Continual Lifelong Learning for Intelligent Agents

Ghada Sokar
Eindhoven University of Technology, The Netherlands
g.a.z.n.sokar@tue.nl

Abstract

Deep neural networks have achieved outstanding performance in many machine learning tasks. However, this remarkable success is achieved in closed and static environments where the model is trained using large training data of a single task and deployed for testing on data with a similar distribution. Once the model is deployed, it becomes fixed and inflexible to new knowledge. This contradicts real-world applications, in which agents interact with open and dynamic environments and deal with non-stationary data. This Ph.D. research aims to propose efficient approaches that can develop intelligent agents capable of accumulating new knowledge and adapting to new environments without forgetting the previously learned ones.

1 Introduction

To consider an agent as truly intelligent, it should be able to continuously learn new knowledge over time, use the previously learned knowledge to help in future learning, and preserve the old knowledge when learning a new one. However, deep neural networks have a non-optimal ability to learn in non-stationary distributions. When the model is optimized to learn new representations, the previously learned knowledge is overwritten. This phenomenon is well known as catastrophic forgetting [McCloskey and Cohen, 1989] and is considered as the main obstacle for providing intelligent agents.

In the last few years, Continual Learning (CL) becomes an active research area that aims to overcome this limitation of classical machine learning and provide agents that can learn a number of tasks sequentially. The most successful methods for mitigating forgetting rely on replaying the data of previous tasks with the current data. However, previous data is not always available in real applications. It also requires additional memory to store previous samples and computation overhead to retrain them. This hinders the agent's ability to fast adapt to new environments and leads to a fast-growing memory footprint. Another direction of research is to expand the network when new tasks arrive. This increases the system size over time which hinders its scalability to a large number of tasks and applicability for embedded devices. The last direction uses a fixed model size and constrains the change in the important weights of previous tasks when the agent learns new ones. Still, these methods have not achieved a satisfactory performance yet.

This research takes an alternative approach to tackle the above-mentioned challenges and limitations by taking inspiration from *human-learning*. People learn throughout life and the brain is efficient in accumulating more and more knowledge without catastrophic forgetting. Fundamental observations from neuroscience showed that brain neurons encode information in a sparse and distributed way [Attwell and Laughlin, 2001] and at any point, only 1 to 4% of the neurons are active [Lennie, 2003]. This is unlike classical training in which all the parameters are optimized to learn each task in the sequence and produce dense representations which likely interfere with each other. Our second inspiration from human-learning is that when we expose to new situations, we select the relevant knowledge only from the past to adapt instead of using all known knowledge. This also contradicts current methods in which each task uses all previously learned knowledge regardless of its relevance for it.

Inspired by these observations, this research aims at providing efficient CL approaches with four research goals: First, reflect the sparse brain activity in the CL paradigm to mitigate forgetting in fixed-capacity models without extensive retain of previous data or adding extra computation and memory overhead. Second, perform a selective transfer and study its role in the forward and backward transfer. Third, disentangle from each task the generic representation that would be useful for future tasks. Fourth, maximize the reusability of previous knowledge without forgetting.

2 Contributions

To address our first goal, we proposed a new brain-inspired method named SpaceNet [Sokar *et al.*, 2021c]. In this work, we harness the significant redundancy of deep neural networks [Denil *et al.*, 2013] and utilize the model capacity efficiently instead of expanding the model with each new task. We train each task from scratch using *sparse connections*. In addition, motivated by the brain sparse activity, we encourage each task to learn *sparse representations*. To fulfill this goal, we proposed a new sparse training algorithm with dynamic sparsity to train each task. During training, the distribution of the sparse connections is adaptively changed and compacted in the most important neurons for the current

task. This results in generating semi-distributed sparse representation for each task. This representation has two key advantages: First, it reduces the interference between tasks, which consequently reduces forgetting. Second, it leaves free neurons for future tasks which increases the scalability of the model. Hence, SpaceNet accounts for both previous and future tasks using a fixed-capacity model; avoiding adding any extra memory or computation overhead to train new tasks or remember old ones.
To assess SpaceNet, we evaluated and illustrated its success in the following aspects compared with the state-of-the-art: (1) the average accuracy across all tasks after learning the whole sequence, (2) the backward transfer to estimates forgetting, (3) the memory reduction, (4) the role of sparse representation in reducing forgetting, and (5) the ability to identify the important neurons corresponding to each task.

To address our second goal, we proposed self-attention meta-learner for continual learning (SAM) [Sokar *et al.*, 2021b]. This work focuses on two requirements that are not widely addressed in the state-of-the-art. First, the necessity of having a good quantity of prior generic knowledge to promote learning new tasks. Second, selective transfer of relevant parts only from previous knowledge to learn each of the future tasks instead of using the whole knowledge. SAM divides the network into two parts. The first part is a shared sub-network meta-trained to learn prior knowledge that can generalize to out-of-domain tasks. The shared sub-network is incorporated with a self-attention mechanism to select the relevant representation for each input task from the prior knowledge. The second part of the network contains a specific branch consists of a few layers for each task to capture the specific discriminative representation. Each task builds this specific branch on top of the *relevant selected sparse* representation from the prior knowledge.
To assess SAM, we evaluated the average accuracy across all tasks and the forward transfer. We showed that SAM achieves a better performance than the state-of-the-art methods. We also analyzed the role of each of our proposed components in decreasing forgetting and increasing positive transfer. Lastly, we demonstrate that popular existing CL methods gain a performance boost when they are combined with our framework.

To address our third goal, we present a new method, named learning Invariant Representation for CL (IRCL) [Sokar *et al.*, 2021a]. In this work, we proposed a new pseudo-rehearsal based method in which we use a unified network for classification and image generation. We harness the conditional generative sub-network to disentangle the invariant representation during learning each task. We showed that this representation is less prone to forgetting which increases the performance of the CL system and reduces the negative transfer. We also illustrated the role of this representation in reducing the number of required pseudo-samples for replay.

3 Conclusion and Future Research

In this work, we have proposed some brain-inspired methods for continual learning. We show the effectiveness of sparse representation in reducing catastrophic forgetting without replay. We also show the importance of selective transfer in increasing forward transfer and reducing negative backward transfer. Finally, we illustrate the effectiveness of extracting the invariant representation which is less prone to forgetting.

In our future work, we intend to continue studying the connection between neuroscience and machine learning; aiming to closing the gap between them. In particular, our goal is to address the following research points. First, study the relation between previously learned knowledge and the current task to reuse some of the sparse connections that are already allocated instead of allocating new ones. This would reduce the memory footprint and increase the scalability of the agent. Second, allow learning using fewer training samples as the agent become more knowledgeable. Third, maximize the usage of previous knowledge when learning a new relevant one and learn only the residual specific representation. Finally, study how the replay is occurred in the brain without using the raw images and reflect this in the CL paradigm.

We envision that by taking these steps, we could push towards providing intelligent agents capable of learning continuously to address real-world problems.

Acknowledgments

I sincerely thank my supervisors Decebal Constantin Mocanu and Mykola Pechenizkiy for their guidance in this research.

References

[Attwell and Laughlin, 2001] David Attwell and Simon B Laughlin. An energy budget for signaling in the grey matter of the brain. *Journal of Cerebral Blood Flow & Metabolism*, 21(10):1133–1145, 2001.

[Denil *et al.*, 2013] Misha Denil, Babak Shakibi, Laurent Dinh, Marc'Aurelio Ranzato, and Nando de Freitas. Predicting parameters in deep learning. In *Proceedings of the 26th International Conference on Neural Information Processing Systems-Volume 2*, pages 2148–2156, 2013.

[Lennie, 2003] Peter Lennie. The cost of cortical computation. *Current biology*, 13(6):493–497, 2003.

[McCloskey and Cohen, 1989] Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. *Psychology of learning and motivation*, 24:109–165, 1989.

[Sokar *et al.*, 2021a] Ghada Sokar, Decebal Constantin Mocanu, and Mykola Pechenizkiy. Learning invariant representation for continual learning. In *Meta-Learning for Computer Vision Workshop at the 35th AAAI Conference on Artificial Intelligence (AAAI-21)*, 2021.

[Sokar *et al.*, 2021b] Ghada Sokar, Decebal Constantin Mocanu, and Mykola Pechenizkiy. Self-attention meta-learner for continual learning. In *20th International Conference on Autonomous Agents and Multiagent Systems, AAMAS*, 2021.

[Sokar *et al.*, 2021c] Ghada Sokar, Decebal Constantin Mocanu, and Mykola Pechenizkiy. Spacenet: Make free space for continual learning. *Neurocomputing*, 439:1–11, 2021.