

On the impact of modern deep-learning techniques to the performance and time-requirements of classification models in experimental high-energy physics

Tuesday, 28 July 2020 15:50 (20 minutes)

Beginning from a basic neural-network architecture, we test the potential benefits offered by a range of advanced techniques for machine learning and deep learning in the context of a typical classification problem encountered in the domain of high-energy physics, using a well-studied dataset: the 2014 Higgs ML Kaggle dataset. The advantages are evaluated in terms of both performance metrics and the time required to train and apply the resulting models. Techniques examined include domain-specific data-augmentation, learning rate and momentum scheduling, (advanced) ensembling in both model-space and weight-space, and alternative architectures and connection methods. Following the investigation, we arrive at a model which achieves equal performance to the winning solution of the original Kaggle challenge, whilst requiring about 2% of the training time and less than 5% of the inference time using much less specialised hardware. Additionally, a new wrapper library for PyTorch called LUMIN is presented, which incorporates all of the techniques studied.

Secondary track (number)

Primary author: STRONG, Giles Chatham (Universita e INFN, Padova (IT))

Presenter: STRONG, Giles Chatham (Universita e INFN, Padova (IT))

Session Classification: Computing and Data Handling

Track Classification: 14. Computing and Data Handling