CIFMA-2024

6th International Workshop on

Cognition: Interdisciplinary Foundations, Models and Applications

5 November 2024

Aveiro, Portugal

Co-located with SEFM 2024

6° CIFMA 2024 - SEFM 2024

University of Aveiro
4/11/2024

On the Morphic Problem of Artificial Neural Networks

Giovanni Galli

University of Teramo

ggalli@unite.it

UNIVERSITÁ DEGLI STUDI DI TERAMO

Agenda

Neuromorphic Engineering

Vector Grounding Problem

Morphic Problem

Neuromorphic Engineering (NE)

- Definition: An interdisciplinary field bridging neuroscience, computer science, and engineering.
- Objective: Design ANNs that mimic biological systems in adaptability, function, and efficiency.

Deep Learning Models and Cognition

- Role of Deep Learning in Cognitive Science
  - Potential and limitations in modeling cognition

- Strengths: Handles vast data, pattern recognition, generalization.
- Limitations: Interpretation, data dependence, and higher-order cognitive tasks.

The Neuromorphic Engineering Project

- Comparison between Biological Neural Networks and ANNs:
  - Adaptation and efficiency vs. artificial imitation.

- Question: Are artificial and biological systems converging?

The Symbol Grounding Problem

- Symbol Grounding Problem: How do symbols get their meanings?

"The symbol grounding problem is the problem of how to make the semantic interpretation of a formal symbol system intrinsic to the system, rather than just parasitic on the meanings in our heads in anything but other meaningless symbols"

(Harnad, 1990)

Solutions: symbolic representations must be grounded bottom-up in nonsymbolic representations of two kinds: (1) iconic representations, which are analogues of the proximal sensory projections of distal objects and events, and (2) categorical representations, which are learned and innate feature detectors that pick out the invariant features of object and event categories from their sensory projections. (Harnad, 1990)

The Vector Grounding Problem

- Vector Grounding Problem: LLMs use vectors as numerical representations of text tokens based on statistical relationships. LLMs struggle to connect their internal representations (vectors) to the real-world. (Coehlo Mollo and Millière, 2023)
- Fine-tuned with Reinforcement Learning from Human Feedback (RLHF), possess the necessary features to overcome the Vector Grounding Problem, as they stand in the requisite causal-historical relations to the world that underpin intrinsic meaning.
- “We also argue that, perhaps unexpectedly, multimodality and embodiment are neither necessary nor sufficient conditions for referential grounding in artificial systems” (Coehlo Mollo and Millière, 2023)

- Current AI Challenges: Lack of intrinsic connection to real-world meanings despite human-like outputs.

Introduction to the Morphic Problem

• Challenge of replicating biological neural systems' adaptability in artificial neural networks (ANNs).

• Importance of Morphic Relation: Aims to capture structure, function, and dynamic correspondences between biological and artificial neural systems.

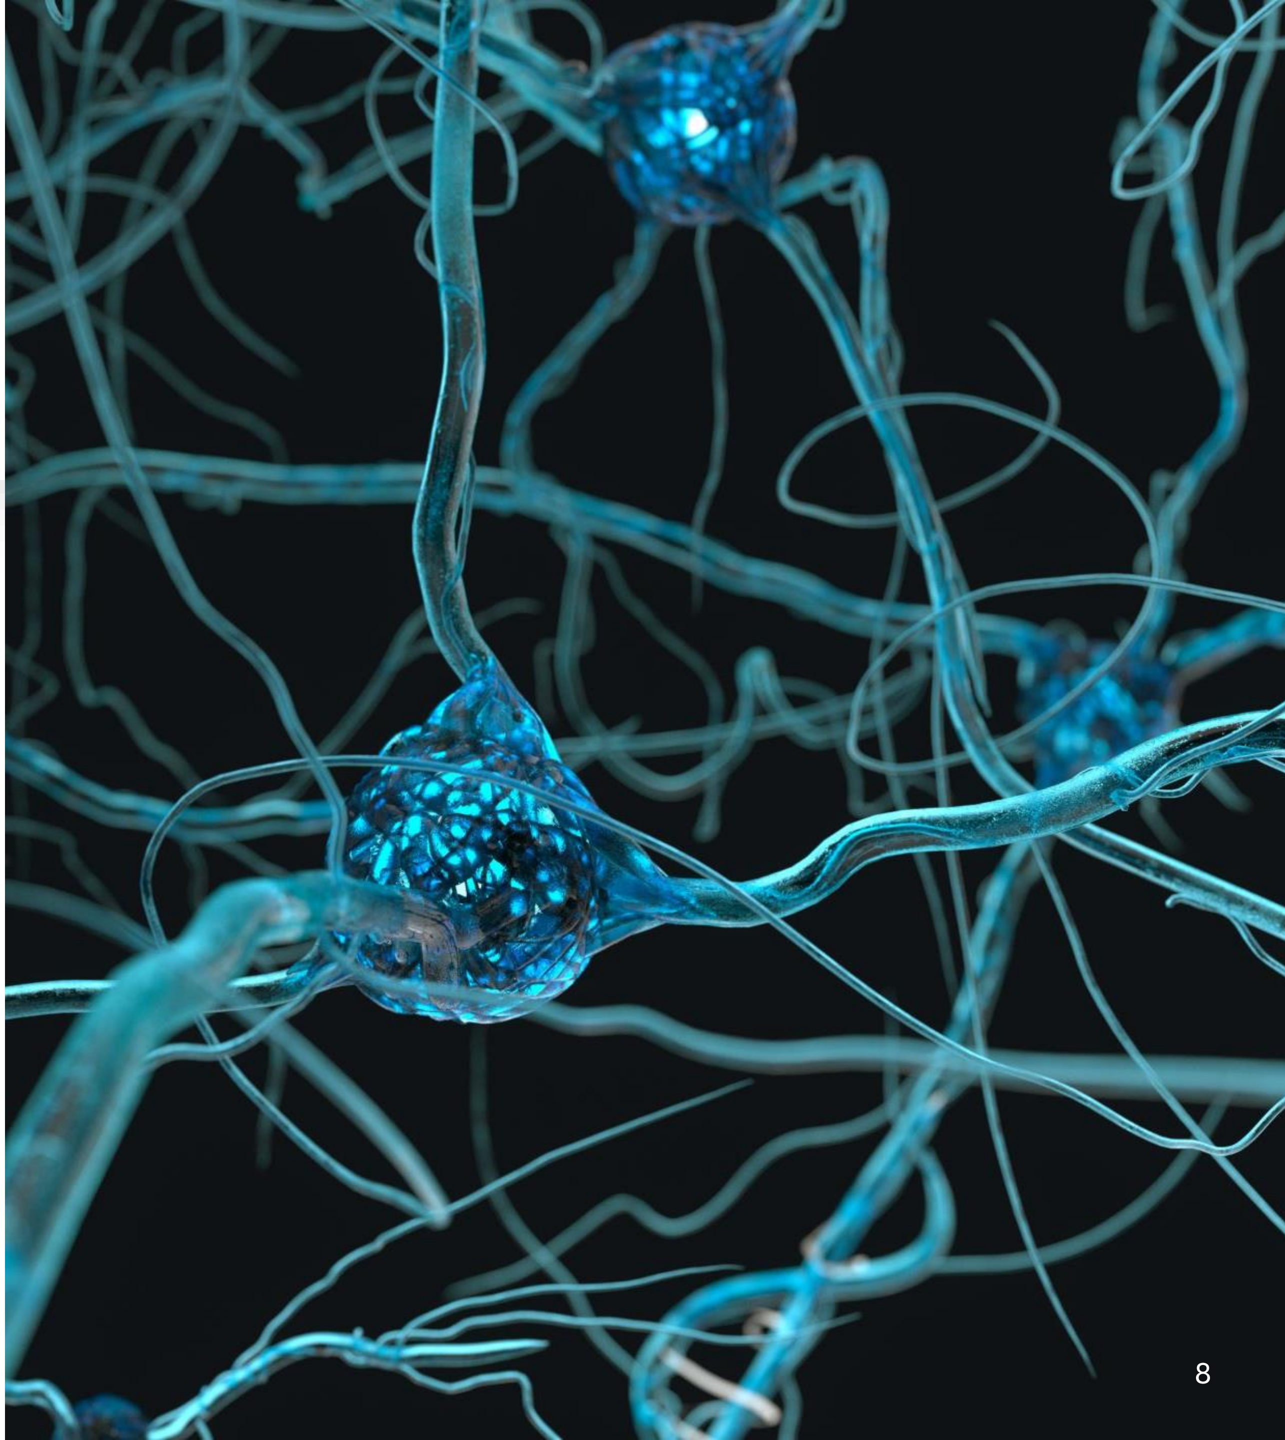

Addressing the Morphic Problem

- Biological Neural Network Attributes: Interaction with the real world, sensory experiences, evolutionary refinement.
- Contrast with ANNs: Training on isolated datasets lacking sensory richness.
- NE Goal: Create embodied cognition and multimodal learning in ANNs.

The Morphic Problem in ANNs (MPANNs)

- It is a generalisation issue of VGP: we want optimised ANNs able to mimic human brain biology, but they are structurally disconnected from the real world.

- Can the morphic program succeed in optimising ANNs without somehow embedding them in a world that allows for meaningful interactions and grounding of representations?

Criteria for Morphic Relations in ANNs

1. Structural Correspondence
2. Functional Equivalence
3. Dynamic Coherence
4. Contextual Integration
5. Robustness and Resilience

Five Key Criteria for MPANNs

- (SC) ANNs must have a structural framework that reflects biological neural systems, including neurons and synapses' basic components and interconnections.
- (FE) ANNs must replicate essential functional operations, like learning through synaptic plasticity and information processing via spike-timing dynamics.
- (DC) Temporal dynamics in ANNs should mirror those in biological systems, particularly in adaptive responses over time to stimuli.
- (CI) Both systems should be integrated into their environments similarly, with neural networks interacting in an ecosystem or framework similar to biological networks.
- (RR) The morphic relation requires robustness and resilience in similar ways, so ANNs should withstand noise, damage, or variability like biological networks.

Methodological Challenges

- Challenges in Comparing DNNs and Human Cognition:
  - Differences in training and experiences.
  - Difficulty in capturing cognitive complexity and environmental interactions.

- Approach: Design more holistic and interactive training environments.

How to deal with MPANNS?

- Foster the morphism

  or

- Cut off the grounding relation

  (Ocham's AI Razor?)

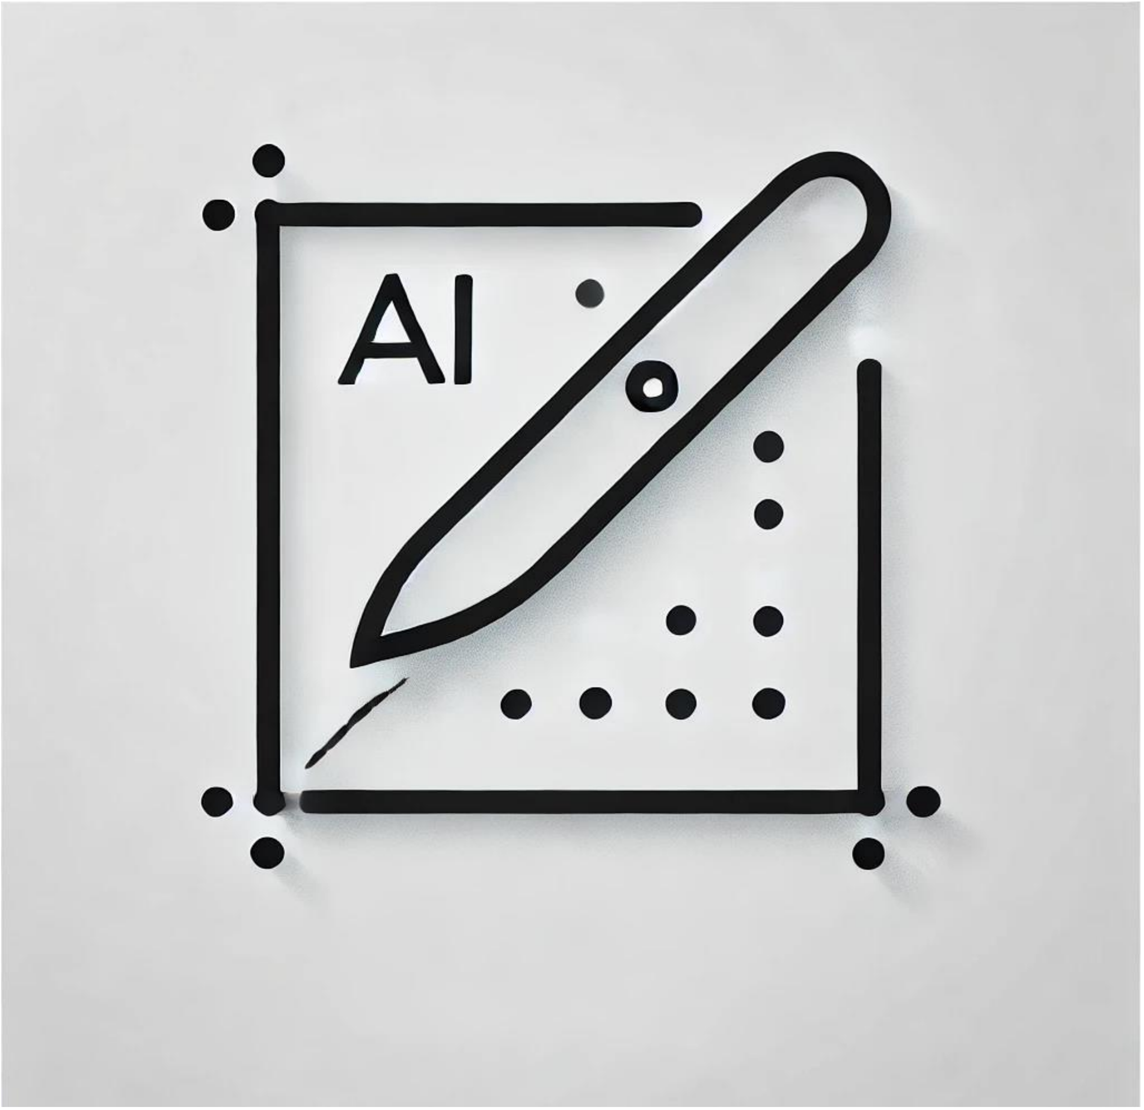

Conclusion

- Morphic Problem: Technical and conceptual challenges.
- Future Outlook for Neuromorphic Engineering: Moving towards more human-like AI

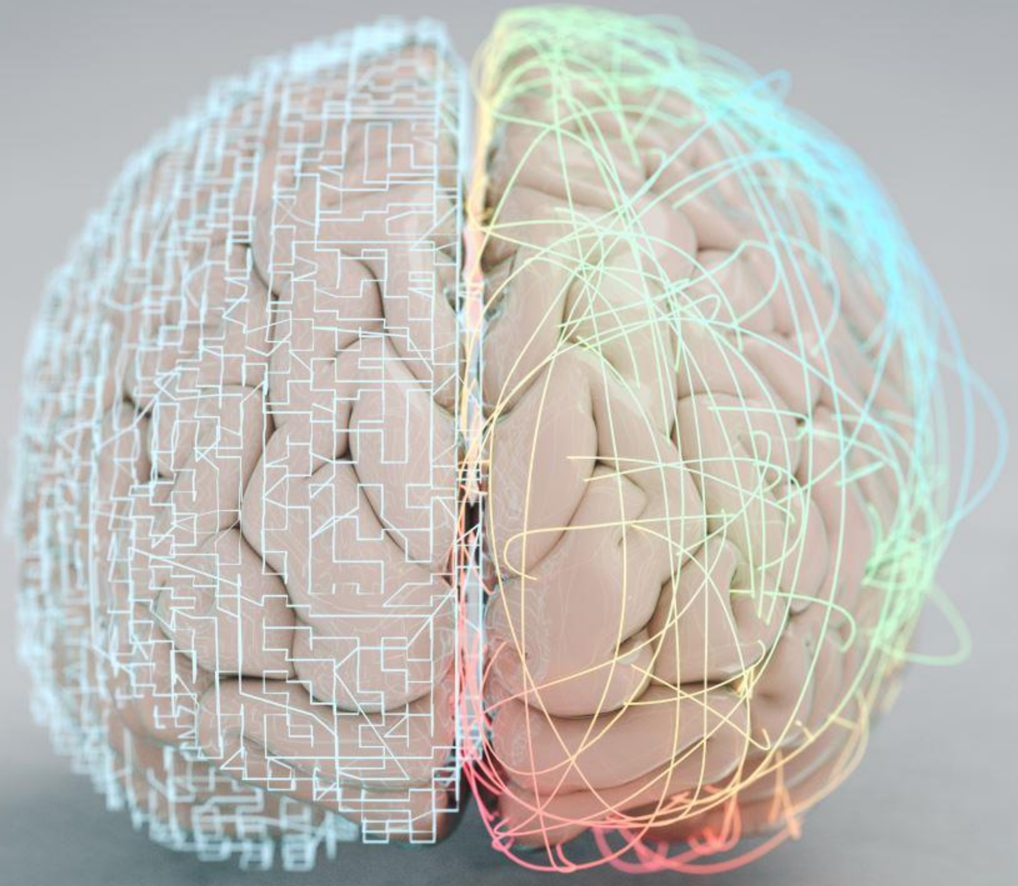

Thank you for listening!

References

Fodor, J. A. and Pylyshyn, Z. W., 1988. Connectionism and cognitive architecture: A critical analysis. Cognition, 28(1-2), pp. 3-71.

Harnad, S., 1990, The Symbol Grounding Problem. Physica D. 42 (1–3): 335–346.

Lake, B. M. and Baroni, M., 2023. Human-like systematic generalization through a meta-learning neural network. Nature, 623, pp. 115–121.

Mandelbaum, E., et al., 2022. Problems and Mysteries of the Many Languages of Thought. Cognitive Science, 46, pp. 1-6.

Marcus, G., 2001. The Algebraic Mind. Cambridge: MIT Press.Marcus, G., 2018. Deep learning: A Critical Appraisal. arXiv:1801.00631v1.

Murty, S., Sharma, P., Andreas, J., and Manning, C., 2023. Grokking of Hierarchical Structure in Vanilla Transformers. In Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), pp. 439–448, Toronto, Canada. Association for Computational Linguistics.

Pinker, S. and Prince, A., 1988. On language and connectionism: Analysis of a parallel distributed processing model of language acquisition. Cognition, 28(1-2), pp. 73-193.

Quilty-Dunn, J., Porot, N., and Mandelbaum, E., 2023. The best game in town: The reemergence of the language-of-thought hypothesis across the cognitive sciences. Behavioral and Brain Sciences, 46, e261.

Tsur, E. E., 2022. Neuromorphic Engineering: The Scientist’s, Algorithm Designer’s, and Computer Architect’s Perspectives on Brain-Inspired Computing. Abingdon: CRC Press.