Just-in-time compute for Neural Architecture Search Benchmarks in Teaching and Research

Project Supervision Team:

Stephen McGough (Stephen.McGough@newcastle.ac.uk)
David Towers (David.Towers@newcastle.ac.uk)

Executive Summary:

Neural Architecture Search (NAS) is a Deep Learning process for identifying the 'best' architecture for a particular problem and dataset. Due to the 'No free lunch' proposition there is no single architecture which will be 'best' in all cases. One of the research directions in NAS is to produce benchmarks where you train thousands of architectures on a dataset recording the performance. This benchmark can then be used to try out different NAS approaches. We need more of these for teaching. But rather than pre-compute all architectures we propose a just-in-time approach where architectures are only trained when needed.

Details:

The intention here is to build a framework where students can propose an architecture (ideally using a standard toolkit such as PyTorch) for a given dataset and the framework will interrogate a database to see if the architecture has been tried before, in which case the previously produced data is returned (including accuracy and training time). If, however, the network/dataset combination has not previously been trained then the framework will submit the architecture to a high-performance computing (HPC) system (e.g., Commet (Newcastle) or Bede (N8)) to train the architecture on the dataset. The performance data will then be added to the database and the results returned to the student.

This project will require development skills both in Deep Learning, Distributed Computing, HPC, and Databases. Although primarily designed for teaching the benchmarks produced by this framework will be of great value to those working in NAS.

Timeline:

Week1: Familiarisation with Deep Learning and the HPC environments.
Week 2: Familiarisation with NAS and the process of developing a benchmark.
Week 3: Development of the backend database system.
Week 4: Development of the framework for submitting architectures to HPC.
Week 5: Development of the overall framework to handle both pre- and non-pretested architectures.
Week 6: Bringing all components together for the framework.
Week 7: Testing and evaluation of the framework.
Week 8: Report and documentation, hopefully producing a paper.

Anticipated outcome:

The outcome will be a software framework which can be used by students to more energy-efficiently develop their skills both with Deep Learning and with NAS.

Background Reading:

1) NAS-Bench-101: Towards Reproducible Neural Architecture Search (https://arxiv.org/abs/1902.09635),
2) NAS-Bench-201: Extending the Scope of Reproducible Neural Architecture Search (https://arxiv.org/abs/2001.00326)
3) Insights from the Use of Previously Unseen Neural Architecture Search Datasets (https://ieeexplore.ieee.org/abstract/document/10657037)
4) Neural Architecture Search with Reinforcement Learning (https://arxiv.org/abs/1611.01578)