Transfer Learning: Boosting Machine Learning Efficiency and Efficacy

Branko Kavšek[a b],

[a]University of Primorska,
Faculty of Mathematics, Natural Sciences and Information Technologies,
branko.kavsek@upr.si
[b]Jožef Stefan Institute,
Department for Artificial Intelligence, branko.kavsek@ijs.si

In recent years, transfer learning has emerged as an innovation in machine learning, redefining the boundaries of data efficiency and model generalization. This technique, which involves the transfer of knowledge from one domain to another, has become crucial in addressing large, annotated datasets and in accelerating the training process for new tasks. The objective of this research is to provide a comprehensive examination of transfer learning methodologies, identify their potential in various applications, and highlight the challenges and frontiers of current research.

We begin by introducing the foundational concepts underlying transfer learning, differentiating between the paradigms of inductive, transductive, and unsupervised transfer learning. By dissecting the process of transferring and adapting features, models, or tasks, we provide insight into how prior knowledge can be used to enhance learning in a novel context.

We continue by critically analyzing the spectrum of transfer learning applications, ranging from natural language processing (NLP), where pre-trained models like BERT [1] have set new benchmarks, to computer vision, where models pre-trained on ImageNet [2] have demonstrated remarkable adaptability. We explore the efficacy of transfer learning in domains burdened with limited data availability, such as medical imaging and remote sensing.

To quantify the impact of transfer learning, we present a meta-analysis of benchmark datasets and tasks, providing empirical evidence that supports the enhanced performance and reduced training times achieved by transfer learning strategies. The analysis reveals when transfer learning offers the most significant advantages, and under what circumstances it may be less effective.

However, transfer learning is not without challenges. We discuss the critical issues of negative transfer, where inappropriate knowledge transfer can degrade performance, the difficulty of identifying relevant source and target domains, and the computational burdens of fine-tuning large pre-trained models.

Finally, we take a look at the future of transfer learning research, emphasizing novel approaches such as few-shot learning, domain adaptation techniques, and cross-lingual transfer learning.

Keywords: Transfer Learning, Domain Adaptation, Pre-trained Models, Model Generalization, Machine Learning Efficiency.

References

[1] Jacob Devlin, Ming-Wei Chang, Kenton Lee, Kristina Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding", In Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics (2019)

[2] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li and Li Fei-Fei, "ImageNet: A large-scale hierarchical image database", In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Miami, Florida, pages 248-255 (2009)