Over-Parameterization and Generalization in Audio Classification

Khaled Koutini¹ Hamid Eghbal-zadeh^{1 2} Florian Henkel¹ Jan Schlüter¹ Gerhard Widmer^{1 2}

Abstract

Convolutional Neural Networks (CNNs) have been dominating classification tasks in various domains, such as machine vision, machine listening, and natural language processing. In machine listening, while generally exhibiting very good generalization capabilities, CNNs are sensitive to the specific audio recording device used, which has been recognized as a substantial problem in the acoustic scene classification (DCASE) community. In this study, we investigate the relationship between over-parameterization of acoustic scene classification models, and their resulting generalization abilities. Specifically, we test scaling CNNs in width and depth, under different conditions. Our results indicate that increasing width improves generalization to unseen devices, even without an increase in the number of parameters.

1. Introduction

An important part of machine auditory perception is to assign tags to audio clips, based on the acoustic activities and events present in the clip. This task has various applications such as content-based multimedia information retrieval, context-aware smart devices, and monitoring and surveillance systems. Convolutional Neural Networks (CNNs) have dominated this area for tasks such as Acoustic Scene Classification (ASC) and tagging (Hershey et al., 2017). Building on successful architectures from the vision domain, the CNN architectures used in the audio domain are usually trained using image-like features such as spectrograms that are extracted from the given audio waveforms. Hence, many deep architectures from vision, such as ResNet (Liu et al., 2021), have been adapted to become state-of-the-art models in audio tasks, e.g., by regularizing their Receptive Field (RF) (Koutini et al., 2019), and Effective Receptive Field (ERF) (Luo et al., 2016; Koutini et al., 2021).

¹Institute of Computational Perception, Johannes Kepler University Linz, Austria ²LIT AI Lab. Correspondence to: Khaled Koutini <Khaled.koutini@jku.at>.

Figure 1. Test accuracy of a Damped ResNet with scaled up variants in width and depth with/without increasing the Receptive Field (RF), with respect to the number of parameters, as explained in Section 4.

Several studies have shown that scaling up the number of parameters of neural networks improves their performance in different domains, such as vision (Neyshabur et al., 2019; Tan & Le, 2019) and language (Kaplan et al., 2020). Tan & Le (2019) investigate the effect of increasing width and depth of CNNs in image classification, and further introduce a high-performance parameter-efficient architecture that exploits the balance between width and depth. Belkin et al. (2019) and Nakkiran et al. (2020) study the double descent phenomenon and show that when increasing some models' complexities pass the point they can perfectly fit the training data, they enter the over-parameterized regime, where test set performance improves and surpasses the best under-parameterized setting; even in the presence of limited label noise in training data. Golubeva et al. (2021) explore the relation between width and generalization. They discover that by increasing width, generalization can be improved even when the number of parameters is kept fixed by introducing sparsity. Although the aforementioned works have shown interesting connections between over-parameterization and

generalization, this connection has not yet been studied in the audio domain, and the challenges of designing effective over-parameterized models for acoustic tasks are still unknown. This becomes an important aspect since the ultimate goal in many applications is to build models that can generalize to new devices, while keeping the complexity low in order to make them suitable for mobile and edge devices.

In this work, we study the characteristics of over-parameterized CNNs in audio classification. We focus on an important generalization aspect: generalization to unseen devices. This is an important and known problem in the acoustic domain, and several international challenges in this area have been dedicated to it.¹ More specifically, we wish to answer the double question of what scaling approaches (if any) achieve better generalization to unseen devices, and whether the added parameters are the cause of generalization, or a side-effect of scaling. To that end, we perform a set of experiments scaling up two architectures that were previously shown to perform well on the ASC task, using four strategies: by 1) width; 2) depth; 3) depth without increasing the RF; and 4) width without increasing the number of parameters. The results show that the width of the network plays the biggest role in its generalization to unseen devices, even when the number of parameters are kept fixed. In contrast, increasing the depth can even hurt generalization. Finally, we will show how these results are connected to the ERF of the scaled up architectures.

2. Scaling up CNNs for Audio Classification

Tan & Le (2019) observed that, for image classification, scaling up the width and depth of the network, or the input resolution, improves a model’s accuracy. However, model depth and its relation to generalization, is a different story in deep models using spectrograms for audio tasks. Koutini et al. (2021) showed that scaling up the depth of a network without restricting its Receptive Field (RF), negatively impacts the generalization capabilities of different CNN architectures. Further, they proposed several techniques to regularize the RF of models in order to increase the depth, and improve generalization. Golubeva et al. (2021) argue that in over-parameterized regimes, widening the network while maintaining the number of parameters results in models with better generalization.

In Section 4 we first investigate the effect of scaling up the depth; both with, and without changing the maximum RF of the base network. Furthermore, we compare the performance of scaled-up models; when their width and depth have been increased.

¹<http://dcase.community/challenge2020/task-acoustic-scene-classification>

Figure 2. Spectrograms of the same audio recorded with two different devices.

3. Aspects of Generalization in Audio

One of the main problems in audio classification and tagging, is the inability of models to generalize to new devices. Using different recording devices at test time introduces a non-trivial domain shift that causes significant performance degradation in tasks such as ASC and tagging. Different recording devices may vary in terms of their level of amplitude, the dynamic range, as well as their response to certain frequencies (cf. Figure 2). The recording device introduces a nonlinear transformation on the input signal, and thereby to the spectrograms calculated from it. We believe that tasks such as ASC under device mismatch can be a practical test bed for evaluating the generalization capabilities of CNNs. Hence, in our empirical study, we focus on this task. We use the TAU Urban Acoustic Scenes 2020 Mobile (Heittola et al., 2020) dataset for our analysis, since it is a well-established benchmark for ASC, and additionally, it provides device annotations. The dataset comprises of 64 hours of audio recorded in 10 different environmental scenes – which are the classification targets – such as airport, shopping mall, bus etc. The audio is recorded in 12 European cities using 3 different microphones: Device A (“Soundman OKM II Klassik/studio A3”) 40 hours, Device B (Samsung Galaxy S7), Device C (iPhone SE) 3h each. Additionally, 6 different microphones are simulated, 3 hours each. We choose the training/testing setup to test (a) Generalization to unseen cities (b) Generalization to unseen devices. We ensure that each (city, class) pair appears only in either the train or the test set. We train using only recordings from Device A, and test both on recordings from Device A (referred to as “Seen”) and on recordings from the other devices (referred to as “Unseen”). In both cases, we only test on unseen cities for an acoustic scene.

4. Creating Over-Parameterized Models

In our experiments, we use a baseline architecture, namely, a variant of the ResNet (Liu et al., 2021) proposed in (Koutini et al., 2019), which is the state-of-the-art in low-complexity ASC under device mismatch, on the DCASE benchmark (2020). The configuration of the network is detailed in Table 1, where the width, i.e., number of channels, of each layer is controlled by W , and D , K control the depth of the network. Using this architecture, a baseline network

Table 1. Baseline architecture and scaling factors

REPEAT	CHANNELS	BLOCK	CONFIG
	W	INPUT	5×5 STRIDE=2
1	W	R	$3 \times 3, 1 \times 1, P$
1	W	R	$3 \times 3, 3 \times 3, P$
1	W	R	$3 \times 3, 3 \times 3$
1	W	R	$3 \times 3, 3 \times 3, P$
D	$2 \times W$	R	$3 \times 3, 3 \times 3$
K	$4 \times W$	R	$1 \times 1, 1 \times 1$
CLASSIFIER $4 \times W \rightarrow 10$ CLASSES			
GLOBAL MEAN POOLING			

P: 2×2 MAX POOLING.

R: RESIDUAL, THE INPUT IS ADDED TO THE OUTPUT

with $W = 32, D = 1, K = 1$ can already fit the training data with 99.99% accuracy. Furthermore, this baseline can fit the training data with random labels to 85.51% accuracy. The same network with $W = 64$ fits the training data with random labels to 100% accuracy, indicating model over-parameterization (Zhang et al., 2017). In addition, we compare the network in each setup to the *Frequency-Damped* version of our baseline architecture (Koutini et al., 2021). Frequency-damping reduces the Effective Receptive Field (ERF) (Luo et al., 2016) of the network over the frequency dimension, by down-scaling the convolutional weights away from the center of the RF along the frequency dimension, as explained in Appendix B. Frequency-Damped networks have an inductive bias towards limiting the RF, which is beneficial for audio tasks and have shown success in the acoustic domain (Koutini et al., 2021). We refer to this variant as *Damped*, and use it to further test the importance of task-specific bias/regularization under different over-parameterization settings.

We scale the baseline networks by increasing the width of the CNN, choosing W (Table 1) in $\{32, 64, 128, 192, 256, 384, 512\}$. Results are discussed in Section 5.1. For comparison, we scale the networks by increasing the depth, adding residual blocks with 3×3 convolutions (increasing D in Table 1). This will result in networks with more parameters, and bigger RFs over the input. However, since such increasing of depth has been shown to cause overfitting on this task (Koutini et al., 2021), we additionally investigate increasing the depth without increasing the RF of the networks, by adding residual blocks with 1×1 convolutions (increasing K in Table 1). The depth scaling results are presented in Section 5.2.

We also study whether the observed performance improvement is due to increased over-parameterization, or due to widening of the network. To that end, we use the same setup as Golubeva et al. (2021), and randomly mask out

Figure 3. The effect of the width on accuracy.

weights of the network at initialization, prioritizing those layers with more parameters, and increase sparsity of the weights. We refer to this approach as *static* random pruning. As a result, we maintain the same number of parameters as the network before widening. We use the number of weights of the network with $W = 64$ as a base (since it can fit the random labels to 100% accuracy), and sparsify the larger networks to have the same number of weights. We further compare this method to *iteratively* and *selectively* removing the lowest-magnitude weights while training. Frankle et al. (2020) show that the latter is expected to have better performance than random static pruning (Golubeva et al., 2021), as well as other more complex methods for pruning at initialization, such as SynFlow (Tanaka et al., 2020) and SNIP (Lee et al., 2019). Although iterative magnitude pruning allows the network to have more parameters during (parts of the) training, at evaluation all the networks have the same number of parameters. We present the results of sparse wide networks in Section 5.3.

5. Results and Discussion

Figure 1 shows a summary of the accuracy as a function of the number of parameters for the different scaling methods for the Damped version. As can be seen, models scaled up in their parameters by increasing width have significantly better generalization on unseen devices compared to models that are scaled up using depth (with, or without increasing the RF). Furthermore, the generalization capabilities of models scaled up with width and depth are on par with the best-case scenarios.

5.1. Increasing the Width

Figure 3 (right) shows that the performance on the seen device improves initially by increasing the network’s width. Increasing W to values larger than 128 seems to have minimal effect on the performance. Looking at the results on unseen devices (left), we can observe that widening both architectures significantly improves their accuracy. Compared to the baseline ($W = 32$, which can fit the training data to 99.99% accuracy), setting $W = 512$ results in 10% increase

Figure 4. The effect of the depth on accuracy.

Figure 5. The effect of the depth (without increasing RF) on Accuracy.

in accuracy on unseen devices for both architectures. Additionally, the Damped model, which encodes the task-specific bias, outperforms the base network in all settings.

5.2. Increasing the Depth

Figure 4 shows the performance of the networks on seen and unseen devices when depth is increased. In both cases, the experiments show that increasing the depth negatively affects the performance. This effect is larger on the base network compared to the damped variant. Koutini et al. (2021) explain this by the increase in the RF of the network, which is not suitable for audio classification tasks. Figure 5 shows that scaling the depth with 1×1 convolutions – which does not change the maximum RF of the network – can give small improvements on the seen device, this effect does not translate to the unseen devices. Moreover, adding these layers worsens the performance of the damped model. Looking at Figures 7 and 9, we can explain this by the observation that adding additional 1×1 convolutions increases the ERF, countering the effect of damping (also see Appendix C.1).

5.3. Increasing the Width with Sparsification

Figure 6 shows the effect of widening the networks while keeping the number of parameters equivalent to the model with $W = 64$. Using static random pruning at initialization, we see initial performance improvements with increasing width. However, these gains diminish on wider networks, in accordance with the findings of Golubeva et al. (2021).

 Figure 6. The effect of the *sparse* width on the accuracy. All the networks have the same number of parameters. Static: random masking at initialization. Itr: iterative magnitude pruning.

 Figure 7. The ERF of different scaled variants of the Damped ResNet. The figures are cropped to the Max RF of the network with $D = 1$ for visibility.

Surprisingly, the performance of the damped version drops on the seen device in iterative pruning. The ERF plots shown in Figure 13 provide a possible explanation: the damped weights are more likely to be pruned first, since they have a lower magnitude, resulting in a smaller ERF. On the contrary, the performance suffers a minimal drop on the unseen devices, even for the model with $W = 512$ that has 98.4% sparsity.

6. Conclusion

We study different methods to scale up CNNs in ASC, and test their ability to generalize to unseen devices. Our experiments show that scaling up the width of CNNs plays the biggest role in their generalization to unseen devices, even if no gains are observed on the already seen devices. We further show that widening the network without increasing parameters can still improve generalization on unseen devices. Additionally, the results suggest that depth scaling can lead to worse performance on unseen devices. Finally, we show a possible interpretation of the results in terms of the ERF of the network.

Acknowledgements

This work has been supported by the COMET-K2 Center of the Linz Center of Mechatronics (LCM) funded by the Austrian Federal Government and the Federal State of Upper Austria. The LIT AI Lab is financed by the Federal State of Upper Austria.

Software and Data

The source code is available on <https://github.com/kkoutini/cpjkudcase21>.

References

- Belkin, M., Hsu, D., Ma, S., and Mandal, S. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32):15849–15854, 2019. ISSN 0027-8424. doi: 10.1073/pnas.1903070116. URL <https://www.pnas.org/content/116/32/15849>.
- Frankle, J., Dziugaite, G. K., Roy, D. M., and Carbin, M. Pruning neural networks at initialization: Why are we missing the mark? In *International Conference on Learning Representations, ICLR 2021*, 2020. URL <https://arxiv.org/abs/2009.08576>.
- Golubeva, A., Neyshabur, B., and Gur-Ari, G. Are wider nets better given the same number of parameters? In *International Conference on Learning Representations, ICLR 2021*, 2021. URL <https://arxiv.org/abs/2010.14495>.
- Heittola, T., Mesaros, A., and Virtanen, T. Acoustic scene classification in DCASE 2020 Challenge: generalization across devices and low complexity solutions. In *Proceedings of the Detection and Classification of Acoustic Scenes and Events 2020 Workshop (DCASE2020)*, 2020. URL <https://arxiv.org/abs/2005.14623>.
- Hershey, S., Chaudhuri, S., Ellis, D. P. W., Gemmeke, J. F., Jansen, A., Moore, R. C., Plakal, M., Platt, D., Saurous, R. A., Seybold, B., Slaney, M., Weiss, R. J., and Wilson, K. CNN architectures for large-scale audio classification. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 131–135, 2017. doi: 10.1109/ICASSP.2017.7952132. URL <https://arxiv.org/abs/1609.09430>.
- Kaplan, J., McCandlish, S., Henighan, T., Brown, T. B., Chess, B., Child, R., Gray, S., Radford, A., Wu, J., and Amodei, D. Scaling laws for neural language models. *CoRR*, abs/2001.08361, 2020.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In Bengio, Y. and LeCun, Y. (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- Koutini, K., Eghbal-zadeh, H., Dorfer, M., and Widmer, G. The Receptive Field as a Regularizer in Deep Convolutional Neural Networks for Acoustic Scene Classification. In *Proceedings of the European Signal Processing Conference (EUSIPCO)*, A Coruña, Spain, 2019. URL <https://arxiv.org/abs/1907.01803>.
- Koutini, K., Eghbal-zadeh, H., and Widmer, G. Receptive field regularization techniques for audio classification and tagging with deep convolutional neural networks. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, pp. 1–1, 2021. doi: 10.1109/TASLP.2021.3082307. URL <https://arxiv.org/abs/2105.12395>.
- Lee, N., Ajanthan, T., and Torr, P. H. S. Snip: single-shot network pruning based on connection sensitivity. In *International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*, 2019. URL <https://openreview.net/forum?id=B1VZqjAcYX>.
- Liu, Y., Neophytou, A., Sengupta, S., and Sommerlade, E. Cross-modal spectrum transformation network for acoustic scene classification. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 830–834. IEEE, 2021.
- Luo, W., Li, Y., Urtasun, R., and Zemel, R. Understanding the Effective Receptive Field in Deep Convolutional Neural Networks. In *Advances in Neural Information Processing Systems*, pp. 4898–4906, 2016. URL <https://arxiv.org/abs/1701.04128>.
- Nakkiran, P., Kaplun, G., Bansal, Y., Yang, T., Barak, B., and Sutskever, I. Deep double descent: Where bigger models and more data hurt. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*, 2020.
- Neyshabur, B., Li, Z., Bhojanapalli, S., LeCun, Y., and Srebro, N. The role of over-parametrization in generalization of neural networks. In *International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*, 2019. URL <https://arxiv.org/abs/1805.12076>.
- Tan, M. and Le, Q. V. EfficientNet: Rethinking model scaling for convolutional neural networks. In *Proceedings of the International Conference on Machine Learning*, pp. 6105–6114, 2019. URL <https://arxiv.org/abs/1905.11946>.

Tanaka, H., Kunin, D., Yamins, D. L., and Ganguli, S. Pruning neural networks without any data by iteratively conserving synaptic flow. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://arxiv.org/abs/2006.05467>.

Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*, 2017. URL <https://arxiv.org/abs/1611.03530>.

Zhang, H., Cissé, M., Dauphin, Y. N., and Lopez-Paz, D. Mixup: Beyond empirical risk minimization. In *International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*, 2018. URL <https://arxiv.org/abs/1710.09412>.

A. Experiment Setup

A.1. Preprocessing

The input clips are 10 seconds, and have a sampling rate of 22050 Hz. We apply Short Time Fourier Transform (STFT) on the input audio signals, with a window size of 2048 and stride of 512. We use perceptual weighting to scale the resulting frequencies, and apply mel filter bank with 256 filters. This results in an input for the CNNs of the shape 256×431 .

A.2. Augmentation

We roll the raw audio wave over time, and use *Mix-up* (Zhang et al., 2018) augmentation with $\alpha = 0.3$, as it has been shown to improve performance on this task (Koutini et al., 2021).

A.3. Training

We train for 250 epochs using Adam (Kingma & Ba, 2015) with a max learning rate of 10^{-4} , starting with 20 epochs of warm-up by exponentially ramping up the learning rate, followed by a linear learning rate decay to 10^{-6} . We repeat all the experiments 3 times and report the mean the standard deviation of the last 10 epochs of the runs.

A.4. Sparsification

In Sparsification experiments, we use 350 epochs of training. For static random pruning, we use the method proposed by Golubeva et al. (2021) to remove weights from the layers with the largest number of parameters first. In iterative

pruning, the number of weights removed in each epoch decays exponentially until we reach the desired number of parameters. Since we are doing global pruning, we prevent layer collapse, by excluding the layers which have more than 99% sparsity from further pruning. Without this measure, we observe significant performance drops starting from $W = 192$.

B. Frequency-Damped Convolutions

Damped convolutions (Koutini et al., 2021) encodes an inductive bias for audio spectrograms, by restricting the Effective Receptive Field (ERF) (Luo et al., 2016) of the CNNs over the input spectrograms. In this paper, we use *frequency-damped convolution* (2021) by replacing every convolutional operator in the $O_n = W_n \star Z_{n-1} + B_n$ with $O_n = (W_n \odot C_n) \star Z_{n-1} + B_n$, where $\star$ is the convolution operator, $\odot$ is the element-wise multiplication, Z_n is the output of the n th layer. C_n is a constant matrix, which has the same shape as the filter weights W_n , C_n has a value of 1 in the center and decays away from the center on the frequency dimension. This biases the classification layer to fit the center of its receptive field, and makes fitting the outer region of its maximum receptive field harder. Figure 8 shows the difference in the effective receptive field between a ResNet and its damped version.

Figure 8. The *effective receptive field* of a ResNet with $W = 128$, $D = 1$, $K = 1$ and its damped version.

C. The Effect of Scaling on the Effective Receptive Field

We calculate the *effective receptive field* as proposed by Luo et al. (2016), by calculating the input gradients to a change in the activation of a single spatial pixel of the penultimate layer of the network (We choose the middle pixel in all the illustrated plots). We take the mean of the absolute gradients over all the test data. The plots show the mean of the absolute gradients over the whole unseen test set. The plots are cropped to the maximum receptive field of the CNN (which can be calculated from the filters' sizes and

the strides of the CNN layers) for better visibility.

C.1. Depth

We show the *ERF* of a Frequency-Damped ResNet with $W = 128, D = 1, K = 0$ and depth-scaled version without increasing the maximum receptive field, by stacking 4 more blocks with filters sizes of 1×1 . The figure shows that adding these blocks – although doesn't change the maximum receptive field – increases the effective receptive field, effectively countering the effect of damping.

Figure 9. The *effective receptive field* of a Frequency-Damped ResNet with $W = 128, D = 1, K = 0$ and $K = 4$.

Figure 10. The *effective receptive field* of the Base ResNet with $W = 128, D = 1, K = 0$ and $K = 4$.

C.2. Width

Figures 11 and Figures 12 show that increasing the width of both variants focuses the ERF in the center of the Max RF.

C.3. Width with Pruning

Figure 13 shows that magnitude pruning reduces the ERF size over frequency for the Damped ResNet, since favors the damped weights (since they are down-scaled). This does not happen in random pruning at initialization, Figure 14. On

Figure 11. The *effective receptive field* of a Frequency-Damped ResNet with $W = 32, D = 1, K = 0$ and $W = 512, D = 1, K = 0$.

Figure 12. The *effective receptive field* of the base ResNet with $W = 32, D = 1, K = 0$ and $W = 512, D = 1, K = 0$.

the base ResNet, comparing Figures 15 and Figure 12 shows widened networks have a similar ERF with and without iterative pruning.

Figure 13. The *effective receptive field* of a Frequency-Damped ResNet with $W = 64, D = 1, K = 0$ and $W = 512, D = 1, K = 0$ with iterative magnitude pruning. Both networks have the same number of parameters.

Figure 14. The *effective receptive field* of a Frequency-Damped ResNet with $W = 64, D = 1, K = 0$ and $W = 512, D = 1, K = 0$ with random pruning at initialization. Both networks have the same number of parameters.

Figure 15. The *effective receptive field* of the base ResNet with $W = 64, D = 1, K = 0$ and $W = 512, D = 1, K = 0$ with iterative magnitude pruning. Both networks have the same number of parameters.

Figure 16. The *effective receptive field* of the base ResNet with $W = 64, D = 1, K = 0$ and $W = 512, D = 1, K = 0$ with random pruning at initialization. Both networks have the same number of parameters.

Figure 17. Test accuracy of scaled up variants of the Base ResNet as the number of parameters increase, as explained in Section 4.