



## OPEN ACCESS

## EDITED BY

Jorge F. Mejias,  
University of Amsterdam, Netherlands

## REVIEWED BY

Mary Alexandria Kelly,  
Carleton University, Canada

## \*CORRESPONDENCE

Claudius Gros  
✉ gros07@itp.uni-frankfurt.edu

RECEIVED 04 October 2025

REVISED 04 November 2025

ACCEPTED 10 November 2025

PUBLISHED 24 November 2025

## CITATION

Gros C (2025) From generative AI to the brain:  
five takeaways.  
*Front. Comput. Neurosci.* 19:1718778.  
doi: 10.3389/fncom.2025.1718778

## COPYRIGHT

© 2025 Gros. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](#). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

# From generative AI to the brain: five takeaways

Claudius Gros\*

Institute for Theoretical Physics, Goethe University, Frankfurt, Germany

The big strides seen in generative AI are not based on somewhat obscure algorithms, but due to clearly defined generative principles. The resulting concrete implementations have proven themselves in large numbers of applications. We suggest that it is imperative to thoroughly investigate which of these generative principles may be operative also in the brain, and hence relevant for cognitive neuroscience. In addition, ML research led to a range of interesting characterizations of neural information processing systems. We discuss five examples, the shortcomings of world modeling, the generation of thought processes, attention, neural scaling laws, and quantization, that illustrate how much neuroscience could potentially learn from ML research.

## KEYWORDS

generative AI, cognitive neuroscience, attention, predictive coding, chain of thought, quantization

## 1 Introduction

A multitude of factors contributes to the current rise of generative artificial intelligence (generative AI). Here we focus on two aspects.

- Algorithmic developments can be formulated in many cases in terms of generic generative principles. These generative principles have proven themselves, giving rise to high-performing machine learning architectures. It is an important question whether corresponding principles may operate in the brain.
- In addition to algorithms, insights regarding general working principles and properties of neural-based information processing systems have been attained. Do these apply also to the human brain?

Machine learning (ML) offers a range of conjectures for the workings of our brain, some of which extend or parallel traditional neuroscience frameworks, while others are new. Cognitive neuroscience should accept the challenge and evaluate these conjectures systematically in the context of wet information processing.

A comprehensive overview of potentially relevant cross-relations between ML and the neurosciences is beyond the scope of this perspective. We will focus instead on five key aspects elucidating the importance of paying attention to the concepts that are being developed for generative artificial intelligence. A flurry of new ideas awaits the scrutiny of cognitive neuroscience.

## 2 World modeling is not enough

The two learning principles, “predictive coding” (neuroscience) and “autoregressive language modeling” are both dedicated to the task of building world models, with the

former also having active components (Brodski-Guerniero et al., 2017), operating in addition on distinct scales and modalities (Caucheteux et al., 2023).

For large language models (LLMs), autoregressive language modeling takes the form of next-word predictions. However, ML tells us that world-model building alone is insufficient.

The base- or foundation model, viz the result of word prediction training, does contain the knowledge of the world, as present in the training data. But all it can do is to complete a given input word by word. At this stage, key concepts of relevance for the interaction with users, such as “question” and “answer”, are not yet explicitly encoded.

In the early 2020s, a significant step forward was the realization that the otherwise essentially useless base model can be turned into a cognitive powerhouse via a secondary process, denoted “fine tuning” or “human supervised fine-tuning” (HSFT).

- A core fine tuning objective is to teach the system to generate meaningful responses for a given prompt, and not just engage in text completion.
- Next comes fine tuning of style, political correctness, etc.
- Models may be fine tuned further for specific downstream tasks, specializing the otherwise universal LLM to excel, e.g., in accounting.

It seems likely that equivalent processes would occur in our brains. In ML, the two processes are normally separated, viz performed subsequently. In the brain, world model training and fine-tuning via reinforcements are conceivably active at the same time.

**Takeaway:** ML offers a concrete construction plan for a basic cognitive system: Universal unsupervised world modeling followed by supervised fine tuning. To which extent does the brain follow this recipe?

### 3 Generative principles for human thinking

The autonomous generation of thoughts is considered to be the basis of human intelligence. It is hence remarkable that commercial chatbots started to engage in rudimentary “thinking” by the mid-2020s. The algorithm used is denoted “Chain-of-Thought” (CoT) (Zhang et al., 2025), originally a prompting technique (Wei et al., 2022). It is unclear to which extent human thought processes may be understood within the CoT framework, if at all. The same holds for its generalizations, viz “Chain-of-X” (CoX) (Xia et al., 2024), such as Chain-of-Feedback, Chain-of-Instructions, or Chain-of-Histories. In any case, of interest are the underlying generative principles.

- CoT is one of many possible fine-tuning processes, characterized by a specific objective function.
- The system auto-prompts, appending its own thoughts to the user prompt.

- The response is then generated using the combined prompt: (user input)+(chain of thoughts).

Why are responses substantially better when the LLM thinks for a while? A possible explanation is based on the information bottleneck (IB) framework (Tishby and Zaslavsky, 2015). We recall that the token sequence is



The middle part, the thought processes, can be interpreted to act as an information bottleneck for the cognitive processing between input and output (Lei et al., 2025). This principle can be expressed as an information-theoretical min-max optimization.

- The mutual information between the input and CoT is – minimized–. This means that the self-generated thoughts should abstract from the specific formulation of the input, retaining only the overall content.
- The mutual information between CoT and the output is – maximized–. This because the latent space, namely the thoughts, should be maximally informative with regard to the final output.

The IB view is not just an abstract cookbook. Instead, it has proven itself as a high-performing training algorithm (Lei et al., 2025).

As an alternative to the notion of an information bottleneck, it has been proposed that CoT-type thinking may be seen as an effort to build a composite object, the final response (Zhu et al., 2025). This interpretation allows to leverage state-of-the-art algorithms for diverse object generation, such as GFlowNet (Bengio et al., 2021), which is used widely in synthetic chemistry.

**Takeaway:** The new views of the functionality of thought processes arising within ML research should motivate us to pose one of the most fundamental questions the neurosciences could consider. Could this provide a possible first step toward an understanding of human thinking?

### 4 No attention without self-consistency

A main driver of generative AI is the self-attention mechanism powering the transformer architecture (Vaswani et al., 2017; de Santana Correia and Colombini, 2022). Regarding the brain, we do not touch here the phenomenology of human attention, or the sometimes controversially discussed question how attention should be defined operatively in psychology and in the neurosciences (Hommel et al., 2019; Wu, 2024). Given this caveat, a few comments can be made:

- The details of how attention works on the level of individual neurons are generally not well understood (Moore and Zirnsak, 2017).
- Top-down attention involves the modulation of sensory processing areas by signals generated in higher brain regions.

It is generally assumed that these modulatory processes depend only on the specific attention signals, viz without being coupled to the actual process generating the top-down signal in the first place.

- Bottom-up attention is observed when early areas react to salient features in the sensory input stream. Pop-out features are then forwarded with higher intensity (Connor et al., 2004).

Bottom-up attention could be interpreted as a variant of self-attention, albeit with a reduced dynamic range. The latter because lateral saliency detection evolves only slowly in the adult brain (Hopfinger, 2017).

Operatively, there is a key difference between the current view of top-down and ML attention. In the neurosciences, attention processes operating in early brain regions are investigated separately from the question of how the modulating top-down signals are generated via cognitive control (Badre, 2024), e.g., in the context of cholinergic signaling (Parikh and Bangasser, 2020). No such separation is present in machine learning, for a good reason. Components of larger models develop their own neural language when trained separately (Ludueña and Gros, 2013); models need therefore to be trained in their entirety for the individual components to be able to talk to each other.

The context window of a transformer represents past states, which implies that the self-attention mechanism discussed above operates in the time domain, involving hence working-memory aspects (Hintzman, 1984). This connection has been addressed in the context of modern Hopfield networks (Ramsauer et al., 2020; Ororbia and Kelly, 2023).

**Takeaway:** A full understanding of attention needs to include the self-consistency loop between the generation of attention signals and their subsequent processing.

## 5 Neural scaling laws

Neural scaling laws describe how performance and training times scale with model and/or data size (Kaplan et al., 2020; Hoffmann et al., 2022; Michaud et al., 2023).

As an example consider the relative performance of two fully trained models (A and B), which are identical in all aspects, apart from model sizes. In good approximation the relative performance is then (Neumann and Gros, 2022, 2024)

$$\frac{P_A}{P_A + P_B} \sim \frac{N_A}{N_A + N_B} \quad (1)$$

where  $N_A$  ( $N_B$ ) are the respective numbers of adaptable parameters, and  $P_A$  ( $P_B$ ) the corresponding performances. In addition, larger systems need longer to train. For the training compute  $C$ , viz for the total amount of resources (chips, time, ...) needed to train a model, one finds a quadratic scaling relation,

$$C \sim N^2.$$

This leads to an interesting hypothesis regarding putative limitations to the phylogenetic growth of the brain. Assume we have two humans, H1 and H2, the first with a standard brain size, the second with a brain twice as large. If we take 15 years as the

time to train standard human brains, H2 would need  $2^2 = 4$  times as long, namely 60 years. After 60 years of growing up, H2 would have higher cognitive capabilities than H1, as given by Equation 1. However, ML finds that performance remains somewhat flat during most training, increasing rapidly only at later stages. This implies that H2 would underperform H1 for extended periods, say the first 40–50 years of adolescence. Evolutionary speaking, doubling brain size may hence not be a viable option. Of course other limiting factors, like metabolic costs, may have determined the size of our brains.

**Takeaway:** Given that information processing networks characterize not only modern machine learning architectures, but also the brain, the biological implications of neural scaling deserve to be investigated.

## 6 Quantization

Large models have large numbers of adaptable parameters, which one needs not only to store, but also to keep in working memory, ready for subsequent use. Typical floating point datatypes are 32 bits (or 64 bits for double precision). In order to save memory, and to make operations faster, data sizes have been reduced in recent years (Wei et al., 2024; Gong et al., 2025). Currently, INT4 (4 bits) devices are being rolled out. For INT4, one has just  $2^4 = 16$  possible values. Synaptic weights are hence “quantized”, taking only one out of 16 possible states. Specialized GPUs support the involved operations efficiently. In analogy, synaptic strength is quantized also in the brain (Petersen et al., 1998; Liu et al., 2017), with the exact number of expressed states being debated.

**Takeaway:** The computational consequences of synapse quantization are well understood for artificial neural nets. This knowledge should be readily transferable to their biological counterparts.

## 7 Conclusions

We reviewed five selected concepts contributing to the rapid progress of generative AI. Interestingly, in machine learning literature, their relevance to biological information processing systems is rarely discussed, if at all, with attention being in part an exception (Lindsay, 2020). The scope of this perspective is to raise awareness that a treasure of generative principles may be hidden in ML literature.

## Data availability statement

The original contributions presented in the study are included in the article/supplementary material, further inquiries can be directed to the corresponding author.

## Author contributions

CG: Conceptualization, Writing – original draft, Writing – review & editing.

## Funding

The author(s) declare that no financial support was received for the research and/or publication of this article.

## Acknowledgments

This article benefited from discussions with Christian Fiebach, Ricardo Kienitz, and Búlsu Sándor.

## Conflict of interest

The author declares that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

## References

- Badre, D. (2024). Cognitive control. *Annu. Rev. Psychol.* 76, 167–195. doi: 10.1146/annurev-psych-022024-103901
- Bengio, E., Jain, M., Korablyov, M., Precup, D., and Bengio, Y. (2021). Flow network based generative models for non-iterative diverse candidate generation. *Adv. Neural Inf. Process. Syst.* 34, 27381–27394. doi: 10.5555/3540261.3542358
- Brodski-Guerniero, A., Paasch, G.-F., Wollstadt, P., Özdemir, I., Lizier, J. T., and Wibral, M. (2017). Information-theoretic evidence for predictive coding in the face-processing system. *J. Neurosci.* 37, 8273–8283. doi: 10.1523/JNEUROSCI.0614-17.2017
- Caucheteux, C., Gramfort, A., and King, J.-R. (2023). Evidence of a predictive coding hierarchy in the human brain listening to speech. *Nat. Human Behav.* 7, 430–441. doi: 10.1038/s41562-022-01516-2
- Connor, C. E., Egeth, H. E., and Yantis, S. (2004). Visual attention: bottom-up versus top-down. *Curr. Biol.* 14, R850–R852. doi: 10.1016/j.cub.2004.09.041
- de Santana Correia, A., and Colombari, E. L. (2022). Attention, please! A survey of neural attention models in deep learning. *Artif. Intellig. Rev.* 55, 6037–6124. doi: 10.1007/s10462-022-10148-x
- Gong, R., Ding, Y., Wang, Z., Lv, C., Zheng, X., Du, J., et al. (2025). A survey of low-bit large language models: Basics, systems, and algorithms. *Neural Netw.* 192:107856. doi: 10.1016/j.neunet.2025.107856
- Hintzman, D. L. (1984). Minerva 2: A simulation model of human memory. *Behav. Res. Methods, Instrum. Comp.* 16, 96–101. doi: 10.3758/BF03202365
- Hoffmann, J., Borgeaud, S., Mensch, A., Buchatskaya, E., Cai, T., Rutherford, E., et al. (2022). Training compute-optimal large language models. *arXiv [Preprint]*. arXiv:2203.15556. doi: 10.48550/arXiv.2203.15556
- Hommel, B., Chapman, C. S., Cisek, P., Neyedli, H. F., Song, J.-H., and Welsh, T. N. (2019). No one knows what attention is. *Attent. Percept. Psychophys.* 81, 2288–2303. doi: 10.3758/s13414-019-01846-w
- Hopfinger, J. B. (2017). Introduction to special issue: attention & plasticity. *Cogn. Neurosci.* 8, 69–71. doi: 10.1080/17588928.2017.1284775
- Kaplan, J., McCandlish, S., Henighan, T., Brown, T. B., Chess, B., Child, R., et al. (2020). Scaling laws for neural language models. *arXiv [Preprint]*. arXiv:2001.08361. doi: 10.48550/arXiv.2001.08361
- Lei, S., Cheng, Z., Jia, K., and Tao, D. (2025). Revisiting llm reasoning via information bottleneck. *arXiv [Preprint]*. arXiv:2507.18391. doi: 10.48550/arXiv.2507.18391
- Lindsay, G. W. (2020). Attention in psychology, neuroscience, and machine learning. *Front. Comput. Neurosci.* 14:29. doi: 10.3389/fncom.2020.00029
- Liu, K. K., Hagan, M. F., and Lisman, J. E. (2017). Gradation (approx. 10 size states) of synaptic strength by quantal addition of structural modules. *Philos. Trans. R. Soc. Lond. B. Biol. Sci.* 372:20160328. doi: 10.1098/rstb.2016.0328
- Ludue na, G. A., and Gros, C. (2013). A self-organized neural comparator. *Neural Comput.* 25, 1006–1028. doi: 10.1162/NECO\_a\_00424
- Michaud, E., Liu, Z., Girit, U., and Tegmark, M. (2023). The quantization model of neural scaling. *Adv. Neural Inf. Process. Syst.* 36, 28699–28722. doi: 10.5555/3666122.3667370
- Moore, T., and Zirnsak, M. (2017). Neural mechanisms of selective visual attention. *Annu. Rev. Psychol.* 68, 47–72. doi: 10.1146/annurev-psych-122414-033400
- Neumann, O., and Gros, C. (2022). Scaling laws for a multi-agent reinforcement learning model. *arXiv [Preprint]*. arXiv:2210.00849. doi: 10.48550/arXiv.2210.00849
- Neumann, O., and Gros, C. (2024). Alphazero neural scaling and zipf's law: a tale of board games and power laws. *arXiv [Preprint]*. arXiv:2412.11979. doi: 10.48550/arXiv.2412.11979
- Ororbias, A. G., and Kelly, M. A. (2023). A neuro-mimetic realization of the common model of cognition via hebbian learning and free energy minimization. *Proc. AAAI Symposium Series* 2, 369–378. doi: 10.31219/osf.io/z7c98
- Parikh, V., and Bangasser, D. A. (2020). Cholinergic signaling dynamics and cognitive control of attention. *Behav. Pharmacol. Cholinerg. Syst.* 45, 71–87. doi: 10.1007/7854\_2020\_133
- Petersen, C. C., Malenka, R. C., Nicoll, R. A., and Hopfield, J. J. (1998). All-or-none potentiation at CA3-CA1 synapses. *Proc. Nat. Acad. Sci.* 95, 4732–4737. doi: 10.1073/pnas.95.8.4732
- Ramsauer, H., Schäfl, B., Lehner, J., Seidl, P., Widrich, M., Adler, T., et al. (2020). Hopfield networks is all you need. *arXiv [Preprint]*. arXiv:2008.02217. doi: 10.48550/arXiv.2008.02217
- Tishby, N., and Zaslavsky, N. (2015). “Deep learning and the information bottleneck principle,” in *2015 IEEE Information Theory Workshop (ITW)* (Jerusalem: IEEE), 1–5.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., et al. (2017). Attention is all you need. *arXiv [Preprint]*. arXiv:1706.03762. doi: 10.48550/arXiv.1706.03762
- Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., et al. (2022). Chain-of-thought prompting elicits reasoning in large language models. *Adv. Neural Inf. Process. Syst.* 35, 24824–24837. doi: 10.5555/3600270.3602070
- Wei, L., Ma, Z., Yang, C., and Yao, Q. (2024). Advances in the neural network quantization: a comprehensive review. *Appl. Sci.* 14:7445. doi: 10.3390/app14177445
- Wu, W. (2024). We know what attention is! *Trends Cogn. Sci.* 28, 304–318. doi: 10.1016/j.tics.2023.11.007
- Xia, Y., Wang, R., Liu, X., Li, M., Yu, T., Chen, X., et al. (2024). Beyond chain-of-thought: a survey of chain-of-X paradigms for LLMs. *arXiv [Preprint]*. arXiv:2404.15676. doi: 10.48550/arXiv.2404.15676
- Zhang, Z., Yao, Y., Zhang, A., Tang, X., Ma, X., He, Z., et al. (2025). Igniting language intelligence: the hitchhiker's guide from chain-of-thought reasoning to language agents. *ACM Comp. Surv.* 57, 1–39. doi: 10.1145/3719341
- Zhu, X., Cheng, D., Zhang, D., Li, H., Zhang, K., Jiang, C., et al. (2025). Flowrl: Matching reward distributions for llm reasoning. *arXiv [Preprint]*. arXiv:2509.15207. doi: 10.48550/arXiv.2509.15207

## Generative AI statement

The author(s) declare that no Gen AI was used in the creation of this manuscript.

Any alternative text (alt text) provided alongside figures in this article has been generated by Frontiers with the support of artificial intelligence and reasonable efforts have been made to ensure accuracy, including review by the authors wherever possible. If you identify any issues, please contact us.

## Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.