

AIMH Lab for a Sustainable Bio-Inspired AI

Gabriele Lagani^{1,*}, Fabrizio Falchi¹, Claudio Gennaro¹ and Giuseppe Amato¹

¹ISTI-CNR, via G. Moruzzi, 1, Pisa, 56100, Italy

Abstract

In this short paper, we report the activities of the Artificial Intelligence for Media and Humanities (AIMH) laboratory of the ISTI-CNR related to Sustainable AI. In particular, we discuss the problem of the environmental impact of AI research, and we discuss a research direction aimed at creating effective intelligent systems with a reduced ecological footprint. The proposal is based on bio-inspired learning, which takes inspiration from the biological processes underlying human intelligence in order to produce more energy-efficient AI systems. In fact, biological brains are able to perform complex computations, with a power consumption which is orders of magnitude smaller than that of traditional AI. The ability to control and replicate these biological processes reveals promising results towards the realization of sustainable AI.

Keywords

Bio-Inspired, Neural Networks, Machine Learning, Sustainable AI

1. Introduction

Energy efficiency is crucial to protect our planet from pollution and global warming disasters.

As society becomes more and more sensible to the climate change related issues, the Computer Science academic world is also observing a change of perspective, towards achieving more environmentally-friendly computing solutions. In this scenario, The Artificial Intelligence (AI) community is particularly affected.

In fact, recent research [1] shows that AI software has an extremely relevant impact in terms of carbon footprint, which is not surprising, given the heavy energy demand required to train Deep Neural Network (DNN) models. As a result, the scientific community is striving to find new directions which yield promises in terms of reducing the environmental impact of AI systems, while preserving the astounding results achieved in the previous years, such as in the fields of vision [2], language [3], and game playing/reinforcement learning [4].

We focus on a research directions in particular, which is promising in regard to the goal of Sustainable AI: biologically inspired approaches to machine learning.

This research direction is motivated by the observation that biological brains are capable of solving incredibly complex problems, while still maintaining an energy budget of just 20W [5]. In comparison, traditional DNNs run on GPUs that consume one order of magnitude more

power, while still lagging behind human-level, both in terms of capabilities, and size of the simulated network models.

Brain connections are thought to obey a constrained optimization, such as maximization of information processing capacity (efficiency) while minimizing the energy expenditure [6].

The key of such efficiency in biological neural systems is rooted in two aspects: first, the communication paradigm adopted by biological neurons is based on discrete pulses, called *spikes*, and second, effective learning is achieved by local learning rules based on the *Hebbian* principle “neurons that fire together wire together”, and Spike Time Dependent Plasticity (STDP) [7]. The research community is currently trying to leverage these principles in order to achieve more energy-efficient computational models for AI.

For example, Spiking Neural Network (SNN) models are being used to implement low-power silicon devices that emulate the behavior of biological neurons: these are called *neuromorphic* devices [8]. Alternatively, instead of emulating real neurons on silicon, in our Lab we are contributing to the design of biological hardware (or *bioware*) for AI [9]. By collaborating with neuroscience colleagues, the idea is to use cultures of biological neurons directly to solve the AI task. Interfacing with neuronal cultures occurs via Multi-Electrode Array (MEA) devices [10], although *optogenetic* (i.e. light-based) stimulation is also under exploration [11].

We are also exploring more biologically realistic training algorithms based on the Hebbian principle, applied to DNNs with the goal of improving current learning paradigms in unsupervised or semi-supervised learning scenarios, also in hybrid combinations with traditional approaches based on the biologically implausible *backpropagation* principle [12], [13, 14, 15, 16, 17, 18].

Ital-IA 2023: 3rd National Conference on Artificial Intelligence, organized by CINI, May 29–31, 2023, Pisa, Italy

*Corresponding author.

✉ gabriele.lagani@isti.cnr.it (G. Lagani); fabrizio.falchi@isti.cnr.it (F. Falchi); claudio.gennaro@isti.cnr.it (C. Gennaro); giuseppe.amato@isti.cnr.it (G. Amato)

📞 0000-0003-4739-5778 (G. Lagani); 0000-0001-6258-5313

(F. Falchi); 0000-0002-3715-149X (C. Gennaro); 0000-0003-0171-4315 (G. Amato)

© 2023 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)

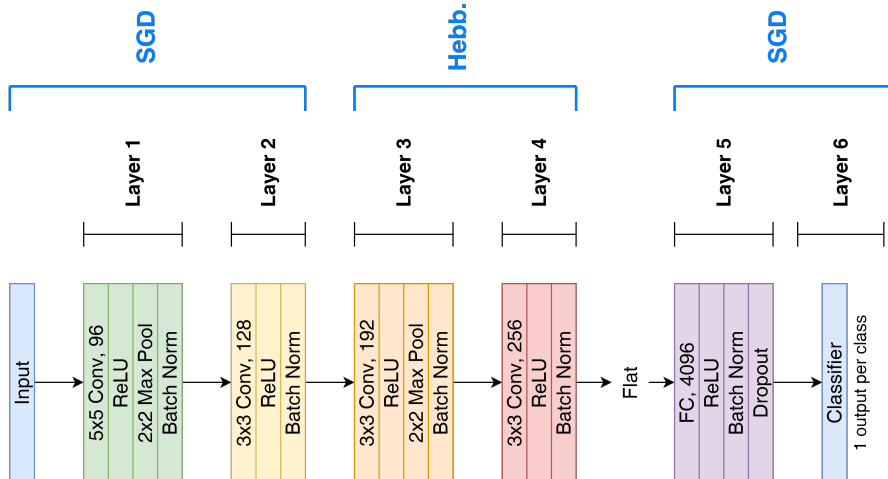


Figure 1: An example hybrid neural network model, where internal layers are trained using Hebbian learning, while the other layers are trained by SGD.

2. Research Themes

Neural networks are said to be biologically inspired since they mimic the behavior of real neurons. However, several processes in state-of-the-art neural networks, including Convolutional Neural Networks (CNNs), are far from the ones found in biological brains [19]. One relevant difference is the training process. In state-of-the-art artificial neural networks, the training process is based on backpropagation and Stochastic Gradient Descent (SGD) optimization. However, studies in neuroscience strongly suggest that this kind of processes does not occur in the biological brain [19]. Rather, learning methods based on STDP or the Hebbian learning rule seem to be more plausible, according to neuroscientists [7]. In [12], we investigated the use of the Hebbian learning rule when training neural networks for image classification by proposing a novel weight update rule for shared kernels in CNNs. We performed experiments using the CIFAR-10 dataset in which we employed Hebbian learning, along with SGD, to train parts of the model or whole networks for the task of image classification, and we discussed their performance thoroughly considering both effectiveness and efficiency aspects.

In [16], we deepened our investigations on Hebbian learning strategies applied to DNN training. We considered two unsupervised learning approaches, Hebbian Winner-Takes-All (HWTA) and Hebbian Principal Component Analysis (HPCA). The Hebbian learning rules were used to train the layers of a CNN in order to extract features that were then used for classification, without requiring backpropagation (backprop). We performed experimental comparisons with state-of-the-art unsu-

pervised (but backprop-based) Variational Auto-Encoder (VAE) training [20]. For completeness, we considered two supervised Hebbian learning variants (Supervised Hebbian Classifiers – SHC, and Contrastive Hebbian Learning – CHL), for training the final classification layer, which were compared to Stochastic Gradient Descent (SGD) training. We also investigated hybrid learning methodologies, where some network layers were trained following the Hebbian approach, and others were trained by backprop. Fig. 1 shows an example of such hybrid neural network models. We tested our approaches on MNIST [21], CIFAR10 and CIFAR100 [22], Tiny ImageNet [23], and ImageNet [24] datasets. Our results suggest that Hebbian learning is generally suitable for training early feature extraction layers, or to retrain higher network layers in fewer training epochs than backprop. Moreover, our experiments show that Hebbian learning outperforms VAE training, with HPCA performing generally better than HWTA. In [15] we further extended our experiments to cover a broader spectrum of competitive learning approaches, beyond HWTA, namely k-WTA and soft-WTA strategies.

In [9] we considered even more biologically realistic models, focusing on the goal of adopting biological neuronal cultures to solve AI tasks. Previous work has shown that it was possible to train neuronal cultures on MEA devices, to recognize very simple patterns. However, this work was mainly focused to demonstrate that it was possible to induce plasticity in cultures, rather than performing a rigorous assessment of their pattern recognition performance. In our paper, we addressed this gap by developing a methodology that allowed us to assess the performance of neuronal cultures on a learn-

ing task. Specifically, we proposed a digital model of the real cultured neuronal networks; we identified biologically plausible simulation parameters that allowed us to reliably reproduce the behavior of real cultures; we used the simulated culture to perform handwritten digit recognition and rigorously evaluate its performance; we also showed that it is possible to find improved simulation parameters for the specific task, which can guide the creation of real cultures.

3. Applications

Our research theme on bio-inspired machine learning found application in two real-world contexts: 1) learning with scarce data, and 2) learning with biological neuronal cultures.

The first scenario is promising in the direction of Sustainable AI, because heavy contributions to the energy footprint of machine learning are given by complex training procedures which require to process huge amounts of data. On the other hand, biological systems appear to be able to learn from little experience. In this perspective, taking inspiration from biological systems holds the promise to revolutionize learning algorithms towards higher sample-efficiency,

In [13, 14, 17] we proposed to address the issue of sample efficiency, in CNNs, with a semi-supervised training strategy that combines Hebbian learning with gradient descent: all internal layers (both convolutional and fully connected) were pre-trained using an unsupervised approach based on Hebbian learning, and the last fully connected layer (the classification layer) was trained using Stochastic Gradient Descent (SGD). In fact, as Hebbian learning is an unsupervised learning method, its potential lies in the possibility of training the internal layers of a CNN without labels. Only the final fully connected layer has to be trained with labeled examples. We realized a machine learning module implementing this strategy, which is available online ¹. We performed experiments on various object recognition datasets, in different regimes of sample efficiency, comparing our semi-supervised (Hebbian for internal layers + SGD for the final fully connected layer) approach with end-to-end supervised backprop training, and with semi-supervised learning based on VAEs. The results showed that, in regimes where the number of available labeled samples was low, our semi-supervised approach outperformed the other approaches in almost all the cases. Further work on Hebbian algorithms [18] allowed us to obtain an extreme performance improvement, up to 50 times in training speed, by leveraging efficient GPU computations. This highlights the promises of bio-inspired solutions as

efficient and sustainable alternatives for complex model training.

In the second scenario, in [9], we developed a simulator of neuronal cultures on MEA devices, which is available online ². This is part of a broader ongoing project, aimed at using biological neuronal cultures for solving AI tasks, with the goal of providing more sustainable solutions for AI. We compared the behavior of the simulated culture with that of biological cultures, tuning the simulation parameters to make the simulated results as close as possible to the real-world data. We validated the simulated culture using digit recognition as a test case. We also found that, by appropriately modifying the simulation parameters, it was possible to further improve performance. We showed that, through simulation, it is possible to obtain insights on the parameters and properties (such as strength and range of excitatory and inhibitory connections) that a neuronal culture should have in order to perform well at a given task. These insights can then be used by neuroscientists in order to develop biological networks, by means of modern cultivation techniques, with the desired properties.

It was observed that simulated neurons were able to develop feature extractors, encoded in their weights, that are reminiscent of edges and shapes found in the patterns from the MNIST [21] digit recognition dataset used during training. Fig. 2 shows a visualization of such feature extractors.

4. Conclusions and Future Work

In conclusion, bio-inspired learning approaches represent a promising direction of future research towards more sustainable AI systems.

Future challenges in the field of bio-inspired learning will involve the practical realization of biological devices to carry out AI computations at scale. In this perspective, the use of MEA devices seems promising, but also costly. On the other hand, we are exploring also the possibility to realize the interfacing with biological cultures via optogenetics, using LED screens and photodetectors as interfaces with the neuronal cultures.

Acknowledgments

This work was partially supported by: Tuscany Health Ecosystem (THE) Project (CUP I53C22000780001), funded by the National Recovery and Resilience Plan (NRPP), within the NextGeneration Europe (NGEU) Program PNRR - M4C2 - Investimento 1.3, Partenariato Esteso PE00000013 - "FAIR - Future Artificial Intelligence

¹<https://github.com/GabrieleLagani/HebbianLearning>

²<https://github.com/GabrieleLagani/SpikingGrid>




Figure 2: Visualization of feature extractors developed by the biologically realistic culture trained on MNIST digits.

Research” - Spoke 1 ”Human-centered AI”, funded by the European Commission under the NextGeneration EU programme.

INAROS (INtelligenza ARTificiale per il mOnitoraggio e Supporto agli anziani) project co-funded by Tuscany Region POR FSE CUP B53D21008060008.

References

- [1] A. Badar, A. Varma, A. Staniec, M. Gamal, O. Magdy, H. Iqbal, E. Arani, B. Zonooz, Highlighting the importance of reducing research bias and carbon emissions in cnns, arXiv preprint arXiv:2106.03242 (2021).
- [2] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2016, pp. 770–778.
- [3] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, Bert: Pre-training of deep bidirectional transformers for language understanding, arXiv preprint arXiv:1810.04805 (2018).
- [4] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, et al., Mastering the game of go with deep neural networks and tree search, nature 529 (2016) 484.
- [5] F. Javed, Q. He, L. E. Davidson, J. C. Thornton, J. Albu, L. Boxt, N. Krasnow, M. Elia, P. Kang, S. Heshka, et al., Brain and high metabolic rate organ mass: contributions to resting energy expenditure beyond fat-free mass, The American journal of clinical nutrition 91 (2010) 907–912.
- [6] L. Yu, Z. Shen, C. Wang, Y. Yu, Efficient coding and energy efficiency are promoted by balanced excitatory and inhibitory synaptic currents in neuronal network, Frontiers in cellular neuroscience 12 (2018) 123.
- [7] W. Gerstner, W. M. Kistler, Spiking neuron models: Single neurons, populations, plasticity, Cambridge university press, 2002.
- [8] C. Gamrat, O. Bichler, D. Roclin, Memristive based device arrays combined with spike based coding can enable efficient implementations of embedded neuromorphic circuits, 2015 IEEE International Electron Devices Meeting (IEDM) 2016 (2015) 4–5.
- [9] G. Lagani, R. Mazziotti, F. Falchi, C. Gennaro, G. M. Cicchini, T. Pizzorusso, F. Cremisi, G. Amato, Assessing pattern recognition performance of neuronal cultures through accurate simulation, in: 2021 10th International IEEE/EMBS Conference on Neural Engineering (NER), IEEE, 2021, pp. 726–729.
- [10] G. W. Gross, E. Rieske, G. Kreutzberg, A. Meyer, A new fixed-array multi-microelectrode system designed for long-term monitoring of extracellular single unit neuronal activity in vitro, Neuroscience letters 6 (1977) 101–105.
- [11] I. Meloni, D. Sachidanandan, A. S. Thum, R. J. Kit-

- tel, C. Murawski, Controlling the behaviour of *drosophila melanogaster* via smartphone optogenetics, *Scientific reports* 10 (2020) 1–11.
- [12] G. Amato, F. Carrara, F. Falchi, C. Gennaro, G. Lagani, Hebbian learning meets deep convolutional neural networks, in: *International Conference on Image Analysis and Processing*, Springer, 2019, pp. 324–334.
- [13] G. Lagani, F. Falchi, C. Gennaro, G. Amato, Hebbian semi-supervised learning in a sample efficiency setting, *Neural Networks* 143 (2021) 719–731.
- [14] G. Lagani, F. Falchi, C. Gennaro, G. Amato, Evaluating hebbian learning in a semi-supervised setting, in: *International Conference on Machine Learning, Optimization, and Data Science*, Springer, 2021, pp. 365–379.
- [15] G. Lagani, F. Falchi, C. Gennaro, G. Amato, Training convolutional neural networks with competitive hebbian learning approaches, in: *International Conference on Machine Learning, Optimization, and Data Science*, Springer, 2021, pp. 25–40.
- [16] G. Lagani, F. Falchi, C. Gennaro, G. Amato, Comparing the performance of hebbian against back-propagation learning using convolutional neural networks, *Neural Computing and Applications* 34 (2022) 6503–6519.
- [17] G. Lagani, D. Bacciu, C. Gallicchio, F. Falchi, C. Gennaro, G. Amato, Deep features for cbir with scarce data using hebbian learning, *arXiv preprint arXiv:2205.08935* (2022).
- [18] G. Lagani, C. Gennaro, H. Fassold, G. Amato, Fasthebb: Scaling hebbian training of deep neural networks to imagenet level, in: *International Conference on Similarity Search and Applications*, Springer, 2022, pp. 251–264.
- [19] R. C. O’Reilly, Y. Munakata, *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*, MIT press, 2000.
- [20] D. P. Kingma, M. Welling, Auto-encoding variational bayes, *arXiv preprint arXiv:1312.6114* (2013).
- [21] Y. LeCun, L. Bottou, Y. Bengio, P. Haffner, et al., Gradient-based learning applied to document recognition, *Proceedings of the IEEE* 86 (1998) 2278–2324.
- [22] A. Krizhevsky, G. Hinton, *Learning multiple layers of features from tiny images*, 2009.
- [23] J. Wu, Q. Zhang, G. Xu, *Tiny ImageNet Challenge*, Technical Report, Stanford University, 2017.
- [24] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, L. Fei-Fei, Imagenet: A large-scale hierarchical image database, in: *2009 IEEE conference on computer vision and pattern recognition*, Ieee, 2009, pp. 248–255.