

EFFICIENT LARGE-SCALE AUDIO TAGGING VIA TRANSFORMER-TO-CNN KNOWLEDGE DISTILLATION

Florian Schmid^{1,2}, Khaled Koutini², Gerhard Widmer^{1,2}

¹Institute of Computational Perception (CP-JKU), ²LIT Artificial Intelligence Lab,
Johannes Kepler University Linz, Austria
florian.schmid@jku.at, khaled.koutini@jku.at

ABSTRACT

Audio Spectrogram Transformer models rule the field of Audio Tagging, outrunning previously dominating Convolutional Neural Networks (CNNs). Their superiority is based on the ability to scale up and exploit large-scale datasets such as *AudioSet*. However, Transformers are demanding in terms of model size and computational requirements compared to CNNs. We propose a training procedure for efficient CNNs based on offline *Knowledge Distillation* (KD) from high-performing yet complex transformers. The proposed training schema and the efficient CNN design based on *MobileNetV3* results in models outperforming previous solutions in terms of parameter and computational efficiency and prediction performance. We provide models of different complexity levels, scaling from low-complexity models up to a new state-of-the-art performance of .483 mAP on *AudioSet*.¹

Index Terms— Audio Tagging, AudioSet, Patchout Audio Transformer, MobileNetV3, Knowledge Distillation

1. INTRODUCTION

Audio Tagging (AT) is the task of assigning one or multiple semantic labels to an audio clip. Until recently, CNNs have dominated the field of AT [1, 2, 3]. CNNs are a well-studied and understood architecture for processing spectrograms and include well-suited inductive biases such as the locality bias, weight sharing and translation equivariance. Additionally, CNNs have a specific receptive field, which can be controlled to optimize a model’s generalization capabilities [4]. Efficient CNN designs have been proposed in the vision domain [5, 6, 7, 8, 9] and proved successful also when applied in the audio domain [2, 10].

Transformers [11] build on global attention and lack corresponding inductive biases resulting, for instance, in a learnable receptive field [12, 10]. While CNNs can learn complex tasks given only a limited amount of data, Transformers show their strength when large datasets are available [12, 13, 14].

Given the scale of *AudioSet* [19], Transformers have recently outperformed CNNs and represent the new state of




Fig. 1. Crosses denote models based on Transformer architecture (Audio-MAE [15], HTS-AT [16], PaSST-S [17], PaSST-S-L [17], AST [18], KD-AST [10]) and circles denote models based on CNNs (PSLA [2], ERANN-1-6 [3], Wavegram-logmel-CNN [1], CNN14 [1], KD-CNN [10]). The gray and green curves present width-scaled MobileNets [7] w/ and w/o ImageNet pre-training using the proposed training schema.

the art in AT [18, 17, 16, 15]. However, Transformers are complex in terms of parameters compared to CNNs, and the global self-attention mechanism scales quadratically with respect to the sequence length, making training and inference slow, and the deployment on edge devices infeasible. In this work, we take the best of both worlds by teaching efficient CNNs from performant Transformer ensembles. Our proposed models have around 10 times fewer parameters and require 100 times fewer multiply-accumulate operations for prediction than state-of-the-art transformers while matching their performance.

The main contributions are (1) to provide highly efficient AT models, both in terms of model size and inference speed. The proposed models outperform existing, more complex models and achieve a new single-model state-of-the-art performance of .483 mAP (AP averaged over classes) on *AudioSet*; and (2) to propose a framework for efficient Knowledge Distillation (KD) from high-performance and complex transformer models to smaller and more efficient CNNs.

¹<https://github.com/feschmid56/EfficientAT/>

2. RELATED WORK

Efficient CNNs for Audio Tagging: Searching for efficient CNN architectures for AT and scaling models has been investigated before [1, 3, 2]. Efficient vision architectures, such as EfficientNets [8, 9] and MobileNets [5, 6, 7] have shown to provide a good performance-complexity trade-off also in the audio domain [1, 2, 10]. The latest versions of MobileNets [7] and EfficientNets [9] are based on inverted residual blocks [6, 8], consisting of a pointwise expansion convolution followed by a depthwise convolution and a pointwise projection convolution operation. Squeeze-and-Excitation layers [20] are integrated into the residual blocks to recalibrate filter responses. By default, both EfficientNets and MobileNets derive the final predictions by global pooling followed by an MLP classifier.

Audio Spectrogram Transformers: Inspired by the Vision Transformer (ViT) [12], transformers capable of processing images have been adapted to the audio domain. Vision and Audio Spectrogram transformers [18, 17, 16, 15] extract overlapping patches with a certain stride and size of the input image, add a positional encoding, and apply transformer layers to the flattened sequence of patches. Transformer layers use a global attention mechanism that leads to computation and memory complexity scaling quadratically with respect to the input sequence.

Knowledge Distillation in Audio Tagging: Knowledge Distillation (KD) [21] is a technique that allows low-complexity models to learn from larger, better-performing teacher models. Besides the common classification loss, a distillation loss based on the teacher’s predictions is minimized by the student. Given student’s logits z_S , teacher’s logits z_T , the labels y , an activation function δ and a weighting coefficient λ , KD can be formulated as the weighted sum of label loss L_1 and distillation loss L_{kd} (see Eq. 1).

$$Loss = \lambda L_1(\delta(z_S), y) + (1 - \lambda) L_{kd}(\delta(z_S), \delta(z_T/\tau)) \quad (1)$$

Learning from the teacher’s soft labels, possibly scaled by a temperature τ , enables the student to exploit rich similarity information between classes established by the teacher models [21]. Based on the results on multiple audio classification tasks, Gong et al. [10] found that CNNs and Transformers are well-suited teachers for each other, and simple KD based on logits works better than more complex attention distillation strategies.

3. EXPERIMENT SETUP

3.1. Dataset

We conduct our experiments on AudioSet [19], a dataset consisting of over 2 million 10-second audio clips sampled from YouTube and labeled with a set of 527 classes. AudioSet is weakly labeled and an audio clip possibly contains more than one label. AudioSet comes with an evaluation set consisting of 20,383 recordings. Since AudioSet needs to be obtained from YouTube, different proportions of the dataset can be

	# PARAMs	mAP
MN-Baseline	4.88M	.401
MN-Baseline pre-trained	4.88M	.417
MN-KD	4.88M	.458
MN-KD pre-trained	4.88M	.470

Table 1. MN performance on AudioSet[19] w/ and w/o KD and pre-training on ImageNet [23].

successfully downloaded. In this regard, our setup is strictly comparable to the dataset used to train *PaSST* models in [17].

3.2. Training Setup

We apply the same preprocessing as in [17]. We use mono audio sampled at 32 kHz and compute Mel features from 25 ms windows with a hop size of 10 ms. Importance sampling based on the label frequency is applied to counter the long tail of infrequent classes.

Models are trained for 200 epochs. In each epoch, we sample 100,000 samples without replacement from the full AudioSet. The learning rate increases to its maximum value 8×10^{-4} within the first 8 epochs and decreases linearly from epoch 80 to epoch 175 to 1% of its maximum. We apply the Adam optimizer and use a batch size of 120.

Mixup [22] with a mixing coefficient of 0.3 is used since it has shown to improve performance on audio classification tasks before [18, 4]. We apply Mixup at the spectrogram level and mix the teacher soft labels accordingly. Since teacher predictions are generated offline on unaugmented data, we do not apply any further data augmentation methods.

4. KNOWLEDGE DISTILLATION

We use the efficient CNN MobileNetV3-Large [7] (abbreviated as *MN* in the following) as our student baseline model. Table 1 shows that *MN-Baseline* falls short of previous CNNs and Transformer models on AudioSet [19]. Pre-training on ImageNet [23] builds a solid weights prior for pattern recognition [16] and improves results, but the model is still inferior to previous AT models. In this section, we explain how to use KD to improve performance drastically.

For KD, we ensemble *PaSST* [17] models with different patch sizes and strides as a teacher. We pre-compute the predictions of 9 different *PaSST* models² on AudioSet to speed up training and form an ensemble by averaging the logits. The *PaSST* ensemble achieves a mAP of .495, which is state of the art for ensembles on AudioSet.

With regard to the components of Eq. 1, we use Binary-Cross-Entropy for both label L_1 and distillation loss L_{kd} and as activation function δ we apply Sigmoid activation. These settings correspond to the natural choices for the AT task. λ and τ are subject to experiments described in Section 6.

²Available at: <https://github.com/kkoutini/PaSST/blob/main/models/passt.py>

Table 1 shows that using KD to teach a pre-trained MN from a PaSST Transformer ensemble improves results substantially by matching the performance of a single state-of-the-art PaSST model while reducing the model size to approximately 6% of PaSST-S, which has 87 million parameters.

5. MODELS AT SCALE

In this section, we vary the model complexity and the spectrogram resolution to obtain models of different sizes and computational complexities. We keep the number of layers constant and scale the model’s width by multiplying the number of channels by a scaling factor α . $\alpha = 1$ is the default MN (shown in Table 1), setting $\alpha < 1$ and $\alpha > 1$ produces models of reduced or increased complexity, respectively. The result is a range of models, starting from low-complexity models, which target edge devices, to larger models that achieve a new state-of-the-art performance on AudioSet [19].

Figures 1 and 2 show that our scaled models without ImageNet [23] pre-training (named ‘MobileNet (width-scaled)’ in the figures) outperform previous AT models in terms of efficiency and performance even without the need for pre-training. However, pre-training a selection of models of varying widths on ImageNet (named ‘MobileNet pre-trained (width-scaled)’ in the figures) demonstrates that the performance of our proposed models, especially the smaller ones, can be further boosted with pre-training.

Parameter Complexity: Figure 1 shows that our scaled models ranging from below 1 million to 68 million parameters compare favorably to other AT models in terms of parameter efficiency. Without pre-training, we need to scale up the model to 16 million and with pre-training to 5 million parameters to reach state-of-the-art performance. The three largest models set a new state-of-the-art AT performance on AudioSet, with the largest achieving a mAP of .483.

Computational Complexity: Complementary to the model size, an important factor to consider about model complexity is the number of multiply-accumulate operations (MACs) required for classifying a single recording. Besides the model width, the resolution of the spectrograms affects the number of MACs. Figure 2 shows how varying hop size (named ‘MobileNet (hop)’ in the figure), the number of mel bands (named ‘MobileNet (mels)’ in the figure) and model width (model sizes corresponding to Figure 1) influences performance and required MACs. MACs are calculated for linear, convolutional and attention layers, as these are the dominating computational factors of the compared models.

Figure 2 shows that our proposed models compare favorably against previous CNN solutions, for instance, PSLA [2] or KD-CNN [10], which are based on EfficientNet-B2 [8]. State-of-the-art transformers, like PaSST [17], require a much higher computational budget compared to CNNs. MN with $\alpha = 1.0$ uses approximately 100 times fewer MACs than PaSST-S while achieving similar performance. Although the




Fig. 2. Comparison of model performances vs. computational effort to predict a single sample. MobileNets with different width (same as Figure 1), spectrogram resolutions (hop size in [10, 15, 20, 25] ms and mels in [40, 64, 128, 256]) are compared to other Audio Tagging models (PSLA [2], CNN14 [1], ResNet38 [1], KD-CNN [10], PaSST-S [17]).

speedup in practice is very hardware-specific, we confirm our theoretical analysis by benchmarking the inference throughput of PaSST and our model on an Nvidia A100 GPU. We use a batch size of 200 and obtain a throughput of 78 clips/sec. for PaSST; for a similarly performing MN ($\alpha = 1.0$) we get 4767 clips/sec., which corresponds to a 61 times speedup.

6. ABLATION STUDY

In this section, we discuss important model design choices and hyperparameters of the training routine. All experiments use the default MN with $\alpha = 1.0$ and the official AudioSet split [19]. In the following tables, cells highlighted in gray mark the final hyperparameter choices our models are based on.

6.1. Knowledge Distillation Hyperparameters

We apply KD with temperatures $\tau = 1$ and $\tau = 3$ to create teacher soft labels and experiment with different weights λ to trade off label and distillation loss (see Eq. 1). Table 2 shows that the best results were obtained using a high distillation loss weight in combination with $\tau = 1$. However, switching off the label loss L_1 completely is counterproductive. Compared to the baseline models in Table 1, using KD leads to a consistent improvement across all configurations.

	$\lambda = 0.5$	$\lambda = 0.3$	$\lambda = 0.1$	$\lambda = 0.0$
$\tau = 1$.454	.460	.464	.454
$\tau = 3$.455	.459	.457	.454

Table 2. Results of offline KD using different values for λ (see Eq. 1) and temperatures $\tau = 1$ or $\tau = 3$ to create the soft labels. Results are based on pre-trained models and a training routine including data augmentation.

6.2. Data Augmentation and Consistent Teaching

Since we are using offline KD, the teacher predictions are generated for the original AudioSet [19] clips before any data augmentation. Additionally, it has been shown [10] that consistent teaching [24] is beneficial. We experiment with the impact of data augmentation and present the results in Table 3. Switching off frequency and time masking leads to a performance improvement, using Mixup [22] has a slight positive performance impact and switching off rolling the waveform over time and applying gain augment does not change results significantly.

SpecAugment Sizes	mAP
Augmented	.464
-Masking	.470
-Masking, -Rolling	.470
-Masking, -Gain Augment	.471
-Masking, -Gain Augment, -Rolling	.470
-Masking, -Mixup	.468

Table 3. Studying the impact of Data Augmentation in KD. *Masking* refers to time and frequency masks as in SpecAugment [25], *Rolling* rolls the waveform over the time axis, *Gain Augment* changes the waveform’s amplitude by ± 7 dB and *Mixup* [22] is applied at spectrogram level.

6.3. Modified Squeeze-And-Excitation

Original Squeeze-and-Excitation (SE) [20] recalibrates the filter responses based on channel weights computed from combined channel statistics. We experiment with applying the same SE mechanism based on frequency statistics to recalibrate frequency bands.

The results in Table 4 show that SE-layers are an important architectural ingredient, in terms of performance. Original channel-wise SE performs best but accounts for approximately 30% of the model’s parameters. We propose frequency-wise SE as a lightweight alternative showing a promising performance-complexity trade-off. Frequency-wise SE only adds insignificantly to the model’s parameters, as it does not scale with the network width, but improves substantially over not using SE layers.

SE dimensions	# PARAMs	mAP
None	3.36M	.453
Channel	4.88M	.470
Frequency	3.37M	.465

Table 4. Comparing the effect of using no, channel-wise or frequency-wise Squeeze-and-Excitation layers in the student network.

6.4. Different Types of Classification Heads

MobileNets come with global channel pooling and an MLP classifier, whereas other popular classification heads in the audio domain include multihead-attention [2] or fully-convolutional architectures [4]. Table 5 shows that 4-headed attention can improve slightly on the MLP classification head, but this comes at the cost of substantially increasing the number of parameters. A fully-convolutional classification head saves around 29% of parameters but results in a performance reduction. The simple MLP classification head gives the best performance-complexity trade-off.

Head Type	# PARAMs	mAP
MLP Classifier	4.88M	.470
Fully-Convolutional	3.48M	.465
Multihead Att. (2 heads)	5.00M	.469
Multihead Att. (4 heads)	7.02M	.471

Table 5. Results of the student network equipped with different types of classification heads.

7. CONCLUSION

The outcome of this paper is an efficient and performant Audio Tagging model that is easy to scale to given resource constraints. Starting from low-complexity models that can operate on resource-constrained platforms, the model can be scaled up to reach a new state-of-the-art performance on AudioSet, while being more efficient in terms of model size and computational effort than previous solutions. Our proposed models exploit the high performance reached by an ensemble of heavy-weight transformers in our KD setup while having the advantage of scaling linearly with increasing input sequence length. For future work, we plan to investigate the performance of our proposed AudioSet pre-trained models on downstream tasks and compare the quality of extracted embeddings to those extracted by Transformer models.

8. ACKNOWLEDGMENT

The computational results presented were achieved in part using the Vienna Scientific Cluster (VSC) and the Linz Institute of Technology (LIT) AI Lab Cluster. The LIT AI Lab is supported by the Federal State of Upper Austria. Gerhard Widmer’s work is supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme, grant agreement No 101019375 (Whither Music?).

9. REFERENCES

- [1] Qiuqiang Kong, Yin Cao, Turab Iqbal, Yuxuan Wang, Wenwu Wang, and Mark D. Plumbley, “Panns: Large-scale pretrained audio neural networks for audio pattern

- recognition,” *IEEE ACM Trans. Audio Speech Lang. Process.*, 2020.
- [2] Yuan Gong, Yu-An Chung, and James R. Glass, “PSLA: improving audio tagging with pretraining, sampling, labeling, and aggregation,” *IEEE ACM Trans. Audio Speech Lang. Process.*, 2021.
- [3] Sergey Verbitskiy, Vladimir B. Berikov, and Viacheslav Vyshegorodtsev, “Eranns: Efficient residual audio neural networks for audio pattern recognition,” *Pattern Recognit. Lett.*, 2022.
- [4] Khaled Koutini, Hamid Eghbal-zadeh, and Gerhard Widmer, “Receptive field regularization techniques for audio classification and tagging with deep convolutional neural networks,” *IEEE ACM Trans. Audio Speech Lang. Process.*, 2021.
- [5] Andrew G. Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam, “Mobilenets: Efficient convolutional neural networks for mobile vision applications,” *CoRR*, 2017.
- [6] Mark Sandler, Andrew G. Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen, “Mobilenetv2: Inverted residuals and linear bottlenecks,” in *CVPR*, 2018.
- [7] Andrew Howard, Ruoming Pang, Hartwig Adam, Quoc V. Le, Mark Sandler, Bo Chen, Weijun Wang, Liang-Chieh Chen, Mingxing Tan, Grace Chu, Vijay Vasudevan, and Yukun Zhu, “Searching for mobilenetv3,” in *ICCV*, 2019.
- [8] Mingxing Tan and Quoc V. Le, “Efficientnet: Rethinking model scaling for convolutional neural networks,” in *ICML*. 2019, Proceedings of Machine Learning Research, PMLR.
- [9] Mingxing Tan and Quoc V. Le, “Efficientnetv2: Smaller models and faster training,” in *ICML*, 2021.
- [10] Yuan Gong, Sameer Khurana, Andrew Rouditchenko, and James R. Glass, “CMKD: cnn/transformer-based cross-model knowledge distillation for audio classification,” *CoRR*, 2022.
- [11] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin, “Attention is all you need,” in *NIPS*, 2017.
- [12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby, “An image is worth 16x16 words: Transformers for image recognition at scale,” in *ICLR*, 2021.
- [13] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou, “Training data-efficient image transformers & distillation through attention,” in *ICML*, 2021.
- [14] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie, “A convnet for the 2020s,” in *CVPR*. 2022, IEEE.
- [15] Po-Yao Huang, Hu Xu, Juncheng Li, Alexei Baevski, Michael Auli, Wojciech Galuba, Florian Metze, and Christoph Feichtenhofer, “Masked autoencoders that listen,” *CoRR*, 2022.
- [16] Ke Chen, Xingjian Du, Bilei Zhu, Zejun Ma, Taylor Berg-Kirkpatrick, and Shlomo Dubnov, “HTS-AT: A hierarchical token-semantic audio transformer for sound classification and detection,” in *ICASSP*, 2022.
- [17] Khaled Koutini, Jan Schlüter, Hamid Eghbal-zadeh, and Gerhard Widmer, “Efficient Training of Audio Transformers with Patchout,” in *Interspeech*, 2022.
- [18] Yuan Gong, Yu-An Chung, and James Glass, “AST: Audio Spectrogram Transformer,” in *Interspeech*, 2021.
- [19] Jort F. Gemmeke, Daniel P. W. Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R. Channing Moore, Manoj Plakal, and Marvin Ritter, “Audio set: An ontology and human-labeled dataset for audio events,” in *ICASSP*, 2017.
- [20] Jie Hu, Li Shen, and Gang Sun, “Squeeze-and-excitation networks,” in *CVPR*, 2018.
- [21] Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean, “Distilling the knowledge in a neural network,” *CoRR*, 2015.
- [22] Hongyi Zhang, Moustapha Cissé, Yann N. Dauphin, and David Lopez-Paz, “mixup: Beyond empirical risk minimization,” in *ICLR*, 2018.
- [23] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei, “Imagenet: A large-scale hierarchical image database,” in *CVPR*, 2009.
- [24] Lucas Beyer, Xiaohua Zhai, Amélie Royer, Larisa Markeeva, Rohan Anil, and Alexander Kolesnikov, “Knowledge distillation: A good teacher is patient and consistent,” *CoRR*, 2021.
- [25] Daniel S. Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D. Cubuk, and Quoc V. Le, “SpecAugment: A simple data augmentation method for automatic speech recognition,” in *Interspeech*, 2019.