Neural Architecture Search and predicting Deep Learning training time

Supervision team
Stephen McGough: Stephen.McGough@newcastle.ac.uk

Research project 1: Neural Architecture Search
Deep Learning has revolutionised the area of Artificial Intelligence allowing problems which were previously considered too difficult to be solved to be solved easily. However, these solutions often require support from highly skilled Deep Learning developers to provide optimal solutions. Meaning that may people lose out on the ability to exploit these technologies well. Neural Architecture Search (NAS) is a branch of AutoML where the system can identify the most optimal Deep Learning network for a given problem (e.g., data set and available compute). Previously we have developed Neural Architecture Search approaches: Bonsai-Net[1] and SpiderNet[2] which offer a novel approach to Neural Architecture Search. We are also advocating that most NAS approaches are too highly tuned to the common datasets (CIFAR-10/100, ImageNet) and do not work as well on other datasets [3]. We are seeking here a bright PhD candidate to work on a range of NAS projects. Of particular interest is someone to work on hardware constrained NAS.

Research project 2: Predicting the computational cost of Deep Learning
It is very difficult to predict how long it will take to train a Deep Learning network – often leading to excessive and unpredicted computational cost (e.g., number of hours for training leading to energy and/or cloud cost). As this cost is so high it would be extremely beneficial to know upfront what this training time would be – to judge if it is worthwhile. In previous work [4], we have made great inroads into understanding the computational cost of each training epoch. This PhD would be to expand this work to look at determining the likely number of epochs required to train the network.

Applicant skills/background
These PhDs would be ideally suited to candidates with strong mathematical and computing backgrounds and a willingness to apply these to novel processes.

References
[1] R. Geada, D. Prangle, and A.S. McGough. "Bonsai-Net: One-shot neural architecture search via differentiable pruners." arXiv preprint arXiv:2006.09264 (2020).
[2] R. Geada, and A.S. McGough. "SpiderNet: Hybrid differentiable-evolutionary architecture search via train-free metrics." In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 1962-1970. 2022.
[3] R. Geada, D. Towers, M. Forshaw, A. Atapour-Abarghouei, and A.S. McGough. "Insights from the use of previously unseen neural architecture search datasets." In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 22541-22550. 2024.
[4] D. Justus, J. Brennan, S. Bonner, and A.S. McGough. "Predicting the computational cost of deep learning models." In 2018 IEEE international conference on big data (Big Data), pp. 3873-3882. IEEE, 2018.