



AudioTagging Done Right: 2nd comparison of deep learning methods for environmental sound classification

Juncheng B Li, Shuhui Qu, Po-Yao Huang, Florian Metze

Carnegie Mellon University

junchen1, poyaoh, fmetze@cs.cmu.edu

Abstract

After its sweeping success in vision and language tasks, pure attention-based neural architectures (e.g. DeiT) [1] are emerging to the top of audio tagging (AT) leaderboards [2], which seemingly obsoletes traditional convolutional neural networks (CNNs), feed-forward networks or recurrent networks. However, taking a closer look, there is great variability in published research, for instance, performances of models initialized with pretrained weights differ drastically from without pretraining [2], training time for a model varies from hours to weeks, and often, essences are hidden in seemingly trivial details.

This urgently calls for a comprehensive study since our 1st comparison [3] is half-decade old. In this work, we perform extensive experiments on AudioSet [4] which is the largest weakly-labeled sound event dataset available, we also did analysis based on the data quality and efficiency. We compare a few state-of-the-art baselines on the AT task, and study the performance and efficiency of 2 major categories of neural architectures: CNN variants and attention-based variants. We also closely examine their optimization procedures. Our open-sourced experimental results¹ provide insights to trade off between performance, efficiency, optimization process, for both practitioners and researchers.

Index Terms: AudioSet, CNNs, ViT, Efficiency, Optimization

1. Introduction

Recently, after seeing tremendous success in language tasks[5], the ML community has been exploring a variety of methods for deploying attention-based architectures, e.g. Vision Transformers (ViT) [6]—in computer vision and other fields, and competitive performances are reported. Recently, AST [2] and PSLA [7] improved the SOTA performance of the AT task² on the AudioSet benchmark by leveraging a suite of improvements including DeiT[1] (distilled ViT) architecture, ImageNet pre-training, data augmentations, and ensemble. However, there is still no clear “winner-takes-all” approach in audio classification tasks that can have the best performance while being efficient.

Audio signals, 1D continuous by nature, require different processing than vision (2D) or language input sequences (discrete). Environmental sounds, compared to well-studied human speech, do not require language model, but are more diverse and span a wide range of frequencies. Thus, the same techniques that worked well in speech are not guaranteed to work out of the box [8]. Plus, the lack of well-defined strongly-labeled data makes environment sounds recognition task not only more challenging, but also understudied so far.

5 years ago, in our first comparison [3], we identified CNN’s superiority over MLPs, and RNNs. CNNs [9, 10] and

¹<https://github.com/lijuncheng16/AudioTaggingDoneRight>

²Audio Tagging (AT) task aims to characterize the acoustic event of an audio stream by selecting a semantic label for it.

its variant CRNNs [11] have become the de facto architecture for acoustic event recognition tasks and have dominated the AudioTagging leaderboard until the rise of ViT [2]. The major difference between convolutional models and attention-based networks (e.g. ViTs) are the locality inductive biases. In a nutshell, the convolutional neural network sweeps through every consecutive pixel with its learned kernel, which is perfect for learning *local features*[3], but cannot see beyond its receptive field; whereas ViT networks skip through, attend between patches to build *global correlations*, and sometimes not even rely on positional information.

In this work, we seek to thoroughly understand the difference between using each type of model on the AT task. We train 4 variants of transformer networks including CNN+Transformer, Vision Transformer (ViT), Transformer, and Conformer on the task of Audio Tagging, compare them with ResNet, CRNN control group. Between these six architectures, we perform analysis on the largest available dataset: Google AudioSet [4]. Our contributions:

1. We systematically compared the performances of different transformer variants between several CNN variants under a permutation of different settings on the same large-scale dataset Audioset (e.g. w/. or w/o. pretraining, lr scheduling). Our experiments suggest that pre-training is *not* always necessary, LR scheduling & data augmentation are always helpful.
2. Our experiments shed light on critical optimization strategies & tradeoffs, e.g. feature size, LR schedule, batch size, momentum, normalization, loss landscape, which have not been thoroughly explained previously.

2. AT Background & Related Works

AudioSet [4] contains 2,042,985 10-second YouTube video clips, summing up to 5,800 hours annotated with 527 types of sound events (weak label³). The same group [17] conducted quality assessments of these labels ranging from 0-100%. The *full* trainset has 2 subsets: classwise *balanced* set (22,176 samples) and *unbalanced* (2,042,985 samples) set, and *eval* set with 20,383 samples [18].

AudioTagging Benchmark: The left columns of Table 1 list state-of-the-art *single models* and training procedures for the AT task trained on the *full* AudioSet and test on the *eval* set⁴. Here, we *do not* include *multi-modal* or *ensembled* models, which would introduce tremendous extra variability, making fair comparison almost impossible⁵. Unlike some previous works, we

³does not specify which second specific event happens

⁴Previous benchmarks are excerpted from the original publications

⁵We consider weight-averaging(WA)[19] an implicit ensemble, hence we report single-checkpoint score w/o WA: AST [2] 0.459 → 0.448, and PSLA [7] 0.444 → 0.439. WavegramCNN in PANNs [10] is also an implicit ensemble, thus leaving out of the comparison.

Table 1: Different training procedures for single-checkpoint Audio-Only models trained on the full AudioSet as of Mar 2022. SSL[12]: pretrained on Self-supervised task using 3.9M (67k hours) proprietary data. To standardize reporting difference in steps VS. epoches, epoch* here means 1 full iteration of the TrainSet. Adam optimizer’s default $\beta_1=0.9, \beta_2=0.999$, the ones without specified β are default ones. LR: learning rate. TimeSpecAug[13]: t, f indicate max length of time and frequency mask, \checkmark means specific setups are unknown. Label enhancement[7] involves altering the original labels. CNN14 [10] achieves 0.442 mean average precision (mAP) using 128 mel bins. We are aware of PaSST [14], due to its high similarity to AST except for dropout/ensemble, we do not separately list it.

Model	Previous approaches				Attention		Our Implementation			
	CNN Variants						see Table 2			
Procedure	PANNs CNN14 ([10])	ResNet [9]	PSLA [7]	ERANNs [15]	Conformer [12]	AST [2]	A1	A2	A3	A4
params	42.2M	23M	13.6M	54.5M	88.1M	88M	see Table 2			
train dataset	1934187	1953082	1953082	1803891	2063949	1953082	1998999		1998999	
eval set	18887	19185	19185	17967	20371	19185	20126		20126	
feature size	64×1001	64×1000	128×1056	128×1280	64×500	128×1024	128×1024		64×400	
pretrained	\checkmark	\checkmark	ImageNet	\checkmark	SSL	ImageNet	ImageNet	\checkmark	ImageNet	\checkmark
epoch*	10	50	30	9	100	5	10	10	10	10
batch size	32	n/a	100	32	640	12	20	400	80	448
optimizer	adam	adam	adam	adam	adam	adam	adam	adam	adam	adam
maxLR	0.001	0.95-0.999	0.95-0.999	0.001	0.9-0.98	0.95-0.999	1.0E-05	4.0E-4	1.0E-05	4.0E-4
LR decay	\checkmark	\checkmark	step	one-cycle	linear	step	step	step	step	step
decay rate	\checkmark	\checkmark	0.5	[16]	3.00E-6	0.5	0.5	0.5	0.5	0.5
decay epochs	\checkmark	\checkmark	5	cyclic	100	2	2	2	2	2
weight decay	\checkmark	5.0E-07	5.0E-07	\checkmark	0.01	5.0E-07	\checkmark	\checkmark	\checkmark	\checkmark
warmup steps	\checkmark	\checkmark	1000	\checkmark	10k	1000	1000	1000	1000	1000
dropout	\checkmark	\checkmark	\checkmark	\checkmark	0.1	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark
databalancing	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark
mixup	\checkmark	\checkmark	\checkmark	modified	\checkmark	\checkmark	0.3	0.3	0.3	0.3
TimeSpecAug	\checkmark	\checkmark	\checkmark	\checkmark	timeonly	\checkmark	t192,f36	t192,f36	t75,f12	t75,f12
label enhance	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark
normalize	\checkmark	\checkmark	x2	\checkmark	\checkmark	x2	x2	x2	x2	x2
train time	3 days	n/a	a week	n/a	n/a	a week	108 Hrs	16.45Hrs	21.0Hrs	8.84Hrs
GPU	V100×1	n/a	titanX×4	n/a	n/a	titanX×4	V100×4	V100×4	V100×2	V100×2
Best mAP	0.431	0.392	0.439	0.450	0.415	0.448	0.430	0.437	0.410	0.411

do not list models trained only using the *balanced* subset since it only accounts for 1% of the AudioSet training set, which are 10-20% mAP worse than models trained on the *full* set.

3. Experiments & Results

We explored many variations with different optimization techniques, data augmentations, choice of regularization, and a reasonable amount of grid search for the hyperparameters. We offer 4 different training procedures with different costs and performance that cover different typical use cases, see table 1.

Procedure A1: aims at reproducing the best SOTA performance. Therefore, it is the longest in terms of training time using the larger feature. **Procedure A2:** is to test whether we could reach SOTA without pre-training, except with a larger batch size and LR.

Procedure A3 & A4: aim at matching SOTA performance using $5.12 \times (\frac{1}{2})$ #Mels, $\frac{1}{2.56}$ less time resolution) smaller features thus a lot faster. It can be trained on-average 6x faster and could be a good setting for exploratory research.

AST/ViT (Table 2, Figure 1(C)). uses pre-trained weights from the DeiT-base-384 [1] model imported from the timm library [20]. In order to preserve the learned positional embedding in the pre-trained DeiT [1] model ($24 \times 24 = 576$ patches), we choose the same 16×16 patch-size for all our AST experiments. We did comparison on stride sizes as shown in Table 3. e.g. A3 procedure use strides of 8 (8 pixels overlap), resulting in ($7 \times 49 = 343$) patches. Shorter strides benefit smaller features

Table 2: Our implementation of architectures with different training procedures mentioned above. A1, A3 contain blanks since ImageNet pretrained weights are only available for ResNet50 and AST/ViT. **Bolded:** the models reported in Table 1

Model	#param	Best mAP		Best mAP	
		A1	A2	A3	A4
AST/ViT	87.9M	0.430	0.274	0.410	0.268
Transformer	28.5M	-	0.230	-	0.209
CNN+Trans	12.1M	-	0.437	-	0.411
Conformer	88.1M	-	0.335	-	0.308
ResNet50	25.6M	0.410	0.399	0.382	0.370
CRNN	10.5M	-	0.429	-	0.406

more than larger features, but there is a catch: Computation cost grows quadratically due to the $\mathcal{O}(p^2)$ attention mechanism, p being the number of patches.

Transformers (Table 2) We implemented pure transformers without convolution or ViT type patches, which takes in logMel spectrogram as input. Since transformer layer is temporal agnostic, we introduce positional encoding to retain the temporal order of inputs. To add positional encoding to the audio model, we adopted the classic positional encoding [5], which is defined as follows, where d represents the dimension of the input, pos is the position in time and i is the dimension index in the input tensor.

$$\mathbf{PE}_{(pos,k)} = \begin{cases} \sin\left(\frac{pos}{10000^{2i/d}}\right) & k = 2i \\ \cos\left(\frac{pos}{10000^{2i/d}}\right) & k = 2i + 1 \end{cases}$$

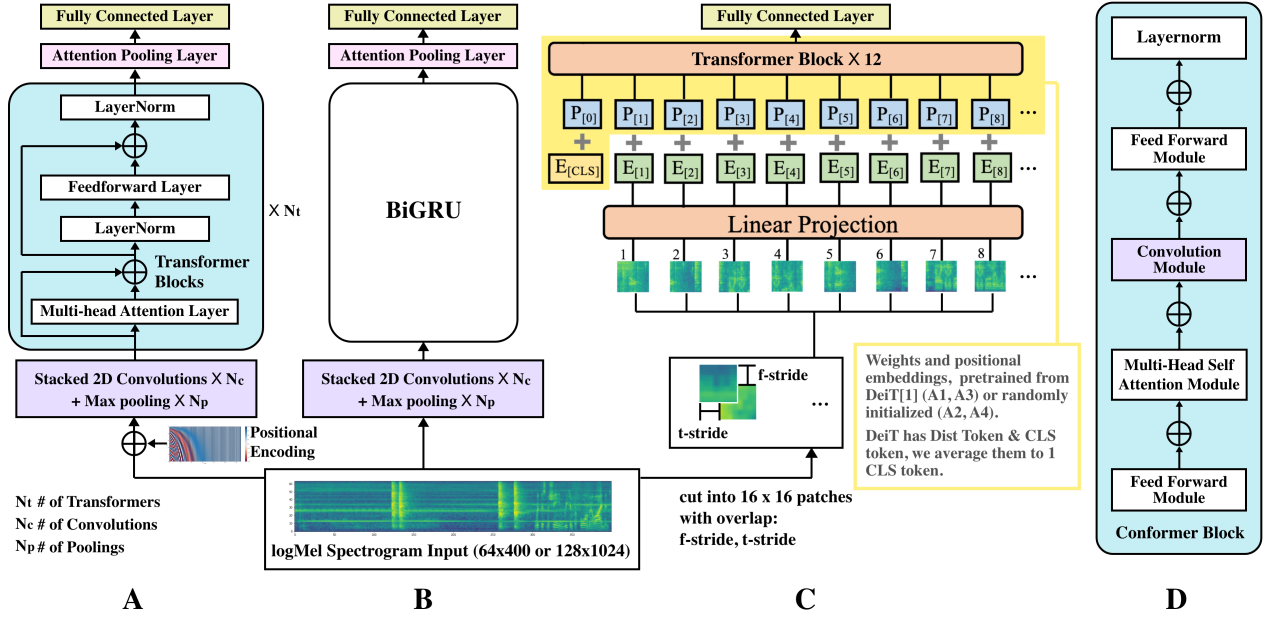


Figure 1: The overall architecture (A): CNN+Transformer (B): TALNet[11](CRNN) (C): AST[2]/DeiT[1](ViT) (D): Conformer

Table 3: time, freq stride influence AST’s mAP and train time, blanks are the experiment could not fit into GPU’s.

Procedure	tstride × fstride	4×4	6×6	8×8	10×10	12×12	16×16
A1 128 × 1024	mAP	-	-	-	0.430	0.429	0.421
	#patches	8192	3570	2048	1212	756	512
	Hrs/Epoch	-	-	-	10.8	6.19	5.53
A3 64 × 400	mAP	0.414	0.408	0.410	0.370	0.320	0.24
	#patches	1485	390	343	195	128	72
	Hrs/Epoch	12.67	2.45	2.1	1.44	1.01	0.8

We added positional encoding to spectrograms before feeding into CNN layers and we scaled up the original input. Given the input x : $n \times d$, we scale up the input by square root of the input dimension: $x = x \cdot \text{round}[\sqrt{d}] + \text{PE}(x)$

The overall architecture is depicted in Figure 1(A) but without the stacked 2D conv + pooling layers. The resulting mAP are not as high as the AST/ViT type. In Table 2, we report the best performing one with $N_t = 8$ layers of Transformer blocks with MultiHeadAttention (MHA) layers (8 attention heads) after testing out $N_t = 2, 4, 6, 8, 10, 12$ with 4, 8, 12 attention heads set-ups. The same as [14], to train faster and generalize better, we implemented dropout layers within the MHA layers. mAP for other setups are too low for meaningful comparisons. Transformers suffer from the same $\mathcal{O}(n^2d)$ cost issue as ViT, n being the input length. Therefore, training for Transformers is quadratically slower than models with the same number of parameters without the attention mechanism.

CRNN (Table 2, Figure 1(B)) We follow the well-tuned TALNet architecture using $N_c = 10$ convolution layers with 3×3 kernels and $N_p = 5$ max-pooling layers in between, resulting in embedding size of 1024 feeding into the biGRU layer. The output of GRU is fed to a fully-connected layer of size 527 to predict frame-wise probabilities. Finally, the frame probabilities is aggregated with a pooling function to make final prediction. We did not tune hyperparameters for TALNet $\sim \mathcal{O}(md^2)$ where m is the stacked convolution output sequence length, d is the representation dimension.

CNN+Transformer (Table 2, Figure 1(A)) The first part of Stacked Conv blocks and Pooling is the same as the CRNN we

implemented, which uses $N_c = 10$ Conv layers and $N_p = 5$ pooling layers. We replace the GRU layer with the aforementioned Transformers⁶ $\sim \mathcal{O}(m^2d)$, and experimented with $N_t = 1, 2, 4$ layers of Transformer layers with 4, 8, 12 attention heads. We use the same positional encoding as our aforementioned Transformer model. Here, we report the best performing $N_t = 2$ layers of transformers with 8 attention heads, followed by the attention pooling layer to aggregate frame probabilities. **ResNet** (Table 2) Following [9], we replace the standard ResNet50’s last layer with attention pooling layer to output 527 classes probabilities. We also implemented ResNet34, but since ResNet architectures are from the same family $\sim \mathcal{O}(knd^2)$, k : kernel size, we only list ResNet50 in Table 2.

Conformer (Table 2) We implemented the same large conformer as [12] $\sim \mathcal{O}(n^2d)$ using 12 conformer blocks with 768D encoder embeddings and 12 attention heads, as this setting was reported to be the best performing setup for conformer architecture. Each conformer block is shown in Figure 1(D).

4. Observations and Discussion

4.1. Data Quality & Data Efficiency

Missing files: In table 1, due to the downloading difference of AudioSet, number of train&test set varies a *whopping* $\pm 5\%$ across previous works, especially the different test size could cause severe fluctuations in final mAP reporting as seen in Figure 2. e.g. one could have downloaded the lower-label-quality test samples that tank their score. We release our test set labels along with our implementation for consistency.

Speedup: As we can see in Table 1, previous approaches benefited from using larger features. In Table 3 we see dramatic speedup using $64(\#mels) \times 400(\text{time})$ logMel spectrogram⁷ VS. $128(\#mels) \times 1024$ filter bank features, this 5.12 times feature size reduction results in more than 6 times speedup. Note that

⁶To avoid repetition, m, n, d are shared among different models

⁷The waveform is downsampled to 16 kHz; frames of 1,024 samples (64 ms) are taken with a hop of 400 samples (25 ms); each frame is Hanning windowed and padded to 4,096 samples before taking the Fourier transform; the filterbank of 64 triangle filters spans a frequency range from 0 Hz to 8 kHz.

mAP < 0.35 not listed, A1: 128x1024, pretrain; A2: 128x1024 w/o pretrain; A3: 64x400 pretrain; A4: 64x400 w/o pretrain;




Figure 2: Performance(mAP) fluctuation on quantiles of data with varying label quality.

half of the pipeline efficiency (for both train & inference) depends on the data efficiency, the rest is then about model efficiency. Using our smaller feature would result in 10^{-1} reduction of data time per sample. Therefore, we highly recommend our A3, A4 procedure for research or inference task, which could easily fit into a single 16GB-GPU with a lot less carbon footprint. Indeed, we trade-off 2 points of mAP for 6 times of speed up. To analyze the performance loss due to feature temporal/frequency resolution loss, we took a closer look at different label quality quantiles of the data. As Figure 2 shows, larger features outperform smaller features in high-quality classes by $2 \pm 0.1\%$ mAP, indicating models benefit, but very limited, from higher freq/time feature resolution.

4.2. Impact of Pretraining

Similar to the conclusion of [21], we find ImageNet pretraining helps certain architectures to converge, but is not a must. As shown in Table 2, CNN-Trans model trained from scratch can outperform AST models that are pre-trained on ImageNet and also be better than the Conformer model pre-trained using SSL in [12]. CRNNs are also competitive. We observe AST models need pre-trained weights to converge to optimal, cold-start weights/positional embeddings from random Gaussian results in a lot lower performance. This is maybe due to huge distribution shift from Gaussian to optimal weight distribution.

4.3. Insights regarding Optimization Process

LR scheduling is crucial for training a high-performance model. The proven strategy is to scale $\mathbf{LR} \propto \mathbf{batch\ size}$ [22] when amplifying batch size, and to fill the GPU-memory till full with large enough batch size, since increasing batch-size linearly speeds up training time. **LR decay** is also important and can be seen as simulated annealing [22]. We view Adam/SGD training as $\frac{dw}{dt} = -\frac{d\mathcal{L}}{dw} + \eta(t)$ where \mathcal{L} is the BCE loss function in our case, summed over all training examples, and w denotes the parameters. $\eta(t)$ denotes Gaussian random noise updating in continuous “time” t towards convergence, which models the effect of estimating the gradient using mini-batches. In [22], they showed that the mean $\mathbb{E}(\eta(t)) = 0$ and variance $\mathbb{E}(\eta(t)\eta(t')) = gF(w)\delta(t-t')$, where $F(w)$ describes the covariance in gradient fluctuations between different parameters. They also proved that the “noise scale” $g = \epsilon(\frac{N}{B} - 1)$, where ϵ is the learning rate (LR), N the training set size and B the batch size. This noise scale controls the magnitude of the random fluctuations in the training dynamics. Intuitively, the initial

noisy phase allows the model to explore a larger fraction of the parameter space without getting trapped in local minima. Once we find a promising region of parameter space, we reduce the noise to fine-tune the parameters. In Table 1 and all our experiments, we observe LR decay outperforms constant LR by at least 2% mAP. We also implemented cyclic schedule [16] used by [15], but the result is suboptimal VS. step decay. Meanwhile, annealing the temperature in a series of discrete steps can sometimes trap the system in a “robust” minimum, in our experiments, this severely impede training attention-based models such as AST, transformers. e.g when we initialize training AST in our A3 procedure with a non-optimal LR: say $1\text{E-}6$, $1\text{E-}4$, the model’s mAP get “stuck” below 0.1. This effect is less severe for non-attention or hybrid models. We postulate the culprit is the *sharp loss landscape* of the attention-based models [23], when the gradient accumulation cannot adapt to changes in the loss landscape, training would be impeded, since we would keep on exploring the wrong direction in parameter space. This also explained why our training also failed when we used *gradient accumulation* step size ≥ 2 . In Table 1, we can see some previous works have tuned the β_1 of Adam to 0.95, this is actually letting gradient updating **twice** as slower than the default 0.9. $m_t = \beta_1 m_{t-1} + (1 - \beta_1)g_t$ where m_t is the first moment of the gradient and g_t is the gradient at t step. This trick might compensate for the aforementioned sharp-loss-landscape phenomenon, but we did not achieve the same performance gain using the trick, hence stick to default. Some previous works advocate *regularization* for better model generalizability by adopting weight decay or dropout, we observe 1-2% mAP performance drop using regularization in our experiments, and therefore did not include them in our procedures. We note [14] used AdamW [24] optimizer. AdamW itself was a fix to weight decay within the adam update iteration, we do not use regularization in this work, hence AdamW is not studied.

4.4. Data Augmentations & Normalization

AST and PSLA [2, 7] applied an unusual normalization (Normalize*) to the feature input: $x = \frac{x-\mu}{2\sigma}$, which results in the normalized feature with $\mathbb{E}(x) = 0$, and $\text{var}(x) = 0.25$. In our experiment, Training AST using input of $\text{var}(x) = 1$ or $\text{var}(x) = 0.0625$ would lead to mAP of 0.09, 0.02 respectively. The latter phenomenon can be explained by vanishing gradients, whereas the former behavior is likely due to the training noise $\eta(t)$ (§4.3) being too large, when input variance is too large, the training would continue exploring suboptimal regions of the parameter space. Previous works in Table 1 reported that data augmentation tricks improve performance across all models/procedures. Table 4 shows their specific impact after ablation. Note that Mixup, TimeSpecAug shift the input distribution’s mean and variance, the higher the mixup coefficient/the longer time/freq mask \rightarrow the lower variance of the input \rightarrow more influence on end performance. We find data balancing helps, but due to the stochasticity(random sampling/shuffling) of the vanilla (no-balancing) train loader, its influence also fluctuates.

Table 4: Data Augmentation’s influence on mAP through ablation, on all models trained with A4 procedure

Augmentation	dataBalancing	Mixup	SpecAug	Normalize*
mAP drop(%)	3.9 ± 2.0	1.5 ± 0.8	1.5 ± 0.3	16.5 ± 9.0

Takeaways: Smaller features: $6 \times$ efficiency VS. 2% mAP loss. Models with Local+Global info \rightarrow efficient & best Performance Attention-based models: More # params, harder/more sensitive to train, sharp loss landscape, but more robust to noise[25]

5. References

- [1] H. Touvron, M. Cord, M. Douze, F. Massa, A. Sablayrolles, and H. Jégou, "Training data-efficient image transformers & distillation through attention," in *International Conference on Machine Learning*. PMLR, 2021, pp. 10 347–10 357.
- [2] Y. Gong, Y.-A. Chung, and J. Glass, "Ast: Audio spectrogram transformer," *arXiv preprint arXiv:2104.01778*, 2021.
- [3] J. Li, W. Dai, F. Metzger, S. Qu, and S. Das, "A comparison of deep learning methods for environmental sound detection," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 126–130.
- [4] J. F. Gemmeke, D. P. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, M. Plakal, and M. Ritter, "Audio set: An ontology and human-labeled dataset for audio events," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 776–780.
- [5] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, ser. NIPS'17. USA: Curran Associates Inc., 2017, pp. 6000–6010. [Online]. Available: <http://dl.acm.org/citation.cfm?id=3295222.3295349>
- [6] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," *arXiv preprint arXiv:2010.11929*, 2020.
- [7] Y. Gong, Y.-A. Chung, and J. Glass, "Psla: Improving audio tagging with pretraining, sampling, labeling, and aggregation," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3292–3306, 2021.
- [8] Y. Wang, "Polyphonic sound event detection with weak labeling," *Thesis*, 2018.
- [9] L. Ford, H. Tang, F. Grondin, and J. Glass, "A Deep Residual Network for Large-Scale Acoustic Scene Analysis," in *Proc. Interspeech 2019*, 2019, pp. 2568–2572. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2019-2731>
- [10] Q. Kong, Y. Cao, T. Iqbal, Y. Wang, W. Wang, and M. D. Plumbley, "Panns: Large-scale pretrained audio neural networks for audio pattern recognition," *arXiv preprint arXiv:1912.10211*, 2019.
- [11] Y. Wang, J. Li, and F. Metzger, "A comparison of five multiple instance learning pooling functions for sound event detection with weak labeling," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 31–35.
- [12] S. Srivastava, Y. Wang, A. Tjandra, A. Kumar, C. Liu, K. Singh, and Y. Saraf, "Conformer-based self-supervised learning for non-speech audio tasks," *arXiv preprint arXiv:2110.07313*, 2021.
- [13] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, "SpecAugment: A simple data augmentation method for automatic speech recognition," *ArXiv*, vol. abs/1904.08779, 2019.
- [14] K. Koutini, J. Schlüter, H. Eghbal-zadeh, and G. Widmer, "Efficient training of audio transformers with patchout," 2021.
- [15] S. Verbitskiy, V. Berikov, and V. Vyshegorodtsev, "Eranns: Efficient residual audio neural networks for audio pattern recognition," *arXiv preprint arXiv:2106.01621*, 2021.
- [16] L. N. Smith and N. Topin, "Super-convergence: Very fast training of neural networks using large learning rates," in *Artificial intelligence and machine learning for multi-domain operations applications*, vol. 11006. International Society for Optics and Photonics, 2019, p. 1100612.
- [17] S. Hershey, D. P. Ellis, E. Fonseca, A. Jansen, C. Liu, R. C. Moore, and M. Plakal, "The benefit of temporally-strong labels in audio event classification," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 366–370.
- [18] S. Hershey, S. Chaudhuri, D. P. W. Ellis, J. F. Gemmeke, A. Jansen, C. Moore, M. Plakal, D. Platt, R. A. Saurous, B. Seybold, M. Slaney, R. Weiss, and K. Wilson, "Cnn architectures for large-scale audio classification," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017. [Online]. Available: <https://arxiv.org/abs/1609.09430>
- [19] P. Izmailov, D. Podoprikin, T. Garipov, D. Vetrov, and A. G. Wilson, "Averaging weights leads to wider optima and better generalization," *arXiv preprint arXiv:1803.05407*, 2018.
- [20] R. Wightman, H. Touvron, and H. Jégou, "Resnet strikes back: An improved training procedure in timm," *arXiv preprint arXiv:2110.00476*, 2021.
- [21] K. He, R. Girshick, and P. Dollár, "Rethinking imagenet pre-training," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 4918–4927.
- [22] S. L. Smith, P.-J. Kindermans, C. Ying, and Q. V. Le, "Don't decay the learning rate, increase the batch size," in *International Conference on Learning Representations*, 2018.
- [23] X. Chen, C.-J. Hsieh, and B. Gong, "When vision transformers outperform resnets without pre-training or strong data augmentations," in *International Conference on Learning Representations*, 2022. [Online]. Available: <https://openreview.net/forum?id=LtKcMgGOeLt>
- [24] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," *arXiv preprint arXiv:1711.05101*, 2017.
- [25] J. B. Li, K. Ma, S. Qu, P.-Y. Huang, and F. Metzger, "Audio-visual event recognition through the lens of adversary," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 616–620.