



AxLSTMs: learning self-supervised audio representations with xLSTMs

Sarthak Yadav^{1,2}, Sergios Theodoridis^{1,3}, Zheng-Hua Tan^{1,2}

¹Department of Electronic Systems, Aalborg University, Denmark

²Pioneer Centre for Artificial Intelligence, Denmark

³Department of Informatics and Telecommunications, National and Kapodistrian University of Athens, Greece

sarthaky@es.aau.dk, sthe@es.aau.dk, zt@es.aau.dk

Abstract

While the transformer has emerged as the eminent neural architecture, several independent lines of research have emerged to address its limitations. Recurrent neural approaches have observed a lot of renewed interest, including the extended long short-term memory (xLSTM) architecture, which reinvigorates the original LSTM. However, while xLSTMs have shown competitive performance compared to the transformer, their viability for learning self-supervised general-purpose audio representations has not been evaluated. This work proposes Audio xLSTM (AxLSTM), an approach for learning audio representations from masked spectrogram patches in a self-supervised setting. Pretrained on the AudioSet dataset, the proposed AxLSTM models outperform comparable self-supervised audio spectrogram transformer (SSAST) baselines by up to 25% in relative performance across a set of ten diverse downstream tasks while having up to 45% fewer parameters.

Index Terms: xLSTM, self-supervised learning, audio representation learning

1. Introduction

In lieu of their excellent generalisation capabilities and domain and data agnostic nature, transformers [1] and their successors have seen widespread adoption. Furthermore, transformers, together with masked predictive modelling, have emerged as a key driving force behind several prominent advancements in the realm of unsupervised and self-supervised representation learning for NLP [2], computer vision [3] and audio and speech processing [4–8]. Recently, however, finding alternatives to scaled dot-product attention has garnered significant interest, with a lot of emphasis on finding sub-quadratic approximations of the attention operation [9] as well as token mixing [10].

Before the advent of transformers, sequence modeling was predominantly done using recurrent neural networks (RNNs). Recurrent models offer several advantages over transformers: they scale linearly with respect to sequence length and they have lower runtime memory requirements since storing the entire key-value (KV) cache is not necessary. State-space models (SSMs) [11], which are a family of sequence models that lie at the intersection of convolutional neural networks, RNNs and classical state spaces, are the most popular amongst recent recurrent approaches. Several variants of SSMs have been proposed, showing competitive performance and scalability versus transformers in several domains, including long sequence modelling [12], computer vision [13] as well as audio [14].

However, prior to transformers and state-space models, LSTMs [15] were the go to neural architecture for sequence modeling. Compared to transformers, LSTMs suffer from several key drawbacks: (i) inability to revise storage decisions, (ii)

compressing information into a scalar cell state, which impacts performance on rare input tokens, and (iii) memory mixing and the resulting lack of parallelizability. Recently, [16] proposed the extended long-short term memory (xLSTM) neural architecture, which revitalises the original LSTM architecture, leveraging the latest techniques and tricks learned from years of transformer and large language modelling research. xLSTM incorporates exponential gating, improved normalization and stabilization techniques while removing traditional memory mixing, resulting in two new fundamental building blocks: sLSTM and mLSTM. xLSTMs have demonstrated on-par or better performance than transformers for large language models [16] as well as patch-based image recognition [17], while also demonstrating superior sequence length extrapolation capabilities. However, the ability of xLSTMs to learn general-purpose audio representations in a self-supervised setting is yet to be thoroughly evaluated. In this work, we propose Audio xLSTMs (AxLSTMs) for learning self-supervised general-purpose audio representations within a masked modelling framework from spectrograms patches. Pretrained on AudioSet [18], AxLSTMs consistently outperform comparable self-supervised audio spectrogram transformer (SSAST) [6] based baselines on ten varied downstream tasks, even matching the performance of more recent and involved masked modeling approaches, such as BEATs [19]. Code and pretrained models are available at github.com/SarthakYadav/axlstm-official.

2. Method

2.1. Prerequisites: xLSTM and the mLSTM block

As previously discussed, xLSTM [16] proposes two new building blocks, sLSTM and mLSTM. In line with Vision-LSTM [17], the parallelizable mLSTM module is the fundamental building block of our approach, and thus is the focus of our discussion. The mLSTM block addresses the issues with the original LSTM cell by utilizing a matrix memory cell $C \in \mathbb{R}^{d \times d}$. mLSTM stores a key-value vector pair $k_t, v_t \in \mathbb{R}^d$ at timestep t , and the relevant value vector is retrieved later using a query $q_{t+\tau} \in \mathbb{R}^d$. Equations 1-9 describe the forward pass for mLSTM module, where C_t is the matrix memory cell, n_t and h_t represent the normalizer state and the hidden state, and i_t, f_t and o_t represent the input, forget and the output gates, respectively. Weights W_q, W_k and W_v are learnable projection matrices for vectors query q , key k and value v , respectively. Since there is no memory mixing in mLSTMs, multiple memory cells and multiple heads are equivalent, and the forward pass can be parallelized. Further, mLSTM uses exponential gating to mitigate the inability of LSTMs to revise storage decisions when a more similar input is found, which impacts retrieval performance. For more information, refer to [16].

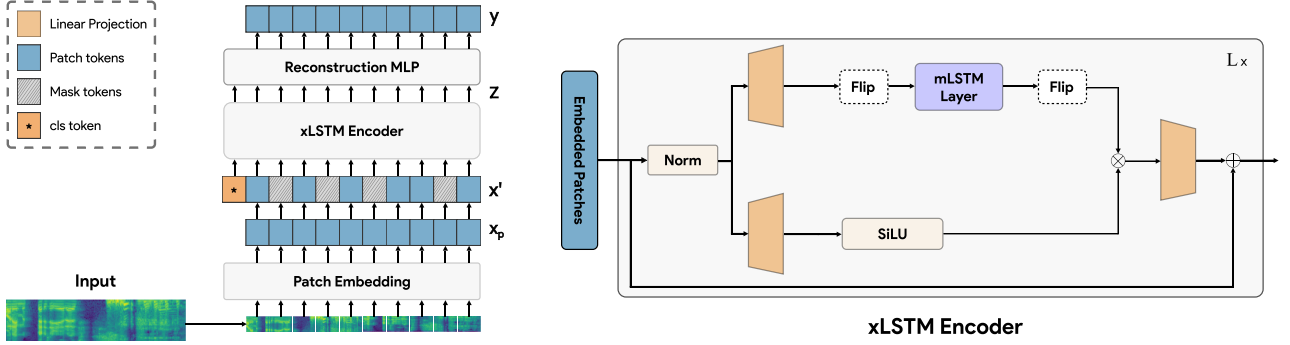


Figure 1: An overview of the proposed AxLSTM approach (left), and the constituent mLSTM blocks (right).

$$C_t = f_t C_{t-1} + i_t v_t k_t^\top, \quad (1)$$

$$n_t = f_t n_{t-1} + i_t k_t, \quad (2)$$

$$h_t = o_t \odot C_t q_t / \max\{|n_t^\top q_t|, 1\}, \quad (3)$$

$$q_t = W_q x_t + b_q \quad (4)$$

$$k_t = \frac{1}{\sqrt{d}} W_k x_t + b_k \quad (5)$$

$$v_t = W_v x_t + b_v \quad (6)$$

$$i_t = \exp(w_i^\top x_t + b_i), \quad (7)$$

$$f_t = \exp(w_f^\top x_t + b_f), \quad (8)$$

$$o_t = \sigma(W_o x_t + b_o) \quad (9)$$

2.2. AxLSTM: Modelling masked patches with xLSTMs

Creating and masking patches: Given an input spectrogram $\mathbf{x} \in \mathbb{R}^{T \times F}$, the axes denoting time and frequency, respectively, we compute non-overlapping patches of shape $t \times f$, yielding $\mathbf{x}_p \in \mathbb{R}^{N \times (t \cdot f)}$ patches. Similar to [7, 20], fixed sinusoidal 2d positional embeddings are added to the patches after projecting them to a $\mathbb{R}^{N \times d_m}$ dimensional space. To remain in line with previous work [6, 20] and to facilitate fine-tuning of the obtained models, a representative class token is then added to the beginning of the sequence. We then proceed to randomly mask 50% of the input patches using an unstructured masking strategy and replace these masked patches with a learnable *mask* token, similar to previous works [6, 14]. Thus, input to the encoder is:

$$\mathbf{x}' = [\text{cls}, \mathbf{x}_p^1, \mathbf{x}_p^2, \dots, \mathbf{x}_p^N] + E_{\text{pos}} \quad (10)$$

Encoding and Reconstruction: These partially masked patches are now fed to the encoder, yielding encoded representations $\mathbf{z} = \text{enc}(\mathbf{x}')$, $\mathbf{z} \in \mathbb{R}^{(N+1) \times d_m}$. The encoder is a stack of mLSTM blocks as shown in Fig. 1, where each block expands the d_m dimensional input by an expansion factor E_f (usually 2 [16, 17]) before projecting it back to d_m . The resulting blocks have fewer parameters than a comparable transformer block. To facilitate bidirectional modeling of the input, a *flip* operation that reverses the order of the input sequence is enabled in even numbered mLSTM blocks, similar to [17]. After encoding, a single hidden layer MLP is used to reconstruct patches from encoded representation \mathbf{z} :

$$\mathbf{y}' = \text{Linear}_{(t \cdot f)}(\sigma(\text{Linear}_{d_m}(\mathbf{z}))), \quad (11)$$

where Linear_d is a parameterized linear projection to dimensions d , and σ denotes the GELU non-linear activation function. Finally, the mean square error (MSE) between the original input spectrogram and the reconstructions, computed only for the masked patches, is used for pretraining. Further details can be found in Section 3.2.

Table 1: Overview of downstream tasks

ID	Name	Description	#Classes	#Hours
BO	Beijing Opera	percussion instrument classification	4	0.3
CD	Crema-D	emotion recognition	6	10
E50	ESC-50	environmental sound classification	50	2.77
LC	LibriCount	counting speakers (classification)	10	8
Mri-S	Mridangam Stroke	classifying Mridangam <i>strokes</i>	10	1.57
Mri-T	Mridangam Tonic	classifying Mridangam <i>tonics</i>	6	1.57
NS-5h	NSynth Pitch 5h	pitch classification	88	5.5
SC-5h	SpeechCommands 5h	keyword spotting	12	6.5
F50K	FSD50K	multilabel audio tagging	200	100
VL	VoxLingua107 Top10	spoken language identification	10	5

SSAST as a Baseline: SSAST [6] is the most directly comparable transformer-based approach to AxLSTMs. As the first patch-based self-supervised learning framework for audio spectrogram transformers (ASTs), it is a well-established method in audio representation learning. While newer transformer-based approaches such as masked autoencoders [7, 21, 22] and self-distilled tokenizers in BEATs [19] introduce additional architectural modifications, SSAST offers a modular framework for evaluating the modeling capabilities of xLSTM v/s the transformer without additional confounding factors. Unlike the original SSAST, which used a combined token prediction and reconstruction based pretraining objective, we simplify the approach by using a reconstruction-only objective.

3. Experiments

3.1. Datasets

Pretraining: For pretraining, we use the AudioSet dataset (AS) [18] which has roughly 2 million 10-second long weakly labeled YouTube clips and over 5000 hours of audio data spanning 527 classes.

Downstream Evaluation: Following recent work [14, 22], we use the following subset of tasks proposed as part of the HEAR benchmark [23] for downstream evaluation: Beijing Opera [23, 25], Crema-D [26], ESC-50 [27], LibriCount [28], Mridangam Stroke and Tonic [29], NSynth Pitch 5h [23, 30], Speech Commands 5h [23, 31], FSD50K [32] and VoxLingua107 [33]. Table 1 provides a brief overview of these tasks.

Table 2: Comparing AxLSTMs with popular self-supervised audio representations. We used pretrained models from cited papers to extract fixed feature vectors and conducted our own downstream experiments. LS, AS, VP, LL stand for LibriSpeech, AudioSet, VoxPopuli and LibriLight datasets, respectively. Original SSAST [6] was trained on AS+LS, whereas we pretrained the directly comparable underlined SSAST baselines (SSAST-Tiny, Small, Base). *includes decoder parameters

Model	Data	#M Params	Music & Pitch				Speech-based tasks				Audio		
			BO	Mri-S	Mri-T	NS-5h	CD	LC	SC-5h	VL	E50	F50K	$s(m)$
Supervised Baselines													
HEAR-Naive [23]	-	-	52.6 \pm 2.4	38.0 \pm 1.3	36.4 \pm 1.9	18.6 \pm 4.4	30.9 \pm 0.8	33.5 \pm 1.1	8.5 \pm 0.4	11.2 \pm 0.5	5.8 \pm 0.2	7.1 \pm 0.2	5.1 \pm 0.7
PaSST-Base [24]	AS	86	94.9 \pm 0.5	96.5 \pm 0.1	87.6 \pm 0.6	23.3 \pm 0.9	61.0 \pm 0.3	60.1 \pm 0.2	66.6 \pm 1.4	25.5 \pm 0.8	94.8\pm0.3	64.2\pm0.1	74.4 \pm 0.4
SSL													
W2V2-large [4]	VP	315.4	93.1 \pm 0.7	93.9 \pm 0.1	77.4 \pm 0.2	42.0 \pm 1.0	66.9 \pm 0.4	62.4 \pm 0.3	87.6 \pm 0.5	53.6 \pm 1.0	60.1 \pm 0.5	34.2 \pm 0.1	74.9 \pm 0.4
WavLM-large [8]	Mix	315.4	96.4\pm0.5	96.8 \pm 0.1	89.5 \pm 0.1	53.7 \pm 0.5	57.2 \pm 0.2	61.1 \pm 0.3	46.2 \pm 0.8	23.7 \pm 0.9	47.9 \pm 0.4	29.0 \pm 0.1	64.8 \pm 0.2
HuBERT-large [5]	LL	315.4	94.1 \pm 0.7	95.3 \pm 0.1	83.5 \pm 0.3	19.3 \pm 0.8	70.7 \pm 0.1	59.9 \pm 0.2	83.2 \pm 0.7	66.1\pm0.9	60.3 \pm 0.4	31.5 \pm 0.1	74.3 \pm 0.3
BEATs-iter3 [19]	AS	90.0	94.0 \pm 0.8	94.7 \pm 0.1	95.8 \pm 0.1	69.4 \pm 0.8	67.3 \pm 0.2	68.0 \pm 0.2	85.2 \pm 0.3	38.5 \pm 1.0	83.7 \pm 0.3	53.6 \pm 0.2	86.7 \pm 0.3
AudioMAE [7]	AS	86.0	93.7 \pm 0.6	89.2 \pm 0.2	86.6 \pm 0.2	64.5 \pm 0.8	68.2 \pm 0.2	42.2 \pm 0.2	28.6 \pm 1.5	29.7 \pm 1.0	60.6 \pm 0.4	37.9 \pm 0.1	63.6 \pm 0.3
MSM-MAE-208 [21]	AS	92.7*	95.7 \pm 0.7	97.3 \pm 0.1	97.9\pm0.1	69.1 \pm 0.5	68.7 \pm 0.2	63.8 \pm 0.5	85.7 \pm 0.3	40.3 \pm 0.6	78.4 \pm 0.6	49.5 \pm 0.1	86.0 \pm 0.2
MWMAE-Tiny [22]	AS	12.6*	93.3 \pm 1.0	97.1 \pm 0.1	97.6 \pm 0.1	68.1 \pm 0.4	64.4 \pm 0.2	65.5 \pm 0.3	77.0 \pm 0.6	28.6 \pm 1.1	71.9 \pm 0.5	43.4 \pm 0.1	80.0 \pm 0.3
MWMAE-Base [22]	AS	92.5*	96.0 \pm 0.5	97.4 \pm 0.1	97.9\pm0.1	69.3 \pm 0.6	73.1\pm0.3	68.8\pm0.2	90.9\pm0.2	44.2 \pm 0.9	81.2 \pm 0.4	51.2 \pm 0.2	90.3\pm0.2
SSAM-Tiny [14]	AS	4.8	93.7 \pm 0.8	97.1 \pm 0.1	94.9 \pm 0.1	62.0 \pm 0.7	61.8 \pm 0.3	59.2 \pm 0.4	74.8 \pm 0.4	27.8 \pm 1.0	70.6 \pm 0.2	41.3 \pm 0.2	75.6 \pm 0.2
SSAM-Small [14]	AS	17.9	94.0 \pm 0.7	97.5 \pm 0.1	96.7 \pm 0.1	66.3 \pm 0.8	67.5 \pm 0.2	60.5 \pm 0.3	83.7 \pm 0.3	39.6 \pm 0.7	78.7 \pm 0.6	48.5 \pm 0.1	83.5 \pm 0.3
SSAM-Base [14]	AS	69.3	93.2 \pm 1.1	97.7\pm0.1	96.9 \pm 0.1	70.5 \pm 0.5	70.3 \pm 0.2	63.5 \pm 0.2	87.9 \pm 0.3	50.4 \pm 0.7	81.0 \pm 0.3	<u>52.2\pm0.1</u>	88.7 \pm 0.3
SSAST Based													
SSAST [6]	Mix	89.0	93.4 \pm 0.9	96.7 \pm 0.1	96.3 \pm 0.1	66.8 \pm 0.7	56.5 \pm 0.2	60.7 \pm 0.3	53.5 \pm 1.3	28.5 \pm 0.9	68.4 \pm 0.4	38.2 \pm 0.1	72.5 \pm 0.2
SSAST-Tiny	AS	5.4	90.4 \pm 0.7	95.7 \pm 0.1	94.3 \pm 0.1	61.2 \pm 0.5	46.9 \pm 0.2	42.7 \pm 0.2	50.6 \pm 1.6	13.8 \pm 1.0	42.4 \pm 0.6	24.6 \pm 0.1	55.6 \pm 0.2
<u>SSAST-Small</u>	AS	21.5	93.2 \pm 0.5	96.2 \pm 0.1	95.0 \pm 0.1	63.8 \pm 0.4	51.6 \pm 0.2	50.0 \pm 0.3	58.3 \pm 1.2	15.6 \pm 0.7	50.1 \pm 0.6	31.6 \pm 0.1	63.0 \pm 0.2
<u>SSAST-Base</u>	AS	85.7	93.1 \pm 0.7	96.6 \pm 0.1	96.2 \pm 0.2	64.6 \pm 0.8	56.0 \pm 0.4	52.9 \pm 0.3	66.1 \pm 1.0	19.2 \pm 0.9	59.6 \pm 0.7	37.5 \pm 0.1	68.7 \pm 0.3
Proposed													
AxLSTM-Tiny	AS	4.3	93.9 \pm 0.7	96.8 \pm 0.1	94.9 \pm 0.1	62.5 \pm 0.6	57.5 \pm 0.2	55.1 \pm 0.5	69.0 \pm 1.6	24.1 \pm 0.6	61.3 \pm 0.6	37.5 \pm 0.2	70.7 \pm 0.2
AxLSTM-Small	AS	16.7	92.9 \pm 1.0	97.4 \pm 0.1	96.6 \pm 0.1	66.6 \pm 0.4	65.0 \pm 0.2	60.3 \pm 0.3	80.5 \pm 0.4	36.5 \pm 0.7	75.5 \pm 0.4	46.5 \pm 0.1	81.1 \pm 0.3
AxLSTM-Base	AS	65.6	93.6 \pm 0.9	97.5 \pm 0.1	97.5 \pm 0.1	71.4 \pm 0.8	68.7 \pm 0.2	63.2 \pm 0.3	85.1 \pm 0.2	43.5 \pm 0.6	79.2 \pm 0.6	<u>51.0\pm0.1</u>	86.6 \pm 0.2

3.2. Implementation details

Spectrogram features: We extract log-scaled mel spectrogram features with a window size of 25 ms, a hop size of 10 ms and $F = 80$ mel-spaced frequency bins. A sampling rate of 16000 Hz is used for all audio clips.

Pretraining: All the proposed models are pretrained on randomly cropped 2-second audio clips, resulting in an input spectrogram of shape $[200 \times 80]$. Our default configuration consists of $l = 12$ number of stacked mLSTM blocks with a model feature dimension of $d_m = 192$, the same as those of a ViT-Tiny [20] encoder, but we also evaluate Small ($d_m = 384$, $l=12$) and Base ($d_m=768$, $l=12$) encoder configurations. By default, a patch embedding layer that computes $[4 \times 16]$ shaped non-overlapping patches (along time and frequency, respectively) is used in all our experiments. While [16, 17] suggest a default expansion factor of $E_f = 2$, in early ablations we found that $E_f = 3$ performed the best (Section 4.2). Ablation experiments also revealed that disabling the flip operation performed better than enabling flip in alternating blocks. We build on hyperparameter recommendations from previous work [6, 14]: all models are trained for 100 epochs with a batch size of 1024 and a weight decay of 0.05. AdamW optimizer with a linear warmup for 10 epochs followed by a cosine learning rate decay schedule is used. No data augmentations were used.

Downstream evaluation: Following the HEAR protocol, we extract fixed-sized feature vectors independent of the input audio duration by taking the mean over time across 2-second au-

dio chunks. We then train a single hidden layer MLP classifier with 1024 neurons for each task, using the official *hear-eval-kit* accompanying the HEAR benchmark. All experiments are repeated with 10 different random seeds, and 95% confidence intervals are reported. It is worth noting that due to the number of experiments conducted, we use a restricted hyperparameter grid (same as [14, 22]) compared to the *hear-eval-kit*.

Aggregated Performance Metric: We use the aggregated normalized score as proposed by [22] to compare evaluated approaches across the proposed list of downstream tasks. For a model m , overall score $s(m) \in [0, 100]$ is given as: $s(m) = \frac{1}{|T|} \sum_{t \in T} \frac{x_t(m) - \min_t}{\max_t - \min_t} * 100$, where $x_t(m)$ denotes performance of the model m on task t , and \min_t and \max_t represent the worst and the best performance across all models on the task, thus taking into account the relative performance amongst all evaluated representations.

4. Results

4.1. Comparison with existing works

Table 2 shows how the proposed AxLSTM models fare against recent audio representations. Directly comparable SSAST and AxLSTM feature representations have identical feature embedding sizes, however, feature vector sizes extracted using other referred methods can vary. While sub-optimal, it is infeasible to retrain all pretrained representations to have the same embedding sizes, and our evaluation protocol is in line with

Table 3: Performance impact of expansion factor (E_f)

Model	E_f	# Params	$s(m)$
SSAST-Tiny	-	5.4 M	55.6 \pm 0.2
AxLSTM-Tiny	2	2.9 M	66.9 \pm 0.3
AxLSTM-Tiny	3	4.3 M	68.9\pm0.4
AxLSTM-Tiny	4	5.8 M	63.4 \pm 0.2

recent frameworks for evaluating self-supervised audio representations [23, 34]. SSAST [6] represents the officially released model trained on AudioSet+LibriSpeech, whereas underlined SSAST representations were pretrained by us. We can see that AxLSTM models consistently outperform their transformer-based SSAST counterparts by a considerable margin, while having over 23% fewer parameters, with AxLSTM-Base yielding over 25% relative improvement in aggregate performance (86.6 \pm 0.2 v/s 68.7 \pm 0.3). Despite being based on the older SSAST framework, our AxLSTM-Base model matches the more recent masked modeling-based approaches such as AudioMAE [7], BEATS-iter3 [19] and MSM-MAE-208 [21] in overall performance, while having over 25% fewer parameters. The recent Mamba-based SSAM [14] models outperform comparable AxLSTM models in overall performance, however, it’s worth noting that AxLSTM models outperform SSAM models in music and pitch perception tasks (BO, Mri-T, NSynth-5h) whereas SSAM models shine in Speech-based and Audio classification tasks, performing notably better for emotion recognition (CD) and spoken language identification (VL). Finally, the recent MWMAE-Base [22] model with local-global self-attention outperforms every evaluated approach, including AxLSTM-Base, albeit with notably more parameters. Overall, while the AxLSTM model doesn’t establish a new state-of-the-art, it significantly outperforms SSAST-based transformer baselines while matching the performance of several other newer masked modeling approaches while possessing fewer parameters, thus demonstrating the viability of xLSTMs to learn self-supervised audio representations.

4.2. Ablations

This section covers ablations for investigating the impact of key configurable hyperparameters. Unless stated otherwise, experiments use the Tiny configuration with an expansion factor $E_f = 2$ and flipping enabled in alternating mLSTM blocks.

Expansion Factor: This hyperparameter affects the size and modeling capacity of mLSTM blocks. Table 3 shows that $E_f = 3$ offers better overall performance than the default $E_f = 2$ (suggested in [16, 17]). However, $E_f = 2$ still outperforms the SSAST baseline, offering 20% better performance with over 45% fewer parameters.

Patch Size: The patch size hyperparameter governs the sequence length as well as the time-frequency resolution of the patches that are processed by the underlying sequence modeling blocks. We pretrain and compare SSAST and AxLSTM Tiny configurations with 3 patch sizes: (4, 8), (4, 16), and (8, 16). From Table 4, we can see that AxLSTM scales better with increasing number of patches, with smaller patch sizes continually leading to better performance, whereas SSAST-Tiny experiences a drop in performance at patch size of (4, 8), suggesting that AxLSTMs scale better with sequence length in terms of modeling performance compared to standard transformers.

Bidirectional operation and exponential gating: In Table 5,

Table 4: Patch size ablations with the Tiny configuration

Model	Patch Size	# Patches	$s(m)$
SSAST-Tiny	(8, 16)	125	47.0 \pm 0.3
AxLSTM-Tiny	(8, 16)	125	55.2\pm0.3
SSAST-Tiny	(4, 16)	250	55.6 \pm 0.2
AxLSTM-Tiny	(4, 16)	250	66.9\pm0.3
SSAST-Tiny	(4, 8)	500	53.8 \pm 0.3
AxLSTM-Tiny	(4, 8)	500	68.2\pm0.3

we evaluate the impact of bidirectional modeling (controlled through the *flip* operation) and replacing exponential gating with a sigmoid gate. When done one at a time, disable both flip operation in alternating xLSTM blocks and the exponential gate improve performance compared to when both of these operations are enabled, with the former leading to a greater performance improvement. This suggests that bidirectional modeling is not as important for downstream audio classification as it is for vision [17], which is inline with observations made in previous work [14]. However, disabling both exponential gating and flipping together does not further improve performance, instead it performs worse than just disabling sequence flipping in alternative blocks. These results suggest that changes at the architectural level to accommodate bidirectional modeling and exponential gating might be warranted.

Table 5: Evaluating the impact of exponential gating and flipping on overall performance

Model	E_f	Exp gate	Flip	$s(m)$
AxLSTM-Tiny	2	✓	✓	66.9 \pm 0.3
AxLSTM-Tiny	2	✓		70.1 \pm 0.3
AxLSTM-Tiny	2		✓	68.2 \pm 0.2
AxLSTM-Tiny	3	✓	✓	68.8 \pm 0.4
AxLSTM-Tiny	3	✓		70.7 \pm 0.2
AxLSTM-Tiny	3		✓	69.6 \pm 0.2

5. Conclusion

This work presents AxLSTMs, an approach that evaluates the viability of the recently proposed xLSTM model for learning self-supervised audio representations from masked audio spectrograms. Extensive empirical analysis shows that AxLSTMs significantly outperform comparable SSAST baselines on ten varied downstream audio recognition tasks. AxLSTM-Base, which is our best performing model, offers over 25% better performance than SSAST-B while having up to 23% fewer parameters, and even matches the performance of several newer masked modeling approaches such as MSM-MAE-208 and BEATs. On the other hand, AxLSTM-Tiny models with $E_f = 2$ can offer over 25% relative improvement compared to an SSAST baseline, while having over 45% fewer parameters, which is a great performance/complexity tradeoff. Overall, our results demonstrate the viability of xLSTMs for learning self-supervised audio representations.

6. Acknowledgements

This research is supported by the Pioneer Centre for Artificial Intelligence, Denmark. We also thank Prof. Lars Kai Hansen from DTU, Denmark, for valuable feedback and suggestions.

7. References

- [1] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [2] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 2019.
- [3] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick, "Masked autoencoders are scalable vision learners," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 16 000–16 009.
- [4] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," in *Advances in Neural Information Processing Systems*, vol. 33, 2020, pp. 12 449–12 460.
- [5] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhota, R. Salakhutdinov, and A. Mohamed, "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE Transactions on Audio, Speech, and Language Processing*, pp. 1–1, 2021.
- [6] Y. Gong, C.-I. Lai, Y.-A. Chung, and J. Glass, "Ssast: Self-supervised audio spectrogram transformer," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, 2022, pp. 10 699–10 709.
- [7] P.-Y. Huang, H. Xu, J. Li, A. Baevski, M. Auli, W. Galuba, F. Metzger, and C. Feichtenhofer, "Masked autoencoders that listen," *Advances in Neural Information Processing Systems*, vol. 35, pp. 28 708–28 720, 2022.
- [8] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao *et al.*, "Wavlm: Large-scale self-supervised pre-training for full stack speech processing," *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1505–1518, 2022.
- [9] A. Katharopoulos, A. Vyas, N. Pappas, and F. Fleuret, "Transformers are rnns: Fast autoregressive transformers with linear attention," in *International conference on machine learning*. PMLR, 2020, pp. 5156–5165.
- [10] I. O. Tolstikhin, N. Houlsby, A. Kolesnikov, L. Beyer, X. Zhai, T. Unterthiner, J. Yung, A. Steiner, D. Keysers, J. Uszkoreit *et al.*, "Mlp-mixer: An all-mlp architecture for vision," *Advances in neural information processing systems*, vol. 34, pp. 24 261–24 272, 2021.
- [11] A. Gu, K. Goel, and C. Re, "Efficiently modeling long sequences with structured state spaces," in *International Conference on Learning Representations*, 2022.
- [12] A. Gu and T. Dao, "Mamba: Linear-time sequence modeling with selective state spaces," *arXiv preprint arXiv:2312.00752*, 2023.
- [13] L. Zhu, B. Liao, Q. Zhang, X. Wang, W. Liu, and X. Wang, "Vision mamba: efficient visual representation learning with bidirectional state space model," in *Proceedings of the 41st International Conference on Machine Learning*. JMLR.org, 2024.
- [14] S. Yadav and Z.-H. Tan, "Audio mamba: Selective state spaces for self-supervised audio representations," in *Proc. INTERSPEECH 2024 – 25th Annual Conference of the International Speech Communication Association*, Kos Island, Greece, Sep. 2024.
- [15] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [16] M. Beck, K. Pöppel, M. Spanring, A. Auer, O. Prudnikova, M. K. Kopp, G. Klambauer, J. Brandstetter, and S. Hochreiter, "xLSTM: Extended long short-term memory," in *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- [17] B. Alkin, M. Beck, K. Pöppel, S. Hochreiter, and J. Brandstetter, "Vision-LSTM: xLSTM as generic vision backbone," in *The Thirteenth International Conference on Learning Representations*, 2025.
- [18] 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.
- [19] S. Chen, Y. Wu, C. Wang, S. Liu, D. Tompkins, Z. Chen, W. Che, X. Yu, and F. Wei, "Beats: audio pre-training with acoustic tokenizers," in *Proceedings of the 40th International Conference on Machine Learning*, 2023, pp. 5178–5193.
- [20] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16x16 words: Transformers for image recognition at scale," in *International Conference on Learning Representations*, 2021.
- [21] D. Niizumi, D. Takeuchi, Y. Ohishi, N. Harada, and K. Kashino, "Masked spectrogram modeling using masked autoencoders for learning general-purpose audio representation," in *HEAR: Holistic Evaluation of Audio Representations*. PMLR, 2022, pp. 1–24.
- [22] S. Yadav, S. Theodoridis, L. K. Hansen, and Z.-H. Tan, "Masked autoencoders with multi-window local-global attention are better audio learners," in *The Twelfth International Conference on Learning Representations*, 2024.
- [23] J. Turian, J. Shier, H. R. Khan, B. Raj, B. W. Schuller, C. J. Steinmetz, C. Malloy, G. Tzanetakis, G. Velarde, K. McNally, M. Henry, N. Pinto, C. Noufi, C. Clough, D. Herremans, E. Fonseca, J. Engel, J. Salamon, P. Esling, P. Manocha, S. Watanabe, Z. Jin, and Y. Bisk, "HEAR: Holistic Evaluation of Audio Representations," in *Proceedings of the NeurIPS 2021 Competitions and Demonstrations Track*. PMLR, Jul. 2022, pp. 125–145, ISSN: 2640-3498.
- [24] K. Koutini, J. Schlüter, H. Eghbal-zadeh, and G. Widmer, "Efficient Training of Audio Transformers with Patchout," in *Proc. Interspeech 2022*, 2022, pp. 2753–2757.
- [25] M. Tian, A. Srinivasamurthy, M. Sandler, and X. Serra, "A study of instrument-wise onset detection in beijing opera percussion ensembles," in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2014, pp. 2159–2163.
- [26] H. Cao, D. G. Cooper, M. K. Keutmann, R. C. Gur, A. Nenkova, and R. Verma, "Crema-d: Crowd-sourced emotional multimodal actors dataset," *IEEE transactions on affective computing*, vol. 5, no. 4, pp. 377–390, 2014.
- [27] K. J. Piczak, "ESC: Dataset for Environmental Sound Classification," in *Proceedings of the 23rd Annual ACM Conference on Multimedia*. ACM Press, 2015.
- [28] F.-R. Stöter, S. Chakrabarty, B. Edler, and E. A. Habets, "Classification vs. regression in supervised learning for single channel speaker count estimation," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018.
- [29] A. Anantapadmanabhan, A. Bellur, and H. A. Murthy, "Modal analysis and transcription of strokes of the mridangam using non-negative matrix factorization," in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2013.
- [30] J. Engel, C. Resnick, A. Roberts, S. Dieleman, D. Eck, K. Simonyan, and M. Norouzi, "Neural audio synthesis of musical notes with wavenet autoencoders," 2017.
- [31] P. Warden, "Speech commands: A dataset for limited-vocabulary speech recognition," 2018.
- [32] E. Fonseca, X. Favory, J. Pons, F. Font, and X. Serra, "Fsd50k: an open dataset of human-labeled sound events," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2021.
- [33] B. Kim, M. Ghei, B. Pardo, and Z. Duan, "Vocal imitation set: a dataset of vocally imitated sound events using the audioset ontology," in *DCASE*, 2018, pp. 148–152.
- [34] S.-w. Yang, P.-H. Chi, Y.-S. Chuang, C.-I. Jeff Lai, K. Lakhota, Y. Y. Lin, A. T. Liu, J. Shi, X. Chang, G.-T. Lin, T.-H. Huang, W.-C. Tseng, K.-t. Lee, D.-R. Liu, Z. Huang, S. Dong, S.-W. Li, S. Watanabe, A. Mohamed, and H.-y. Lee, "SUPERB: Speech Processing Universal PERFORMANCE Benchmark," in *Proc. Interspeech 2021*, 2021, pp. 1194–1198.