



Contribution ID: 63

Type: **PhD student talk**

Neural networks on a budget: Speed without sacrificing performance

Thursday, 25 June 2020 14:55 (12 minutes)

As machine-learning and deep-learning techniques become more and more ingrained in academic research tasks, two problems that researchers invariably face are:

- 1) How to train and apply models in a timely manner without sacrificing performance;
- 2) And how to do so with only a limited hardware-budget.

In my recent paper [\[1\]](#), I explore how recent techniques for neural network training and architecture can be used to match state-of-the-art performance on a HEP benchmark dataset (2014 Kaggle HiggsML) in just 14 minutes on a standard laptop CPU, as opposed to the previous time of 12 hours on a top-of-the range GPU (~100 minutes accounting for hardware advancements between 2013 and 2017).

Developed partly alongside the above referenced investigation, LUMIN [\[2\]](#) is an open-source PyTorch wrapper and data-analysis framework designed specifically for HEP usage, providing implementations of many recent and state-of-the-art methods for deep learning. The main aim is to allow HEP researchers to apply easily optimal approaches to standard tasks in HEP (i.e. classification & regression), whilst being modular enough to be adapted to more novel and investigative problems (e.g. studying the applicability of HEP-specific graph-nets to top-tagging).

This presentation would include: a quick overview of the techniques studied; and a call for contributors to help grow LUMIN to meet better the requirements of the HEP community.

[\[1\]](#) Strong 2020, <https://iopscience.iop.org/article/10.1088/2632-2153/ab983a>

[\[2\]](#) LUMIN <https://lumin.readthedocs.io/>

Primary author: STRONG, Giles (LIP)

Presenter: STRONG, Giles (LIP)

Session Classification: Session 3