



ON THE CHOICE OF THE OPTIMAL TEMPORAL SUPPORT FOR AUDIO CLASSIFICATION WITH PRE-TRAINED EMBEDDINGS

Aurian Quelenec, Michel Olvera, Geoffroy Peeters, Slim Essid

► To cite this version:

Aurian Quelenec, Michel Olvera, Geoffroy Peeters, Slim Essid. ON THE CHOICE OF THE OPTIMAL TEMPORAL SUPPORT FOR AUDIO CLASSIFICATION WITH PRE-TRAINED EMBEDDINGS. ICASSP, IEEE, Apr 2024, Séoul, South Korea. hal-04360221

HAL Id: hal-04360221

<https://hal.science/hal-04360221>

Submitted on 21 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

ON THE CHOICE OF THE OPTIMAL TEMPORAL SUPPORT FOR AUDIO CLASSIFICATION WITH PRE-TRAINED EMBEDDINGS

Aurian Queleennec, Michel Olvera, Geoffroy Peeters, Slim Essid

LTCI - Télécom Paris, Institut Polytechnique de Paris

© 2023 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

ABSTRACT

Current state-of-the-art audio analysis systems rely on pre-trained embedding models, often used off-the-shelf as (frozen) feature extractors. Choosing the best one for a set of tasks is the subject of many recent publications. However, one aspect often overlooked in these works is the influence of the duration of audio input considered to extract an embedding, which we refer to as Temporal Support (TS). In this work, we study the influence of the TS for well-established or emerging pre-trained embeddings, chosen to represent different types of architectures and learning paradigms. We conduct this evaluation using both musical instrument and environmental sound datasets, namely OpenMIC, TAU Urban Acoustic Scenes 2020 Mobile, and ESC-50. We especially highlight that Audio Spectrogram Transformer-based systems (PaSST and BEATs) remain effective with smaller TS, which therefore allows for a drastic reduction in memory and computational cost. Moreover, we show that by choosing the optimal TS we reach competitive results across all tasks. In particular, we improve the state-of-the-art results on OpenMIC, using BEATs and PaSST without any fine-tuning.

Index Terms— audio embeddings, acoustic scene classification, instrument recognition, temporal support, transformers

1. INTRODUCTION

In the audio processing domain, the lack of vast task-specific annotated datasets has been a persistent challenge. This scarcity has led to the development of models that are pre-trained on large-scale and diverse audio datasets, enabling them to generate very effective representations [1, 2]. These pre-trained models are subsequently either fine-tuned or used as frozen feature extractors when applied to smaller task-specific datasets [3, 4, 5].

Notably, pre-trained audio models are typically trained using either Supervised Learning (SL) [6, 7, 8] or Self Supervised Learning (SSL) techniques [9, 10, 11, 12]. While Convolution Neural Network (CNN) architectures were once the prevailing choice, the newer Audio Spectrogram Transformer (AST) architectures [13] have proven to be more effective than CNNs in numerous scenarios. To assess the effectiveness and the generalization ability of such models across diverse downstream tasks, various benchmarking frameworks and challenges have been proposed [14, 15, 16, 9], aiding in the selection of the most suitable frozen pre-trained model to develop audio analysis systems.

However, one aspect that has not been systematically addressed in model evaluation is the impact of the duration of audio input considered for the extraction of an embedding, which we refer to as Temporal Support (TS). Remarkably, the subject of TS has received limited consideration in current research, with only one study, to our knowledge, addressing it [17]. In this study, the authors show that altering the TS of a PaSST model from 160 ms to 640 ms improves the results on the DCASE 2016 Task 2 dataset [18]. Shedding light on the influence of the TS in model performance, we conduct an in-depth evaluation of a selection of well-established and current state-of-the-art pre-trained models.

We selected the models to study based on four criteria: i) data used for training: we chose models that were pre-trained using AudioSet [1], a dataset rich in both musical content and environmental sounds; ii) training paradigm: we considered SL vs SSL approaches; iii) type or architecture: CNN vs AST variants; and iv) levels of complexity: choosing both lightweight and computationally intensive models. This resulted in the selection of BYOL-A [9], PaSST [8], and BEATs [10]. Needless to say, the selection was also motivated by the recent success of these models in various audio classification tasks.

Our evaluation of these frozen pre-trained audio embeddings extends across three distinct audio classification tasks: Instrument Recognition (IR) using OpenMIC [19], Audio Scene Classification (ASC) employing TAU Urban Acoustic Scenes 2020 Mobile [20], and Environmental Sound Classification (ESC) utilizing ESC-50 [21]. We chose these datasets as reference results exist for them in previous works with most of the selected embeddings [17, 15, 9]. We adopted a simple and uniform classification protocol across these datasets for all pre-trained models to ensure a fair and robust basis for comparative analysis.

The contributions of this paper are as follows. Firstly, we provide valuable guidelines for the selection of pre-trained audio embeddings with regard to the influence of TS. Notably, we highlight the efficacy of AST models such as PaSST and BEATs, even when operating with smaller TS across all tasks. Reducing TS substantially decreases memory and computational overhead at inference time. Secondly, through the selection of the optimal TS, we improve the state-of-the-art results on the OpenMIC dataset using BEATs and PaSST without resorting to any fine-tuning procedure. Finally, we explored the utility of intermediate layer representations when using PaSST and BEATs, enhancing overall results.

2. METHOD

The aim is to use a frozen pre-trained model, denoted as $f(\cdot)$, to project audio instances into embeddings. These embeddings are then used to make class predictions using a simple classification probe denoted as $g(\cdot)$.

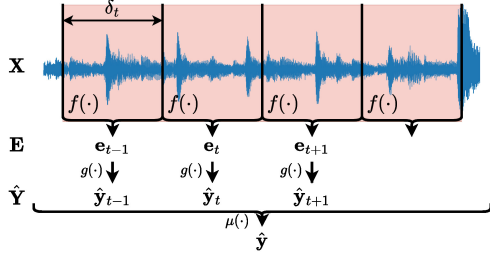


Fig. 1: Projection of an audio instance \mathbf{X} into an embedding \mathbf{E} with pre-trained model $f(\cdot)$, using a Temporal Support δ_t . This is followed by the label prediction facilitated by a classification probe $g(\cdot)$ and aggregation operation $\mu(\cdot)$.

2.1. Embedding projection using $f(\cdot)$

Given an audio instance $\mathbf{X} = (x_n)_{n \in \{1, \dots, N\}}$, where N represents the number of audio samples, we project \mathbf{X} into an embedding space using the pre-trained model $f(\cdot)$. This results in an embedding sequence $\mathbf{E} = (\mathbf{e}_t)_{t \in \{1, \dots, T\}}$, where T corresponds to the number of embedding vectors, and $\mathbf{e}_t \in \mathbb{R}^m$ with m indicating the dimensionality of the embedding. Typically, generating a single embedding entails feeding a segment of \mathbf{X} as input of $f(\cdot)$. We denote by Temporal Support, δ_t , the duration (in seconds) of this segment. The number of embeddings T is a function of N , and the value of δ_t . Precisely, T can be calculated as follows: $T = \lfloor \frac{N}{\delta_t \times s_r} \rfloor$, where s_r is the sample rate of \mathbf{X} . To illustrate, if we consider $\delta_t = 1$ and an audio instance of 10 s, we obtain 10 embeddings $(\mathbf{e}_t)_{t \in \{1, \dots, 10\}}$. For a given embedding system, δ_t should not be chosen to be too small and is always assumed to be significantly greater than the length of the analysis window used to compute the input representation (typically a spectrogram).

2.2. Label prediction using $g(\cdot)$

To predict labels from the embedding representation \mathbf{E} , we employ a classification probe $g(\cdot)$ resulting in $\hat{\mathbf{Y}} = g(\mathbf{E}) = (\hat{y}_t)_{t \in \{1, \dots, T\}}$. This yields a label prediction \hat{y}_t for each time step of the embedding representation. Yet, we place ourselves in a so-called *weakly labeled data* setting where for each audio instance \mathbf{X} we only have a global label $\mathbf{y} \in \{0, 1\}^L$ with L representing the number of classes. Here, $y_l = 1$ indicates the presence of the class l in the audio \mathbf{X} . Since the projection head yields a label prediction for each time step of \mathbf{E} , it is necessary to aggregate these “local” predictions \hat{y}_t to get the final estimated label $\hat{\mathbf{y}}$, as illustrated in Figure 1. We consider two alternative aggregation methods: a learnable aggregation operation, commonly referred to as attention [22], and simple average pooling, also known as mean temporal integration. We denote the aggregation function by $\mu(\cdot)$, such that $\hat{\mathbf{y}} = \mu(g(\mathbf{E})) \in [0, 1]^L$, where $\hat{\mathbf{y}}$ is the final clip-level label prediction for an instance \mathbf{X} . Specifically, we refer to $\mu(\cdot)$ as $\mu_m(\cdot)$ for mean-based aggregation and $\mu_a(\cdot)$ for attention-based aggregation.

3. PRE-TRAINED EMBEDDINGS $f(\cdot)$

The selection of the following embeddings was guided by the criteria outlined in the introduction.

BYOL-A [9] is a SSL model that extends the BYOL [24] method tailored to learn representations that capture foreground acoustic events and sound texture details. Its encoder relies on a CNN architecture

trained on AudioSet. The training data consists of 1-second long audio segments, corresponding to $\delta_t = 1$.

PaSST [8], falls under the AST category [13] and draws inspiration from the popular Vision Transformers (ViT) models [25]. Unlike BYOL-A, PaSST is trained in a supervised fashion on AudioSet. To reduce memory and computational costs, it employs patchout masking during training, enabling classification with incomplete patch sequences. The training segments in this case are 10 seconds long, *i.e.* $\delta_t = 10$ s.

BEATs [10] similarly to PaSST, adopts an AST architecture, but uses a SSL iterative training procedure. Optimization involves a tokenizer model for generating pseudo-labels and knowledge distillation. BEATs is also trained on the AudioSet dataset using full-length audio instance, hence with $\delta_t = 10$ s.

It is important to bear in mind that, regardless of the TS that was used for training the embedding networks, one may always exploit them to extract embeddings with a different temporal support. To clarify this with an example: although BEATs was initially trained with $\delta_t = 10$, we are interested in examining its behavior when, for a specific downstream classification task, we opt for a different temporal support (*e.g.*, $\delta_t = 3$) for extracting these embeddings. These embeddings will then serve as input features for training the classifier.

4. EXPERIMENTAL SETUP

4.1. Datasets

We employ three datasets for our study: OpenMIC for musical instrument recognition, TAU Urban for audio scene classification, and ESC-50 for environmental sound classification, which are described in the following.

OpenMIC [19] is sourced from the Free Music Archive [26], featuring a broad and diverse set of musical genres and instruments. This dataset comprises 20 possible instrument classes, with each of the 20,000 10-second audio clip annotated with at least one instrument label at the clip-level, making it suitable for multi-label classification. In our experiments, we use the provided training and testing splits and further partition the training set, keeping 15% of the samples for validation purposes as in [5, 27]. To handle missing labels within OpenMIC, we compute the loss exclusively for the observed labels, following the approach outlined in [5, 27]. For evaluation, we report the instrument-wise Mean Average Precision (mAP) scores, as done in prior works such as [27, 8, 23].

TAU Urban Acoustic Scenes 2020 Mobile [20] This dataset comprises 10-second audio recordings captured in various urban environments from 12 European cities. It covers the following scenes: airport, shopping mall, metro station, pedestrian street, public square, street traffic, tram, bus, metro, and park. Designed for multi-class classification, it contains a total of 23,040 samples. We used the official training and testing splits and further divided the training set by reserving 30% of the samples for validation following the approach used in [8]. We use classification accuracy as a metric for this dataset.

ESC-50 [21] is a dataset intended for multi-class audio classification of environmental sounds. It consists of 2,000 5-second audio clips categorized into 50 classes with 40 examples per class. These classes can be grouped into five broader categories: Animals, Natural soundscapes\water sounds, Human\non-speech sound, Interior\Domestic sounds, and Exterior\urban noises. The audio samples are sourced from the Freesound.org project. We use a 5-fold cross-validation setting as it is customary for this dataset [14]. We use classification accuracy as a metric for this dataset.

Model	δ_t	OpenMIC		TAU Urban		ESC-50		Emb. Size	#Param. $f(\cdot)$
		$\mu_m(\cdot)$	$\mu_a(\cdot)$	$\mu_m(\cdot)$	$\mu_a(\cdot)$	$\mu_m(\cdot)$	$\mu_a(\cdot)$		
BYOL-A v2	1	0.792 ± 0.001	0.797 ± 0.003	52.5 ± 1.4	50.6 ± 1.7	69.1 ± 1.4	68.7 ± 1.1	3072	6.3M
PaSST		0.851 ± 0.001	0.860 ± 0.002	63.3 ± 0.4	62.0 ± 0.5	93.1 ± 0.2	93.0 ± 0.4	768	87M
BEATs		0.852 ± 0.001	0.865 ± 0.001	67.5 ± 0.2	61.0 ± 4.3	93.2 ± 0.1	93.4 ± 0.4	$48 \cdot 768$	90M
BYOL-A v2	3	0.805 ± 0.001	0.804 ± 0.005	53.9 ± 0.9	52.3 ± 0.9	71.2 ± 1.1	72.6 ± 1.0	3072	6.3M
PaSST		0.866 ± 0.001	0.865 ± 0.000	65.0 ± 0.4	64.5 ± 0.5	95.7 ± 0.1	95.0 ± 0.1	768	87M
BEATs		0.862 ± 0.000	0.866 ± 0.002	66.8 ± 0.2	64.9 ± 1.4	95.4 ± 0.1	93.4 ± 0.3	$144 \cdot 768$	90M
BYOL-A v2	5	0.806 ± 0.002	0.808 ± 0.003	53.8 ± 1.1	53.6 ± 0.9	72.8 ± 1.8	74.0 ± 1.1	3072	6.3M
PaSST		0.866 ± 0.001	0.868 ± 0.001	66.5 ± 0.5	65.9 ± 1.0	96.8 ± 0.2	96.6 ± 0.2	768	87M
BEATs		0.869 ± 0.002	0.869 ± 0.001	67.5 ± 0.2	65.4 ± 2.6	96.1 ± 0.0	95.7 ± 0.3	$248 \cdot 768$	90M
BYOL-A v2	10	0.803 ± 0.001	0.805 ± 0.002	52.4 ± 1.5	54.7 ± 0.8	-	-	3072	6.3M
PaSST		0.861 ± 0.001	0.857 ± 0.001	66.7 ± 0.5	66.9 ± 0.4	-	-	768	87M
BEATs		0.866 ± 0.000	0.867 ± 0.000	67.5 ± 0.3	67.2 ± 1.1	-	-	$496 \cdot 768$	90M
Results from papers									
ResAtt [23]	10	0.860	-	-	-	-	-	2048	-
PaSST-S [8]	10/5	0.843	-	75.6	-	96.8	-	768	87M
BEAT _{iter3+} [10]	5	-	-	-	-	98.1	-	$248 \cdot 768$	90M

Table 1: Results for frozen pre-trained embeddings with varying δ_t and two temporal aggregations $\mu_m(\cdot)$ and $\mu_a(\cdot)$. Metrics include mAP for OpenMIC, and accuracy for TAU Urban and ESC-50.

4.2. Training Setup

We applied the methodology described in Section 2 to each dataset employing different pre-trained models $f(\cdot)$ and varying TS. For $g(\cdot)$, we used a linear layer to map the embedding to the number of classes of the task, followed by either $\mu_m(\cdot)$ or $\mu_a(\cdot)$ for temporal aggregation of predictions. We selected $\mu_a(\cdot)$ as described in [5] for OpenMIC, and we utilized the attention mechanism commonly described in the DCASE Challenge [28] for TAU Urban and ESC-50. For training the classification probe, we used the Adam optimizer, a learning rate of $1e^{-4}$ with a binary cross entropy loss for the IR task, and a learning rate¹ of $1e^{-3}$ for ASC and ESC, both using cross entropy loss. We used a batch size of 128 for IR and a batch size of 32 for the other tasks. While batch size differences were not intended, our preliminary experiments revealed no performance changes linked to this variation. We kept the model with the best validation loss. Each experiment was conducted 5 times, and we report the averaged results along with the standard deviation. Regarding the pre-trained models, we used publicly available weights: BYOL-A version 2 for BYOL-A, PaSST-S for PaSST, and *iter3+* weights for BEATs.

5. RESULTS

In this section, we conduct a performance evaluation of the selected embeddings. To this end, we extract embeddings with audio segments of varying durations: $\delta_t = 1$ s, $\delta_t = 3$ s, $\delta_t = 5$ s and $\delta_t = 10$ s, denoted by δ_1 , δ_3 , δ_5 and δ_{10} , respectively. The comprehensive results of this evaluation are presented in Table 1.

5.1. Influence of the Temporal Support

First, it is evident that pre-trained AST models consistently outperform BYOL-A across all δ_t configurations, which is to be expected due to the inherent differences in the number of parameters between

¹We used a fixed learning rate for every classification probe as a strategic trade-off. While it may not optimize performance for every embedding type, it strikes a balance that ensures fairness and simplifies comparisons.

these models. A more interesting finding is that pre-trained models are not necessarily more effective when used with the largest possible δ_t . For instance, all models achieve the best performance on the OpenMIC dataset for δ_5 . BEATs behaves similarly on the TAU Urban dataset, since results for δ_1 , δ_5 and δ_{10} yielded the same results. While BYOL-A achieves its highest score for δ_3 on TAU Urban, PaSST achieves its best scores when using either δ_5 or δ_{10} . Conversely, for ESC-50, all pre-trained models achieve the best results for the largest possible temporal support, which is δ_5 . Overall, this highlights the importance of selecting the optimal δ_t value for each pre-trained model, depending on the specific task, in order to attain the best possible performance. This is of paramount importance for pre-trained AST models, given the computational cost associated with larger TS. Indeed, let L denote the sequence length of an attention layer within an AST model. It’s worth noting that the computation complexity grows quadratically, specifically in $\mathcal{O}(L^2)$ as the sequence length L increases. Consequently, opting for a smaller TS implies processing sequences of much smaller length in each attention layer, which reduces drastically the computational burden during inference. For instance, one should extract embeddings with δ_1 rather than δ_{10} with BEATs on TAU Urban due to the significantly reduced computational inference cost while maintaining equivalent performance. Furthermore, it is noteworthy to highlight that we improve the state-of-the-art performance on OpenMIC with BEATs and PaSST for every value of δ_t (with the exception of δ_{10} for PaSST) without any fine-tuning effort.

Discussion. When it comes to AST models, adjusting their TS can be likened to altering the receptive field of a CNN. Indeed, in CNNs the receptive field is defined as the input area that contributes to the resulting embedding at a given position. In contrast, AST models rely on the attention mechanism, where the entire input contributes to the final embedding (especially on short sequences as it is the case here), effectively determined by the input’s length, *i.e.* δ_t . Remarkably, prior research [27] demonstrated that regularizing the receptive field of a CNN for a specific task can significantly improve performance. Therefore, selecting the optimal TS for pre-trained AST models may be as advantageous as fine-tuning the receptive field of a CNN model for a given task.

5.2. Aggregation method

Our analysis of Table 1 reveals two distinct trends regarding the performance of the aggregation function $\mu(\cdot)$. For the IR task, consistent with [5], we observe that $\mu_a(\cdot)$ is more effective, particularly for small δ_t values, which results in an increased number of time steps in \hat{Y} . In contrast, for ASC and ESC tasks, $\mu_m(\cdot)$ aggregation has the ascendancy in most cases. There are few exceptions, notably BYOL-A with δ_{10} on TAU Urban, and δ_3 and δ_5 on ESC-50. This is an interesting finding, as this aggregation approach is far more efficient, which was not expected. Indeed, given that the dynamics of target sound classes vary (regardless of the dataset), with some class activations being localized in time while other classes being active for longer periods, we expected to see $\mu_a(\cdot)$ be the best aggregation overall. Even if we would need to further investigate the causes that can explain this behavior, we believe that in the case of the multi-class nature of the problems, where events overlap in time, yet the model yields only one class, the training of the attention mechanism may be hindered, especially in this regime of rather small datasets.

5.3. Comparing results to previous works

With equal TS, δ_{10} , and without fine-tuning, we achieve a higher score on OpenMIC with PaSST compared to [8]. This shows that fine-tuning large pre-trained models on smaller datasets may not always be beneficial. This phenomenon is akin to *forgetting* which may limit the generalization ability of the fine-tuned variant compared to the original one. On the other hand, on the TAU Urban dataset, we obtain scores with PaSST that are below those obtained in [8], where fine-tuning is used. For BYOL-A, our results on ESC-50 deviate from those reported in [9] and [15]. This divergence can be partly attributed to differences in methodology. In [9], the authors normalize the embeddings before training the probe $g(\cdot)$ and optimize the learning rate to maximize the scores. In [15], the authors employed a meticulously tuned support vector machine classifier in conjunction with data augmentation, which is expected to be more effective than the shallow probe we use across all tasks. Such improvements and careful optimizations are beyond the scope of this study whose sole focus is on the influence of the temporal support with a neutral, lightweight downstream classification strategy kept fixed across all tasks.

5.4. Exploiting outputs of all embedding layers

This last experiment looks into exploiting the outputs of all intermediate layers of the embedding network as opposed to solely relying on the last one. Indeed, different types of information are encoded across these layers which may be beneficial to the classification task. This practice aligns with common conventions in the

Speech SSL domain, where benchmarking pre-trained embedding models often involves leveraging multiple layers. In particular, the SUPERB benchmark [16] suggests using a weighted sum of the outputs from each layer as the final embedding to be exploited by the downstream classifier $g(\cdot)$, using $z = \sum_{l=1}^{12} \alpha_l z^{(l)}$, where $z^{(l)}$ is the output of layer l . The weights α_l are trained jointly with the downstream classifier, while the representations $z^{(l)}$ are merely the intermediate outputs of the embedding network without fine-tuning. This is because fine-tuning such networks is costly and often sub-optimal, with catastrophic forgetting easily compromising performance, as explained earlier.

Hence, we explore the impact of combining the outputs from all layers of the pre-trained AST models, PaSST and BEATs, across the three datasets for different TS. This yields a total of 11 intermediate representations along with the final layer embedding that are used (as a weighted combination) as the final embedding for both PaSST and BEATs. We select the most effective temporal aggregation method for each dataset, namely $\mu_a(\cdot)$ for OpenMIC and $\mu_m(\cdot)$ for TAU Urban and ESC-50. The results of this analysis are presented in Table 2. In general, we observe only marginal performance improvements when leveraging outputs from all layers considering the augmentation of the dimensionality of the used features, with no noticeable influence of TS on the results. Nevertheless, this approach enables us to improve the state-of-the-art mAP score on OpenMIC, achieving a remarkable 0.87, and outperforming [8] on ESC-50 without the need of fine-tuning the pre-trained model.

6. CONCLUSION

In this study, we have conducted a comprehensive experimental analysis to investigate the influence of Temporal Support on audio embeddings extracted with BYOL-A, PaSST, and BEATs, across tasks including Instrument Recognition, Audio Scene Classification, and Environmental Sound Classification. Our findings underscore the crucial importance of selecting the optimal Temporal Support based on both the pre-trained model and the task at hand. This is particularly compelling as operating Audio Spectrogram Transformer pre-trained models with smaller Temporal Support significantly reduces their inference computational cost. With adequate Temporal Support, these pre-trained models obtain competitive results across all tasks, all without the need for fine-tuning. Moreover, we improve the state-of-the-art result on the OpenMIC dataset. Notably, by using a weighted combination encompassing all layer outputs from PaSST and BEATs, we further improved over our prior results. Future works could delve into understanding how the weights of the last layers correlate with the Temporal Support as well as explore more pre-trained models and tasks.

7. ACKNOWLEDGEMENT

This work was supported by the Audible project, funded by French BPI. This work was performed using HPC resources from GENCI-IDRIS.

8. REFERENCES

- [1] Jort F Gemmeke, Daniel PW 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 *Proc. ICASSP*. IEEE, 2017, pp. 776–780.

Model	δ_t	OpenMIC	TAU Urban	ESC-50
PaSST	1	0.860 \pm 0.001	63.0 \pm 0.4	93.3 \pm 0.2
	3	0.868 \pm 0.001	65.2 \pm 0.4	96.2 \pm 0.0
	5	0.869 \pm 0.002	67.1 \pm 0.6	97.0 \pm 0.0
	10	0.868 \pm 0.002	67.9 \pm 0.4	-
BEATs	1	0.864 \pm 0.002	67.2 \pm 0.2	93.6 \pm 0.1
	3	0.867 \pm 0.001	67.4 \pm 0.3	95.4 \pm 0.1
	5	0.870 \pm 0.002	68.5 \pm 0.3	96.0 \pm 0.1
	10	0.870 \pm 0.002	68.2 \pm 0.3	-

Table 2: Results for PaSST and BEATs using a learned weighted sum of the output of all layers.

- [2] Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushal Lakhotia, Ruslan Salakhutdinov, and Abdelrahman Mohamed, "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE ACM Trans. Audio Speech Lang. Process.*, vol. 29, pp. 3451–3460, 2021.
- [3] Salah Zaiem, Robin Algayres, Titouan Parcollet, Slim Essid, and Mirco Ravanelli, "Fine-tuning strategies for faster inference using speech self-supervised models: A comparative study," in *Proc. ICASSP*. IEEE, 2023, pp. 1–5.
- [4] Florian Angulo, Slim Essid, Geoffroy Peeters, and Christophe Mietlicki, "Cosmopolite sound monitoring (cosmo): A study of urban sound event detection systems generalizing to multiple cities," in *Proc. ICASSP*. IEEE, 2023, pp. 1–5.
- [5] Siddharth Gururani, Mohit Sharma, and Alexander Lerch, "An attention mechanism for musical instrument recognition," in *Proc. ISMIR*, 2019, pp. 83–90.
- [6] Shawn Hershey, Sourish Chaudhuri, Daniel PW Ellis, Jort F Gemmeke, Aren Jansen, R Channing Moore, Manoj Plakal, Devin Platt, Rif A Saurous, Bryan Seybold, et al., "Cnn architectures for large-scale audio classification," in *Proc. ICASSP*. IEEE, 2017, pp. 131–135.
- [7] 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.*, vol. 28, pp. 2880–2894, 2020.
- [8] Khaled Koutini, Jan Schlüter, Hamid Eghbal-zadeh, and Gerhard Widmer, "Efficient training of audio transformers with patchout," in *Proc. Interspeech*. IEEE, 2022, pp. 2753–2757.
- [9] Daisuke Niizumi, Daiki Takeuchi, Yasunori Ohishi, Noboru Harada, and Kunio Kashino, "Byol for audio: Exploring pre-trained general-purpose audio representations," *IEEE ACM Trans. Audio Speech Lang. Process.*, vol. 31, pp. 137–151, 2023.
- [10] Sanyuan Chen, Yu Wu, Chengyi Wang, Shujie Liu, Daniel Tompkins, Zhuo Chen, and Furu Wei, "Beats: Audio pre-training with acoustic tokenizers," in *Proc. ICML*. PMLR, 2022, vol. 202, pp. 5178–5193.
- [11] Aurora Linh Cramer, Ho-Hsiang Wu, Justin Salamon, and Juan Pablo Bello, "Look, listen, and learn more: Design choices for deep audio embeddings," in *Proc. IEEE ICASSP*. IEEE, 2019, pp. 3852–3856.
- [12] Aaqib Saeed, David Grangier, and Neil Zeghidour, "Contrastive learning of general-purpose audio representations," in *Proc. ICASSP*. IEEE, 2021, pp. 3875–3879.
- [13] Yuan Gong, Yu-An Chung, and James R. Glass, "AST: audio spectrogram transformer," in *Proc. Interspeech*. IEEE, 2021, pp. 571–575.
- [14] J. Turian, J. Shier, H. R. Khan, B. Raj, B. W. Schuller, C. J. Steinmetz, C. Malloy, G. Tzanetakis, G. Velarde, K. McNally, et al., "Hear: Holistic evaluation of audio representations," in *NeurIPS 2021 Competitions and Demonstrations Track*. PMLR, 2022, vol. 176, pp. 125–145.
- [15] Luyu Wang, Pauline Luc, Yan Wu, Adrià Recasens, Lucas Smaira, Andrew Brock, Andrew Jaegle, Jean-Baptiste Alayrac, Sander Dieleman, Joao Carreira, and Aaron van den Oord, "Towards learning universal audio representations," in *Proc. ICASSP*. IEEE, 2022, pp. 4593–4597.
- [16] Shu-Wen Yang, Po-Han Chi, Yung-Sung Chuang, Cheng-I Jeff Lai, Kushal Lakhotia, Yist Y. Lin, Andy T. Liu, Jiatong Shi, Xuankai Chang, Guan-Ting Lin, Tzu-Hsien Huang, Wei-Cheng Tseng, Ko-tik Lee, et al., "SUPERB: speech processing universal performance benchmark," in *Proc. INTERSPEECH*. ISCA, 2021, pp. 1194–1198.
- [17] Khaled Koutini, Shahed Masoudian, Florian Schmid, Hamid Eghbal-zadeh, Jan Schlüter, and Gerhard Widmer, "Learning general audio representations with large-scale training of patchout audio transformers," in *HEAR*. PMLR, 2021, vol. 166, pp. 65–89.
- [18] Annamaria Mesaros, Toni Heittola, Emmanouil Benetos, Peter Foster, Mathieu Lagrange, Tuomas Virtanen, and Mark D. Plumbley, "Detection and classification of acoustic scenes and events: Outcome of the dcase 2016 challenge," *IEEE ACM Trans. Audio Speech Lang. Process.*, vol. 26, pp. 379–393, 2018.
- [19] Eric Humphrey, Simon Durand, and Brian McFee, "Openmic-2018: An open data-set for multiple instrument recognition," in *Proc. ISMIR*, 2018, pp. 438–444.
- [20] Annamaria Mesaros, Toni Heittola, and Tuomas Virtanen, "A multi-device dataset for urban acoustic scene classification," in *Proc. DCASE*, 2018, pp. 9–13.
- [21] Karol J. Piczak, "ESC: dataset for environmental sound classification," in *Proc. Multimedia*. ACM, 2015, pp. 1015–1018.
- [22] Brian McFee, Justin Salamon, and Juan Pablo Bello, "Adaptive pooling operators for weakly labeled sound event detection," *IEEE ACM Trans. Audio Speech Lang. Process.*, vol. 26, pp. 2180–2193, Jul. 2018.
- [23] Zhi Zhong, Masato Hirano, Kazuki Shimada, Kazuya Tateishi, Shusuke Takahashi, and Yuki Mitsufuji, "An attention-based approach to hierarchical multi-label music instrument classification," in *Proc. ICASSP*. IEEE, 2023, pp. 1–5.
- [24] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H. Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Ávila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, and Michal Valko, "Bootstrap your own latent - A new approach to self-supervised learning," in *Proc. NeurIPS*, 2020, pp. 6–12.
- [25] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou, "Training data-efficient image transformers & distillation through attention," in *Proc. ICML*. PMLR, 2021, vol. 139, pp. 10347–10357.
- [26] Michaël Defferrard, Kirell Benzi, Pierre Vandergheynst, and Xavier Bresson, "FMA: A dataset for music analysis," in *Proc. ISMIR*, 2017, pp. 316–323.
- [27] 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.*, vol. 29, pp. 1987–2000, 2021.
- [28] Francesca Ronchini, Samuele Cornell, Romain Serizel, Nicolas Turpault, Eduardo Fonseca, and Daniel P. W. Ellis, "Description and analysis of novelties introduced in DCASE task 4 2022 on the baseline system," in *Proc. DCASE*, 2022, pp. 3–4.