

Prediction: An Algorithmic Principle Meeting Neuroscience and Machine Learning Halfway

Younes Bouhadjar^{1,†}, Caterina Moruzzi^{2,*,†} and Melika Payvand^{3,†}

¹ Institute of Neuroscience and Medicine (INM-6), & Institute for Advanced Simulation (IAS-6), & JARA BRAIN Institute Structure-Function Relationships (INM-10), & Peter Grünberg Institute (PGI-7,10), Jülich Research Centre and JARA, Jülich, Germany & RWTH Aachen University, Aachen, Germany

² Department of Philosophy, University of Konstanz, Universitätsstraße 10, 78464 Konstanz, Germany

³ Institute of Neuroinformatics, University of Zurich and ETH Zurich, Winterthurerstrasse 190, 8057 Zurich, Switzerland

Abstract

In this paper, we support the relevance of the collaboration and mutual inspiration between research in Artificial Intelligence and neuroscience to create truly intelligent and efficient systems. In contrast to the traditional top-down and bottom-up strategies designed to study and emulate the brain, we propose an alternative approach where these two strategies are met halfway, defining a set of algorithmic principles. We present *prediction* as a core algorithmic principle and advocate for applying the same approach to identify other neural principles which can constitute core mechanisms of new Machine Learning frameworks.

Keywords

Prediction, Neuroscience, Reasoning, Algorithmic Principles, Computation, Bottom-up, Top-down

1. Introduction

The impressive accomplishments by Machine Learning (ML) research in the last decade substantiate the argument that, given enough computing power and scale, almost any kind of task can be successfully achieved with just statistical correlation of data [1, 2, 3, 4]. Yet, leveraging of this type of computation does not seem to be enough in order to solve two hurdles that research in Artificial Intelligence (AI) still has to confront: (i) *Functional*: the current ML models lack the ability to achieve abstraction and generalization, e.g., classification of images in unfamiliar environments, or making predictions on out-of-distribution data [5, 6, 7]; (ii) *Technological*: ML models typically require a lot of data to train and consume a substantial amount of energy [8].

In this paper, we support the claim that by exploiting the knowledge we have about the brain, as a proof of the existence of an efficient intelligent machine, we can provide insight to ML

International Workshop on Human-Like Computing - International Joint Conference on Learning & Reasoning 2022

*Corresponding author.

† These authors contributed equally.

✉ y.bouhadjar@fz-juelich.de (Y. Bouhadjar); caterina.moruzzi@uni-konstanz.de (C. Moruzzi); melika@ini.uzh.ch (M. Payvand)

🌐 <https://younesbouhadjar.github.io> (Y. Bouhadjar); <https://caterinamoruzzi.weebly.com> (C. Moruzzi);

<https://services.ini.uzh.ch/people/melika> (M. Payvand)

🆔 0000-0003-4367-8236 (Y. Bouhadjar); 0000-0002-9728-3873 (C. Moruzzi); 0000-0001-5400-067X (M. Payvand)



© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

computational models for solving these problems.

Functionally, the shift from behaviorism to cognitivism and the study of mental processes was aided and inspired by the emerging field of computer science. Many researchers shared the vision to undertake the project of building *Thinking Machines*. Conversely, the inspiration for many discoveries in the AI research field comes from the observation of underlying computational mechanisms in the animal and human brain [9, 10, 11, 12]. The first computational model of an artificial neuron was developed in 1943 by McCulloch and Pitts, and it was based on their understanding of the structure and functionality of a biological neuron [13, 14]. More than fifty years later, the creation of deep neural networks was inspired by the hierarchical structure of the brain [15]. The observation of reward prediction brain processes was the basis for the development of reinforcement learning techniques [16]. **Technologically**, many of the solutions used by the current AI hardware for more efficient implementations are inspired by and/or compatible with the solutions found by the biological counterparts. For example, pruning the neural network parameters [17], reducing parameter precision [18], reducing the precision of the activation functions [19], and exploiting the sparsity of the activations [20] and network parameters [21].

The collaboration and mutual inspiration between AI and brain studies can bear fruit yet another time, by helping the AI research to tackle the functional and technological problems it is facing. To achieve this aim, in section 2 we propose an alternative strategy to more traditional top-down and bottom-up approaches to the emulation of the human brain. We propose to meet these approaches halfway and start from an *algorithmic principles* level connecting the two. In section 3, we identify **prediction** - an important mechanism which had already received attention in the AI and cognitive field - as one of these principles, crucial in order to endow machines with more human-like capacities.

2. Algorithmic Principles: A Middle Layer

Figure 1 illustrates the articulation of mental processes and functions which stem from a *mind*, i.e., a functional entity that supports intelligent behavior [22]. This functional entity needs to have a substratum, which can either be biological or artificial. In the quest for building human-level intelligence, the most straightforward strategy would seem that of emulating the functioning of the organic brain. In performing this attempt, two opposite approaches are traditionally followed: **top-down** or **bottom-up** [23].

Top-down strategies start from the analysis of the brain's behavior and its higher cognitive functionalities, e.g. *learning*, *reasoning*, and *memory*. Bottom-up approaches attempt instead at understanding and replicating the organizing principles of the brain, such as *modularity*, *spatial awareness*, *self-organization*, and *temporal coding*, or, at an even lower abstraction level, the biological structure and physiological features of the brain. Both approaches may face criticism, though. Top-down approaches are subject to interpretation biases, as framing the behavior of the system-under-analysis consequently influences the choice of strategy followed for reproducing it [24]. On the other hand, bottom-up approaches are comparatively more resource-expensive, and lack high-level functional guidance.

As an alternative strategy to top-down and bottom-up approaches towards the emulation of

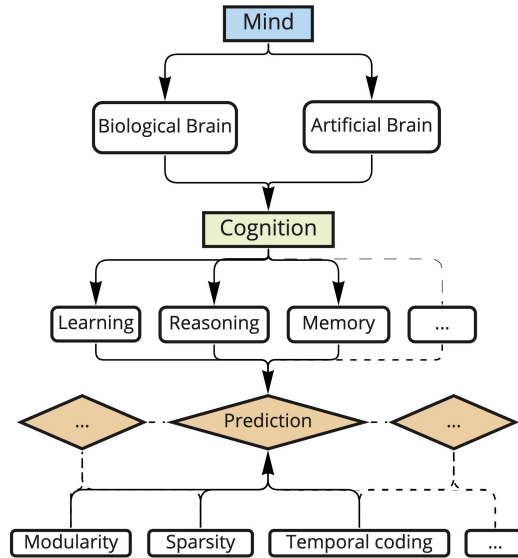


Figure 1: *Articulation of Mental Processes and Functions.* In the diagram, prediction is presented as one of the algorithmic principles that bridge cognitive faculties and organizing principles.

the human brain, we propose to meet these approaches halfway, and start from an *algorithmic principles level* connecting the two. The intermediate layer is defined by using insights from both top-down and bottom-up studies, and it expresses the higher-order cognitive functions of the brain, while simultaneously explaining the biological functions and neural system’s organizing principles. Our approach is reminiscent of Marr’s three levels [23], but advocates for a new interpretation that can be accessible to practitioners in different fields.

At this stage, our proposal is speculative in nature: it is intended as an invitation to engage in a discussion about how an alternative third way, between top-down and bottom-up perspectives on learning and cognition, could facilitate collaborations between different fields of research aimed at endowing machines with human reasoning skills. We leave to further research a detailed consideration of how this approach would be implemented in practice.

Focusing on the algorithmic principles layer, as suggested by this paper, will help to identify how the structure of the brain might give rise to its cognitive functionality. It helps neuroscientists to abstract away from the many details of biology to find the organizing principles of intelligence, and provides insights to ML researchers to move towards more human-like intelligence. In the remaining part of the paper, we focus on **prediction**, as an exemplary mechanism that governs the interaction between top-down expectations and bottom-up sensory input and which can help neuroscience and ML research meet halfway.

3. Prediction as a Primary Computation Performed by the Brain

We define prediction as an inference of future states learned by means of prior experiences, which can result in the fast detection and adaptation to changes in the environment. Prediction has been identified by many as a requirement for planning, decision-making, motor and cognitive

control, counterfactual reasoning, and the improvement of behavior on the basis of experiences [25, 26, 27, 28, 29, 30, 24]. To discuss this principle with a simple example, let us say you are thirsty. You get up from the couch, enter the kitchen and open the fridge, because you *predict* we will find a fresh bottle of water there. The brain predicts and prepares actions on the basis of prior experiences: your action of opening the fridge, and the cognitive and motor control aspects responsible for this action, were primed by the expectation to find the water bottle in there.

The neuronal structure responsible for this functionality in the animal cortex has been extensively studied. The cortex consists of a repeating modular structure with hierarchical connectivity of a population of neurons. The majority of these neurons are pyramidal neurons characterized by a distinct multi-compartment structure receiving inputs from different directions. The bottom-up sensory information is received in compartments closer to the neuron body (known as the basal dendrites), and the top-down feedback information is received in compartments further away from the body (known as the apical dendrites) [31]. The spatio-temporal features of the sensory input is detected by the *coincident* activation of the neighboring synapses in the dendrites. In case of sufficient activation, a dendritic action potential (dAP) is generated [32, 33, 34]. The activation of a dAP results in a long-lasting depolarization of the neuron body, or soma. The depolarized state of these neurons make them more likely to fire as a response to an input.

Recent studies show that the dAPs can encode contextual information representing prior experiences or top-down expectations [35, 36]. The contextual information allows pyramidal neurons to predict the most likely next outcome [37, 38] or compare top-down expectations from higher-level areas with bottom-up sensory signals [39]. In case of a prediction error, the cells fire in a non-sparse manner, signaling a “mismatch”. This activity derives the learning and in turn reduces the error [40].

Pyramidal neurons are thus powerful computational units that not only passively sum up incoming input, but can also signal predictions as a result of the dendritic activity. As these neurons are grouped within different sensory and hierarchical areas, they allow for simultaneous predictions of different features, which can be of low-level details, such as predicting sensory information, or of high-level details, such as predicting concepts or thought processes. As you open the fridge to grab the bottle of water, the visual areas of the cortex predict the shape and color of the bottle, the somatosensory areas predict tactile sensations, and other areas predict more complex features such as the weight of the bottle. These predictions are unified by higher-level areas into an abstract concept.

In the light of top-down studies, prediction emerged as a building block of cognition. Using this knowledge has helped to guide bottom-up studies toward an interpretation of neuronal mechanisms that were otherwise obscure. This, in turn, solidifies the concept of prediction and presents it as an algorithmic principle implemented by the brain.

We advocate for applying the same approach to identify other neural principles. These principles can constitute core mechanisms of new machine learning frameworks, bridging the gap in the performance between machines and brains. In the quest toward the goal of emulating human-level intelligence in artificial systems, what is needed is a shared platform where methodologies of different disciplines can meet and cross-pollinate each other. Algorithmic principles, we argue, may constitute such a platform.

Acknowledgments

The idea for this paper was developed during the 2022 edition of “The CapoCaccia Workshops toward Neuromorphic Intelligence”. We would like to thank the organizers of the workshop, as well as the Office of Naval Research for supporting the participation to it through the ONR Global grant (grant number: N62909-22-1-2029).

References

- [1] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al., Language models are few-shot learners, *Advances in neural information processing systems* 33 (2020) 1877–1901.
- [2] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, M. Chen, Hierarchical text-conditional image generation with clip latents, *arXiv preprint arXiv:2204.06125* (2022).
- [3] S. Reed, K. Zolna, E. Parisotto, S. G. Colmenarejo, A. Novikov, G. Barth-Maron, M. Gimenez, Y. Sulsky, J. Kay, J. T. Springenberg, et al., A generalist agent, *arXiv preprint arXiv:2205.06175* (2022).
- [4] D. Silver, T. Hubert, J. Schrittwieser, I. Antonoglou, M. Lai, A. Guez, M. Lanctot, L. Sifre, D. Kumaran, T. Graepel, et al., Mastering chess and shogi by self-play with a general reinforcement learning algorithm, *arXiv preprint arXiv:1712.01815* (2017).
- [5] K. Kawaguchi, L. P. Kaelbling, Y. Bengio, Generalization in deep learning, *arXiv preprint arXiv:1710.05468* (2017).
- [6] J. Miüller, R. Weibel, J.-P. Lagrange, F. Salgé, Generalization: state of the art and issues, *GIS and Generalization* (2020) 3–17.
- [7] C. Moruzzi, Artificial creativity and general intelligence, *Journal of Science and Technology of the Arts* 12 (2020) 84–99.
- [8] A. Mehonic, A. J. Kenyon, Brain-inspired computing needs a master plan, *Nature* 604 (2022) 255–260.
- [9] J. Fan, L. Fang, J. Wu, Y. Guo, Q. Dai, From brain science to artificial intelligence, *Engineering* 6 (2020) 248–252.
- [10] D. Hassabis, D. Kumaran, C. Summerfield, M. Botvinick, Neuroscience-inspired artificial intelligence, *Neuron* 95 (2017) 245–258.
- [11] N. Savage, Marriage of mind and machine, *Nature* 571 (2019) S15–S17.
- [12] S. Ullman, Using neuroscience to develop artificial intelligence, *Science* 363 (2019) 692–693.
- [13] S. Chakraverty, D. M. Sahoo, N. R. Mahato, Mcculloch–pitts neural network model, in: *Concepts of Soft Computing*, Springer, 2019, pp. 167–173.
- [14] F. Rosenblatt, The perceptron: a probabilistic model for information storage and organization in the brain., *Psychological review* 65 (1958) 386.
- [15] Y. LeCun, Y. Bengio, G. Hinton, Deep learning, *nature* 521 (2015) 436–444.
- [16] W. Dabney, Z. Kurth-Nelson, N. Uchida, C. K. Starkweather, D. Hassabis, R. Munos, M. Botvinick, A distributional code for value in dopamine-based reinforcement learning, *Nature* 577 (2020) 671–675.

- [17] S. Han, H. Mao, W. J. Dally, Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding, arXiv preprint arXiv:1510.00149 (2015).
- [18] Edge TPU (2019). URL: <https://cloud.google.com/edge-tpu/>.
- [19] S. Gupta, A. Agrawal, K. Gopalakrishnan, P. Narayanan, Deep learning with limited numerical precision, in: International conference on machine learning, PMLR, 2015, pp. 1737–1746.
- [20] C. Gao, A. Rios-Navarro, X. Chen, S.-C. Liu, T. Delbruck, EdgeDRNN: Recurrent neural network accelerator for edge inference, IEEE Journal on Emerging and Selected Topics in Circuits and Systems 10 (2020) 419–432.
- [21] C. Gao, T. Delbruck, S.-C. Liu, Spartus: A 9.4 TOp/s FPGA-based LSTM accelerator exploiting spatio-temporal sparsity, arXiv preprint arXiv:2108.02297 (2021).
- [22] J. E. Laird, C. Lebiere, P. S. Rosenbloom, A standard model of the mind: Toward a common computational framework across artificial intelligence, cognitive science, neuroscience, and robotics, Ai Magazine 38 (2017) 13–26.
- [23] D. Marr, Vision: A computational investigation into the human representation and processing of visual information, MIT press, 2010.
- [24] M. György Buzsáki, The brain from inside out, Oxford University Press, 2019.
- [25] A. Bubic, D. Y. Von Cramon, R. I. Schubotz, Prediction, cognition and the brain, Frontiers in human neuroscience 4 (2010) 25.
- [26] A. Clark, Whatever next? predictive brains, situated agents, and the future of cognitive science, Behavioral and brain sciences 36 (2013) 181–204.
- [27] K. Friston, The free-energy principle: a unified brain theory?, Nature reviews neuroscience 11 (2010) 127–138.
- [28] J. Hawkins, S. Blakeslee, On intelligence: How a new understanding of the brain will lead to the creation of truly intelligent machines, Macmillan, 2007.
- [29] C. Moruzzi, Climbing the ladder: How agents reach counterfactual thinking, in: ICAART 2022: 14th International Conference on Agents and Artificial Intelligence, 2022, pp. 555–560.
- [30] R. S. Sutton, Learning to predict by the methods of temporal differences, Machine learning 3 (1988) 9–44.
- [31] N. Spruston, Pyramidal neurons: dendritic structure and synaptic integration, Nature Review Neuroscience 9 (2008) 206–221.
- [32] G. J. Stuart, N. Spruston, Dendritic integration: 60 years of progress, Nature Neuroscience 18 (2015) 1713–1721.
- [33] G. Major, M. Larkum, J. Schiller, Active properties of neocortical pyramidal neuron dendrites, Annual Review Neuroscience 36 (2013) 1–24.
- [34] M. London, M. Häusser, Dendritic computation, Annual Review Neuroscience 28 (2005) 503–532.
- [35] N. Takahashi, C. Ebner, J. Sigl-Glöckner, S. Moberg, S. Nierwetberg, M. E. Larkum, Active dendritic currents gate descending cortical outputs in perception, Nature Neuroscience 23 (2020) 1277–1285. URL: <https://doi.org/10.1038/s41593-020-0677-8>. doi:10.1038/s41593-020-0677-8.
- [36] C. J. Gillon, J. E. Pina, J. A. Lecoq, R. Ahmed, Y. N. Billeh, S. Caldejon, P. Groblewski, T. M. Henley, I. Kato, E. Lee, J. Luviano, K. Mace, C. Nayan, T. V. Nguyen, K. North,

- J. Perkins, S. Seid, M. T. Valley, A. Williford, Y. Bengio, T. P. Lillicrap, B. A. Richards, J. Zylberberg, Learning from unexpected events in the neocortical microcircuit (2021). URL: <https://doi.org/10.1101/2021.01.15.426915>. doi:10.1101/2021.01.15.426915.
- [37] J. Hawkins, S. Ahmad, Why neurons have thousands of synapses, a theory of sequence memory in neocortex, *Frontiers in Neural Circuits* 10 (2016) 23. URL: <https://doi.org/10.3389/fncir.2016.00023>. doi:10.3389/fncir.2016.00023.
- [38] Y. Bouhadjar, D. J. Wouters, M. Diesmann, T. Tetzlaff, Sequence learning, prediction, and replay in networks of spiking neurons, *arXiv preprint arXiv:2111.03456* (2021).
- [39] M. Larkum, A cellular mechanism for cortical associations: an organizing principle for the cerebral cortex, *Trends in Neurosciences* 36 (2013) 141–151. URL: <https://doi.org/10.1016/j.tins.2012.11.006>. doi:10.1016/j.tins.2012.11.006.
- [40] A. Payeur, J. Guerguiev, F. Zenke, B. A. Richards, R. Naud, Burst-dependent synaptic plasticity can coordinate learning in hierarchical circuits, *Nature Neuroscience* 24 (2021) 1010–1019. URL: <https://doi.org/10.1038/s41593-021-00857-x>. doi:10.1038/s41593-021-00857-x.