

# From Xception to NEXception: New Design Decisions and Neural Architecture Search

Hadar Shavit<sup>a</sup>, Filip Jatełnicki<sup>b</sup>, Pol Mor-Puigventós<sup>c</sup> and Wojtek Kowalczyk<sup>d</sup>

Leiden Institute of Advanced Computer Science (LIACS), Leiden University, Niels Bohrweg 1, 2333CA, The Netherlands  
{h.shavit, f.o.jatelnicki, p.mor.puigventos}@umail.leidenuniv.nl, w.j.kowalczyk@liacs.leidenuniv.nl

Keywords: Deep Learning, ConvNeXt, Xception, Image Classification, ImageNet, Computer Vision.

Abstract: In this paper, we present a modified Xception architecture, the NEXception network. Our network has significantly better performance than the original Xception, achieving top-1 accuracy of 81.5% on the ImageNet validation dataset (an improvement of 2.5%) as well as a 28% higher throughput. Another variant of our model, NEXception-TP, reaches 81.8% top-1 accuracy, similar to ConvNeXt (82.1%), while having a 27% higher throughput. Our model is the result of applying improved training procedures and new design decisions combined with an application of Neural Architecture Search (NAS) on a smaller dataset. These findings call for revisiting older architectures and reassessing their potential when combined with the latest enhancements. Our code is available at <https://github.com/hadarshavit/NEXception>.

## 1 INTRODUCTION

There are multiple deep-learning-based approaches to tackle the image classification problem. In the last decade, the main attention was put into Transformers and convolutional-based architectures. Most of the recent findings in the convolutional neural networks field focused on improving the performance of the ResNet architecture (Liu et al., 2022; Wightman et al., 2021). In this paper we investigate how similar modifications can affect other convolutional architectures, specifically, the Xception model. In the following sections, we present NEXception, several Xception-based models that reach state-of-the-art level accuracies. Our models are the result of running Neural Architecture Search (NAS) experiments on the CIFAR-100 dataset (Krizhevsky et al., 2009) and an improved training procedure based, among other techniques, on new optimization methods and data augmentation. Our search space for experiments consists of variants of network architectures with different sizes of convolutional layers, activation functions, modern normalization and pooling methods, together with other recently introduced designs. As a final result of our experiments, we create three variants of the NEXception network. All of them outperform

the Xception model in the image classification task in terms of accuracy and inference throughput. Comparing our NEXception-TP model to the recently published ConvNeXt-T (Liu et al., 2022), our network reaches higher throughput (1428±9 vs. 1125±5 images/second) while having a similar accuracy.

## 2 RELATED WORK

### 2.1 Recent Research

During the last decade, numerous architectures have been proposed for computer vision with convolutional neural networks as their central building block. AlexNet (Krizhevsky et al., 2017) trained on ImageNet with striking results, compared to the then state-of-the-art models, achieving top-1 and top-5 test set error rates of 37.5% and 17.0% respectively. This was the moment when state-of-the-art solutions progressed from pattern recognition to deep learning. Thereafter, remarkable progress has been made almost yearly, starting with GoogLeNet (Szegedy et al., 2014) with a more efficient (due to the *Inception* module) and deeper architecture. Later on, ResNet (He et al., 2015) was introduced, with residual (*skip*) connections that allowed for even deeper networks. In 2016, both of those last contributions were merged generating *Extreme Inception*, shown in Xception (Chollet, 2017). This time, the *Inception* module

<sup>a</sup> <https://orcid.org/0000-0001-6709-9955>

<sup>b</sup> <https://orcid.org/0000-0002-1717-167X>

<sup>c</sup> <https://orcid.org/0000-0002-4843-3732>

<sup>d</sup> <https://orcid.org/0000-0002-6973-1341>

was replaced by the Xception module which used a *depthwise separable convolution layer* as its basic building block. In the following year, SENet (Hu et al., 2018) boosted previous networks, proposing *SE-Inception* and *SE-ResNet*. SENet proposed the “Squeeze-and-Excitation” (SE) block focusing on the depth dimension, recalibrating channel-wise feature responses. EfficientNet (Tan and Le, 2019) and EfficientNetV2 (Tan and Le, 2021) suggested a proper scaling of existing architectures achieving superior performances.

In recent years, the machine vision community adopted the Transformers architecture, originally developed for Natural Language Processing (NLP) (Devlin et al., 2019) by introducing the Vision Transformer (ViT) (Dosovitskiy et al., 2021). Since this installation of ViT, various improvements have been introduced including the Data-efficient image Transformers (DeiT) (Touvron et al., 2021), Swin Transformer (Liu et al., 2021b) and the recent Neighbourhood Attention Transformer (Hassani et al., 2022).

In addition to the improvements in the macro-level architectures, other micro-level improvements were introduced. While ReLU was widely employed a few years ago, newer activation functions were published, for instance, the Gaussian Error Linear Unit (GELU) (Hendrycks and Gimpel, 2016).

Moreover, new training procedures were exploited. While in the past basic Stochastic Gradient Descent (SGD) was used to train state-of-the-art models, today there are new variants of gradient-based optimizers such as RAdam (Liu et al., 2020) AdamP (Heo et al., 2021) and LAMB (You et al., 2020) with sophisticated learning rate schedules such as cosine decay (Loshchilov and Hutter, 2017). Furthermore, data augmentation techniques such as RandAugment (Cubuk et al., 2020), RandErase (Zhong et al., 2020), Mixup (Zhang et al., 2018) and CutMix (Yun et al., 2019) greatly improved the accuracy of neural networks.

## 2.2 Xception

The Xception neural network was introduced by Chollet (2017). This architecture implements the *depthwise separable convolution* operation. These convolutions consist of two parts: a depthwise convolution followed by a pointwise convolution. We refer to them as separable convolutional layers. The three parts of the Xception architecture are:

**Entry Flow.** First, a stem of two convolutional layers of increasing sizes, followed by the first layers of the model, and then, three downsampling blocks. Each of these blocks has two separable convolutional layers

with a kernel size of 3 combined with a *Max Pooling* layer. Each block has a skip connection with a  $1 \times 1$  convolution with stride 2.

**Middle Flow.** The central unit contains 8 Xception blocks. Each block has three separable convolutional layers with a kernel size of 3 and stride 1. The method applied does not reshape the input size. For this reason, the size of the features map remains  $19 \times 19 \times 728$  through this part of the network. In addition, there is a residual identity connection around every block.

**Exit Flow.** The closure section starts with one downsampling block, like the ones in the entry flow, followed by two separable convolutions. Lastly, there is a classification head with a global average pooling and fully connected layer(s).

## 2.3 Neural Architecture Search

Neural Architecture Search (NAS) is a collection of methods for automating the design of neural network architectures (Elsken et al., 2019). This can be done using automated search in a pre-defined configuration space using automated algorithm configuration methods such as Bayesian Optimization (Hutter et al., 2011; van Stein et al., 2019; Jin et al., 2019) or Evolutionary Algorithms (Liu et al., 2021a). The usage of NAS methods has grown significantly, which can be observed in recent works like EfficientNetV2 (Tan and Le, 2021), which used NAS to improve EfficientNet.

## 3 NEXception

In this section, we present and explain our reasoning behind the chosen techniques for our search space, inspired by many recent design decisions, including ConvNeXt by Liu et al. (2022) and the re-study of ResNet by Wightman et al. (2021), and extending those ideas with other innovations.

The search space is built with the PyTorch library (Paszke et al., 2019) and *timm* (Wightman, 2019). We apply SMAC (Lindauer et al., 2022) automated algorithm configurator to find a good configuration of improvements. Due to the considerable training time of a full network on ImageNet, we test the configurations with a reduced network, with four blocks in its main part instead of eight, and on a smaller dataset (CIFAR-100 (Krizhevsky et al., 2009))<sup>1</sup>, with only one downsampling block in the entry flow. This allows the trials of as many configurations as possible

<sup>1</sup><https://www.cs.toronto.edu/~kriz/cifar.html>

within 3 days on a single RTX 3090 Ti. The search space containing the multiple combinations of parameters has more than fifty thousand different possible configurations. We optimize the architecture to maximize accuracy.

Our search space consists of various kernel sizes (3, 5, 7, 9), stem types (convolutional stem or patchify stem), different pooling types (max pooling, convolutional downsampling layer or blur pool), whether to implement bottleneck in the middle flow, or to add Squeeze-and-Excitation at the end of each block. Moreover, we experiment with various positions and types of activation functions (ReLU, GELU, ELU or CELU) and with different positions and types of normalization methods (batch normalization or layer normalization).

We performed several preliminary experiments to find the optimal training procedure. However, we find that existing ones perform better on the final models. Therefore, we use similar training procedures to the ones created by the authors of Wightman et al. (2021) for the ResNet network. For more details about training procedure parameters, see Table 2 and the comparison between Xception and NEXcepTion architectures in Tables 5 and 6, all in the Appendix section.

### 3.1 Training Procedures

**Stochastic Depth.** The original Xception network performs regularization by adding a Dropout layer before the classification layer. Stochastic depth (Huang et al., 2016) changes the network depth during training by randomly bypassing groups of blocks and using the entire network for inference. Consequently, training time is reduced substantially and accuracy is improved introducing regularization into the network.

**Optimizer.** In our paper, we choose *Layer-wise Adaptive Moments optimizer for Batch training* (LAMB) optimizer inspired by You et al. (2020). As stated by Wightman et al. (2021), LAMB optimizer increases the efficiency and performance of the network, in comparison to other common optimizers, like AdamW in Liu et al. (2022), LAMB performs more accurate updates of the learning rate.

**Data Augmentation.** While the original Xception model was trained without any data augmentation methods, newer training procedures utilize multiple techniques, which improve generalization. In our NEXcepTion model, we apply Rand Augment (Cubuk et al., 2020) that performs a few random transformations, Mixup (Zhang et al., 2018) and CutMix (Yun et al., 2019), which merge images, see Table 2 for specific values.

**Learning Rate Decay.** Similarly to recent models such as DeiT (Touvron et al., 2021), we adopt cosine annealing (Loshchilov and Hutter, 2017) with warmup epochs. This method initially sets a low learning rate value, which gradually increases during the warmup epochs. Then, the learning rate is gradually reduced using the cosine function to achieve rapid learning.

### 3.2 Structural Changes

**“Soft” Patchify Stem.** Patchify layers are characterized by large kernel sizes and non-overlapping convolutions (by setting the stride and the kernel size to the same value). Inspired by this design, we add a  $2 \times 2$  patchify layer to the search space, which we consider a “soft” patch, different from the aggressive  $16 \times 16$  solution proposed by Dosovitskiy et al. (2021) in the Transformer schema and the  $4 \times 4$  from ConvNeXt (Liu et al., 2022). We use the initial block with kernel  $2 \times 2$  and stride 2 to match the original Xception network and to fit the output size. This stem is adapted to the reduced resolution of the input images, similarly to the efficient configuration introduced by Cordonnier et al. (2020).

**Bottleneck.** The idea of inverted bottleneck was introduced by Sandler et al. (2018) and has been prevalent in modern attention-based architectures, significantly improving performance. The Xception architecture does not feature a bottleneck and has a constant number of channels through the middle flow of the network. In the NEXcepTion architecture, we introduce a bottleneck in the middle flow blocks, as proposed by Liu et al. (2022), see Figure 1.

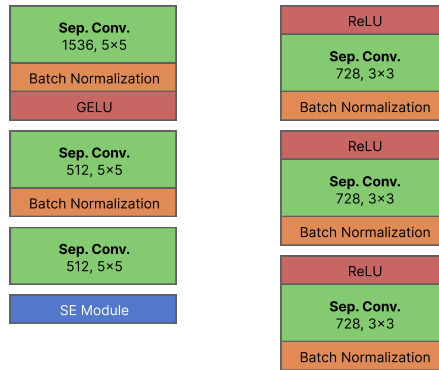


Figure 1: NEXcepTion block (left) and Xception block (right).

**Larger Kernels.** Inspired by Liu et al. (2022), among others, we pick larger kernels for our experiments, and we achieve the best accuracy with their size set to 5. Combining this idea with bottleneck blocks and the reduced resolution allows using bigger kernels with-

out excessive increase in the computational demand.

**Squeeze-and-Excitation Block.** Squeeze-and-Excitation block (SE block) from Hu et al. (2018) improves channel interdependencies with an insignificant decrease in efficiency by recalibrating the feature responses channel-wise, creating superior feature maps. SE blocks provide significant performance improvement and are easy to include in existing networks as specified by Hu et al. (2018).

**Fewer Activations and Normalizations.** Similarly to Liu et al. (2022), we employ less activation layers than in the original Xception network. Fewer activation layers is a distinctive property of the state-of-the-art Transformer blocks and, by replicating this concept, we can achieve higher accuracy.

Moreover, what is also inherent to Transformers architectures, is fewer normalization layers than in typical convolution-based solutions. It is important to mention that in the original Xception architecture, all convolutional layers are followed by batch normalization.

**Activation Function.** Concerning neuron activations, GELU (Hendrycks and Gimpel, 2016) is used in modern Transformer architectures like BERT (Devlin et al., 2019) and recent convolutional-based architectures like ConvNeXt (Liu et al., 2022). Despite ReLU’s simplicity and efficiency, we decide to experiment with different activation functions, inspired by the survey performed by Dubey et al. (2021). Based on our search, we achieve the best results with the GELU activation function.

**Standardizing the Input.** The original Xception model uses an input size of  $299 \times 299$ . We found that standardizing the input size to  $224 \times 224$ , as in He et al. (2015), makes the training faster on Nvidia Tensor Cores. To compensate for the lower resolution, we make the network wider by adding more channels.

**Blur Pooling.** Inspired by the solution from Zhang (2019), we integrate a blurring procedure before subsampling the signal. By introducing this anti-aliasing technique, our network generalizes better and achieves higher accuracy.

## 4 NEXception VARIANTS

As a result of our experiments, we produce a configuration of a downsized network.

We prepare two different NEXception variants, adapted to the complexities of the main “Tiny” and “Small” recent state-of-the-art models. This allows us to compare them to recent models with similar features. Additionally, we construct NEXception-TP with pyramid-like architecture. All the variants use

the methods described in the previous section and the NEXception block presented on Figure 1.

**NEXception-T.** This model exploits all the methods described in the previous section, see Table 5 in the Appendix. It has 24.5M parameters and 4.7 GFLOPs. The motivation for it is to have similar FLOPs and a number of parameters to the recent state-of-the-art models such as ConvNeXt (Liu et al., 2022) and Swin Transformers’ (Liu et al., 2021b) tiny models, for instance in ConvNeXt-T and Swin-T.

**NEXception-S.** This architecture is a wider variant with 8.5 GFLOPs and 43.4M parameters. The motivation for it is to have a model with similar FLOPs to the original Xception network (Chollet, 2017).

**NEXception-TP.** While the Xception architecture and the two other variants have an isotropic architecture with a constant resolution through the middle flow, other architectures such as ResNet and ConvNeXt have a pyramid-like architecture. Such architecture incorporates a few stages in its middle flow and the resolution decreases from stage to stage, hence, the name “Pyramid”. We use the ConvNeXt architecture and replace the ConvNeXt blocks with NEXception blocks, as well as substituting Layer Normalization with Batch Normalization and adding one more block in the second phase to have a comparable number of FLOPs. This Pyramid NEXception model is trained with the same training procedure as NEXception. Our motivation is to check how the NEXception blocks’ performance changes with the pyramid architecture, as the pyramid ConvNeXt has significantly higher accuracy than the isotropic ConvNeXt (82.1% vs. 79.7%). This variant has 4.5GFLOPs and 26.6M parameters.

## 5 RESULTS

### 5.1 DeepCAVE Analysis

We first present a hyperparameter importance analysis of our NAS process, which can be seen in Figure 4. We also measured the importance of the stem shape, the pooling procedure and the SE module, obtaining an importance of less than 0.1 for those features. We calculate the Local Hyperparameter Importance (LPI) using DeepCave (Sass et al., 2022). We can see that the most important hyperparameter is the block type, for instance, most of the improvement comes from shifting to a bottleneck block. Changing the positions of the normalizations and the kernel sizes of the convolutions also has an impact on the performance. We can see that the activation function type has a relatively small impact on the accuracy of the model.




Figure 2: FLOPs and Accuracy comparison of the NEXception variants (in bold), with other contemporary convolutional or Transformer networks with similar features. The size of the bubbles corresponds to the number of parameters. More details can be found in Table 1, in the Appendix.

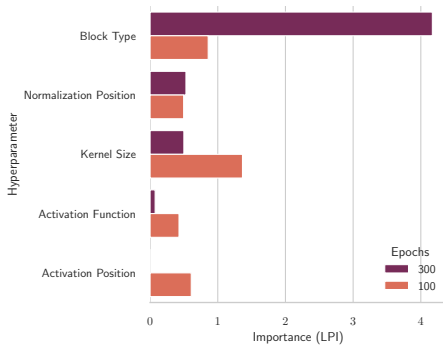


Figure 4: Local Hyperparameter Importance (LPI). Depending on the number of epochs, the influence of the selected methods differs for the final result.

## 5.2 Image Classification

We train our networks on the widely used ImageNet-1K image classification benchmark<sup>2</sup> (Russakovsky et al., 2015). We run our experiments on a single node of a local cluster. A node has 4 Nvidia RTX 2080 Ti GPUs and 2 Intel Xeon Gold 6126 2.6GHz with 12 cores CPUs. We run 3 repetitions for each of the NEXception variants. Each training of the NEXception-T variant takes 100 hours on average and for the NEXception-TP 89 hours on average. Finally, the biggest model we test, the NEXception-S, takes 150 hours on average. The top-1 accuracy of the NEXception-T model is  $81.6 \pm 0.08$ , for NEXception-TP is  $81.7 \pm 0.07$ , for NEXception-S is  $82.0 \pm 0.07$ .

<sup>2</sup><https://image-net.org/>

We compare our three variants with the original Xception network (Chollet, 2017), with the convolutional neural networks ConvNeXt (Liu et al., 2022), EfficientNetV1 (Tan and Le, 2019) and EfficientNetV2 (Tan and Le, 2021). We also compare our networks with the Transformer-based models Vision Transformer (ViT) (Dosovitskiy et al., 2021), Data Efficient Transformer (DeiT) (Touvron et al., 2021) and Swin Transformer (Liu et al., 2021b). The reported accuracies of these models belong to the papers cited next to their names. We calculate the throughputs using the *timm* library (Wightman, 2019), on a single RTX 2080 Ti with a batch size of 256, mixed precision and channels last, on the ImageNet validation dataset during 30 repetitions. For the EffNet-B4 and EffNet-B5 the calculations are made using a batch size of 128, due to GPU memory issues.

We evaluate the models using *timm*. For the NEXception set models, we use our own trained weights. For isotropic ConvNeXt the trained weights are from Liu et al. (2022). For Neighbourhood Attention Transformer (NAT), the trained weights are from Hassani et al. (2022). The results are presented in Figures 2 and 3. Additionally, in the Appendix, we offer the results values in Table 1 and we evaluate the robustness of the NEXception architectures in Table 3.

All variants of NEXception have higher accuracy than Xception, as well as higher throughput. The NEXception-TP model has significantly higher throughput than the other compared models with similar accuracy.




Figure 3: Throughput and Accuracy comparison of the NEXcepTion variants (in bold) with other contemporary convolutional or Transformer networks with similar features. The size of the bubbles corresponds to the number of parameters exploited. More details can be found in Table 1, in the Appendix.

## 6 CONCLUSIONS

In this work, we implement the backbone of the existing Xception architecture, with some modifications and improved training. We show that on the ImageNet classification task it is possible to achieve significantly higher accuracy than with the original architecture. Our findings strengthen the work published recently by Liu et al. (2022), in which ConvNeXt was presented. While ConvNeXt only showed results on modernizing ResNet, we generalize their findings to another convolutional architecture.

We also present a NAS method that combines both the usage of applying modern design decisions to existing architectures and automated algorithm configuration for neural architecture search. This method can be used to apply those modifications to other architectures. Nevertheless, it could be useful to generalize the usage of NAS to enhance existing architectures with other existing networks to confirm this idea.

When it comes to the obtained results, we provide three variants of the NEXcepTion network, and all reach higher accuracy and throughput than Xception. Our NEXcepTion-T outperforms the original Xception, using half of the FLOPs and a similar number of parameters.

In comparison to ConvNeXt, NEXcepTion-TP reaches similar accuracy with higher throughput, as reported in Section 5. We note that ConvNeXt’s pyramid compute ratio gives better results both in terms of accuracy and inference throughput, as using the NEXcepTion block with this compute ratio performs

better than using Xception’s compute ratio. In addition, we can see that the NEXcepTion block has less impact from the compute ratio than the ConvNeXt block as the difference between the isotropic NEXcepTion-T and the pyramid NEXcepTion-TP is only 0.2 while the difference between ConvNeXt-T and isotropic ConvNeXt-S is 2.4 (Liu et al., 2022).

Finally, we also check the generalization of our models testing their performance on robustness datasets and comparing them to other state-of-the-art models. For all datasets, the NEXcepTion set obtains better results than Xception and is frequently above the rest of the architectures, see Table 3 in the Appendix section.

Overall, this work can inspire future research to use algorithm configuration libraries like SMAC (Lindauer et al., 2022), RayTune (Liaw et al., 2018) or KerasTuner (O’Malley et al., 2019) as they require only the definition of a base model and a configuration space.

Another future research direction can be to perform an in-depth importance analysis of architectural designs such as the ones we use, similarly to checking the hyperparameter’s importance as in van Rijn and Hutter (2018) for traditional machine learning models.

## ACKNOWLEDGEMENTS

This work was performed using the compute resources from the Academic Leiden Interdisciplinary Cluster Environment (ALICE) provided by Leiden University. We thank Andrius Bernatavicius, Shima Javanmardi and the participants of the Advances in Deep Learning 2022 class in LIACS for the valuable discussions and feedback.

## REFERENCES

- Chollet, F. (2017). Xception: Deep learning with depthwise separable convolutions. In *CVPR*, pages 1251–1258.
- Cordonnier, J., Loukas, A., and Jaggi, M. (2020). On the relationship between self-attention and convolutional layers. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Cubuk, E. D., Zoph, B., Shlens, J., and Le, Q. (2020). Randaugment: Practical automated data augmentation with a reduced search space. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H., editors, *Advances in Neural Information Processing Systems*, volume 33, pages 18613–18624. Curran Associates, Inc.
- Devlin, J., Chang, M., Lee, K., and Toutanova, K. (2019). BERT: pre-training of deep bidirectional transformers for language understanding. In Burstein, J., Doran, C., and Solorio, T., editors, *NAACL-HLT*, pages 4171–4186. Association for Computational Linguistics.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissensborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., and Houlsby, N. (2021). An image is worth 16x16 words: Transformers for image recognition at scale. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Dubey, S. R., Singh, S. K., and Chaudhuri, B. B. (2021). A comprehensive survey and performance analysis of activation functions in deep learning. *arXiv preprint arXiv:2109.14545*.
- Elsken, T., Metzen, J. H., and Hutter, F. (2019). *Neural Architecture Search*, pages 63–77. Springer International Publishing, Cham.
- Hassani, A., Walton, S., Li, J., Li, S., and Shi, H. (2022). Neighborhood attention transformer. *arXiv preprint arXiv:2204.07143*.
- He, K., Zhang, X., Ren, S., and Sun, J. (2015). Deep residual learning for image recognition. *CVPR*, pages 770–778.
- Hendrycks, D., Basart, S., Mu, N., Kadavath, S., Wang, F., Dorundo, E., Desai, R., Zhu, T., Parajuli, S., Guo, M., et al. (2021a). The many faces of robustness: A critical analysis of out-of-distribution generalization. In *ICCV*, pages 8340–8349.
- Hendrycks, D. and Gimpel, K. (2016). Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*.
- Hendrycks, D., Zhao, K., Basart, S., Steinhardt, J., and Song, D. (2021b). Natural adversarial examples. In *CVPR*, pages 15262–15271.
- Heo, B., Chun, S., Oh, S. J., Han, D., Yun, S., Kim, G., Uh, Y., and Ha, J. (2021). Adamp: Slowing down the slowdown for momentum optimizers on scale-invariant weights. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Hu, J., Shen, L., and Sun, G. (2018). Squeeze-and-excitation networks. In *CVPR*, pages 7132–7141.
- Huang, G., Sun, Y., Liu, Z., Sedra, D., and Weinberger, K. Q. (2016). Deep networks with stochastic depth. In *European conference on computer vision*, pages 646–661. Springer.
- Hutter, F., Hoos, H. H., and Leyton-Brown, K. (2011). Sequential model-based optimization for general algorithm configuration. In *Proc. of LION-5*, page 507–523.
- Jin, H., Song, Q., and Hu, X. (2019). Auto-keras: An efficient neural architecture search system. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1946–1956. ACM.
- Krizhevsky, A., Hinton, G., et al. (2009). Learning multiple layers of features from tiny images. Technical report, University of Toronto.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2017). ImageNet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90.
- Liaw, R., Liang, E., Nishihara, R., Moritz, P., Gonzalez, J. E., and Stoica, I. (2018). Tune: A research platform for distributed model selection and training. *arXiv preprint arXiv:1807.05118*.
- Lindauer, M., Eggensperger, K., Feurer, M., Biedenkapp, A., Deng, D., Benjamins, C., Ruhkopf, T., Sass, R., and Hutter, F. (2022). Smac3: A versatile bayesian optimization package for hyperparameter optimization. *Journal of Machine Learning Research*, 23(54):1–9.
- Liu, L., Jiang, H., He, P., Chen, W., Liu, X., Gao, J., and Han, J. (2020). On the variance of the adaptive learning rate and beyond. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Liu, Y., Sun, Y., Xue, B., Zhang, M., Yen, G. G., and Tan, K. C. (2021a). A survey on evolutionary neural architecture search. *IEEE transactions on neural networks and learning systems*.
- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., and Guo, B. (2021b). Swin transformer: Hierarchical vision transformer using shifted windows. *ICCV*, pages 9992–10002.
- Liu, Z., Mao, H., Wu, C., Feichtenhofer, C., Darrell, T., and Xie, S. (2022). A convnet for the 2020s. *CVPR*, pages 11966–11976.
- Loshchilov, I. and Hutter, F. (2017). SGDR: stochastic gradient descent with warm restarts. In *5th International Conference on Learning Representations, ICLR 2017*,

- Toulon, France, April 24-26, 2017, Conference Track Proceedings. OpenReview.net.
- O'Malley, T., Bursztein, E., Long, J., Chollet, F., Jin, H., Invernizzi, L., et al. (2019). Kerastuner. <https://github.com/keras-team/keras-tuner>.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. (2019). Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A. C., and Fei-Fei, L. (2015). ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252.
- Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., and Chen, L.-C. (2018). Mobilenetv2: Inverted residuals and linear bottlenecks. In *CVPR*, pages 4510–4520.
- Sass, R., Bergman, E., Biedenkapp, A., Hutter, F., and Lindauer, M. (2022). Deepcave: An interactive analysis tool for automated machine learning. *arXiv preprint arXiv:2206.03493*.
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., and Rabinovich, A. (2014). Going deeper with convolutions.
- Tan, M. and Le, Q. (2021). Efficientnetv2: Smaller models and faster training. In *International Conference on Machine Learning*, pages 10096–10106. PMLR.
- Tan, M. and Le, Q. V. (2019). Mixconv: Mixed depthwise convolutional kernels. In *30th British Machine Vision Conference 2019, BMVC 2019, Cardiff, UK, September 9-12, 2019*, page 74. BMVA Press.
- Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., and Jégou, H. (2021). Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning*, pages 10347–10357. PMLR.
- van Rijn, J. N. and Hutter, F. (2018). Hyperparameter importance across datasets. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery; Data Mining, KDD '18*, page 2367–2376, New York, NY, USA. Association for Computing Machinery.
- van Stein, B., Wang, H., and Bäck, T. (2019). Automatic configuration of deep neural networks with parallel efficient global optimization. In *2019 International Joint Conference on Neural Networks (IJCNN)*, pages 1–7.
- Wang, H., Ge, S., Lipton, Z., and Xing, E. P. (2019). Learning robust global representations by penalizing local predictive power. *Advances in Neural Information Processing Systems*, 32.
- Wightman, R. (2019). Pytorch image models. <https://github.com/rwightman/pytorch-image-models>.
- Wightman, R., Touvron, H., and Jegou, H. (2021). Resnet strikes back: An improved training procedure in timm. In *NeurIPS 2021 Workshop on ImageNet: Past, Present, and Future*.
- You, Y., Li, J., Reddi, S. J., Hseu, J., Kumar, S., Bhojanapalli, S., Song, X., Demmel, J., Keutzer, K., and Hsieh, C. (2020). Large batch optimization for deep learning: Training BERT in 76 minutes. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Yun, S., Han, D., Oh, S. J., Chun, S., Choe, J., and Yoo, Y. (2019). Cutmix: Regularization strategy to train strong classifiers with localizable features. In *ICCV*, pages 6023–6032.
- Zhang, H., Cissé, M., Dauphin, Y. N., and Lopez-Paz, D. (2018). mixup: Beyond empirical risk minimization. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Zhang, R. (2019). Making convolutional networks shift-invariant again. In *International conference on machine learning*, pages 7324–7334. PMLR.
- Zhong, Z., Zheng, L., Kang, G., Li, S., and Yang, Y. (2020). Random erasing data augmentation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 13001–13008.

## APPENDIX

The Appendix of the paper, including a detailed architectures description, detailed results, and a robustness analysis, is available at <https://github.com/hadarshavit/NEXception/blob/main/appendix.pdf>.