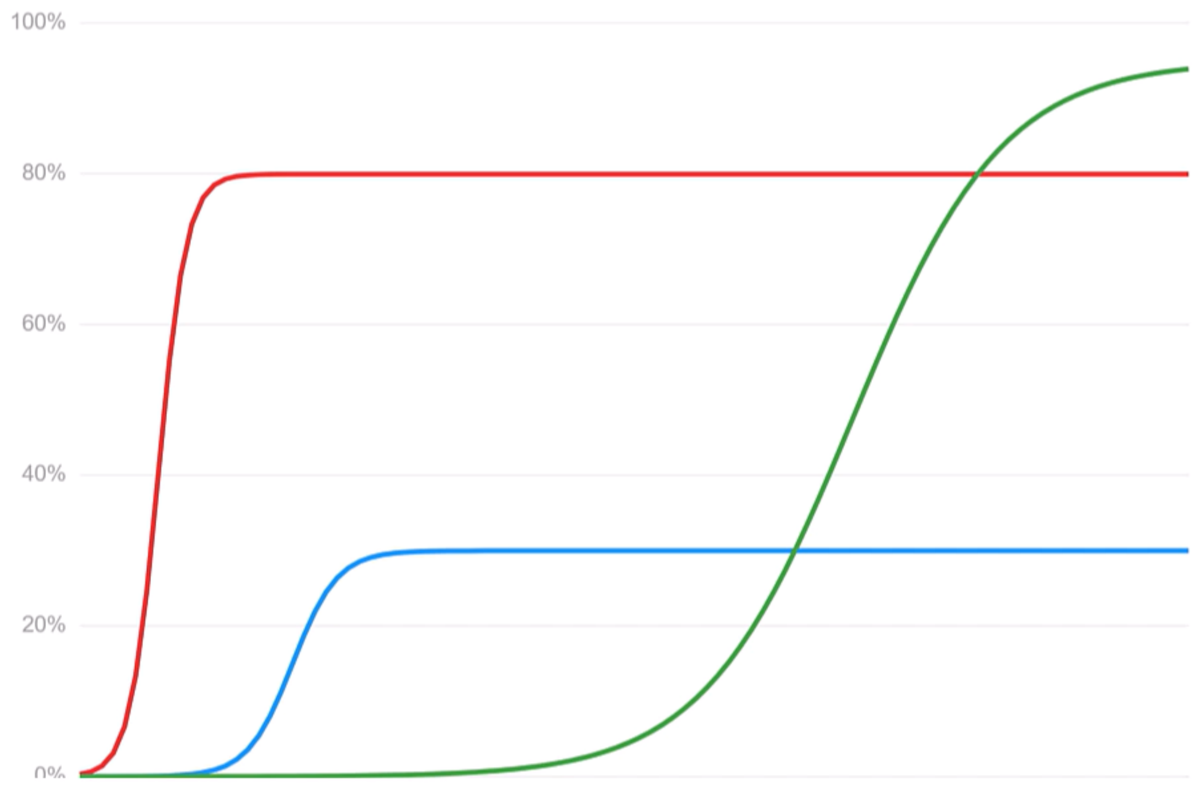

Grokking: From Memorization to Understanding

- Training accuracy ON
- Test accuracy ON
- Memorization ON
- Fast heuristic ON
- True understanding ON

Well-scaffolded Teaching 0%

Memorization

"How are you?" *I am fine, thank you*
"How is your mother?" *Huh?*

Fast Heuristic

Rule: *add -ed for past tense walked, talked ✓*
But: *goed, runned ✗*

True Understanding

Handles anything new—even sentences never heard before

Accuracy Over Training dashed = train, solid = test

Pattern Strength Over Training what's happening inside the learner?

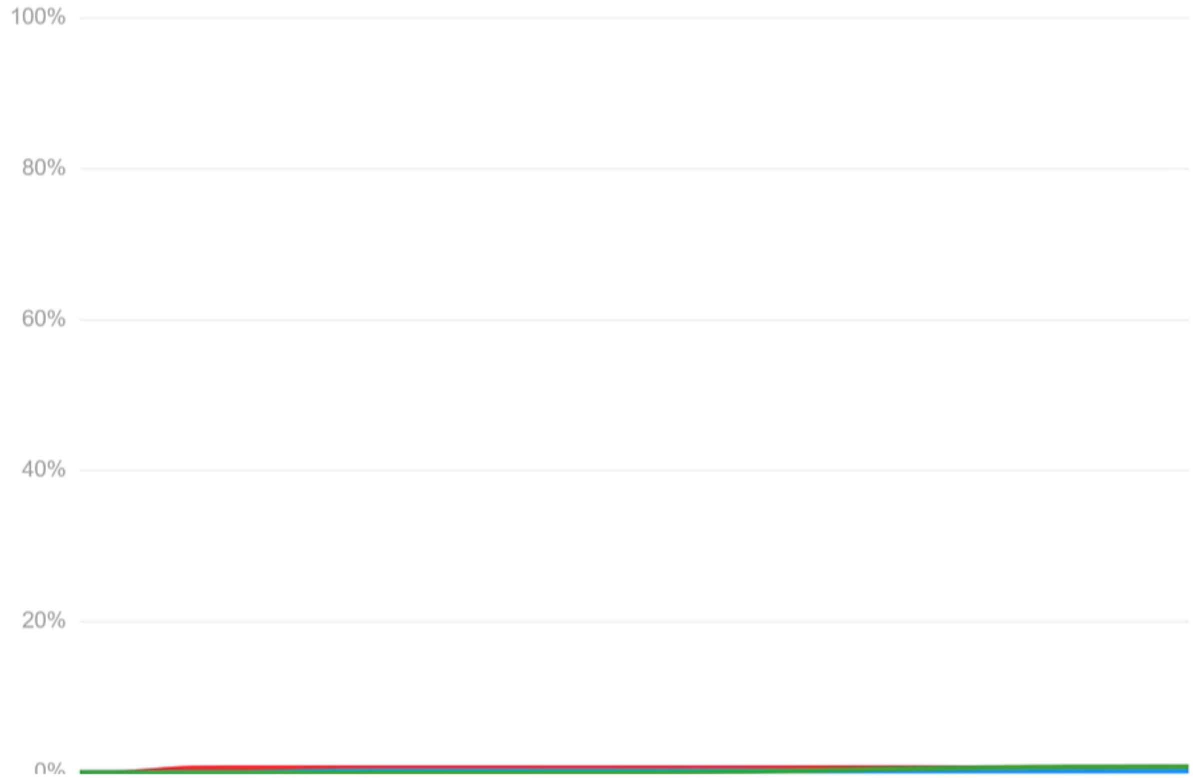

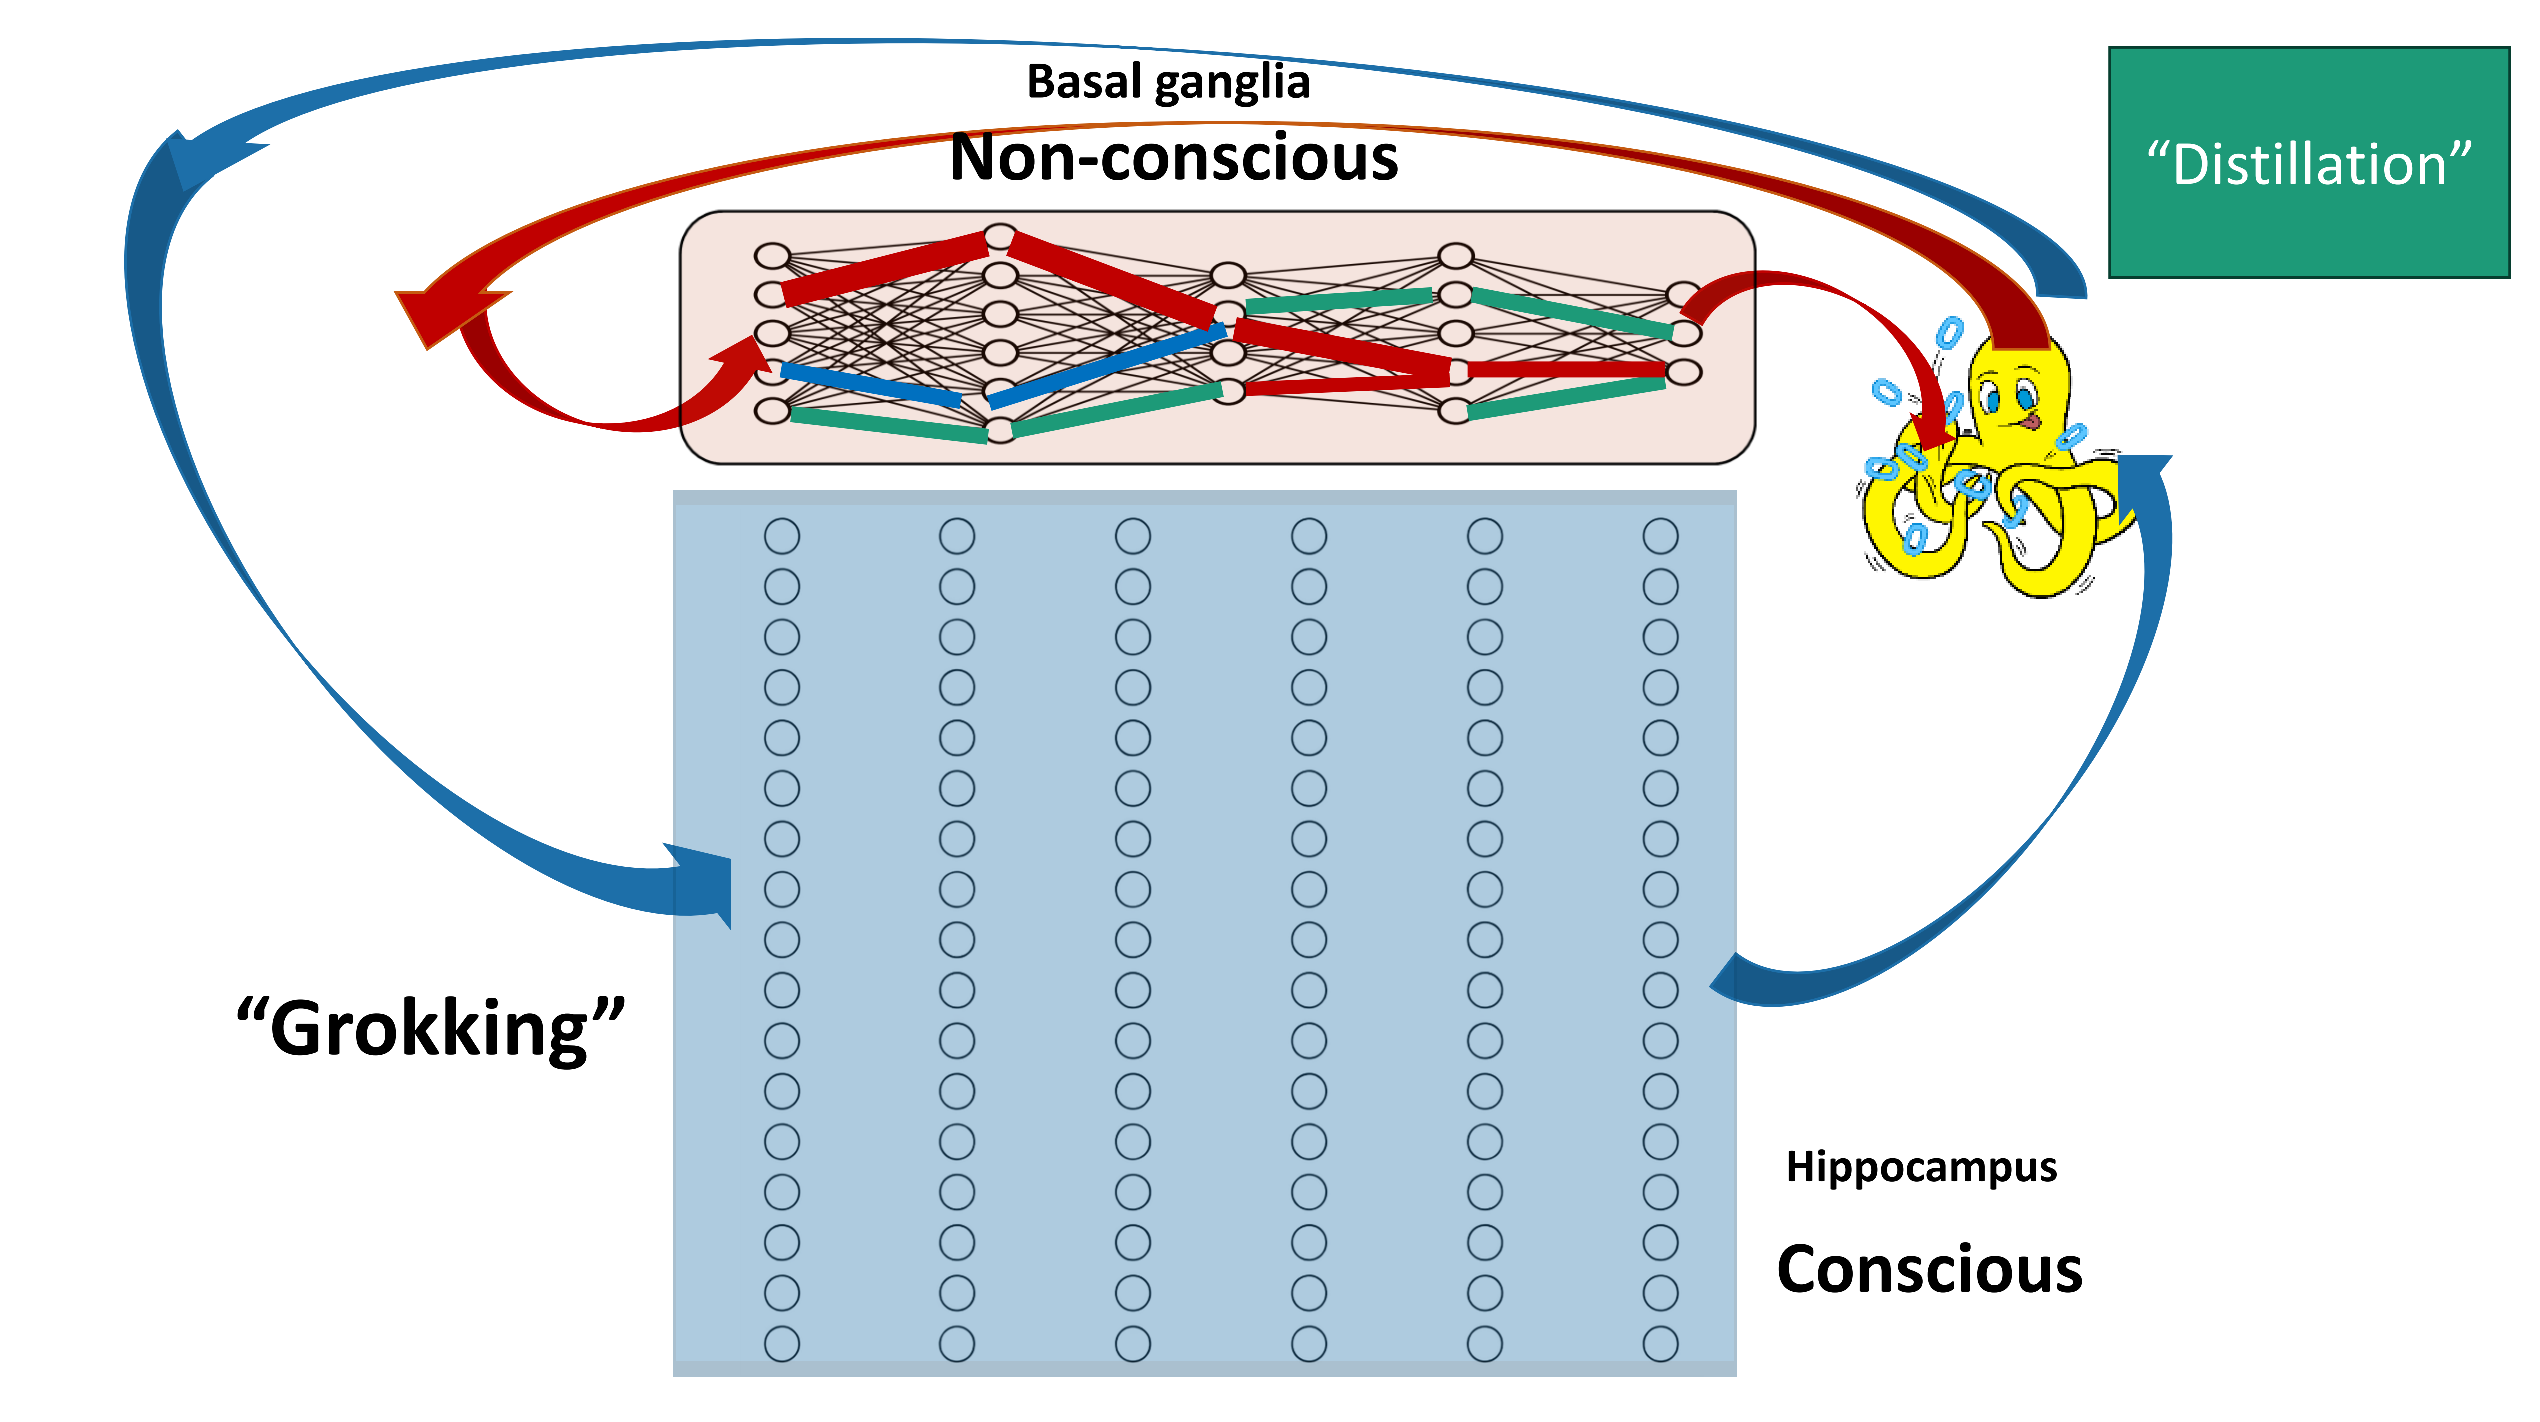
Basal ganglia
Non-conscious
"Distillation"
"Grokking"
Hippocampus
Conscious

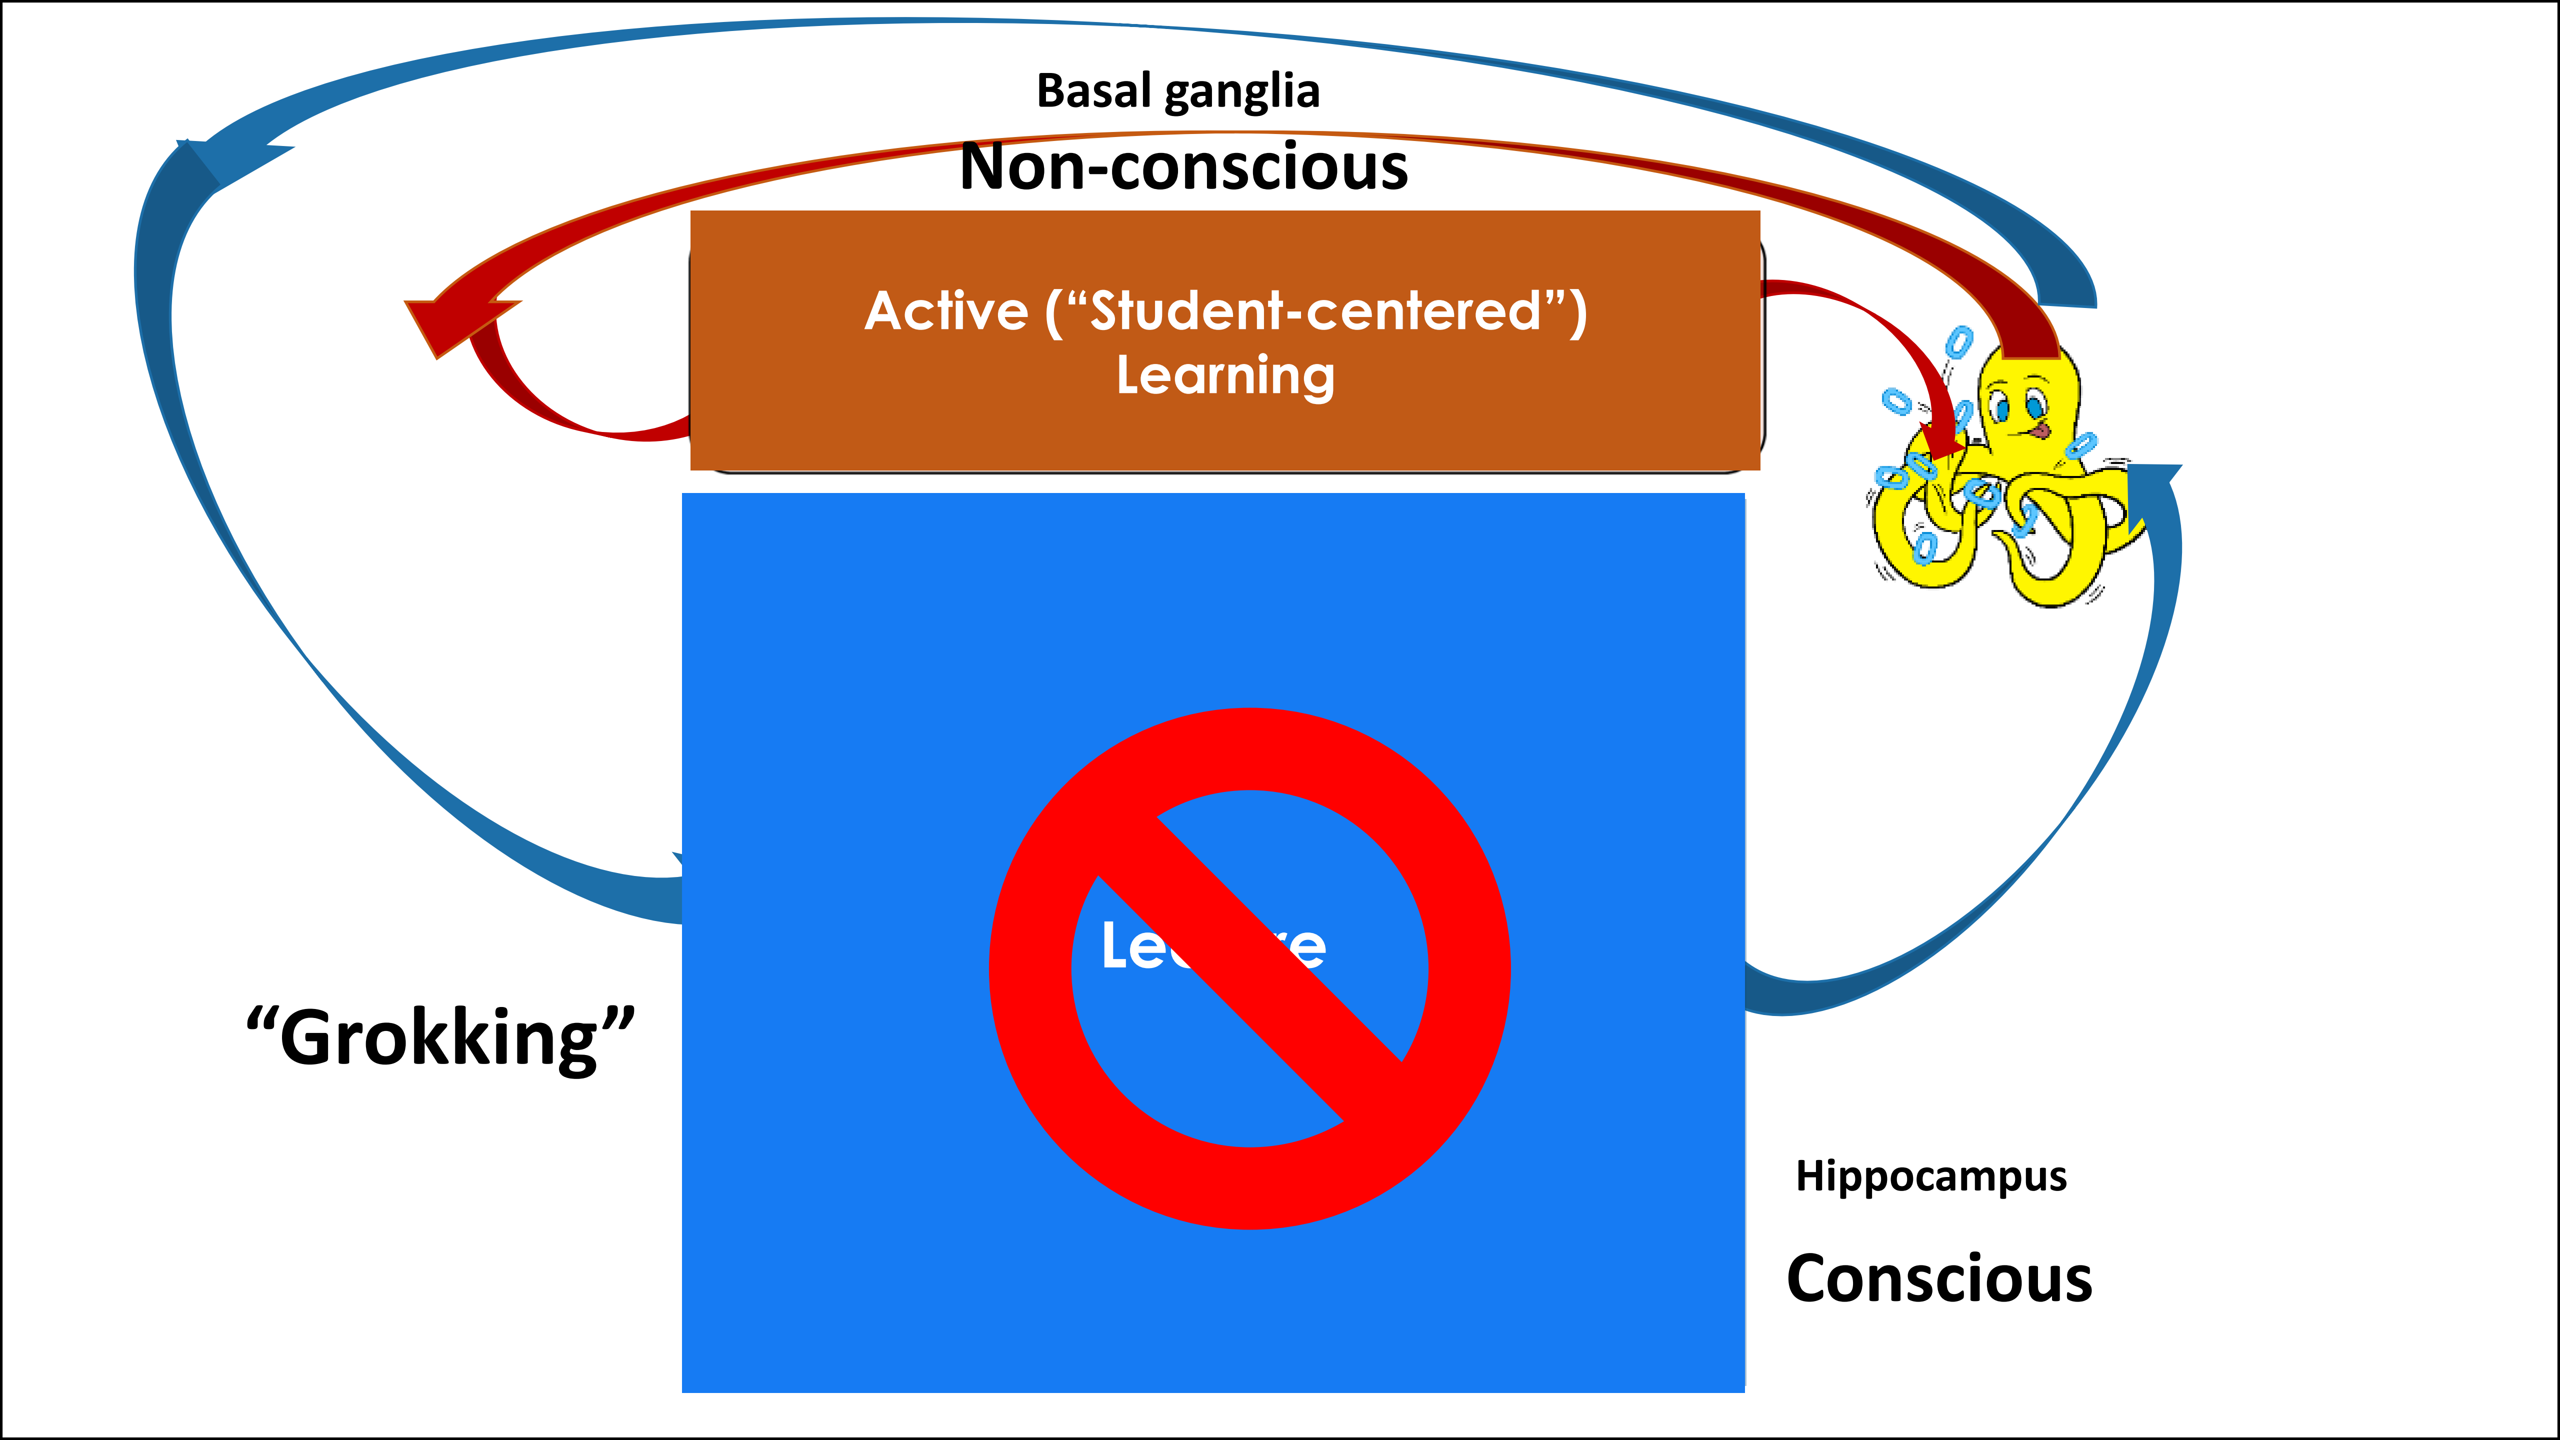
Basal ganglia
Non-conscious
Active (“Student-centered”)
Learning
“Grokking”
Hippocampus
Conscious

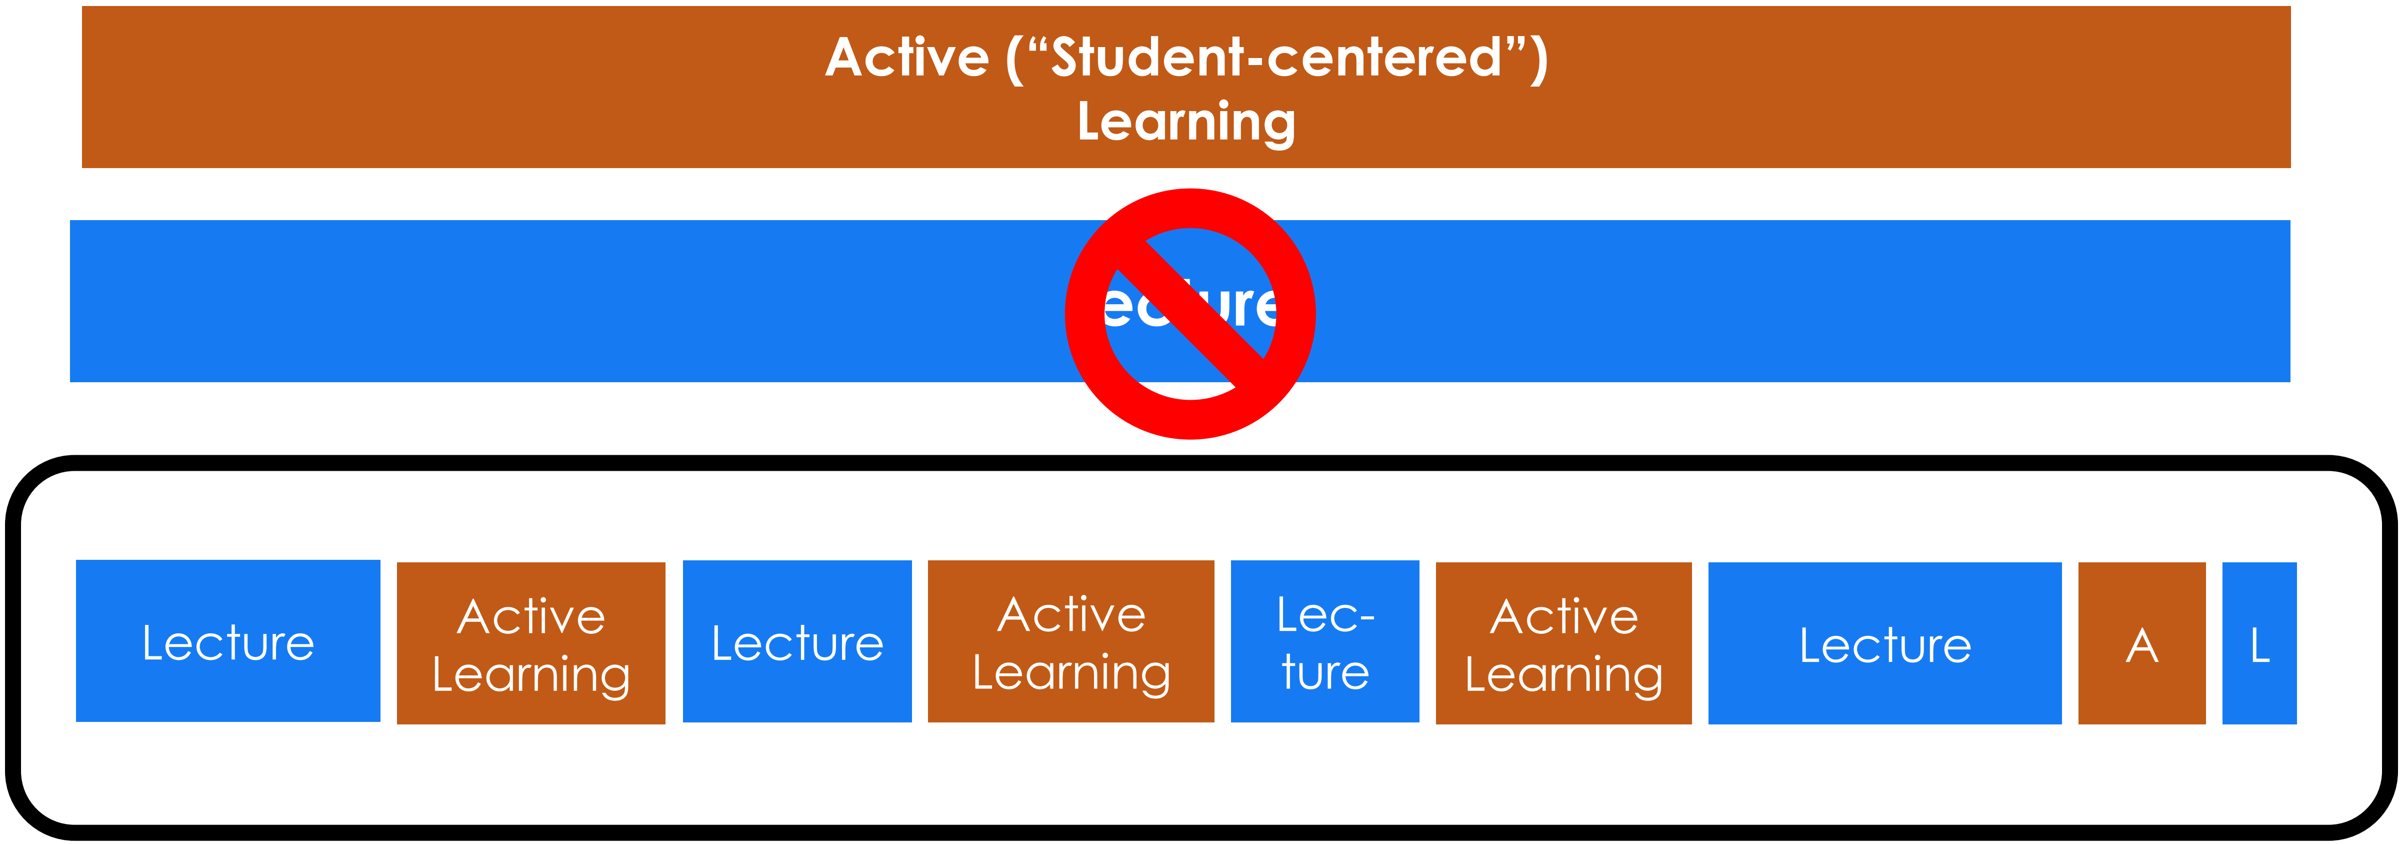

Direct instruction

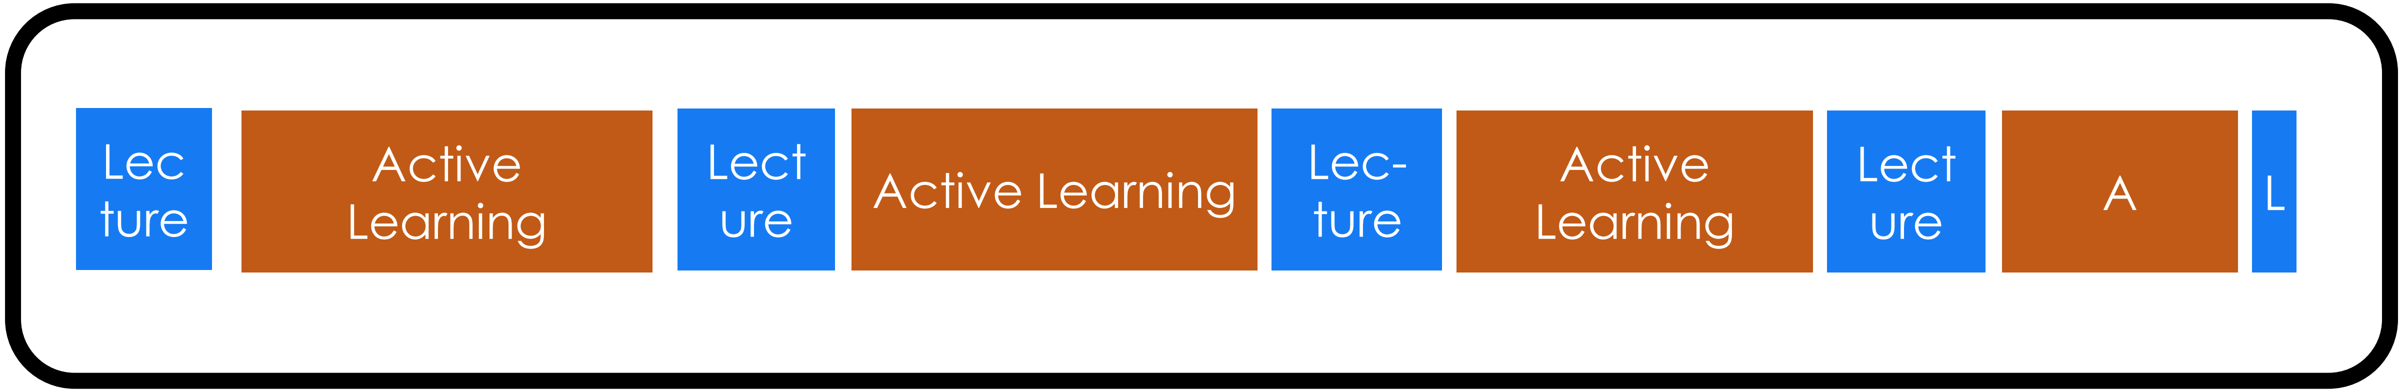

Direct instruction

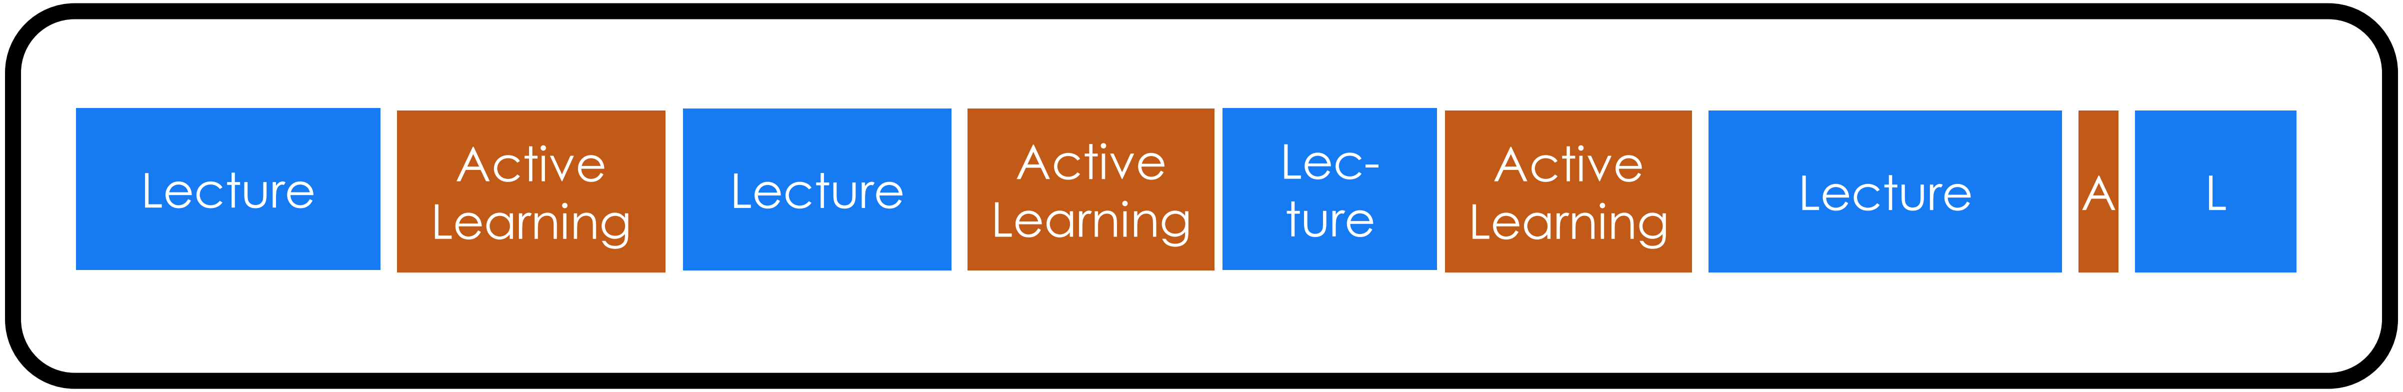

Direct instruction

The "Activity Ratchet"

The Flynn Effect

Flynn, James R. "The Mean IQ of Americans: Massive Gains 1932 to 1978." *Psychological Bulletin 95, no. 1 (1984): 29.*

- Dutton, Edward, Dimitri van der Linden, and Richard Lynn. "The Negative Flynn Effect: A Systematic Literature Review." *Intelligence 59 (2016): 163-69.*
- Bratsberg, Bernt, and Ole Rogeberg. "Flynn Effect and Its Reversal Are Both Environmentally Caused." *Proceedings of the National Academy of Sciences* 115, no. 26 (2018):
- Dworak, Elizabeth M., William Revelle, and David M. Condon. "Looking for Flynn Effects in a Recent Online U.S. Adult Sample: Examining Shifts within the SAPA Project." *Intelligence* 98 (2023).

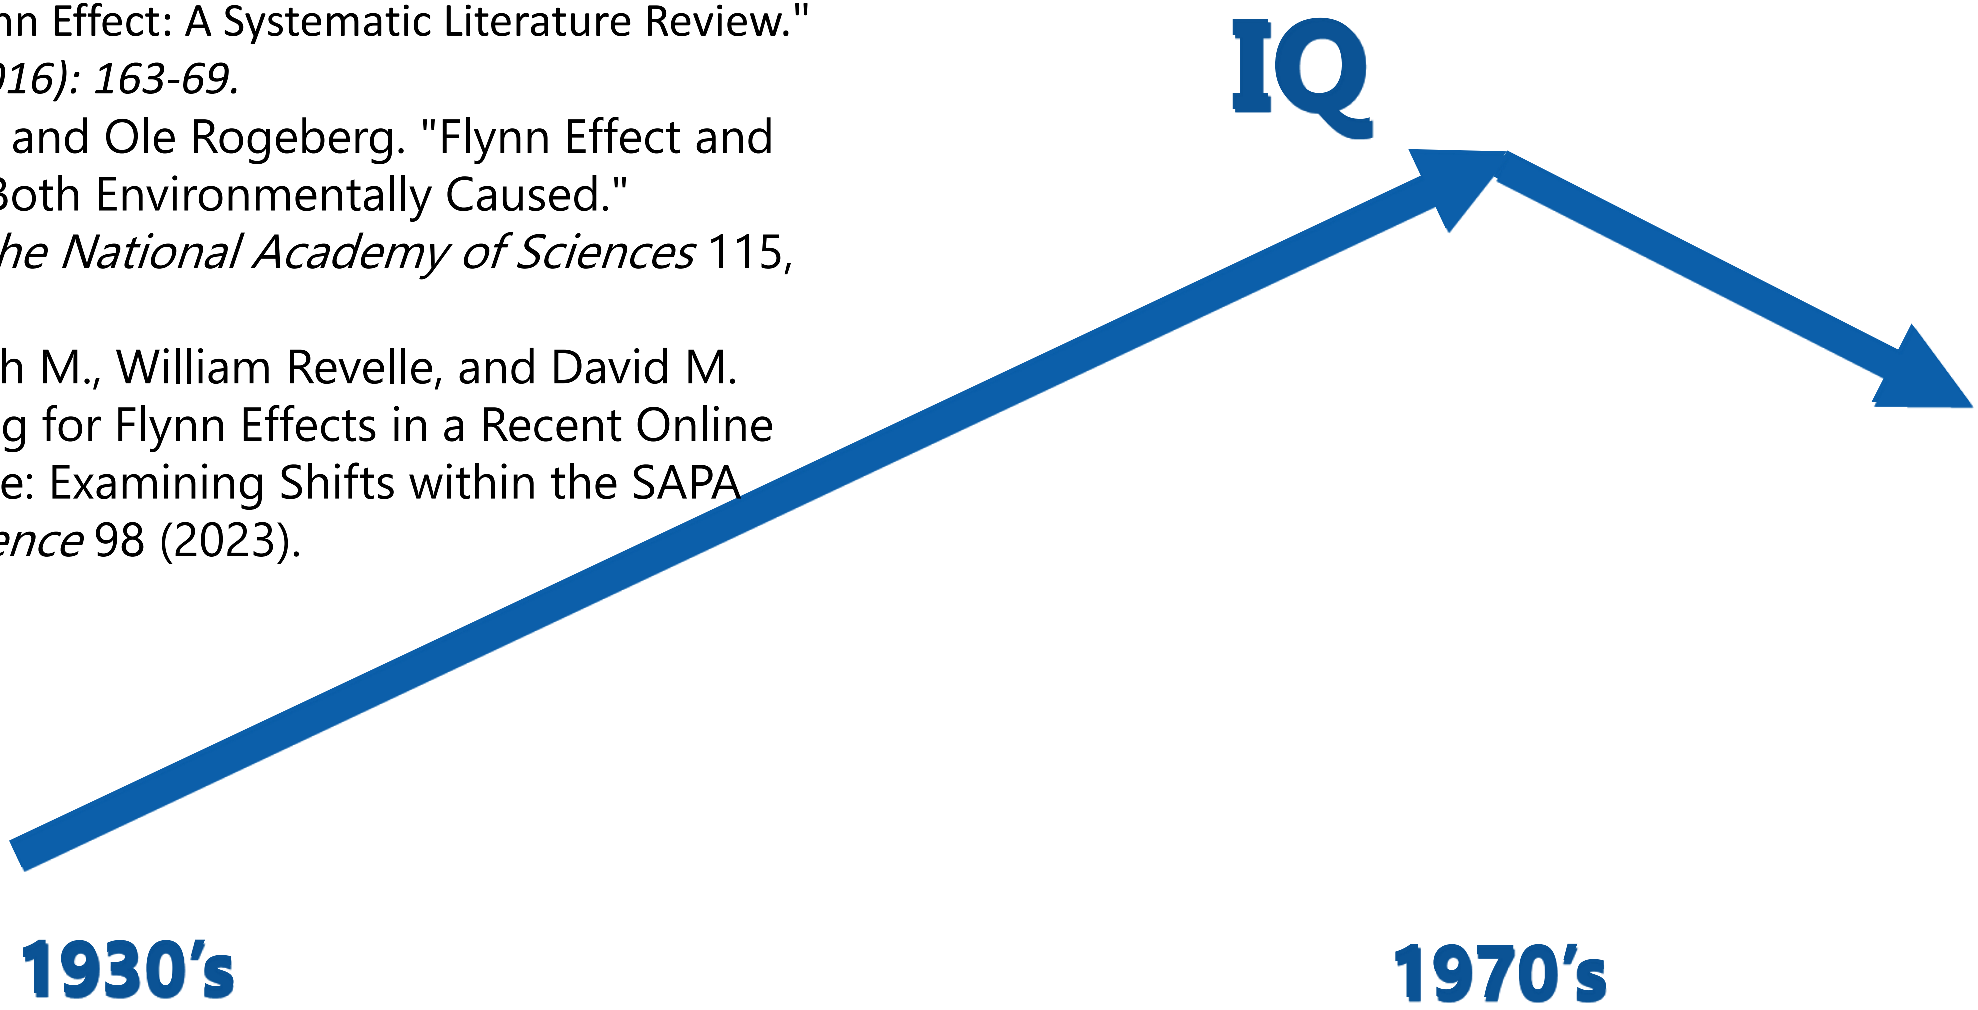

What will our future hold?

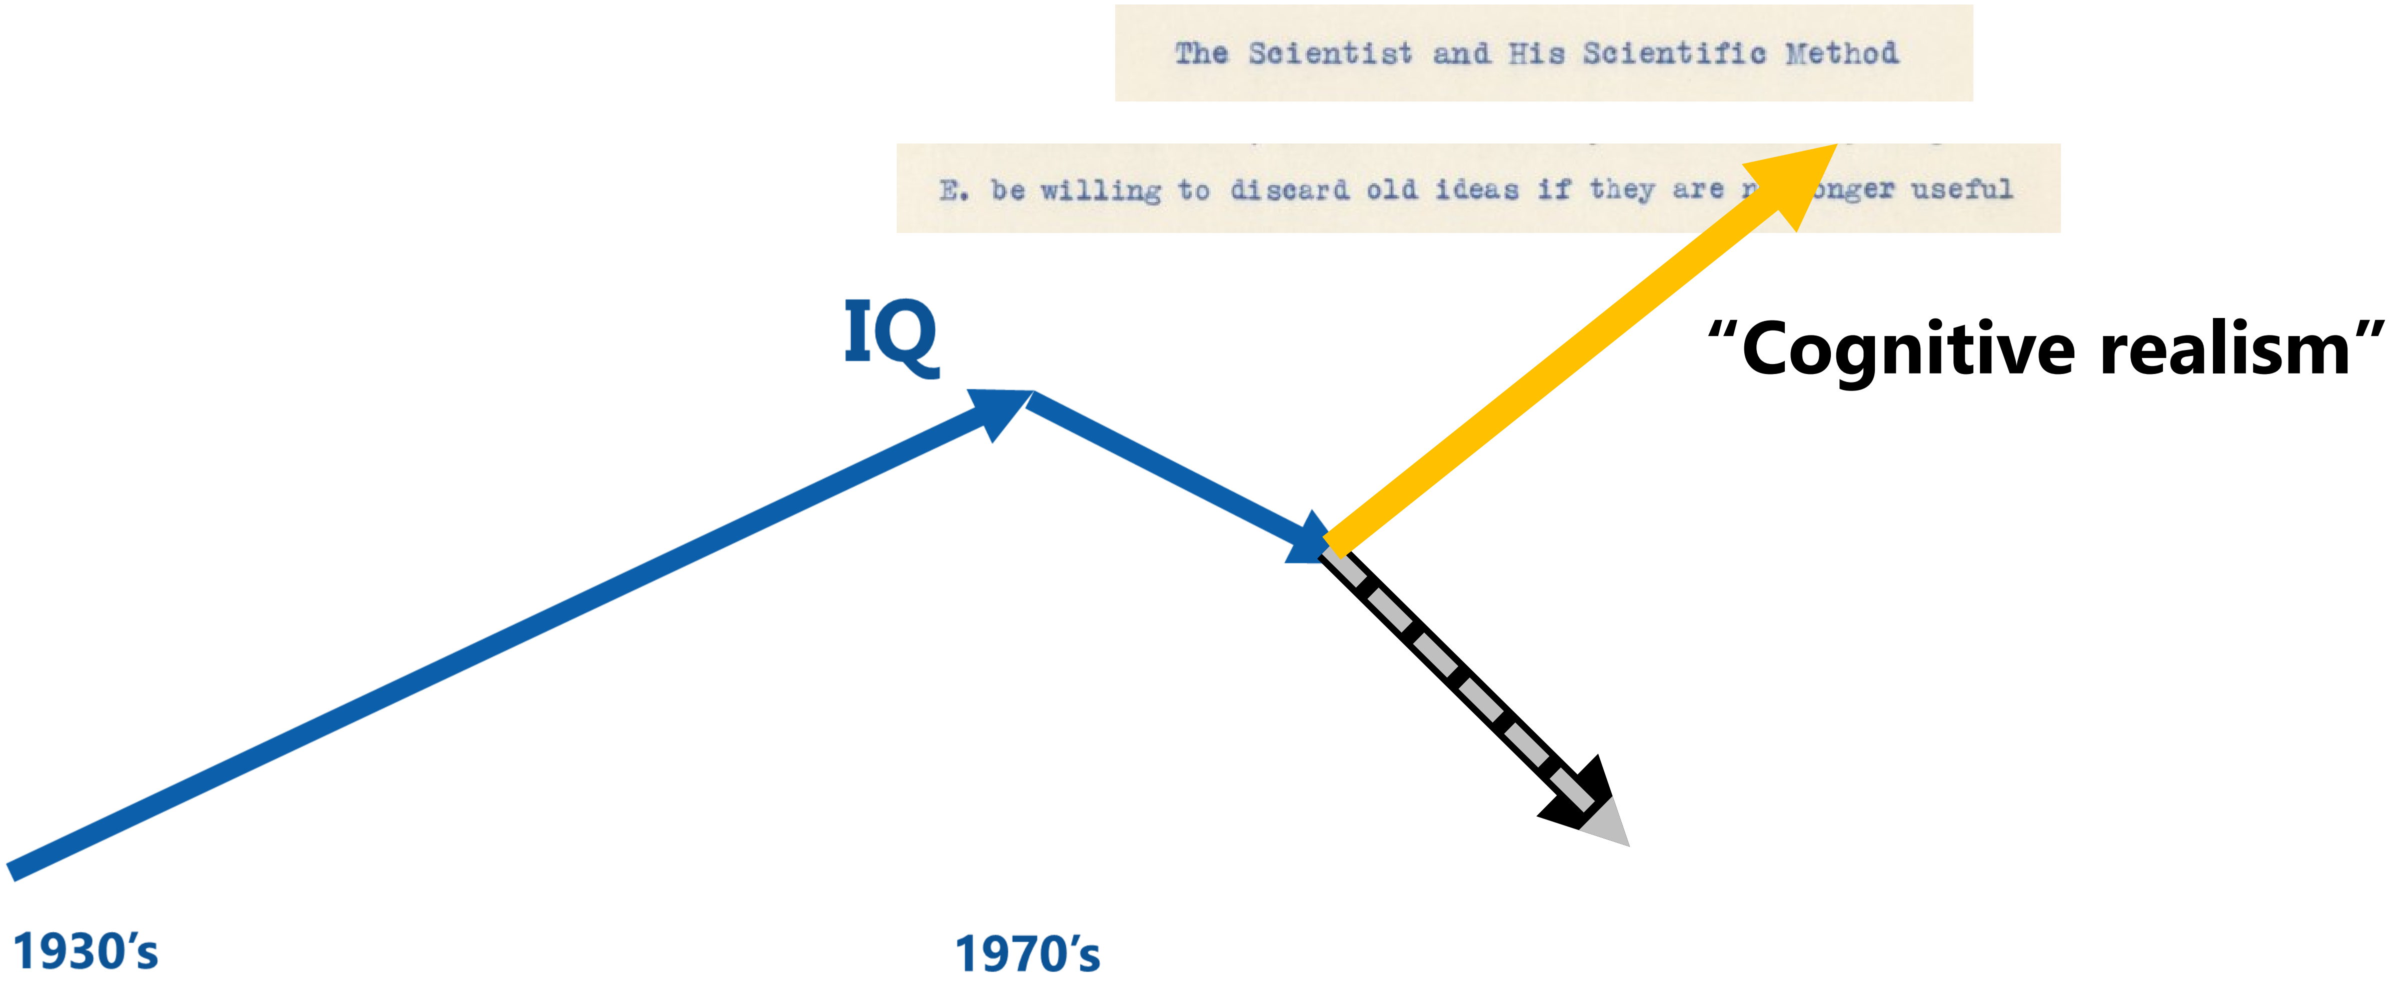

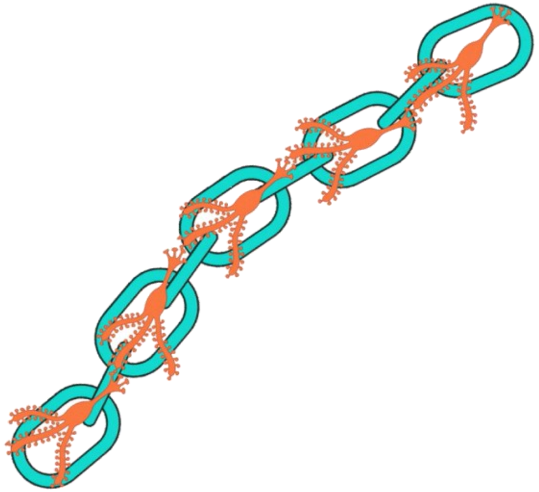

"Cognitive realism" tells us:

- **Build the links.** *Foundational knowledge is vital.*
- **Test frequently.** *Frequent low-stakes tests with no access to AI.*
- **Spice *enhances*—it can't carry the meal**

Carefully designed explicit instruction

The main dish

Lecture	Active Learning	Lecture	Active Learning	Lec-ture	Active Learning	Lecture	A	L

Student-centered learning is important…

Great teaching matters too.

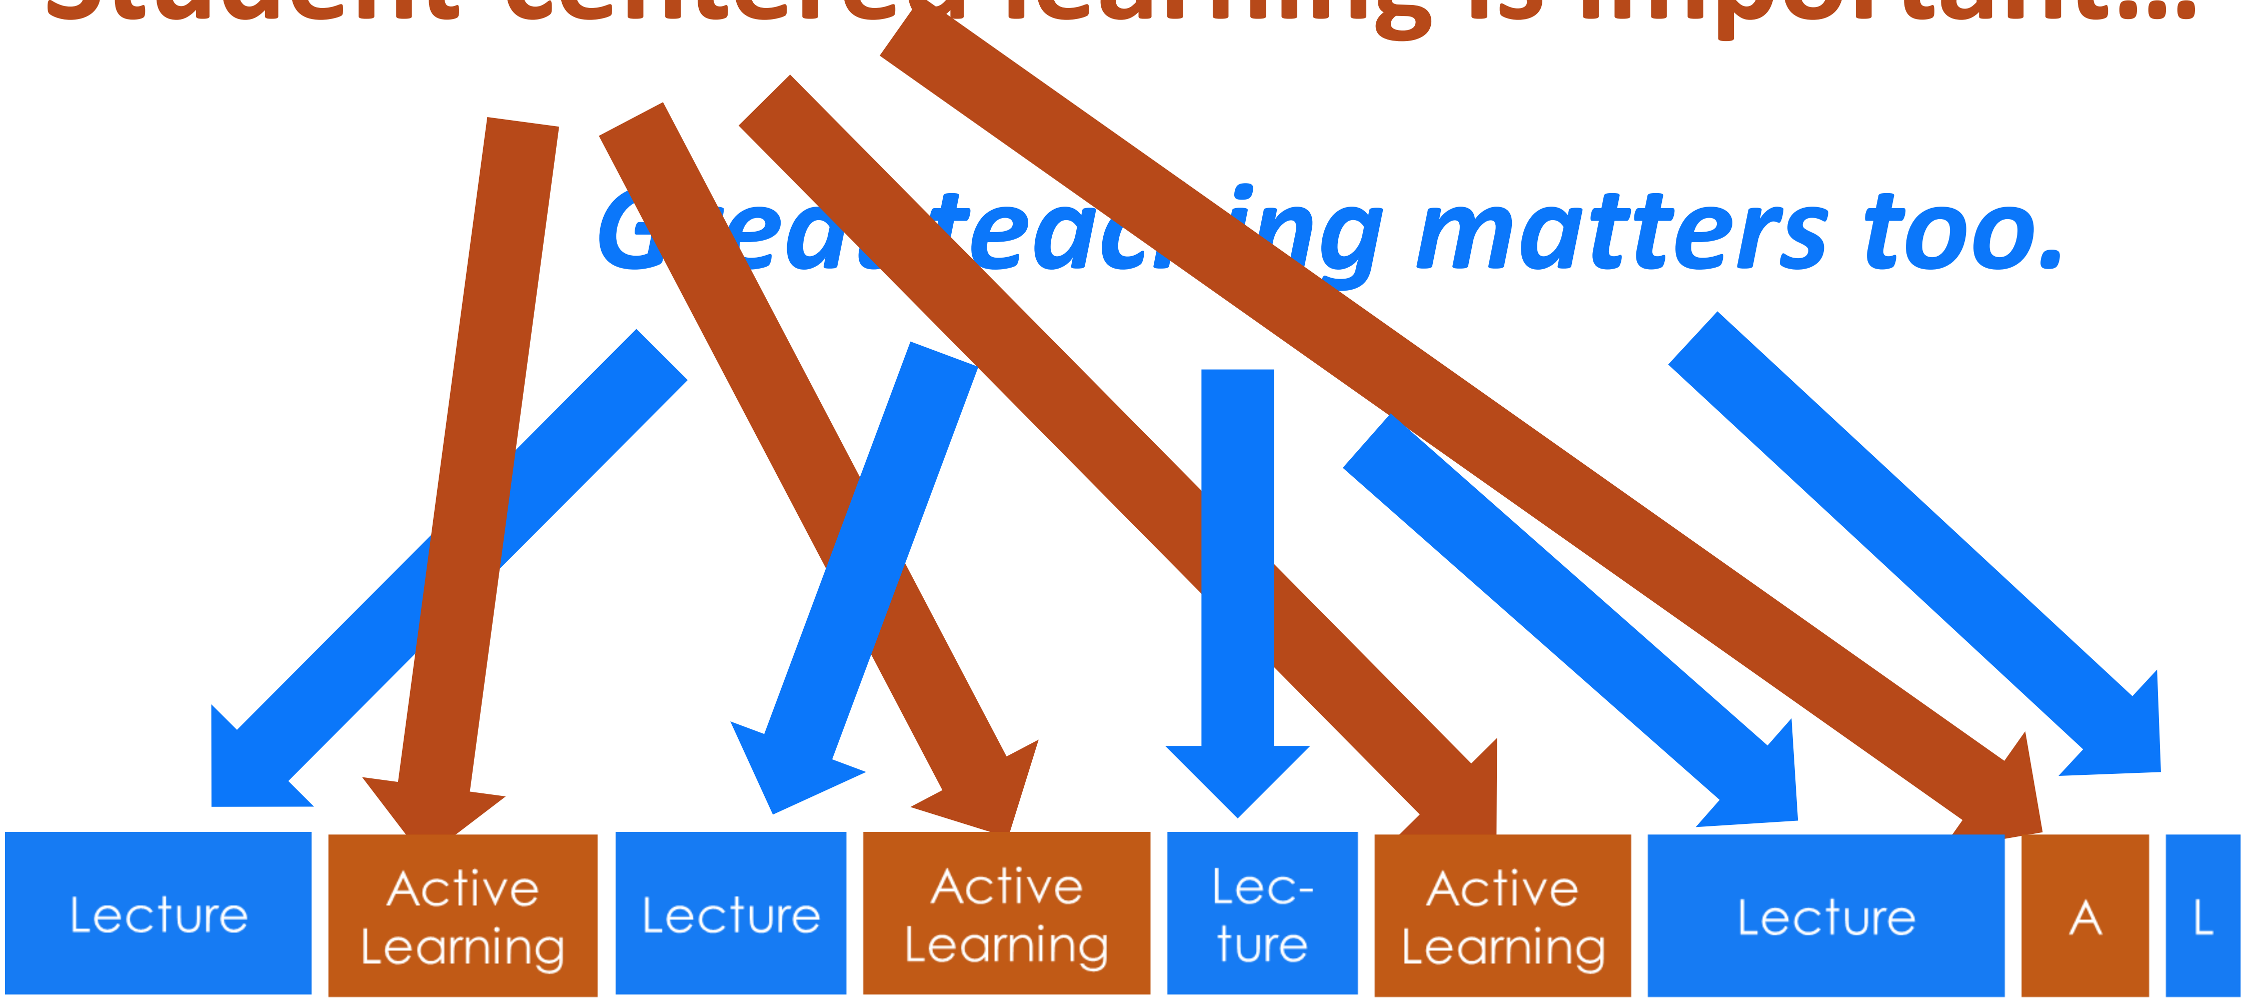

“Cognitive realism”