

HiggsML challenge with optimized DNN

Wednesday 17 October 2018 10:15 (20 minutes)

A lot of work done in advancing the performance of deep-learning approaches often takes place in the realms of image recognition - many papers use famous benchmark datasets, such as Cifar or Imagenet, to quantify the advantages their idea offers. However it is not always obvious, when reading such papers, whether the concepts presented can also be applied to problems in other domains and still offer improvements.

One such example of another domain is the task of event classification in high-energy particle-collisions, such as those which occur at the LHC. In this presentation, a classifier trained on publicly available physics data (from the HiggsML Kaggle challenge) is used to test the domain transferability of several recent Machine-Learning concepts.

A system utilising relatively recent concepts, such as cyclical learning-rate schedules and data-augmentation, is found to slightly outperform the winning solution of the HiggsML challenge, whilst requiring less than 10% of the training time, no further feature engineering, and less specialised hardware

Application

Author: STRONG, Giles Chatham (LIP Laboratorio de Instrumentacao e Fisica Experimental de Part)

Presenter: STRONG, Giles Chatham (LIP Laboratorio de Instrumentacao e Fisica Experimental de Part)

Session Classification: External To Atlas