

R&D defense Arun Rajendra Prabhu

Date: Tuesday, 11.06.2019

Time: 15:30

Room: C181

Title:

Neural Scrutiny: A comparative study of neural network architecture generation techniques

Abstract:

Deep Convolutional Neural Networks (CNN's) based techniques have outperformed, classical computer vision algorithms on a variety of computer vision tasks like image classification, object detection, semantic segmentation, etc. Work by [1] illustrates the importance of CNN architectures on the performance of deep learning based object detectors. Hence coming up with proficient CNN architectures is essential for improving the performances of such sophisticated deep learning based computer vision techniques. The development of many state-of-the-art CNN architectures has been fostered by the competitions such as ImageNet Large Scale Visual Recognition Challenge (ILSVRC)[2]. Though these network architectures are themselves capable of learning from data (images), they still need to be designed manually by a team of engineers and experts in the field. This is a highly iterative process and involves a lot of trial and error experiments. Hence arriving at an optimal network is a time-intensive endeavor.

Neural Architecture Search or NAS, in short, is a relatively new research area which intends to eliminate the involvement of human experts in this time-intensive and repetitive activity. Work by [3],[4] have shown very promising results where the machine-generated network architectures outperformed the previous human designed architectures. However one of the dissuading factors of these techniques was the use of a gargantuan amount of computational resources. From then on, the field of NAS has gained a lot of popularity. The sheer amount of publications on the topic make it difficult to distinguish a promising NAS technique from the rest.

Neural Scrutiny investigates the existing work in the field of automated neural architecture search and provides a curated list of promising NAS techniques. Practically relevant metrics like open source implementation, reasonable computational requirements, etc are used to make the distinction between promising techniques and the rest. In addition to this, it also provides an unbiased empirical performance evaluation of the 4 most promising NAS techniques according to our analysis, which are DARTS[5], DeepArchitect[6], CGP-CNN[7] and Autokeras[8]. This evaluation will be done using multiple datasets of varying difficulty, image sizes, sample sizes, color spaces, etc. The datasets shortlisted for this are CIFAR-10[9], Devanagari[10], Fashion MNIST[11], MNIST[12], STL-10[13] and SVHN[14]. The experiments that are devised for this evaluation, judge the NAS techniques based on two factors,

1. The quality of the best-found architecture.
2. The quality of the overall search process.

To provide a relative ranking between these shortlisted NAS techniques, this work makes use of many standard metrics like validation and test accuracy, precision, recall, F1 score, etc. In addition to these, this work also introduces a metric called Network

Betterment Score which quantifies the quality of the overall search process. Neural Scrutiny also introduces a Fraunhofer lines style visualization technique, which can be used to better perceive the correlation between the training and validation accuracies of the different generated architectures during the search process. The result of this grueling evaluation process was that DARTS[63] beat the competition by a huge margin, this was followed by DeepArchitect[72] in second, Autokeras[47] came in third and CGP-CNN[92] was fourth.

References:

- [1] Jonathan Huang, Vivek Rathod, Chen Sun, Menglong Zhu, Anoop Korattikara, Alireza Fathi, Ian Fischer, Zbigniew Wojna, Yang Song, Sergio Guadarrama, and Kevin Murphy. Speed/accuracy trade-offs for modern convolutional object detectors. CoRR, abs/1611.10012, 2016. URL <http://arxiv.org/abs/1611.10012>.
- [2] Jia Deng, Wei Dong, Richard Socher, Li jia Li, Kai Li, and Li Fei-fei. Imagenet: A large-scale hierarchical image database. In In CVPR, 2009.
- [3] Barret Zoph and Quoc V. Le. Neural architecture search with reinforcement learning. CoRR, abs/1611.01578, 2016. URL <http://arxiv.org/abs/1611.01578>.
- [4] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V. Le. Learning transferable architectures for scalable image recognition. CoRR, abs/1707.07012, 2017. URL <http://arxiv.org/abs/1707.07012>.
- [5] Hanxiao Liu, Karen Simonyan, and Yiming Yang. DARTS: differentiable architecture search. CoRR, abs/1806.09055, 2018. URL <http://arxiv.org/abs/1806.09055>.
- [6] Renato Negrinho and Geoffrey J. Gordon. Deeparchitect: Automatically designing and training deep architectures. CoRR, abs/1704.08792, 2017.
- [7] Masanori Suganuma, Shinichi Shirakawa, and Tomoharu Nagao. A genetic programming approach to designing convolutional neural network architectures. CoRR, abs/1704.00764, 2017. URL <http://arxiv.org/abs/1704.00764>.
- [8] Haifeng Jin, Qingquan Song, and Xia Hu. Efficient neural architecture search with network morphism. CoRR, abs/1806.10282, 2018. URL <http://arxiv.org/abs/1806.10282>.
- [9] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar-10 (canadian institute for advanced research). URL <http://www.cs.toronto.edu/~kriz/cifar.html>.
- [10] Shailesh Raj Acharya, Ashok Kumar Pant, and Prashanna Kumar Gyawali. Deep learning based large scale handwritten devanagari character recognition. 2015 9th International Conference on Software, Knowledge, Information Management and Applications (SKIMA), pages 1–6, 2015.
- [11] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. CoRR, abs/1708.07747, 2017. URL <http://arxiv.org/abs/1708.07747>.
- [12] Yann Lecun, Lon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. In Proceedings of the IEEE, pages 2278–2324, 1998.
- [13] Adam Coates, Andrew Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In Geoffrey Gordon, David Dunson, and Miroslav Dudk, editors, Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics, volume 15 of Proceedings of Machine Learning Research, pages 215–223, Fort Lauderdale, FL, USA, 11–13 Apr 2011. PMLR. URL <http://proceedings.mlr.press/v15/coates11a.html>.
- [14] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y. Ng. Reading digits in natural images with unsupervised feature learning. In NIPS Workshop on Deep Learning and Unsupervised Feature Learning 2011, 2011. URL http://ufldl.stanford.edu/housenumbers/nips2011_housenumbers.pdf