



# Towards Attention-based Contrastive Learning for Audio Spoof Detection

Chirag Goel, Surya Koppiseti, Ben Colman, Ali Shahriyari, Gaurav Bharaj

Reality Defender Inc., USA

{chirag,surya,ben,ali,gaurav}@realitydefender.com

## Abstract

Vision transformers (ViT) have made substantial progress for classification tasks in computer vision. Recently, Gong et al. '21, introduced attention-based modeling for several audio tasks. However, relatively unexplored is the use of a ViT for audio spoof detection task. We bridge this gap and introduce ViTs for this task. A vanilla baseline built on fine-tuning the SSAST (Gong et al. '22) audio ViT model achieves sub-optimal equal error rates (EERs). To improve performance, we propose a novel attention-based contrastive learning framework (SSAST-CL) that uses cross-attention to aid the representation learning. Experiments show that our framework successfully disentangles the bonafide and spoof classes and helps learn better classifiers for the task. With appropriate data augmentations policy, a model trained on our framework achieves competitive performance on the ASVSpooF 2021 challenge. We provide comparisons and ablation studies to justify our claim.

**Index Terms:** spoof detection, contrastive learning, attention-based modeling, ASVSpooF challenge

## 1. Introduction

Vision Transformers (ViTs) (Dosovitskiy et al. [1]) have emerged as the state-of-the-art method for multiple computer vision [2] and audio (Gong et al. [3]) tasks. Audio ViTs can *self-attend* to extract optimal patch embeddings and are capable of learning long-range global context within spectrograms [4, 3]. Inspired by the success of audio ViTs in sound and speech classification, here we introduce ViTs for the audio spoof detection task.

Similar to [3], the standard approach is to fine-tune a pre-trained audio ViT for a given downstream classification task using cross-entropy. Pre-training is done on a large dataset, such as AudioSet [5] or LibriSpeech [6], and fine-tuning on a smaller task-specific dataset. However, this standard approach does not empirically work well for audio spoof detection (see Section 4.1), and results in high equal error rate (EER). Pretraining on a large audio dataset such as AudioSet-2M [5] is a key requirement for audio ViTs to perform well [4, 3]. However, such *bonafide-only* datasets are not optimal as they don't contain spoof samples; spoof detection is an out-of-distribution downstream task.

The *de-facto* dataset for audio spoof detection is ASVSpooF 2021 logical access (LA) challenge [7]. Its training set is the same as ASVSpooF2019 [8] and contains clean studio quality data. Its test set, on the other hand, contains data corrupted by codecs and transmission channel artifacts. Also, when finetuning an audio ViT for spoof detection, the ASVSpooF19 LA training dataset falls under the small data regime.<sup>1</sup> We thus require an

appropriate training and augmentation strategy that helps achieve robustness against data corruption in the test set.

Supervised contrastive learning (CL) methods [9] that use Siamese networks (Koch et al. [10]) have emerged as an algorithm class that helps learn efficient data representations. Such approaches learn a representation space from generic input features by pulling samples from the same class closer together and pushing samples from different classes apart, even on limited data. Thus, they are able to train on small datasets, learn better classification margins, and are more robust to data corruptions than cross-entropy classifiers [9]. Inspired by these merits, we pursue a CL approach. We introduce a cross-attention branch into the training framework and propose a novel loss formulation that measures the (dis-)similarity between the self and cross-attention representations to separate the bonafide and spoof classes. We observe that the proposed attention-based CL framework, with appropriate data augmentations, is able to learn *discriminative representations* that disentangle the bonafide samples from the spoofed ones. Further, our approach shows a significant gain in performance over the baseline cross-entropy classifier. To summarize, the main contributions of our work are:

1. We propose a two-stage contrastive learning framework to train an audio ViT for the spoof detection task.
2. We consider Siamese training for representation learning and introduce a cross-attention branch into the training framework to learn discriminative representations for bonafide and spoof classes; a suitable loss function is formulated.
3. A MLP classifier trained on the learned representations outperforms the ASVSpooF 2021 challenge baselines [7] and competes with the best-performing models.

Finally, while in this paper we focus on audio spoof detection, the proposed CL framework is not constrained by design to our specific problem, and is a general-purpose framework. In the future, we plan to explore the framework for other downstream tasks, such as language identification and acoustic scene classification.

## 2. Related works

**Attention-based modeling.** Self-attention mechanism relates different patches within an input image using attention weights and encodes efficient representation of an image. It is a building block in ViTs, and has facilitated state-of-the-art (SOTA) classification performance on multiple computer vision tasks [1, 11, 12]. Cross-attention [13, 14] extends the idea to image pairs and uses attention matrices to encode an aggregate representation for the input pair based on the similarity in their patches. Due to its ability to align features from different inputs, cross-attention has been used for domain adaptation [14] and multi-modal feature

<sup>1</sup>Contains 25,380 samples, which is about 1% of AudioSet-2M [5].




Figure 1: SSAST-CL: A two-stage contrastive learning framework to train the SSAST model [3] for audio spoof detection. In Stage I, we employ Siamese training with weight-sharing across two multi-head self-attention (MH-SA) and one multi-head cross-attention (MH-CA) branches. Model weights are learned using a contrastive loss which measures the (dis-)similarity between the self and cross-attention representations ( $r_1^{SA}$ ,  $r_2^{SA}$ ,  $r_{12}^{CA}$ ). In Stage II, a MLP classifies the learned representations as bonafide or spoof.

fusion [15, 16]. ViTs with self-attention only [3] have achieved SOTA performance on multiple audio/speech classification tasks, however use of both self and cross-attention in audio ViTs, especially for audio spoof detection, is unexplored.

**Contrastive learning for audio classification.** In CL, samples are fed into the network in form of data pairs in order to contrast between them when learning the representation spaces of each class. A typical CL formulation maximizes the similarity of samples in an input pair when they belong to the same class, and their dissimilarity otherwise [17, 18]. In [17, 18], CL was used to derive audio representations for a variety of speech and non-speech classification tasks, including speaker identification, acoustic scene classification, and music recognition.

Xie *et al.* [18] focus on detecting studio quality spoofed audio and proposed a CL algorithm where a Siamese CNN is trained on the ASVspoof19 LA. On the other hand, we propose a CL formulation that uses cross-attention to train a ViT for spoof detection. Further, it is unclear whether a Siamese network can perform well on the ASVspoof 2021 LA, which contains impairments from transmission codecs [7]. To this end, we use appropriate data augmentations that help achieve robustness against such codec impairments.

**Spoof detection.** Early works have focused on handcrafted features, such as cepstral coefficients [19], with classical methods such as Gaussian mixture models [20] (also see [8] and references therein). Recent methods use CNNs that operate on 2D features extracted from audio signals [21, 22, 23], or directly map raw waveforms to their labels in an end-to-end fashion [24, 25]. Introduction of attention mechanism into CNN models has further enhanced the single model performance for the ASVspoof21 challenge [26], particularly those employing temporal and/or spatial attention. For this challenge, the works [27] and [28] employ wav2vec2.0 [29] large audio model as a feature extractor and achieve significantly better EER. Among other contributors, the improvement comes from wav2vec2.0 - a deep network of CNN and transformer blocks. While it is clear from SOTA that the introduction of attention layers in CNNs gave improved performance, it is unclear whether a pure attention-based model can be used for spoof detection. We bridge this gap and study the use of ViTs for audio spoof detection.

### 3. Method

We introduce an audio ViT to learn efficient representations for the spoof detection task. We consider the self-supervised audio spectrogram transformer (SSAST) [3], a ViT pretrained on the AudioSet [5] and LibriSpeech [6] datasets. We propose a CL-based two-stage training framework, summarized in Fig. 1. In Stage I, we adapt the SSAST to train in a Siamese manner (we call it SSAST-CL), in order to learn discriminative representations for the bonafide and spoof classes. In Stage II, we train an MLP using the SSAST-CL backbone to classify the learned representations as bonafide or spoof.

#### 3.1. Stage I - representation learning

In Stage I, the goal is to learn discriminative representations for the bonafide and spoof classes. To this end, we propose the SSAST-CL framework which adapts the SSAST architecture in [3] to a three-branch Siamese training (see Fig. 1). The framework takes a pair of data  $(x_1, x_2)$  as input and feeds them to SSAST-CL, whose weights are shared across the three branches. Two branches in the SSAST-CL use self-attention to compute representations  $r_1^{SA}$  and  $r_2^{SA}$ , whereas the third branch uses cross-attention to compute an aggregate representation  $r_{12}^{CA}$  for the input pair.

In the self-attention branch for  $x_1$  (and similarly for  $x_2$ ), the transformer blocks encode the intermediate representations  $R_1^{SA'}$  using the query, key and value matrices  $(Q_1, (K_1, V_1))$  from the same  $x_1$  branch. Whereas in the cross-attention branch, the transformer blocks encode the intermediate representations  $R_{12}^{CA'}$  using the query  $Q_1$  from the  $x_1$  branch and key-value matrices  $(K_2, V_2)$  from the  $x_2$  branch (as shown in Fig. 1). By way of this design, the final representation  $r_{12}^{CA}$  becomes an aggregate representation of the input pair because it captures the information in  $x_2$  that is relevant to  $x_1$ .

When computing the attention matrices in the transformer blocks, the query, key, and value are assigned learnable weights  $W_q$ ,  $W_k$ , and  $W_v$  respectively. These weights are shared across the three branches, as is the case with all other learnable parameters in SSAST-CL. Finally, we add a linear projection MLP at the end of SSAST-CL to upsample all the attention-based repre-




Figure 2:  $t$ -SNE embeddings for the (a) vanilla WCE baseline [3] and (b-c) the proposed SSAST-CL solution on ASVSpooF 2021 dataset.

representations ( $\mathbf{r}_1^{SA}, \mathbf{r}_2^{SA}, \mathbf{r}_{12}^{CA}$ ) for the contrastive loss calculations.

The model weights are optimized using a novel contrastive loss formulation which measures the (dis-)similarity between the self and cross-attention representations to separate the bonafide and spoof classes. Similar to [30], we use the cosine distance metric, given by  $\cos(\mathbf{a}, \mathbf{b}) = (\mathbf{a}^\top \mathbf{b}) / \max(\|\mathbf{a}\| \cdot \|\mathbf{b}\|, \epsilon)$ ,  $\epsilon > 0$ , to measure the similarity between the representations<sup>2</sup>. We then define the contrastive loss  $L_{con}$  as:

$$L_{con} = L_{SA} + \alpha L_{CA}, \text{ where ,}$$

$$L_{SA} = -1 \begin{cases} \log(\cos(\mathbf{r}_1^{SA}, \mathbf{r}_2^{SA})), & c(x_1) = c(x_2) \\ \log(1 - \cos(\mathbf{r}_1^{SA}, \mathbf{r}_2^{SA})), & c(x_1) \neq c(x_2) \end{cases},$$

$$L_{CA} = -1 \begin{cases} \log(\cos(\mathbf{r}_1^{SA}, \mathbf{r}_{12}^{CA})), & c(x_1) = c(x_2) \\ \log(1 - \cos(\mathbf{r}_1^{SA}, \mathbf{r}_{12}^{CA})), & c(x_1) \neq c(x_2) \end{cases} \quad (1)$$

In Eq (1),  $c(x_1)$  and  $c(x_2)$  are the (bonafide or spoof) classes to which  $x_1$  and  $x_2$  belong,  $L_{SA}$  denotes the self-attention loss,  $L_{CA}$  denotes the cross-attention loss, and  $\alpha \in [0, 1]$  a weighting parameter. The self-attention loss  $L_{SA}$  operates on representations  $\mathbf{r}_1^{SA}$  and  $\mathbf{r}_2^{SA}$  from the self-attention branches. It maximizes the similarity between  $\mathbf{r}_1^{SA}$  and  $\mathbf{r}_2^{SA}$  if the input pair  $(x_1, x_2)$  belong to the same class, and their dissimilarity otherwise. The cross-attention loss  $L_{CA}$  operates on the self and cross-attention representation pair  $(\mathbf{r}_1^{SA}, \mathbf{r}_{12}^{CA})$ , as they have the same query  $Q_1$  for each transformer block.  $L_{CA}$  maximizes the similarity between  $\mathbf{r}_1^{SA}$  and its cross-attention counterpart  $\mathbf{r}_{12}^{CA}$  if the input pair  $(x_1, x_2)$  belongs to the same class, and their dissimilarity otherwise. While the self-attention loss term directs the model to learn representations that separate the sample classes, the cross-attention loss term serves as a regularizer by pushing the class-specific representations away from the aggregate representations computed by the cross-attention branch.

A traditional Siamese training network only computes the self-attention representations  $\mathbf{r}_1^{SA}$  and  $\mathbf{r}_2^{SA}$  and measures their (dis-)similarity to separate the bonafide and spoof classes. Here, we introduce an additional cross-attention branch into the training framework, in order to help the model learn more discriminative representations. After the training is complete, the projection MLP at the end of Stage I is discarded, as is commonplace in contrastive learning methods [31].

### 3.2. Stage II - classifier

In Stage II, we train an MLP using weighted cross-entropy to classify the representations from Stage I as bonafide or spoof. For a given input, the model weights from Stage I are frozen and a self-attention branch is used to compute the representation. Note that the two self-attention branches are identical due to weight sharing.

<sup>2</sup>Similarity is measured after feeding the representations  $\mathbf{r}_1^{SA}, \mathbf{r}_2^{SA}, \mathbf{r}_{12}^{CA}$  through the projection MLP (see Fig. 1)

Table 1: EER performance of SSAST-CL on ASVSpooF 2021 LA evaluation set. Vanilla WCE is the policy recommended in [3]. WCE-updated is the same as vanilla WCE, but with our suggested augmentation policy. Smaller EERs are better.

Training policy	Augmentations	EER
Vanilla WCE	Mixup, SpecAugment, Noise	19.48
WCE-updated	WavAugment, Pitch Shift, Time Stretch, RawBoost, Narrowband FIR	8.96
<b>SSAST-CL (Ours)</b> (self-attention only)	<b>WavAugment, Pitch Shift, Time Stretch, RawBoost, Narrowband FIR</b>	<b>5.64</b>
<b>SSAST-CL (Ours)</b> (self and cross-attention, $\alpha = 0.2$ )	<b>WavAugment, Pitch Shift, Time Stretch, RawBoost, Narrowband FIR</b>	<b>4.74</b>

### 3.3. Data augmentations

We need a suitable data augmentation policy to support our training framework in three ways: prevent overfitting, handle speaker variability, and achieve robustness to telephony codec impairments. The following augmentations are used: pitch-shift, time-stretch, time and frequency masking from WavAugment [32], linear and non-linear convolutive noise and impulsive signal dependent additive noise from RawBoost [33], and the narrowband frequency impulse response (FIR) filters suggested in [34]. All augmentations are applied on-the-fly during the model training.

## 4. Experiments and results

**Audio pre-processing.** Raw audio waveforms of length 6-seconds are used to create log-mel spectrograms of size 128 mel-frequency bins and 512 timebins, computed using PyTorch Kaldi [35] with 25ms Hanning window and 10ms overlap. Longer waveforms are cut off at the end, whereas shorter waveforms are repeat padded by concatenating the original signal with its time-inverted version until the length is 6 seconds.

**Sampling and batching.** For the contrastive learning in Stage I, data pairs  $(x_1, x_2)$  are created as follows: The sample  $x_1$  is picked up in sequence from the training dataset, covering each datapoint once over a training epoch. For each  $x_1$ , we make use of its class information  $c(x_1)$  to select the pairing sample  $x_2$  such that positive and negative pairs are created with equal probability of 0.5. When picking up  $x_2$  from the spoof class, we assign equal probability of 0.5 to the text-to-speech (TTS) and voice conversion (VC) subclasses. Once a pair is picked, each of the data augmentations from Section 3 are applied to  $x_1$  with a probability of 0.8, and a randomly selected subset of data augmentations is applied to  $x_2$ . The above procedure is repeated until a batch of 64 pairs is created. During Stage II, we pick batches of size 64 sequentially from the dataset; to each sample, all data augmentations are applied with 0.8 probability.

**Training policy.** For both Stage I and II, we use Adam optimizer with a learning rate of  $10^{-4}$ , and an exponential rate decay scheduler with  $\gamma = 0.95$  for every 5 epochs. The model is trained for 50 epochs for each stage. For Stage I, the epoch checkpoint reporting the least validation loss is chosen. For Stage II, the epoch checkpoint reporting the smallest validation EER is chosen.

**Implementation of SSAST-CL.** For details on the SSAST-CL implementation, please see the supplementary material.

We now present numerical studies to demonstrate the improvements in EER when using our CL framework. We compare with the state of the art and ablate on the data augmentations.

#### 4.1. Impact of data augmentation, contrastive learning and cross-attention

In Table 1, we compare the performance of the proposed contrastive learning framework against baselines that finetune the pretrained SSAST using weighted cross-entropy (WCE). The vanilla WCE, which follows the training and augmentation policy recommended in [3], reports a high EER of 19.48 on ASVSpooF 2021 LA evaluation set. This is likely because the augmentations in vanilla WCE do not account for codec impairments. When our augmentation policy from the ablation studies in Table 3 is used, the resulting WCE-updated policy reports an EER of 8.96, which is marginally better than the 9.26 reported by the best-performing baseline B03 in ASVSpooF 2021 [7]. When the training framework is *additionally* replaced with our SSAST-CL framework, the EER improves significantly to 4.74.

Table 1 also demonstrates the impact of introducing cross-attention. We see that the EER improves from 5.64 to 4.74 when we set a weight of  $\alpha = 0.2$  to the cross-attention loss term  $L_{CA}$ . The t-SNE plots in Fig. 2 suggest that the proposed SSAST-CL framework better disentangles the bonafide and spoof classes when compared to the vanilla WCE baseline. There is also an improvement from introducing cross-attention; although the visual differences are subtle in this case due to the high reduction in dimensionality from 192 to 2 when plotting the t-SNEs.

While the impact of SSAST-CL is evident from Table 1 and Fig. 2, we note that our framework is not fully optimized to report the best EERs. A grid search on the training hyperparameters, including the optimizer, learning rate, batch size, and cross-attention weight, will likely boost the EERs further.

Table 2: EER comparison with SOTA on ASVSpooF21 LA evaluation set. (\*) estimate based on authors’ description of model.

Model	DA	EER	Parameters (million)	GMACs <sup>3</sup>
Wave2Vec2.0+ AASiST [27]	RawBoost (On-the-fly)	0.82	317.84	1050
LCNN-LSTM [34]	RS Mixup + FIR (On-the-fly)	2.21	0.5*	21.77*
ResNet-LDE [36]	RIR+Background noise + Frequency Masking	3.68	11.5*	100.74*
SSAST-CL (Ours)	RawBoost + WavAugment + Pitch Shift + Time Stretch + Narrowband FIR (On-the-fly)	4.74	5.99	126.75
RawNet2 [33]	RawBoost (On-the-fly)	5.31	25.43	59.45
ResNet-34 [37]	Cellular, Telephony, and VoIP Codecs (Offline)	5.18	21.5	623.89
ASVSpooF21 B03 [7]	None	9.26	0.27	10.75

#### 4.2. Comparison with ASVSpooF21 top-performing models

In Table 2, we compare the performance of the proposed SSAST-CL system against the top-performing single system models on the ASVSpooF 2021 LA evaluation dataset. Firstly, our system

<sup>3</sup><https://github.com/B06901052/DeepSpeed/tree/superb-challenge>

comprehensively outperforms the challenge’s best-performing baseline B03 [7]. We see that an audio ViT, with appropriate training and data augmentations, can indeed achieve competitive performance on the audio spoof detection task. Note that a vanilla WCE finetuning on the SSAST model, as suggested in [3], results in worse EER than the challenge baseline (see Table 1). Secondly, when positioned against the best-performing models, our system reports comparable EERs while being significantly smaller in size than most of them. The LCNN-LSTM [34] is the only *lightweight* system that reports a smaller EER than us. Lastly, our augmentations are much simpler than in the ResNet-LDE [36] and ResNet-34 [37] systems, as we do not use external noise or impulse response datasets for the augmentations.

Table 3: Impact of augmentations on the model performance for ASVSpooF2021 evaluation data. All experiments additionally use time-stretch, pitch-shift, time masking, and frequency masking.

Model	RawBoost	Narrowband FIR	EER
SSAST-CL (Ours)	✗	✗	8.69
	✗	✓	6.88
	✓	✗	5.34
	✓	✓	4.74

#### 4.3. Ablation on data augmentations

In Table 3, we ablate over data augmentation combinations for the proposed SSAST-CL framework. We note that the RawBoost and FIR augmentations are both crucial for the model to perform well. These augmentations differ in design but both help capture the telephony codec artifacts, although to different extents. Telephony impairments are known to be present in the ASVSpooF 2021 LA dataset. The remaining augmentations in our policy, namely, time masking, frequency masking, pitch shift, and time stretch, are now commonplace in spoof detection modeling. These are essential because they account for speaker variability and prevent the model from overfitting [32].

## 5. Conclusion

In this work, we answered the question “Can we leverage ViTs for the audio spoof detection task?” We conducted investigations using SSAST-CL (our adaption of the SSAST model [3] for contrastive learning). A vanilla finetuning of the pretrained SSAST with cross-entropy loss gave sub-optimal performance. To learn more discriminative audio representations, the proposed SSAST-CL framework used Siamese training with a cross-attention branch and a novel contrastive loss formulation. An MLP was later used to classify the learned representations as bonafide or spoof. We empirically showed that the SSAST-CL framework successfully disentangles the bonafide and spoof classes, and it helps learn better classifiers for the task at hand. The introduction of cross-attention, along with suitable augmentations, has allowed our system to achieve competitive performance on the ASVSpooF 2021 challenge.

Several research directions are enabled as a result of our work. A joint training of the two stages in our framework, using a multi-task loss formulation, could likely improve the model performance and also reduce the training time. Another direction would be to build an importance sampling/training policy where hard-to-learn samples (for example, of the voice conversion type) are prioritized. The proposed contrastive learning framework can also be extended to other downstream audio tasks where limited training data is available, such as emotion recognition [38] and language identification [39].

## 6. References

- [1] 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 *Proc. ICLR*, 2021.
- [2] A. Hatamizadeh, H. Yin, J. Kautz, and P. Molchanov, “Global context vision transformers,” *arXiv preprint arXiv:2206.09959*, 2022. [Online]. Available: <https://arxiv.org/abs/2206.09959>
- [3] Y. Gong, C.-I. Lai, Y.-A. Chung, and J. Glass, “Ssast: Self-supervised audio spectrogram transformer,” in *Proc. AAAI Conference on Artificial Intelligence*, 2022, pp. 10 699–10 709.
- [4] Y. Gong, Y.-A. Chung, and J. Glass, “AST: Audio Spectrogram Transformer,” in *Proc. Interspeech*, 2021, pp. 571–575.
- [5] J. F. Gemmeke, D. P. W. 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,” *Proc. IEEE ICASSP*, pp. 776–780, 2017.
- [6] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: An asr corpus based on public domain audio books,” *IEEE ICASSP*, pp. 5206–5210, 2015.
- [7] J. Yamagishi, X. Wang, M. Todisco, M. Sahidullah, J. Patino, A. Nautsch, X. Liu, K. A. Lee, T. Kinnunen, N. Evans *et al.*, “Asvspoof 2021: accelerating progress in spoofed and deepfake speech detection,” in *ASvspoof 2021 Workshop-Automatic Speaker Verification and Spoofing Countermeasures Challenge*, 2021.
- [8] A. Nautsch, X. Wang, N. W. D. Evans, T. H. Kinnunen, V. Vestman, M. Todisco, H. Delgado, M. Sahidullah, J. Yamagishi, and K.-A. Lee, “Asvspoof 2019: Spoofing countermeasures for the detection of synthesized, converted and replayed speech,” *IEEE Transactions on Biometrics, Behavior, and Identity Science*, vol. 3, pp. 252–265, 2021.
- [9] P. Khosla, P. Teterwak, C. Wang, A. Sarna, Y. Tian, P. Isola, A. Maschinot, C. Liu, and D. Krishnan, “Supervised contrastive learning,” in *Proc. NeurIPS*, 2020, pp. 1–13.
- [10] G. Koch, R. Zemel, R. Salakhutdinov *et al.*, “Siamese neural networks for one-shot image recognition,” in *ICML deep learning workshop*, vol. 2, no. 1, 2015.
- [11] Y. Qian, M. Lin, X. Sun, Z. Tan, and R. Jin, “Entroformer: A transformer-based entropy model for learned image compression,” in *Proc. ICLR*, 2022, pp. 1–15.
- [12] L. Yuan, Q. Hou, Z. Jiang, J. Feng, and S. Yan, “Volo: Vision outlooker for visual recognition,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–13, 2022.
- [13] C.-F. Chen, Q. Fan, and R. Panda, “Crossvit: Cross-attention multi-scale vision transformer for image classification,” in *Proc. ICCV*, 2021, pp. 347–356.
- [14] T. Xu, W. Chen, P. Wang, F. Wang, H. Li, and R. Jin, “Cdtrans: Cross-domain transformer for unsupervised domain adaptation,” in *Proc. ICLR*, 2022, pp. 1–14.
- [15] N. Li, S. Liu, Y. Liu, S. Zhao, and M. Liu, “Neural speech synthesis with transformer network,” in *Proc. AAAI Conference on Artificial Intelligence*, 2018, pp. 6706–6713.
- [16] R. Hu and A. Singh, “Unit: Multimodal multitask learning with a unified transformer,” in *Proc. ICCV*, 2021, pp. 1419–1429.
- [17] A. Saeed, D. Grangier, and N. Zeghidour, “Contrastive learning of general-purpose audio representations,” in *Proc. IEEE ICASSP*, 2021, pp. 3875–3879.
- [18] Y. Xie, Z. Zhang, and Y. Yang, “Siamese network with wav2vec feature for spoofing speech detection,” in *Proc. Interspeech*, 2021, pp. 4269–4273.
- [19] M. Sahidullah, T. H. Kinnunen, and C. Hanilçi, “A comparison of features for synthetic speech detection,” in *Proc. Interspeech*, 2015, pp. 2087–2091.
- [20] M. Todisco, H. Delgado, and N. W. D. Evans, “Constant q cepstral coefficients: A spoofing countermeasure for automatic speaker verification,” *Comput. Speech Lang.*, vol. 45, pp. 516–535, 2017.
- [21] Y. Gao, T. Vuong, M. Elyasi, G. Bharaj, and R. Singh, “Generalized spoofing detection inspired from audio generation artifacts,” in *Proc. Interspeech*, 2021, pp. 4185–4188.
- [22] N. Zeghidour, O. Teboul, F. de Chaumont Quiry, and M. Tagliasacchi, “Leaf: A learnable frontend for audio classification,” in *Proc. ICLR*, 2021, pp. 1–16.
- [23] G. Lavrentyeva, S. Novoselov, A. Tseren, M. Volkova, A. Gorlanov, and A. Kozlov, “Stc antispoofing systems for the asvspoof2019 challenge,” in *Proc. Interspeech*, 2019, pp. 1033–1037.
- [24] H. Tak, J. Patino, M. Todisco, A. Nautsch, N. W. D. Evans, and A. Larcher, “End-to-end anti-spoofing with rawnet2,” in *Proc. IEEE ICASSP*, 2020, pp. 6369–6373.
- [25] G. Hua, A. B. J. Teoh, and H. Zhang, “Towards end-to-end synthetic speech detection,” *IEEE Signal Processing Letters*, vol. 28, pp. 1265–1269, 2021.
- [26] X. Liu, X. Wang, M. Sahidullah, J. Patino, H. Delgado, T. H. Kinnunen, M. Todisco, J. Yamagishi, N. W. D. Evans, A. Nautsch, and K.-A. Lee, “Asvspoof 2021: Towards spoofed and deepfake speech detection in the wild,” *ArXiv*, vol. abs/2210.02437, 2022. [Online]. Available: <https://arxiv.org/abs/2210.02437>
- [27] H. Tak, M. Todisco, X. Wang, J.-w. Jung, J. Yamagishi, and N. Evans, “Automatic speaker verification spoofing and deepfake detection using wav2vec 2.0 and data augmentation,” in *Proc. Odyssey*, 2022, pp. 112–119.
- [28] Y. Eom, Y. Lee, J. S. Um, and H. Kim, “Anti-spoofing using transfer learning with variational information bottleneck,” in *Proc. Interspeech*, 2022, pp. 3568–3572.
- [29] A. Baeovski, Y. Zhou, A. Mohamed, and M. Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” in *Proc. NeurIPS*, 2020, pp. 12 449–12 460.
- [30] K. R. Prajwal, R. Mukhopadhyay, V. P. Namboodiri, and C. Jawahar, “A lip sync expert is all you need for speech to lip generation in the wild,” in *Proc. ACM International Conference on Multimedia*, 2020, pp. 484–492.
- [31] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, “A simple framework for contrastive learning of visual representations,” *Proc. ICML*, pp. 1–11, 2020.
- [32] E. Kharitonov, M. Rivièrè, G. Synnaeve, L. Wolf, P.-E. Mazaré, M. Douze, and E. Dupoux, “Data augmenting contrastive learning of speech representations in the time domain,” in *Proc. IEEE Spoken Language Technology Workshop*, 2020, pp. 215–222.
- [33] H. Tak, M. R. Kamble, J. Patino, M. Todisco, and N. W. D. Evans, “Rawboost: A raw data boosting and augmentation method applied to automatic speaker verification anti-spoofing,” in *Proc. IEEE ICASSP*, 2021, pp. 6382–6386.
- [34] A. Tomilov, A. F. Svishchev, M. Volkova, A. Chirkovskiy, A. S. Kondratev, and G. Lavrentyeva, “Stc antispoofing systems for the asvspoof2021 challenge,” *2021 Edition of the Automatic Speaker Verification and Spoofing Countermeasures Challenge*, pp. 61–67, 2021.
- [35] M. Ravanelli, T. Parcollet, and Y. Bengio, “The pytorch-kaldi speech recognition toolkit,” in *Proc. IEEE ICASSP*, 2019, pp. 6465–6469.
- [36] T. Chen, E. el Khoury, K. Phatak, and G. Sivaraman, “Pindrop labs’ submission to the asvspoof 2021 challenge,” *2021 Edition of the Automatic Speaker Verification and Spoofing Countermeasures Challenge*, pp. 90–93, 2021.
- [37] A. Cohen, I. Rimon, E. Aflalo, and H. H. Permuter, “A study on data augmentation in voice anti-spoofing,” *Speech Commun.*, vol. 141, pp. 56–67, 2021.
- [38] A. Satt, S. Rozenberg, R. Hoory *et al.*, “Efficient emotion recognition from speech using deep learning on spectrograms,” in *Proc. Interspeech*, 2017, pp. 1089–1093.
- [39] W. Cai, Z. Cai, X. Zhang, X. Wang, and M. Li, “A novel learnable dictionary encoding layer for end-to-end language identification,” in *Proc. IEEE ICASSP*, 2018, pp. 5189–5193.