



TADA: Training-free Attribution and Out-of-Domain Detection of Audio Deepfakes

Adriana Stan^{1,2}, David Combei¹, Dan Oneata², Horia Cucu²

¹Technical University of Cluj-Napoca, Romania

²University “Politehnica” Bucharest, Romania

adriana.stan@com.utcluj.ro, david.combei@cs.utcluj.ro,
dan.oneata@gmail.com, horia.cucu@upb.ro

Abstract

Deepfake detection has gained significant attention across audio, text, and image modalities, with high accuracy in distinguishing real from fake. However, identifying the exact source—such as the system or model behind a deepfake—remains a less studied problem. In this paper, we take a significant step forward in audio deepfake model attribution or source tracing by proposing a training-free, green AI approach based entirely on k-Nearest Neighbors (kNN). Leveraging a pre-trained self-supervised learning (SSL) model, we show that grouping samples from the same generator is straightforward—we obtain an 0.93 F1-score across five deepfake datasets. The method also demonstrates strong out-of-domain (OOD) detection, effectively identifying samples from unseen models at an F1-score of 0.84.

We further analyse these results in a multi-dimensional approach and provide additional insights. All code and data protocols used in this work are available in our open repository: <https://github.com/adrianastan/tada/>.

Index Terms: audio deepfake, model attribution, source tracing, checkpoint attribution, out-of-domain detection

1. Introduction

The rise of deepfake audio has raised concerns about authenticity and security, prompting progress in detection methods to distinguish real from synthetic audio [1–8]. However, an equally critical task is *model attribution* or *source tracing*—identifying the specific generative model behind a deepfake. Unlike detection, which only confirms if audio is synthetic, attribution traces its origin, offering insights into the generator system and enabling targeted countermeasures. We define a *generator* as a complete system comprising both the acoustic model and the vocoder. To emphasize that while different generators may share the same or somewhat similar architectures, it is the model weights that determine the exact source of the audio, for the evaluation section of this work we will predominantly use the term **checkpoint attribution**.

The task of model attribution or source tracing for spoofed or deepfake audio has only recently gained significant attention in the research community. Müller et al. [9] use either signal-based features (e.g., duration, jitter, pitch) or features derived from neural networks to perform clustering on both in-domain and out-of-domain attacks within the ASV19 dataset [10]. Zhu et al. [11] employ ResNet architectures for multi-task classification, focusing on waveform generator, conversion method, and speaker representation. All data used is in-domain, with a speaker-based split. An important contribution to audio deepfake source tracing is the work of Klein et al. [12], which combines various front-end architectures, such as ResNet, self-

supervised learning, and processes them through AASIST [13] or the Whisper encoder [14]. They propose multiple classification tasks, including identifying the acoustic model, vocoder, and input type (speech, text, or bonafide) to distinguish between TTS and VC models. However, they do not address the out-of-distribution (OOD) problem. The AASIST architecture is also applied in [15] on top of SSL features for in- and out-of-domain classification with the ADD2023 dataset. The OOD detection utilizes a novel similarity metric. In terms of interpretable features and attributes, Phukan et al. [16] investigate the ability of various pre-trained models to capture prosodic information for deepfake attribution. In [17], Yan et al. propose a rejection threshold adaptation method for out-of-domain (OOD) detection. They convert closed-set classifiers into open-set classifiers and adaptively predict class-specific thresholds using Gaussian kernels. Xie et al. [18] are the first to explore audio neural codec attribution, introducing a novel dataset of this kind. They propose methods for source tracing and OOD detection over this dataset using AASIST-derived architectures.

Building on previous work, we propose a novel method to identify the sources of generated audio across multiple datasets. Our approach is a training-free, green AI solution that requires no additional modules and leverages a simplified version of a pre-trained SSL model. The **contributions** of our work can be summarized as follows: (i) we introduce the first system for **model attribution across 5 multilingual audio deepfake datasets**, which also incorporates robust **out-of-domain detection** capabilities; (ii) we propose a **highly accurate, ultra-light, computationally efficient, training-free approach** based on k-Nearest Neighbors (kNN), leveraging early features from a pre-trained self-supervised learning (SSL) model; (iii) we conduct an **extensive performance analysis**, demonstrating that the SSL features also capture attributes that pertain to the generative architectures and speaker identities; (iv) we highlight the challenges in **cross-dataset attribution**, revealing significant limitations in generalization across diverse models.

2. Methodology

2.1. Audio Deepfake Datasets

Table 1 provides an overview of the audio datasets chosen for evaluating our work, with a brief description presented below. We emphasize that we selected those datasets where information about the generation method is available in the metadata or can be inferred from the dataset documentation and/or related publications.

ASVspoof 2019 (ASV19) [10] serves as a benchmark for anti-spoofing research and has been extensively used for deepfake detection. We use all 3 subsets: train, dev and eval. **ASVspoof 2021 (ASV21)** [19] expands on previous ASV19 and

Table 1: An overview of the audio deepfake datasets used in this work. The Checkpoints column indicates the number of individual generative systems listed in the datasets’ metadata.

Dataset		Langs.	Checkpoints	Spks	Utts	Duration
ASV19	[10]	en	17	78	63k	3.1±2.9
ASV21	[19]	en	97	14	49k	2.9±2.4
ASV5	[22]	en	32	1159	410k	9.5±4.5
MLAAD v5	[25]	38	82	N/A	154k	8.2±9.0
TIMIT	[24]	en	15	9	20k	3.1±2.3

includes codec manipulate samples, as well samples coming from the Voice Conversion Challenges. As most of the data is already contained in ASV19, we only use the samples from the DF subsection which originate from the Voice Conversion Challenges 2018 [20] and 2020 [21]. We also discard the data which was altered by codecs. **ASVspooF 5 (ASV5)** [22,23] is the latest dataset used in the ASVspooF Challenge. We use the complete train and test and select only the data which had no codec applied to it from the evaluation subset. **TIMIT-TTS (TIM)** [24] combines the TIMIT corpus with text-to-speech (TTS) technology to create a collection of synthetic speech samples. We select the CLEAN subset from it. **Multi-Language Audio Anti-Spoofing Dataset (MLAAD) v5** [25] represents an essential resource as it also contains multispeaker, multilingual checkpoints, as well as some checkpoint overlap with other datasets (as we will show in the following sections). MLAAD does not include details on the vocoders and speakers.

For all datasets we discard the bonafide samples. A total of 741k fake samples coming from 243 (presumably) different text-to-speech synthesis (TTS) or voice cloning (VC) model checkpoints are used.

2.2. Checkpoint Attribution System

The previous works of Wang et al. [26] and Pascu et al. [5] found that large SSL models are better at discriminating between real and fake speech samples. Yet, it may be the case that “smarter” built models may not require as many parameters to incorporate similar information. Therefore we analyse one of the latest wav2vec-based models: `w2v-bert-2.0` [27].¹ The model is lighter (600m parameters) and showed increased performance over various speech classification tasks. As proxy, we select the best deepfake detection model from [5], i.e. `wav2vec2-xls-r-2b`.² Following the approach of Pascu et al. [5], we extract audio features from the frozen SSL model and compute their temporal average. Consistent with our green AI objective, we adopt a simple k-Nearest Neighbors (kNN) model based on Euclidean distance. This model serves to investigate whether the SSL-derived features possess innate properties suitable for the task of audio deepfake model attribution. We experimented with various values of k and identified $k = 21$ as the optimal trade-off between accuracy and computational complexity. Due to space constraints we only report the macro F1-scores as objective measure, and refer the readers to our repository for more numeric results.

3. Evaluation and Analysis

3.1. SSL Model and Layer Selection

The first evaluation step is to verify that the SSL models used for deepfake detection can also perform accurate model attribution. We compare the performance of two selected SSL models,

as well as their intermediate representations (layer-wise) [6,28]. And test their ability to perform the task with different numbers of support samples.³ Tables 2 and 3 depict the F1-scores obtained for each intermediate layer of the two models (columns) when selecting different numbers of samples from each checkpoint to fit the kNN (rows). The number of neighbours is fixed to $k = 21$, and all checkpoints are seen in the support set, i.e. in-domain attribution. There are a total of 243 classes (checkpoints) to predict. We observe very high F1-scores even from the early layers of both models, with as few as 50 samples per model checkpoint: `w2v-bert-2.0` model at layer 4 achieves an F1-score of 0.76, and the `wav2vec2-xls-r-2b` model at layer 9 attains the same F1-score. When 80% of the checkpoint samples are used for fitting the kNN, the F1-scores increase to 0.91 for both models as early as the second layer. Compared to the results of Müller et al. [9] (97.10% accuracy) and Phykan et al. [16] (98.91% accuracy), our kNN-based approach achieves perfect accuracy (100%) and F1-score (1.0) for ASV19 over three 80:20 random splits.

The final rows in Tables 2 and 3 present the average F1-score across different sample sizes for each layer. Notably, the `w2v-bert-2.0` model outperforms the `wav2vec2-xls-r-2b` model, with layer 4 yielding the highest average F1-score (0.76). In line with our green AI focus, we select the `w2v-bert-2.0` model for all subsequent experiments, and extract the features from the best layer. As a result, we use 121m parameters ($\approx 6\%$ of the complete `wav2vec2-xls-r-2b` model) and attain an average 14 msec inference time for 10 second-long audio chunks at 1.13 GB VRAM usage on a Tesla V100 GPU.

3.2. Checkpoint Attribution Analysis

A generative model may be part of multiple deepfake detection datasets. For example, both MLAAD and TIMIT contain samples generated by VITS. However, the VITS model has multiple checkpoints and it is unclear which ones were used by the two datasets. Since the deepfake datasets do not provide this information, we treat each model id listed in the datasets as an independent source of synthetic audio. To gain a deeper understanding of the large and diverse dataset we are using, we conduct an exploratory analysis of its samples. We go back to our kNN approach and examine the 21 closest neighbours of each sample. We count how many of these neighbours originate from a different checkpoint. Figure 1 illustrates these results. The darker shades indicate a higher number of non-target samples. Due to the large number of checkpoints, we only list the datasets in the figure (the full matrix is available in the code repository). Additionally, we note that the matrix is not necessarily symmetrical, as the closest samples in the immediate vicinity may not be the same. For clarity, we omit the main diagonal values (same checkpoint), as they are not relevant to this analysis.

The first observation is that there is very little confusion between the samples pertaining to the different datasets. The single exception is for MLAAD and TIMIT. Upon further investigation, we find that these confusions primarily occur for the VITS and Speedy Speech checkpoints. In the case of TIMIT Speedy Speech, the confusion is notably high, with 46% of the neighbours coming from the MLAAD checkpoint. This suggests that the dataset developers may have used the same checkpoint to generate the samples. Additionally, for TIMIT and MLAAD, more non-target neighbours appear for checkpoints

¹<https://huggingface.co/facebook/w2v-bert-2.0>

²<https://huggingface.co/facebook/wav2vec2-xls-r-2b>

³To emphasize that there is no training involved in instance-based models, we use the *support set* notion instead of *training set*.

Table 2: Mean F1-scores \uparrow across 3 random seeds for the kNN ($k = 21$) checkpoint attribution model with **wav2vec2-xls-r-2b** features. Results for all intermediate layers (columns) using different number of support samples to fit the kNN (rows).

	Layer																																															
Samp	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47	48
1	0.25	0.33	0.31	0.41	0.41	0.43	0.48	0.48	0.48	0.50	0.49	0.50	0.48	0.46	0.47	0.44	0.42	0.44	0.42	0.41	0.37	0.37	0.33	0.34	0.32	0.33	0.32	0.30	0.27	0.32	0.34	0.36	0.36	0.39	0.40	0.44	0.44	0.41	0.45	0.40	0.37	0.32	0.22	0.22	0.23	0.22		
10	0.38	0.46	0.52	0.57	0.60	0.63	0.66	0.67	0.68	0.67	0.66	0.67	0.66	0.65	0.64	0.63	0.63	0.62	0.61	0.60	0.58	0.57	0.51	0.52	0.51	0.51	0.48	0.49	0.48	0.50	0.49	0.54	0.55	0.58	0.60	0.61	0.62	0.62	0.63	0.60	0.58	0.53	0.46	0.42	0.40	0.38	0.38	0.36
50	0.51	0.61	0.66	0.70	0.73	0.73	0.74	0.76	0.75	0.72	0.72	0.71	0.70	0.70	0.69	0.69	0.68	0.68	0.65	0.64	0.63	0.63	0.62	0.63	0.62	0.61	0.63	0.63	0.64	0.67	0.68	0.71	0.73	0.73	0.74	0.74	0.71	0.71	0.67	0.62	0.58	0.56	0.54	0.55	0.53			
100	0.58	0.66	0.71	0.74	0.75	0.75	0.76	0.76	0.77	0.76	0.76	0.77	0.77	0.78	0.77	0.76	0.76	0.75	0.75	0.74	0.72	0.71	0.69	0.68	0.68	0.68	0.66	0.66	0.67	0.67	0.69	0.71	0.73	0.74	0.77	0.77	0.78	0.78	0.76	0.75	0.72	0.68	0.64	0.63	0.62	0.60		
500	0.70	0.75	0.79	0.78	0.79	0.80	0.80	0.79	0.81	0.80	0.80	0.79	0.79	0.80	0.82	0.80	0.81	0.82	0.81	0.81	0.80	0.78	0.76	0.75	0.75	0.75	0.74	0.73	0.75	0.75	0.77	0.78	0.80	0.81	0.83	0.83	0.84	0.84	0.83	0.82	0.80	0.77	0.75	0.74	0.73	0.72	0.71	
1000	0.72	0.76	0.80	0.82	0.83	0.84	0.85	0.85	0.85	0.85	0.84	0.83	0.82	0.83	0.84	0.85	0.83	0.83	0.83	0.83	0.82	0.80	0.78	0.77	0.77	0.76	0.76	0.77	0.78	0.79	0.81	0.82	0.83	0.85	0.85	0.86	0.86	0.85	0.85	0.83	0.81	0.78	0.77	0.77	0.76	0.75	0.73	
80%	0.89	0.91	0.92	0.93	0.93	0.93	0.93	0.93	0.92	0.92	0.91	0.90	0.89	0.86	0.86	0.83	0.82	0.80	0.77	0.75	0.73	0.71	0.70	0.70	0.69	0.68	0.67	0.69	0.68	0.70	0.71	0.74	0.76	0.79	0.82	0.86	0.88	0.87	0.87	0.85	0.82	0.78	0.76	0.75	0.75	0.73		
AVG	0.58	0.64	0.67	0.71	0.72	0.73	0.74	0.75	0.75	0.75	0.74	0.74	0.74	0.73	0.73	0.72	0.72	0.71	0.70	0.69	0.68	0.67	0.64	0.63	0.62	0.63	0.62	0.61	0.60	0.62	0.62	0.64	0.66	0.68	0.69	0.72	0.72	0.73	0.74	0.72	0.71	0.67	0.64	0.60	0.58	0.57	0.57	0.56

Table 3: Mean F1-scores \uparrow across 3 random seeds for the kNN ($k = 21$) checkpoint attribution model with **w2v-bert-2.0** features. Results for all intermediate layers (columns) using different numbers of support samples in the kNN (rows).

	Layer																								
Samp	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	
1	0.20	0.31	0.40	0.42	0.35	0.34	0.35	0.32	0.36	0.32	0.30	0.26	0.25	0.28	0.30	0.34	0.38	0.39	0.39	0.37	0.35	0.29	0.24	0.18	
10	0.35	0.48	0.59	0.59	0.52	0.52	0.51	0.50	0.54	0.51	0.47	0.40	0.39	0.41	0.44	0.44	0.51	0.57	0.58	0.58	0.55	0.54	0.49	0.42	0.30
50	0.52	0.67	0.76	0.76	0.69	0.68	0.70	0.69	0.71	0.66	0.59	0.49	0.49	0.54	0.58	0.66	0.74	0.74	0.74	0.73	0.71	0.67	0.61	0.45	
100	0.62	0.76	0.82	0.82	0.75	0.73	0.76	0.77	0.76	0.71	0.63	0.53	0.53	0.59	0.64	0.73	0.80	0.80	0.79	0.78	0.74	0.68	0.52		
500	0.74	0.85	0.88	0.88	0.83	0.81	0.84	0.84	0.84	0.79	0.70	0.61	0.61	0.68	0.72	0.81	0.87	0.87	0.87	0.86	0.85	0.82	0.77	0.61	
1000	0.80	0.87	0.90	0.89	0.85	0.84	0.87	0.86	0.86	0.81	0.73	0.64	0.64	0.71	0.76	0.83	0.89	0.89	0.88	0.87	0.85	0.81	0.68		
80%	0.87	0.91	0.92	0.93	0.89	0.88	0.88	0.90	0.89	0.85	0.77	0.71	0.71	0.76	0.80	0.87	0.92	0.91	0.91	0.90	0.88	0.86	0.76		
AVG	0.59	0.69	0.75	0.76	0.70	0.69	0.70	0.70	0.71	0.66	0.60	0.52	0.52	0.57	0.61	0.68	0.74	0.74	0.74	0.73	0.71	0.68	0.63	0.50	

using LJSpeech as the synthesis identity. The LJSpeech checkpoints also contribute to intra-dataset confusions. For TIMIT, the only checkpoints that do not exhibit non-target samples in the nearby vicinity are those from the multi-speaker models. This observation may indicate that the **w2v-bert-2.0** features are encoding the speaker information to a large extent.

For ASV21, the separate regions in the plot are a result of the different sources of the checkpoints, i.e. VCC2018 and VCC2020 audio samples. We still observe a considerable number of non-target neighbours, as well as a certain symmetry among them. Both challenges involved two main tasks (parallel and non-parallel or cross-lingual voice conversion), and most teams submitted the same conversion system for both tasks. Additionally, many submissions were very similar in terms of the architecture and/or vocoder. The fact that the **w2v-bert-2.0** model was able to capture these similarities is quite surprising. However, it may be attribute that these non-target neighbours pertain to the same speaker identities used for voice conversion.

In ASV5 we have variants of the same TTS architecture: GlowTTS (checkpoints A01, A02, A03); GradTTS (A04, A05, A06); and ToucanTTS (A09, A10, A21, A22).⁴ Based on informal listening tests, these checkpoints generate different speaker identities. Yet, the **w2v-bert-2.0** features capture neighbouring samples within the same speech generation architecture. This suggests that we may also perform architecture attribution using kNN, which is another promising indication of the model’s capabilities.

With the confusions observed for the LJSpeech samples, we wonder if the entire checkpoint attribution is not in fact a speaker attribution system—given that most checkpoints use different speaker identities. Therefore, we first do a contracted checkpoint attribution over the systems trained only on LJSpeech—a single speaker dataset. A total of 20 checkpoints were selected (see our code repository). The same kNN with $k = 21$ is fitted. We split the data into 80:20 ratio and obtain an 0.74 F1-score for the attribution. If this were a mere speaker attribution, we would have seen a much lower performance. We also examine the multi-speaker models, which, from a checkpoint perspective represent the same source. We select the multispeaker models from TIMIT and MLAAD, and separate them based on speaker and language (a different label for each identity). We obtain an 0.72 F1-score for this attribution.

⁴Samples altered by adversarial attacks are not considered.




Figure 1: Percent of samples (colour shades) in the 21-neighbour vicinity of the **w2v-bert-2.0** features-based kNN which do not pertain to the target class. Main diagonal values (same checkpoint) and individual checkpoint names are omitted for readability.

This is in contrast to the 0.99 F1-score we obtain when not splitting the multi-speaker checkpoints. So it appears that although the speaker attributes are captured by the representations, other architecture-indicative signal characteristics are inherent to the **w2v-bert-2.0** features.

3.3. Unseen Checkpoint Detection

Speech generation is an ever-evolving field, with new systems being introduced regularly. Framing the task of deepfake source tracing as a closed classification problem is insufficient for practical, end-user applications. To address this scenario, we investigate how our method can also be applied to detect unseen or out-of-domain (OOD) checkpoints [29]. Here, we define out-of-domain data as checkpoints for which several samples are available, but no labels have yet been assigned to them.

Vaze et al. [30] showed that out-of-domain classification can be efficiently performed when a robust closed-set classifier is available. We already have such a classifier, and we simulate out-of-distribution detection by randomly setting aside four checkpoints from each of our five datasets (i.e., 20 checkpoints out of a total of 243). Half of these checkpoints are used for validation, and half for evaluation, with no overlap between the sets. The remaining in-domain class samples are then split into support, validation, and test subsets with an 80:10:10 ratio. We fit the kNN model over the support set of in-domain samples. Next, we compute the average Euclidean distance of all valida-

Table 4: Mean and standard deviation of the F1-scores \uparrow reported across 3 random seeds for the kNN-based **OOD detection**. We use different numbers of neighbours (rows) and report results for the complete (All) and individual subsets (columns).

k	ASV19	ASV21	ASV5	MLAAD	TIMIT	All
1	0.87±0.02	0.58±0.12	0.86±0.06	0.78±0.11	0.70±0.07	0.88±0.03
5	0.88±0.04	0.58±0.10	0.84±0.06	0.74±0.14	0.69±0.07	0.86±0.04
21	0.86±0.03	0.58±0.08	0.81±0.07	0.71±0.15	0.66±0.06	0.84±0.04

Table 5: F1-scores \uparrow for the kNN-based **acoustic model and vocoder classification** with different datasets and data splits. The splits in MLAAD3 are those reported in [12]. AM stands for acoustic model, and Voc stands for vocoder attribution. The Split column refers to the use of the same checkpoints for fitting and evaluation (ID) or setting aside one (OOD-1) or half of the checkpoints (OOD-H) from the support set. MLD3 = MLAAD v3; MLD5 = MLAAD v5; and TIM = TIMIT.

Support set	Test data	Split	AM	Voc
0 MLD3 baseline [12]			0.82	0.93
1 MLD3 train	MLD3 test	ID	0.89	0.90
2 MLD3 train+dev	MLD3 test	ID	0.96	0.96
3 ASV5+MLD5+TIM	ASV5+MLD5+TIM	ID	0.99±0.01	N/A
4 MLD5	ASV5+TIM	OOD	0.30±0.00	N/A
5 ASV5+MLD5+TIM	ASV5+MLD5+TIM	OOD-1	0.55±0.19	N/A
6 ASV5+MLD5+TIM	ASV5+MLD5+TIM	OOD-H	0.48±0.10	N/A

tion data points to their k closest neighbours and establish an OOD detection threshold based on the Equal Error Rate (EER) of this subset. The same average distance is then computed for the test samples. The final OOD decision is made relative to the threshold derived from the validation set.

The mean and standard deviation of the F1-score across three different OOD partitions are presented in Table 4. The results are shown in terms of the number of neighbours (rows), with individual dataset results, as well as results for the entire test set (columns). For the OOD evaluation on individual datasets, we compare the OOD samples from each dataset against all in-domain samples from all datasets, meaning that a single OOD detector is trained, and its predictions are processed accordingly. We observe a relatively high average F1-score of 0.84 ($k = 21$) across all datasets, with some variations among them. Best results are for $k = 1$, but in practice this may not be a feasible option due to improper labelling. The performance for individual datasets appears strongly correlated with the results from Section 3.2: a higher number of non-target neighbours for a dataset’s checkpoints translates into a worse OOD detection performance.

For the OOD checkpoints, the previous section showed that as little as 10 random labelled samples in the support set would yield an F1-score of 0.59 for in-domain attribution.

3.4. Beyond checkpoint attribution

Previous work has also explored the provision of more explainable decisions or attributes for speech generation checkpoints [12, 31]. We follow the tasks outlined in [12] and use their evaluation protocol for MLAAD. It is important to note that Klein et al. [12] used version 3 of MLAAD, and for a direct comparison, we revert to that version in this experiment. Two files from the test set are discarded—they are very long (>90 seconds) and crashed the inference, and we do not use the genuine files for attribution purposes. The same front-end and kNN classifier ($k = 21$) were employed. The results are presented in

Table 5, rows 1 and 2, and clearly surpass the baseline (row 0).

Based on these results, it might be tempting to assume that identifying the acoustic model and vocoder is a trivial task. However, the evaluation in [12] is limited to intra-dataset in-domain scenarios. As indicated by our results in Section 3.2 this is essentially a larger grouping of the intra-dataset samples. To further assess the robustness of our approach, we proceed to test *cross-dataset* and *out-of-domain* acoustic model attribution. Vocoder attribution is less reliable due to the information not being clearly marked in the datasets’ metadata. We select the following architectures common to ASV5, MLAAD v5, and TIMIT: FastPitch, GlowTTS, Tacotron, Tacotron2, VITS, and xTTS. First, we compute the F1-score using a random 80:20 train-test split ratio of the samples. The results are shown in Table 5 row 3. While the performance is notably high, this is the same in-domain attribution, and does not capture the challenge of cross-dataset or out-of-domain attribution.

Next, for a true out-of-domain acoustic model attribution, we use the checkpoints from MLAAD v5 to build the kNN support set, and test on TIMIT and ASV5. The results, shown in row 4 of Table 5 indicate a significant drop in performance (F1-score of 0.30). The majority of samples are classified as xTTS, which is the dominant class in the support set (33k samples). To address the class imbalance, we experimented with under-sampling approaches and Condensed kNN [32], but the results remain similar. So it may be that MLAAD is not sufficiently diverse to generalise to the other datasets. We lastly employ a leave-N-out strategy, grouping all the samples from the datasets into six acoustic models, and randomly selecting one or more checkpoints from each architecture for testing. The results show that performance remains low, with an F1-score of 0.55 when a single checkpoint is set aside for testing (row 5, OOD-1) and 0.48 when half of the checkpoints are set aside (row 6, OOD-H). It is also important to mark the wide variance in the last two rows, which means that different unseen checkpoints pose different issues to the attribution system.

As a result, although checkpoint attribution appears to be a rather simple task for our training-free `w2v-bert-2.0` features, unravelling the underlying architecture or vocoder for the model may require alternative classification strategies, as well as most-likely more fine-grained audio datasets.

4. Conclusions

In this work we showed that using a powerful pretrained SSL model alleviates the need for any additional, complex back-end modules to accurately predict the source of a generated speech sample. We also showed that early layers of the SSL models are sufficiently informative for kNN-based checkpoint attribution (0.93 macro F1-score). The same setup can be successfully applied to detect samples from out-of-domain checkpoints, unseen in the kNN support set (0.84 F1-score). As a result, our framework is lightweight and training-free, in line with green AI endeavours.

We argue that our results support the exploitation of pretrained models as more than just naive feature extractors. They should rather be considered as a fundamental source of information that needs to be further understood, perhaps in a more traditional manner (e.g. feature analysis). As future work, we would like to understand how minor checkpoint changes may impact this very accurate attribution. For example, would a few-shot adaptation process for a TTS model change its attribution? How about the outputs from earlier epochs of the training?

Acknowledgement. This work was co-funded by EU Horizon project AI4TRUST (No. 101070190), and by the Romanian Ministry of Research, Innovation and Digitization project DLT-AI SECSPP (ID: PN-IV-P6-6.3-SOL-2024-2-0312).

5. References

- [1] X. Wang and J. Yamagishi, “A comparative study on recent neural spoofing countermeasures for synthetic speech detection,” in *Proc. Interspeech*, 2021.
- [2] P. Kawa, M. Plata, M. Czuba, P. Szymański, and P. Syga, “Improved deepfake detection using Whisper features,” in *Proc. Interspeech*, 2023.
- [3] G. Channing, J. Sock, R. Clark, P. Torr, and C. S. de Witt, “Toward robust real-world audio deepfake detection: Closing the explainability gap,” *arXiv preprint arXiv:2410.07436*, 2024.
- [4] D.-T. Truong, R. Tao, T. Nguyen, H.-T. Luong, K. A. Lee, and E. S. Chng, “Temporal-channel modeling in multi-head self-attention for synthetic speech detection,” *arXiv preprint arXiv:2406.17376*, 2024.
- [5] O. Pascu, A. Stan, D. Oneata, E. Oneata, and H. Cucu, “Towards generalisable and calibrated audio deepfake detection with self-supervised representations,” in *Proc. Interspeech*, 2024.
- [6] J. M. Martín-Doñas, A. Álvarez, E. Rosello, A. M. Gomez, and A. M. Peinado, “Exploring self-supervised embeddings and synthetic data augmentation for robust audio deepfake detection,” in *Proc. Interspeech*, 2024.
- [7] Z. Pan, T. Liu, H. B. Sailor, and Q. Wang, “Attentive merging of hidden embeddings from pre-trained speech model for anti-spoofing detection,” in *Proc. Interspeech*, 2024.
- [8] Y. Chen, J. Yi, J. Xue, C. Wang, X. Zhang, S. Dong, S. Zeng, J. Tao, Z. Lv, and C. Fan, “RawBMamba: End-to-end bidirectional state space model for audio deepfake detection,” in *Proc. Interspeech*, 2024.
- [9] N. Müller, F. Diekmann, and J. Williams, “Attacker attribution of audio deepfakes,” in *Proc. Interspeech*, 2022.
- [10] A. Nautsch, X. Wang, N. 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, no. 2, 2021.
- [11] T. Zhu, X. Wang, X. Qin, and M. Li, “Source tracing: Detecting voice spoofing,” in *Proc. APSIPA*, 2022.
- [12] N. Klein, T. Chen, H. Tak, R. Casal, and E. Khoury, “Source tracing of audio deepfake systems,” in *Proc. Interspeech*, 2024.
- [13] J.-w. Jung, H.-S. Heo, H. Tak, H.-j. Shim, J. S. Chung, B.-J. Lee, H.-J. Yu, and N. Evans, “AASIST: Audio anti-spoofing using integrated spectro-temporal graph attention networks,” in *Proc. ICASSP*, 2022.
- [14] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, “Robust speech recognition via large-scale weak supervision,” *arXiv preprint arXiv:2212.043562*, 2022.
- [15] Y. Xie, R. Fu, Z. Wen, Z. Wang, X. Wang, H. Cheng, L. Ye, and J. Tao, “Generalized source tracing: Detecting novel audio deepfake algorithm with real emphasis and fake dispersion strategy,” in *Proc. Interspeech*, 2024, pp. 4833–4837.
- [16] O. C. Phukan, D. Singh, S. R. Behera, A. B. Buduru, and R. Sharma, “Investigating prosodic signatures via speech pre-trained models for audio deepfake source attribution,” *arXiv preprint arXiv:2412.177962*, 2024.
- [17] X. Yan, J. Yi, J. Tao, Y. Chen, H. Gu, G. Li, J. Zhou, Y. Ren, and T. Xu, “Reject threshold adaptation for open-set model attribution of deepfake audio,” *arXiv preprint arXiv:2412.014252*, 2024.
- [18] Y. Xie, X. Wang, Z. Wang, R. Fu, Z. Wen, S. Cao, L. Ma, C. Li, H. Cheng, and L. Ye, “Neural codec source tracing: Toward comprehensive attribution in open-set condition,” *arXiv preprint arXiv:2501.06514*, 2025.
- [19] X. Liu, X. Wang, M. Sahidullah, J. Patino, H. Delgado, T. Kinnunen, M. Todisco, J. Yamagishi, N. Evans, A. Nautsch, and K. A. Lee, “ASVspoof 2021: Towards spoofed and deepfake speech detection in the wild,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 31, pp. 2507–2522, 2023.
- [20] J. Lorenzo-Trueba, J. Yamagishi, T. Toda, D. Saito, F. Villavicencio, T. Kinnunen, and Z. Ling, “The voice conversion challenge 2018: Promoting development of parallel and nonparallel methods,” *arXiv preprint arXiv:1804.04262*, 2018.
- [21] Z. Yi, W.-C. Huang, X. Tian, J. Yamagishi, R. K. Das, T. Kinnunen, Z.-H. Ling, and T. Toda, “Voice conversion challenge 2020: Intra-lingual semi-parallel and cross-lingual voice conversion,” in *Joint Workshop for the Blizzard Challenge and Voice Conversion Challenge 2020*, 2020.
- [22] X. Wang, H. Delgado, H. Tak, J. weon Jung, H. jin Shim, M. Todisco, I. Kukanov, X. Liu, M. Sahidullah, T. H. Kinnunen, N. Evans, K. A. Lee, and J. Yamagishi, “ASVspoof 5: Crowdsourced speech data, deepfakes, and adversarial attacks at scale,” in *Automatic Speaker Verification Spoofing Countermeasures Workshop (ASVspoof 2024)*, 2024.
- [23] X. Wang, H. Delgado, H. Tak, J. weon Jung, H. jin Shim, M. Todisco, I. Kukanov, X. Liu, M. Sahidullah, T. Kinnunen, N. Evans, K. A. Lee, J. Yamagishi, M. Jeong, G. Zhu, Y. Zang, Y. Zhang, S. Maiti, F. Lux, N. Müller, W. Zhang, C. Sun, S. Hou, S. Lyu, S. L. Maguer, C. Gong, H. Guo, L. Chen, and V. Singh, “ASVspoof 5: Design, collection and validation of resources for spoofing, deepfake, and adversarial attack detection using crowd-sourced speech,” *arXiv preprint arXiv:2502.08857*, 2025.
- [24] D. Salvi, B. Hosler, P. Bestagini, M. C. Stamm, and S. Tubaro, “TIMIT-TTS: A text-to-speech dataset for multimodal synthetic media detection,” *IEEE Access*, vol. 11, 2023.
- [25] N. M. Müller, P. Kawa, W. H. Choong, E. Casanova, E. Gölge, T. Müller, P. Syga, P. Sperl, and K. Böttinger, “MLAAD: The multi-language audio anti-spoofing dataset,” in *Proc. IJCNN*, 2024.
- [26] X. Wang and J. Yamagishi, “Can large-scale vocoded spoofed data improve speech spoofing countermeasure with a self-supervised front end?” in *Proc. ICASSP*, 2024.
- [27] S. Communication, L. Barrault, Y.-A. Chung, M. C. Meglioli, D. Dale, N. Dong, M. Duppenhaler, P.-A. Duquenne, B. Ellis, H. ElSahar, J. Haaheim, J. Hoffman, M.-J. Hwang, H. Inaguma, and C. K. et al., “Seamless: Multilingual expressive and streaming speech translation,” *arXiv preprint arXiv:2312.05187*, 2023.
- [28] Q. Zhang, S. Wen, and T. Hu, “Audio deepfake detection with self-supervised XLS-r and SLS classifier,” in *Proc. ACM Multimedia*, 2024.
- [29] J. Lu, Y. Zhang, Z. Li, Z. Shang, W. Wang, and P. Zhang, “Detecting unknown speech spoofing algorithms with nearest neighbors,” in *Proceedings of IJCAI Workshop on Deepfake Audio Detection and Analysis*, 2023.
- [30] S. Vaze, K. Han, A. Vedaldi, and A. Zisserman, “Open-set recognition: A good closed-set classifier is all you need?” in *Proc. ICLR*, 2022.
- [31] M. Chhibber, J. Mishra, H. Shim, and T. H. Kinnunen, “An explainable probabilistic attribute embedding approach for spoofed speech characterization,” in *Proc. ICASSP*, 2025.
- [32] P. Hart, “The condensed nearest neighbor rule (corresp.),” *IEEE Transactions on Information Theory*, vol. 14, no. 3, pp. 515–516, 1968.