



ID de Contribution: 52

Type: Non spécifié

Enabling Empirically and Theoretically Sound Algorithmic Alignment

mercredi 20 avril 2022 09:00 (1h 30m)

Neural networks that are able to reliably execute algorithmic computation may hold transformative potential to both machine learning and theoretical computer science. On one hand, they could enable the kind of extrapolative generalisation scarcely seen with deep learning models. On another, they may allow for running classical algorithms on inputs previously considered inaccessible to them.

Both of these promises are shepherded by the neural algorithmic reasoning blueprint, which I have recently proposed in a position paper alongside Charles Blundell. On paper, this is a remarkably elegant pipeline for reasoning on natural inputs which carefully leverages the tried-and-tested power of deep neural networks as feature extractors. However, in practice, its success rests on successful realisation of algorithmic alignment – neural networks that are capable of executing algorithmic computations in a high-dimensional latent space, ideally in a manner that extrapolates outside of the training set.

In this talk, I will outline some of the recent work we have done on strengthening algorithmic alignment from both an empirical and a theoretical point of view. These efforts include a dataset of algorithmic reasoning tasks, to be used as a bootstrapping basis, as well as formalising the theoretical connection between (graph) neural networks and dynamic programming using the tools of category theory and abstract algebra.

Orateur: VELICKOVIC, Petar (Deepmind / University of Cambridge)

Classification de Session: Representation Learning workshop