Abstract

Deep learning models in recent years have proven to be extremely proficient in many different domains such as computer vision, language processing, and cheminformatics. However, a deep learning model is usually restricted to a specific, stationary setting and cannot be trivially adapted to new conditions. On the contrary, animals and humans acting in the real world have to constantly readjust to new conditions, solve sequences of heterogenous tasks and calibrate response times depending on the situation. The constant change specific to the real world is one of the major obstacles to building more generally applicable artificial intelligence agents that can be used for a wide variety of problems. In this thesis, we look into different ways of adapting previously trained networks to new tasks and settings.

This thesis consists of five papers and is divided into two parts. The first part of this thesis contains Publications [I-III] which focus on the problem of continual learning. This research area deals with building algorithms capable of training on a shifting data distribution while reusing past knowledge to efficiently learn new tasks. It is an important problem, since artificial neural networks experience catastrophic forgetting, i.e. after a data shift the model loses its performance on the previous tasks. Publication [I] proposes a continual reinforcement learning benchmark built using a sequence of robotic manipulation tasks, along with a set of crucial metrics. Publication [II] extends the previous paper by focusing on the issue of transfer, i.e. how to reuse the knowledge from the past in order to efficiently learn on the new data. We identify key architectural and algorithmic aspects of transfer and propose a set of guidelines that lead to more effective approaches. In Publication [III], we propose a continual learning method with guarantees on forgetting, showing that performance drop can be bounded from above using interval arithmetic.

The second part of this thesis focuses on methods for adapting existing models to different settings and paradigms. In Publication [IV], we show how to efficiently adapt a deep learning model with a static computation path, so that it dynamically chooses the computation path depending on the example. We do this by introducing early exit classifiers, which allow us to control the processing time of each example. In Publication [V] we use normalizing flows to introduce conditioning factors to pre-trained unsupervised generative models, allowing for a more controlled generation process.

All the Publications listed here were presented at top A* conferences and I am the first author of all of them. They were written as a result of national as well as international cooperation. Additionally, I published several works on different topics in deep learning, interned at established research institutions, served as a reviewer for top-ranking journals and conferences, and organized workshops and summer schools on machine learning.

Keywords: deep learning, adaptation, continual learning, reinforcement learning