Continual learning of neural networks

Supervisors

dr hab. inż. Tomasz Trzciński, prof. PW (IDEAS NCBR)
dr hab. inż. Paweł Morawiecki, prof. IPI PAN

Description

Even with recent advancements in image, text, and sound processing powered by neural networks, tailoring these models to suit dynamic data conditions presents a noteworthy challenge. Continual learning, a field that addresses the alteration of data characteristics used in training models over time, plays a crucial role in this regard. The most significant challenge within this discipline is catastrophic forgetting. This phenomenon occurs when a model trained sequentially on two datasets loses its performance capabilities on the first dataset as it gets trained on the second. The detrimental effect of this challenge is that the network tends to forget the previously acquired knowledge, significantly limiting its practical application.

This doctoral project aims to develop methods for training deep neural networks, directly addressing this issue of forgetfulness. The principal objective of the research is to establish a strategy to handle the 'forgetting' problem effectively, thereby enhancing the efficiency and performance of the network when exposed to new data, while still retaining knowledge from older data. To attain this, the project will engage in an in-depth exploration of various training methods, their strengths, and limitations, in the context of continual learning.

Additionally, it will critically analyze existing learning mechanisms and algorithms, aiming to modify and extend them, if necessary, to better tackle catastrophic forgetting. One of the approaches this project may consider is the implementation of memory augmentation strategies, including architectural modifications, or devising new learning rules that are more resilient to forgetting.

Furthermore, the project also aspires to explore new realms of application for continual learning. This involves identifying fields and sectors where this approach can be beneficial, such as real-time data processing, IoT devices, autonomous systems, dynamic environments, or in any setting where the system needs to adapt to the changing data continually [1, 2].

Overall, the intention is to make a significant contribution to the field of continual learning, providing valuable insights, developing practical methodologies, and opening up new areas of application that can help to unlock the full potential of deep neural networks in a world that is increasingly characterized by dynamic and changing data conditions.

Requirements

- MSc degree in computer science or related field,
- Good knowledge of deep learning, including practical experience with programming in Python and relevant libraries (PyTorch, TensorFlow, Keras)
- Advanced skills in written and spoken English
- Publication track record in major CS/ML/CV venues (e.g. CVPR, ICML, NeurIPS) **is a plus**

References

[1] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(12):2935–2947, 2018.

[2] Maciej Wolczyk, Karol J. Piczak, Bartosz Wójcik, Lukasz Pustelnik, Pawel Morawiecki, Jacek Tabor, Tomasz Trzcinski, and Przemyslaw Spurek. Continual learning with guarantees via weight interval constraints. In *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 23897–23911. PMLR, 2022.