



# NetScore: Towards Universal Metrics for Large-Scale Performance Analysis of Deep Neural Networks for Practical On-Device Edge Usage

Alexander Wong<sup>1,2(✉)</sup>

<sup>1</sup> Waterloo Artificial Intelligence Institute, University of Waterloo, Waterloo, ON, Canada

a28wong@uwaterloo.ca

<sup>2</sup> DarwinAI Corp., Waterloo, ON, Canada

**Abstract.** Much of the focus in the design of deep neural networks has been on improving accuracy, leading to more powerful yet highly complex network architectures that are difficult to deploy in practical scenarios, particularly on edge devices such as mobile and other consumer devices given their high computational and memory requirements. As a result, there has been a recent interest in the design of quantitative metrics for evaluating deep neural networks that accounts for more than just model accuracy as the sole indicator of network performance. In this study, we continue the conversation towards universal metrics for evaluating the performance of deep neural networks for practical on-device edge usage. In particular, we propose a new balanced metric called **NetScore**, which is designed specifically to provide a quantitative assessment of the balance between accuracy, computational complexity, and network architecture complexity of a deep neural network, which is important for on-device edge operation. In what is one of the largest comparative analysis between deep neural networks in literature, the NetScore metric, the top-1 accuracy metric, and the popular information density metric were compared across a diverse set of 60 different deep convolutional neural networks for image classification on the ImageNet Large Scale Visual Recognition Challenge (ILSVRC 2012) dataset. The evaluation results across these three metrics for this diverse set of networks are presented in this study to act as a reference guide for practitioners in the field. The proposed NetScore metric, along with the other tested metrics, are by no means perfect, but the hope is to push the conversation towards better universal metrics for evaluating deep neural networks for use in practical on-device edge scenarios to help guide practitioners in model design for such scenarios.

**Keywords:** Performance analysis · Deep neural networks · On-device · Edge usage

Supported by Natural Sciences and Engineering Research Council of Canada (NSERC), the Canada Research Chairs program, Nvidia, and DarwinAI.

## 1 Introduction

There has been a recent urge in both research and industrial interests in deep learning [4], with deep neural networks demonstrating state-of-the-art performance in recent years across a wide variety of applications. In particular, deep convolutional neural networks [5,6] has been shown to outperform other machine learning approaches for visual perception tasks ranging from image classification [19] to object detection [22] and segmentation [11]. One of the key driving factors behind the tremendous recent successes in deep neural networks has been the availability of massive computing resources thanks to the advances and proliferation of cloud computing and highly parallel computing hardware such as graphics processing units (GPUs). The availability of this wealth of computing resources has enabled researchers to explore significantly more complex and increasingly deeper neural networks that has resulted in significant performance gains over past machine learning methods. For example, in the realm of visual perception, the depth of deep convolutional neural networks with state-of-the-art accuracies have reached hundreds of layers, hundreds of millions of parameters in size, and billions of calculations for inferencing.

While the ability to build such large and complex deep neural networks has led to a constant increase in accuracy, the primary metric for performance widely leveraged for evaluating networks, it has also created significant barriers to the deployment of such networks for practical edge device usage. The practical deployment bottlenecks associated with the powerful yet highly complex deep neural networks in research literature has become even more visible in recent years due to the incredible proliferation of mobile devices, consumer devices, and other edge devices and the increasing demand for machine learning applications in such devices. As a result, the design of deep neural networks that account for more than just accuracy as the sole indicator of network performance and instead strike a strong balance between accuracy and complexity has very recently become a very hot area of research focus, with a number of different deep neural network architectures designed specifically with efficiency in mind [14,18,26,28,33,34,36].

One of the key challenges in designing deep neural networks that strikes a strong balance between accuracy and complexity for practical usage lies in the difficulties with assessing how well a particular network architecture is striking that balance. As previous mentioned, using accuracy as the sole metric for network performance does not provide the proper indicators of how efficient a particular network is in practical scenarios such as deployment on mobile devices and other consumer devices. As a result, there has been a recent interest in the design of quantitative metrics for evaluating deep neural networks that accounts for more than just model accuracy. In particular, it is generally desirable to design such metrics in a manner that is as hardware vendor agnostic as possible so that different network architectures can be compared to each other in a consistent manner. One of the most widely cited metrics in research literature for assessing the performance of deep neural networks that accounts for both accuracy and architectural complexity is the information density metric

proposed by [1], which attempts to measure the relative amount of accuracy captured within one of the most basic building blocks of a deep neural network: a parameter. More specifically, the information density ( $D(\mathcal{N})$ ) of a deep neural network  $\mathcal{N}$  is defined as the accuracy of the deep neural network (denoted by  $a(\mathcal{N})$ ) divided by the number of parameters needed for representing it (denoted by  $p(\mathcal{N})$ ),

$$D(\mathcal{N}) = \frac{a(\mathcal{N})}{p(\mathcal{N})} \quad (1)$$

While highly effective for giving a good general idea of the balance between accuracy and architectural complexity (which also acts as a good indicator for memory requirements), the information density metric does not account for the fact that, depending on the design of the network architecture, the architecture complexity does not necessarily reflect the computational requirements for performing network inference (e.g., MobileNet [14] has more parameters than SqueezeNet [18] but has lower computational requirements for network inference). Therefore, the exploration and investigation towards universal performance metrics that account for accuracy, architectural complexity, and computational complexity is highly desired as it has the potential to improve network model search and design.

In this study, we continue the conversation towards universal metrics for evaluating the performance of deep neural networks for practical usage. In particular, we propose a new balanced metric called **NetScore**, which is designed specifically to provide a quantitative assessment of the balance between accuracy, computational complexity, and network architecture complexity of a deep neural network. This paper is organized as follows. Section 2 describes the proposed NetScore metric and the design principles around it. Section 3 presents and discusses experimental results that compare the NetScore, information density, and top-1 accuracy across 60 different deep convolutional neural networks for image classification on the ImageNet Large Scale Visual Recognition Challenge (ILSVRC 2012) dataset [25], making this one of the largest comparative studies between deep neural networks.

## 2 NetScore: Design Principles

The proposed NetScore metric (denoted here as  $\Omega$ ) for assessing the performance of a deep neural network  $\mathcal{N}$  for practical usage can be defined as:

$$\Omega(\mathcal{N}) = 20 \log \left( \frac{a(\mathcal{N})^\alpha}{p(\mathcal{N})^\beta m(\mathcal{N})^\gamma} \right) \quad (2)$$

where  $a(\mathcal{N})$  is the accuracy of the network,  $p(\mathcal{N})$  is the number of parameters in the network,  $m(\mathcal{N})$  is the number of multiply-accumulate (MAC) operations performed during network inference, and  $\alpha$ ,  $\beta$ ,  $\gamma$  are coefficients that control the influence of accuracy, architectural complexity, and computational complexity of the network on  $\Omega$ . A number of design principles were taken into consideration in the design of the proposed NetScore metric, which is described below.

## 2.1 Model Accuracy Representation

In the NetScore metric, the obvious incorporation of the model accuracy  $a(\mathcal{N})$  of the network  $\mathcal{N}$  into the metric is in the numerator of the ratio, as an increase in accuracy should naturally lead to an increase in the metric, similar to the information density metric [1]. We further introduce a coefficient  $\alpha$  in the proposed NetScore metric to provide better control over the influence of model accuracy on the overall metric. In particular, we set  $\alpha = 2$  to better emphasize the importance of model accuracy in assessing the overall performance of a network in practical usage, as deep convolutional neural networks that have unreasonably low model accuracy remain unusable in practical scenarios, regardless how small or fast the network is. In this study, the unit used for  $a(\mathcal{N})$  is in percent top-1 accuracy on the ILSVRC 2012 dataset [25].

## 2.2 Model Architectural and Computational Complexity Representations

Taking inspiration from the information density metric [1], we represent the architectural complexity of a deep neural network by the number of parameters  $p(\mathcal{N})$  in the network  $\mathcal{N}$  and incorporate it in the denominator of the ratio. As such, the architecture complexity of the network is inversely proportional to the metric  $\Omega$ , where an increase in architectural complexity results in a decrease in  $\Omega$ . In addition, we incorporate the computational complexity of the deep neural network as an additional factor in the denominator of the ratio to be taken into consideration for assessing the overall performance of a network for practical usage, which is particularly important in operational scenarios such as inference on mobile devices and other consumer devices where computational power is limited. To represent the computational complexity of the network  $\mathcal{N}$  in a manner that is relatively hardware vendor agnostic, thus enabling a more consistent comparison between networks, we chose to leverage the number of multiply-accumulate (MAC) operations necessary for performing network inference. Given that the computational bottleneck associated with performing network inference on a deep neural network is predominantly in the computation of MAC operations, the number of MAC operations  $m(\mathcal{N})$  is a good proxy for the computational complexity of the network. By incorporating both architectural and computational complexity, the proposed NetScore metric can better quantify the balance between accuracy, memory requirements, and computational requirements in practical usage. Furthermore, we introduce two coefficients ( $\beta$  and  $\gamma$ , respectively) to provide better control over the influence of architectural and computational complexity on the overall metric. In particular, we set  $\beta = 0.5$  and  $\gamma = 0.5$  since, while architectural and computational complexity are both very important factors to assessing the overall performance of a network in practical scenarios, the most important metric remains the model accuracy given that, as eluded to before, networks with unreasonably low model accuracy are not useful in practical scenarios regardless of size and speed.

Given these coefficients, NetScore is in the units of squared percentage accuracy per root parameter per root MAC operation, and represents the capacity of a network architecture to utilize its full learning and computing capacity.

### 2.3 Logarithmic Scaling

One of the difficulties with comparing the overall performance of different deep neural networks with each other is their great diversity in their model accuracy, architectural complexity, and computational complexity. This makes the dynamic range of the performance metric quite large and unwieldy for practitioners to compare for model search and design purposes. To account for this large dynamic range, we take inspiration from the field of signal processing; in particular, the logarithmic scale commonly used to express the ratio between one value of a property to another. In the proposed NetScore metric, we transform the ratio between the model accuracy property ( $a(\mathcal{N})$ ) and the model architectural and computational complexity ( $p(\mathcal{N})$  and  $m(\mathcal{N})$ ) into the logarithmic scale to reduce the dynamic range to within a more readily interpretable range.

## 3 Experimental Results and Discussion

To get a better sense regarding the overall performance of the huge wealth of deep convolutional neural networks introduced in research literature in the context of practical usage, we perform a large-scale comparative analysis across a diverse set of 60 different deep convolutional neural networks designed for image classification using the following quantitative performance metrics: (i) top-1 accuracy, (ii) information density, and (iii) the proposed NetScore metric. The dataset of choice for the comparative analysis in this study is the ImageNet Large Scale Visual Recognition Challenge (ILSVRC 2012) dataset [25], which consists of 1000 different classes. To the best of the author’s knowledge, this comparative analysis is one of the largest in research literature and the hope is that the results presented in this study can act as a reference guide for practitioners in the field.

The set of deep convolutional neural networks being evaluated in this study are: AlexNet [19], AmoebaNet-A (4, 50) [23], AmoebaNet-A (6, 190) [23], AmoebaNet-A (6, 204) [23], AmoebaNet-B (3, 62) [23], AmoebaNet-B (6, 190) [23], AmoebaNet-C (4, 50) [23], AmoebaNet-C (6, 228) [23], CondenseNet (G = C = 4) [16], CondenseNet (G = C = 8) [16], DenseNet-121 (k = 32) [17], DenseNet-169 (k = 32) [17], DenseNet-161 (k = 48) [17], DenseNet-201 (k = 32) [17], DPN-131 [2], GoogleNet [31], IGC-L100M2 [35], IGC-L16M16 [35], IGC-L100M2 [35], Inception-ResNetv2 [30], Inceptionv2 [32], Inceptionv3 [32], Inceptionv4 [30], MobileNetv1 (1.0-224) [14], MobileNetv1 (1.0-192) [14], MobileNetv1 (1.0-160) [14], MobileNetv1 (1.0-128) [14], MobileNetv1 (0.75-224) [14], MobileNetv2 [26], MobileNetv2 (1.4) [26], NASNet-A (4 @ 1056) [38], NASNet-A (6 @ 4132) [38], NASNet-B (4 @ 1536) [38], NiN [20], OverFeat [27], PNASNet-5 (4, 216) [21], PolyNet [37], PreResNet-152 [13], PreResNet-200 [13], PyramidNet-101 (alpha = 250) [9], PyramidNet-200

(alpha = 300) [9], PyramidNet-200 (alpha = 450) [9], ResNet-152 [12], ResNet-50 [12], ResNet-101 [12], ResNeXt-101, SENet [15], ShuffleNet (1.5) [36], ShuffleNet (x2) [36], SimpleNet [10], SqueezeNet [18], SqueezeNetv1.1 [18], SqueezeNext (1.0-23v5) [7], SqueezeNext (2.0-23) [7], SqueezeNext (2.0-23v5) [7], TinyDarkNet [24], VGG16 [29], Xception [3], ZynqNet [8].

In this study, the units used for  $p(\mathcal{N})$  and  $m(\mathcal{N})$  for two of the quantitative performance metrics (information density and the proposed NetScore metric) are in M-Params (millions of parameters) and G-MACs (billions of MAC operations), respectively, given that most modern deep convolutional neural networks are within those architectural and computational complexity ranges.

### 3.1 Top-1 Accuracy


The top-1 accuracies across 60 different deep convolutional neural networks for the ILSVRC 2012 dataset is shown in Fig. 1. It can be clearly observed that significant progress has been made in the design of deep convolutional neural networks for image classification over the past six years, with the difference between the deep convolutional neural network with the highest top-1 accuracy in this study (i.e., AmoebaNet-C (6, 228)) and that of AlexNet exceeding 25%. It is also interesting to see that more recent developments in efficient deep convolutional neural networks such as MobileNetv1, MobileNetv2, and ShuffleNet all have top-1 accuracies that exceed VGG-16, the third largest tested network evaluated in the study that was also the state-of-the-art just four years ago, thus further illustrating the improvements in network design over the past few years.

### 3.2 Information Density

The information densities across 60 different deep convolutional neural networks for the ILSVRC 2012 dataset is shown in Fig. 2. It can be clearly observed that the deep convolutional neural networks that were specifically designed for efficiency (e.g., MobileNetv1, MobileNetv2, ShuffleNet, SqueezeNet, Tiny DarkNet, and SqueezeNext) have significantly higher information densities compared to networks that were designed purely with accuracy as a metric. More specifically, the SqueezeNext (1.0-23v5), Tiny DarkNet, and the SqueezeNet family of networks had the highest information density by a wide margin compared to the other tested deep convolutional neural networks, which can be attributed to their significantly lower architectural complexity in terms of number of network parameters. Another notable observation from the results in Fig. 2 is that the dynamic range of the information density metric is quite large across the diverse set of 60 deep convolutional neural networks evaluated in this study.

### 3.3 NetScore

The NetScore across 60 different deep convolutional neural networks for the ILSVRC 2012 dataset is shown in Fig. 3. Similar to the trend observed in Fig. 2,



**Fig. 1.** Top-1 accuracy across 60 different deep convolutional neural networks for the ILSVRC 2012 dataset.

Information Density (higher is better)





Fig. 2. Information density across 60 different deep convolutional neural networks for the ILSVRC 2012 dataset. Units are in %/M-Params.





**Fig. 3.** NetScore across 60 different deep convolutional neural networks for the ILSVRC 2012 dataset.

it can be clearly observed that many of the deep convolutional neural networks that were specifically designed for efficiency have significantly higher NetScores compared to networks that were designed purely with accuracy as a metric. However, what is interesting to observe is that the NetScore ranking amongst these efficient networks are quite different than that when using the information density metric. In particular, the top ranking deep convolutional neural networks with the highest NetScores are SqueezeNext (1.0-23v5), CondenseNet ( $G = C = 8$ ), and MobileNetv2.

The SqueezeNet family of networks, on the other hand, had much lower relative NetScores compared to the aforementioned efficient networks despite having the top two highest information densities. This observation illustrates the effect of incorporating computational complexity to the assessment of deep convolutional neural networks for practical usage, given that while the SqueezeNet family of networks has significantly lower architectural complexities compared to other tested networks, it also is offset by noticeably higher computational complexities compared to other tested efficient networks such as the MobileNetv1, MobileNetv2, SqueezeNext, and ShuffleNet network families.

The proposed NetScore metric, which by no means is perfect, could potentially be useful for guiding practitioners in model search and design and hopefully push the conversation towards better universal metrics for evaluating deep neural networks for use in practical scenarios. Future work includes incorporating additional or alternative factors that are important to assessing architectural and computational complexities of deep neural networks beyond what is being used in the NetScore metric, as well as finding a good balance between these different factors based on relative importance for the deployment of deep neural networks for practical usage in scenarios such as mobile devices and other edge devices.

## References

1. Canziani, A., Paszke, A., Culurciello, E.: An analysis of deep neural network models for practical applications. arXiv preprint [arXiv:1605.07678](https://arxiv.org/abs/1605.07678) (2017)
2. Chen, Y., Li, J., Xiao, H., Jin, X., Yan, S., Feng, J.: Dual path networks. CoRR abs/1707.01629 (2017). <http://arxiv.org/abs/1707.01629>
3. Chollet, F.: Xception: deep learning with depthwise separable convolutions. CoRR abs/1610.02357 (2016). <http://arxiv.org/abs/1610.02357>
4. Cun, Y.L., Bengio, Y., Hinton, G.: Deep learning. *Nature* **521**, 436–444 (2015)
5. Cun, Y.L., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proc. IEEE* **86**, 2278–2324 (1998)
6. Cun, Y.L., Denker, J., Henderson, D., Howard, R., Hubbard, W., Jackel, L.: Handwritten digit recognition with a back-propagation network. In: *Proceedings of the Advances in Neural Information Processing Systems (NIPS)* (1989)
7. Gholami, A., et al.: Squeezenext: hardware-aware neural network design. CoRR abs/1803.10615 (2018). <http://arxiv.org/abs/1803.10615>
8. Gschwend, D.: ZynqNet: an FPGA-accelerated embedded convolutional neural network (2016). <https://github.com/dgschwend/zynqnet>

9. Han, D., Kim, J., Kim, J.: Deep pyramidal residual networks. CoRR abs/1610.02915 (2016). <http://arxiv.org/abs/1610.02915>
10. HasanPour, S.H., Rouhani, M., Fayyaz, M., Sabokrou, M.: Lets keep it simple, using simple architectures to outperform deeper and more complex architectures. CoRR abs/1608.06037 (2016). <http://arxiv.org/abs/1608.06037>
11. He, K., Gkioxari, G., Dollar, P., Girshick, R.: Mask R-CNN. In: ICCV (2017)
12. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. CoRR abs/1512.03385 (2015). <http://arxiv.org/abs/1512.03385>
13. He, K., Zhang, X., Ren, S., Sun, J.: Identity mappings in deep residual networks. CoRR abs/1603.05027 (2016). <http://arxiv.org/abs/1603.05027>
14. Howard, A., et al.: Mobilenets: efficient convolutional neural networks for mobile vision applications. arXiv preprint [arXiv:1704.04861](https://arxiv.org/abs/1704.04861) (2017)
15. Hu, J., Shen, L., Sun, G.: Squeeze-and-excitation networks. CoRR abs/1709.01507 (2017). <http://arxiv.org/abs/1709.01507>
16. Huang, G., Liu, S., van der Maaten, L., Weinberger, K.Q.: Condensenet: an efficient densenet using learned group convolutions. CoRR abs/1711.09224 (2017). <http://arxiv.org/abs/1711.09224>
17. Huang, G., Liu, Z., Weinberger, K.Q.: Densely connected convolutional networks. CoRR abs/1608.06993 (2016). <http://arxiv.org/abs/1608.06993>
18. Iandola, F., Han, S., Moskewicz, M., Ashraf, K., Dally, W., Keutzer, K.: Squeezenet: Alexnet-level accuracy with 50x fewer parameters and <0.5 MB model size. arXiv preprint [arXiv:1602.07360](https://arxiv.org/abs/1602.07360) (2016)
19. Krizhevsky, A., Sutskever, I., Hinton, G.: Imagenet classification with deep convolutional neural networks. In: NIPS (2012)
20. Lin, M., Chen, Q., Yan, S.: Network in network. CoRR abs/1312.4400 (2013). <http://arxiv.org/abs/1312.4400>
21. Liu, C., et al.: Progressive neural architecture search. CoRR abs/1712.00559 (2017). <http://arxiv.org/abs/1712.00559>
22. Liu, W., et al.: SSD: single shot multibox detector. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9905, pp. 21–37. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-46448-0\\_2](https://doi.org/10.1007/978-3-319-46448-0_2)
23. Real, E., Aggarwal, A., Huang, Y., Le, Q.V.: Regularized evolution for image classifier architecture search. CoRR abs/1802.01548 (2018). <http://arxiv.org/abs/1802.01548>
24. Redmon, J.: Tiny darknet (2016). <https://pjreddie.com/darknet/tiny-darknet/>
25. Russakovsky, O., et al.: Imagenet large scale visual recognition challenge. Int. J. Comput. Vis. **115**(3), 211–252 (2015)
26. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L.: Mobilenetv 2: inverted residuals and linear bottlenecks. arXiv preprint [arXiv:1704.04861](https://arxiv.org/abs/1704.04861) (2017)
27. Sermanet, P., Eigen, D., Zhang, X., Mathieu, M., Fergus, R., LeCun, Y.: Overfeat: integrated recognition, localization and detection using convolutional networks. CoRR abs/1312.6229 (2013). <http://arxiv.org/abs/1312.6229>
28. Shafiee, M., Li, F., Chwyl, B., Wong, A.: Squishednets: squishing squeezenet further for edge device scenarios via deep evolutionary synthesis. In: NIPS (2017)
29. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. CoRR abs/1409.1556 (2014). <http://arxiv.org/abs/1409.1556>
30. Szegedy, C., Ioffe, S., Vanhoucke, V.: Inception-v4, inception-resnet and the impact of residual connections on learning. CoRR abs/1602.07261 (2016). <http://arxiv.org/abs/1602.07261>
31. Szegedy, C., et al.: Going deeper with convolutions. CoRR abs/1409.4842 (2014). <http://arxiv.org/abs/1409.4842>

32. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z.: Rethinking the inception architecture for computer vision. CoRR abs/1512.00567 (2015). <http://arxiv.org/abs/1512.00567>
33. Wong, A., Shafiee, M.J., Jules, M.S.: muNet: a highly compact deep convolutional neural network architecture for real-time embedded traffic sign classification. CoRR abs/1804.00497 (2018). <http://arxiv.org/abs/1804.00497>
34. Wong, A., Shafiee, M.J., Li, F., Chwyl, B.: Tiny SSD: a tiny single-shot detection deep convolutional neural network for real-time embedded object detection. CoRR abs/1802.06488 (2018). <http://arxiv.org/abs/1802.06488>
35. Zhang, T., Qi, G., Xiao, B., Wang, J.: Interleaved group convolutions for deep neural networks. CoRR abs/1707.02725 (2017). <http://arxiv.org/abs/1707.02725>
36. Zhang, X., Zhou, X., Lin, M., Sun, J.: Shufflenet: an extremely efficient convolutional neural network for mobile devices. CoRR abs/1707.01083 (2017). <http://arxiv.org/abs/1707.01083>
37. Zhang, X., Li, Z., Loy, C.C., Lin, D.: Polynet: a pursuit of structural diversity in very deep networks. CoRR abs/1611.05725 (2016). <http://arxiv.org/abs/1611.05725>
38. Zoph, B., Vasudevan, V., Shlens, J., Le, Q.V.: Learning transferable architectures for scalable image recognition. CoRR abs/1707.07012 (2017). <http://arxiv.org/abs/1707.07012>