

The PartialSpoof Database and Countermeasures for the Detection of Short Fake Speech Segments Embedded in an Utterance

Lin Zhang , *Student Member, IEEE*, Xin Wang , *Member, IEEE*, Erica Cooper , *Member, IEEE*, Nicholas Evans , *Member, IEEE*, and Junichi Yamagishi , *Senior Member, IEEE*

Abstract—Automatic speaker verification is susceptible to various manipulations and spoofing, such as text-to-speech synthesis, voice conversion, replay, tampering, adversarial attacks, and so on. We consider a new spoofing scenario called “*Partial Spoof*” (PS) in which synthesized or transformed speech segments are embedded into a bona fide utterance. While existing countermeasures (CMs) can detect fully spoofed utterances, there is a need for their adaptation or extension to the PS scenario. We propose various improvements to construct a significantly more accurate CM that can detect and locate short-generated spoofed speech segments at finer temporal resolutions. First, we introduce newly developed self-supervised pre-trained models as enhanced feature extractors. Second, we extend our PartialSpoof database by adding segment labels for various temporal resolutions. Since the short spoofed speech segments to be embedded by attackers are of variable length, six different temporal resolutions are considered, ranging from as short as 20 ms to as large as 640 ms. Third, we propose a new CM that enables the simultaneous use of the segment-level labels at different temporal resolutions as well as utterance-level labels to execute utterance- and segment-level detection at the same time. We also show that the proposed CM is capable of detecting spoofing at the utterance level with low error rates in the PS scenario as well as in a related logical access (LA) scenario. The equal error rates of utterance-level detection on the PartialSpoof database and ASVspoof 2019 LA database were 0.77 and 0.90%, respectively.

Index Terms—Anti-spoofing, deepfake, PartialSpoof, self-supervised learning, spoof localization, countermeasure.

I. INTRODUCTION

SPEECH technologies play a crucial role in many aspects of life, e.g., keyword spotting in smart home devices,

Manuscript received 8 April 2022; revised 31 October 2022 and 14 December 2022; accepted 18 December 2022. Date of publication 30 December 2022; date of current version 16 January 2023. This work was supported in part by JST CREST through the Japanese-French Joint National VoicePersonae Project under Grants JPMJCR18A6, JPMJCR20D3, and JPMJFS2136, in part by the ANR under Grant ANR-18-JSTS-0001, in part by MEXT KAKENHI under Grants 21K17775, 21H04906, 21K11951, and 18H04112, and in part by Google AI for Japan Program. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Omid Sadjadi. (Corresponding author: Lin Zhang.)

Lin Zhang, Xin Wang, Erica Cooper, and Junichi Yamagishi are with the National Institute of Informatics, Tokyo 101-8340, Japan (e-mail: zhanglin@nii.ac.jp; wangxin@nii.ac.jp; ecooper@nii.ac.jp; jyamagis@nii.ac.jp).

Nicholas Evans is with the EURECOM, Campus SophiaTech, 06410 Biot, France (e-mail: evans@eurecom.fr).

Digital Object Identifier 10.1109/TASLP.2022.3233236

speaker recognition in online banking, diarization of meeting recordings, and speech recognition for captioning news broadcasts. However, such technologies are also vulnerable to spoofing — synthesized, transformed, or manipulated speech can fool machines and even humans. A number of initiatives and challenges such as ASVspoof [1], [2], [3], [4] have hence been organized to encourage research in countermeasure (CM) solutions, which are needed to protect speech applications and human listeners from spoofing attacks. Several types of spoofing scenarios have been considered and explored, including logical access (LA), DeepFake (DF), and physical access (PA). The LA scenario is designed for text-to-speech (TTS) synthesis and voice conversion (VC) attacks with and without telephony codecs, reflecting a use case of authentication in call centers. The DF scenario is similar to the LA scenario, but takes into account audio under strong compression for media streaming without an authentication process. The PA scenario targets the development of CMs against replay attacks.

In the LA and DF scenarios, entire audio signals are generated using TTS or VC algorithms. Missing in past work is the consideration of scenarios in which synthesized or transformed speech segments are embedded into a bona fide speech utterance such that only a fraction of an utterance is spoofed. There are many possible motives for attackers to take such an approach. To give a few examples, specific words or phrases may be replaced with different ones, and negation words, such as “not” may be generated using TTS or VC, and inserted into an original utterance to completely change the meaning of a given sentence. If an attacker has an audio file containing a phrase such as “Search Google” for a particular person, the attacker can replace the word “Search” with the synthesized phrase “OK” and attempt a presentation attack against a text-dependent automatic speaker verification (ASV) system running on the person’s device. An attacker could also use segments of units smaller than words. If an attacker synthesizes certain vowels and replaces the original vowels with the synthesized ones, he or she can manipulate words such as “bat,” “bet,” “bit,” “bot,” and “but.” If the attacker has knowledge of phonology, he or she can use even smaller units and manipulate the consonants /b/, /g/, and /d/ by synthesizing and replacing only the transitional part of the second formant, which is an acoustic cue for the consonants. With modern speech-synthesis technologies having the ability to

produce high-quality speech resembling a given target speaker’s voice, these types of partial audio manipulations are becoming more likely to occur. We call this new spoofing scenario “*Partial Spoof*” (PS).

We believe that speech utterances containing such short generated spoofed audio segments will likely be difficult to detect using CMs trained for the LA or DF scenarios since they typically use aggregation operations over time, in which case the short segments will have little bearing on the final CM score. Therefore, new CMs for the PS scenario are needed.

How should CMs be implemented for the PS scenario? We propose that CMs for the PS scenario should have two functions. The first involves the simple utterance-level detection of utterances containing any short spoofed segments. This is similar to the standard *spoof vs. bona fide* classification in the LA and DF scenarios. The second function is to detect which segments in an utterance are spoofed. This is to show which specific parts of an audio sample may have been generated using TTS or VC and to improve the explainability of the CMs. We refer to the former as “utterance-level detection” and the latter as “segment-level detection.” While earlier we described examples of spoofing at the phone and word levels, in practice, there is no restriction on the speech unit to be used by attackers. Segment-level detection should be carried out without prior knowledge of the length of the unit used by the attacker and needs to support variable-length segments.

In our previous study [5], we constructed a speech database called “PartialSpoof” designed for the PS scenario and reported initial results for utterance- and segment-level detection. In this preliminary work, segment-level detection was applied in a straightforward manner with a fixed temporal resolution. Whether the detection is per-utterance or per-segment, the features needed for CMs are likely to be similar. Therefore, in our subsequent work [6], we constructed a CM that simultaneously executes detection at the utterance level and (fixed) segment level using multi-task learning. However, these two studies showed that there is much room for improvement.

We thus extend our previous studies to improve the utterance- and segment-level detection of CMs in the PS scenario. We use self-supervised pre-trained models based on wav2vec 2.0 [7] and WavLM [8] as enhanced feature extractors. We also extend the PartialSpoof database. Specifically, we add segment-level labels for various temporal resolutions, instead of only for a fixed temporal resolution. The spoofed audio segments to be embedded by attackers are of variable length; thus, by using these labels, CM models can be trained to execute detection at various temporal resolutions. In the extended PartialSpoof database, there are six different temporal resolutions for these segment-level labels, ranging from as fine as 20 ms to as coarse as 640 ms. We also propose a CM that enables the simultaneous use of the segment-level labels at different temporal resolutions as well as utterance-level labels to execute utterance- and segment-level detection at the same time.

The remainder of this paper is structured as follows. In Section II, we explain the PS scenario in more detail and review relevant topics. In Section III, we explain our PartialSpoof database with new segment labels at various temporal resolutions. In

Section IV, we define the tasks and overview existing CMs for the PS scenario. Then, in Section V, we introduce the proposed CM. We present experiments and results in Section VI, and conclude the paper and discuss future work in Section VII.

II. WHAT IS THE PS SCENARIO?

A. Spoofing in the PS Scenario and Related Scenarios

In the PS scenario, we consider manipulated audio in which generated audio segments are embedded in a bona fide speech utterance and vice versa. The speaker characteristics of the embedded segments are similar to those of the true speaker. However, due to differences in the performance of TTS or VC methods, there are differences in speaker similarity. Although the PS scenario resembles other scenarios in which spoofing and CMs have already been widely studied, there are some crucial differences.

The conventional LA and DF scenarios assume that speech in a single utterance is generated using a single TTS or VC method. In the PS scenario, however, a single speech utterance may contain audio segments generated using more than one TTS or VC method, even if the entire utterance consists only of spoofed segments.

The PS scenario is also closely related to copy-moves and splicing, which are scenarios well-studied for tampering forgery [9], [10]. The copy-move method of audio forgery involves the copying and pasting of segments within the same bona fide audio sample. Splicing forgery involves assembling a speech utterance using spliced segments obtained from other audio recordings¹. Spoofed audio in the PS scenario is hence a special case of splicing forgery in which segments do not come from other bona fide audio recordings but are generated using TTS or VC.

B. New Realistic Threats

Certain companies have started to develop technologies and services that enable users to modify specific segments of a speech recording using TTS or VC without affecting other segments, making the final utterance match as closely as possible to the original, e.g., [13], [14], [15]. Although such technologies and services are desirable for users who would like to manipulate their own speech without re-recording, they increase the possibility of misuse in the form of impersonation and fraud. They may also pose a threat to other speech-based applications such as text-dependent ASV. There is thus a need to develop new CMs that are capable of detecting partially-spoofed utterances.

C. Databases for PS Scenario

During the same time as when we built the initial PartialSpoof database in 2021 [5], another database was also proposed for the PS scenario [16] for which a single multi-speaker TTS system was used to replace a single word within an utterance. This later became a part of the Audio Deep synthesis Detection (ADD)

¹In the TTS field, “unit selection” techniques [11], [12] apply a similar splicing operation on the basis of dynamic time warping to reduce concatenation artifacts.

challenge database [17]. In addition to the limited number of TTS systems, the database has predetermined the target for replacement as a word, which is a form of prior knowledge even on the detection side. In contrast, we consider a spoofing scenario based on variable-length spoofed segments generated using several multi-speaker TTS and VC systems.

III. NEW PARTIALSPOOF DATABASE

With the goal of stimulating research on CM models suitable for the PS scenario, an early version of the PartialSpoof² database [5] containing the spoofed speech needed for model training was built and made available to the research community in 2021. In the PS scenario, attackers may create partially-spoofed audio in various ways. Attackers may insert segments of spoofed speech generated by TTS or VC into a natural utterance to add new content or substitute parts of the original utterance with spoofed speech to replace content. Alternately, instead of starting with a natural utterance and inserting or replacing segments with spoofed speech, they may start with an utterance generated by TTS or VC and insert or replace segments with natural speech, for example, in the case in which the attackers only have a small amount of bona fide data. Thus, it is necessary for the database to contain spoofed speech with different proportions of generated audio segments within a single utterance. We call the ratio of the duration of generated audio segments in an utterance to the total duration of the utterance “intra-speech generated segment ratio.” An efficient way to achieve this is to prepare pairs of speech utterances consisting of one that is entirely generated using TTS or VC and one original bona fide speech utterance then randomly substitute short segments within the pair. This enables us to create a large number of spoofed audio files with different intra-speech generated segment ratios. This is a key concept of the PartialSpoof database and we describe its construction policies and procedures as well as newly generated segment-level labels for a new version of the PartialSpoof database.

A. Construction Policies

Since the methods used by attackers are unknown in practice, these will likely be different from the training data used to train a real system. Therefore, we should assume that TTS/VC methods for generating spoofed audio segments in the training set are mostly different from those in the evaluation set. To achieve this, each subset of the ASVspoof 2019 LA database [18] was used to construct each corresponding subset of the PartialSpoof database. It was also assumed that the attacker could use generated audio segments at a variety of acoustic units, not limited to linguistic units such as words. Hence, we used variable-length speech segments found by voice activity detection (VAD) regardless of the content of the speech.

Finally, we assumed that the attacker could carry out not advanced but basic digital signal processing (DSP) to reduce artifacts of the spoofed audio for the purpose of a reliable assessment of CMs. Since it is unrealistic to assume that advanced

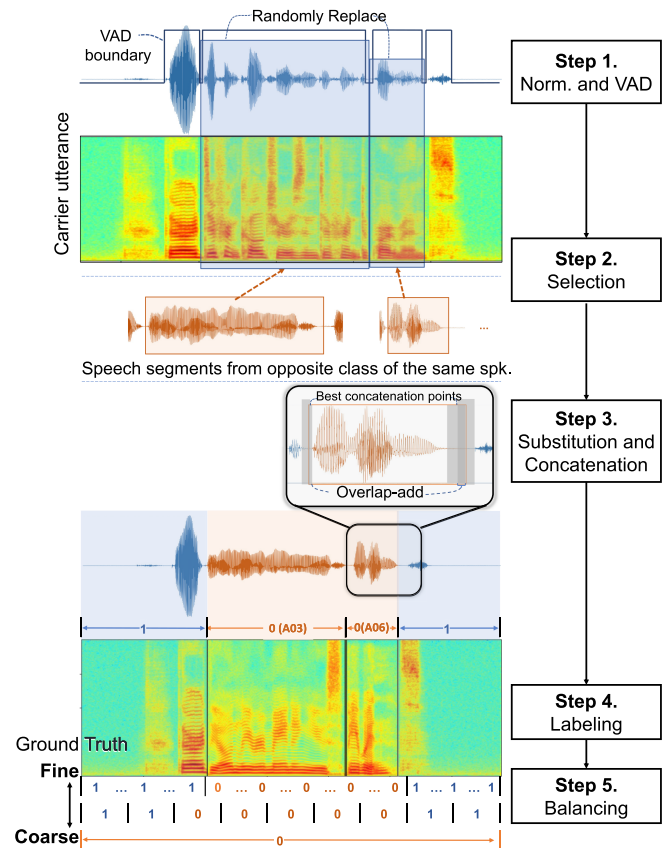


Fig. 1. Procedure to automatically construct our PartialSpoof database. (0 refers to *spoof*, and 1 refers to *bona fide*).

processing will be carried out, the DSP and VAD used by the attacker were assumed to be basic methods, and certain errors and artifacts were assumed to occur naturally.

B. Construction Procedure

The following is the procedure for constructing the PartialSpoof database automatically. As shown in Fig. 1, the entire processing pipeline involves five stages:

Step 1. Normalization and VAD: When substituting segments into another speech utterance, the volume of the segments and utterance may be different. If there are significant volume mismatches, it becomes straightforward to detect the substituted segments, which would interfere with the CM’s ability to learn fine-grained mismatches. Therefore, we first normalized the waveform amplitudes of the original and spoofed utterances contained within the ASVspoof 2019 LA database to -26 dBov according to the ITU-T SV56 standard [19]. Next, to select the variable-length candidate segments, three types of VAD [20], [21], [22] were used, and a majority vote was taken. Segments detected as speech regions by two or more types of VAD within a human speech utterance were considered as candidate regions to be replaced with generated segments. Candidate segments within spoofed utterances were also found in the same manner.

Step 2. Selection: The next step was to determine which segments found by the VAD should be replaced with which

²<https://zenodo.org/record/5766198>

TABLE I

DETAILS OF TRIALS IN PARTIALSPOOF DATASET. NUMBER OF AUDIO FILES, TOTAL DURATION, MAXIMUM NUMBER OF UNIQUE TTS AND VC METHODS WITHIN ONE UTTERANCE, AND STATISTICS OF AUDIO LENGTH FOR EACH SUBSET ARE SHOWN

	Subset	# of Samples	Duration (h)	Max # of TTS/VC	Audio length (s)		
					min	mean	max
<i>bona fide</i>	Train	2,580	2.43	-	1.36	3.39	11.13
	Dev.	2,548	2.48	-	1.28	3.51	11.39
	Eval.	7,355	6.94	-	0.90	3.40	13.01
<i>spooft</i>	Train	22,800	21.82	6	0.60	3.45	21.02
	Dev.	22,296	21.86	6	0.62	3.53	15.34
	Eval.	63,882	60.74	9	0.48	3.42	18.20

segments from the other class of the same speaker. We considered both directions, that is, replacement of a randomly chosen segment from a *bona fide* utterance with a spoofed segment and substitution of a spoofed segment into a *bona fide* segment. The following conditions were used to select the segments: (1) segments are chosen not from the same utterance but from different utterances of the same speaker; (2) the same segment cannot be injected into an utterance more than once; (3) original and substitute segments must be of similar duration.

Step 3. Substitution and concatenation: We then conducted substitution of the speech segments. To concatenate the segments without significant artifacts, we computed time-domain cross-correlation to find the best concatenation points. When searching for these points, parts of the silent regions around the corresponding speech were considered and the waveform overlap-add method was used.

Step 4. Labeling: After the concatenation of the segments, each utterance was annotated with fine-grained segment labels at various temporal resolutions in addition to utterance-level labels. Speech frames generated by TTS/VC and those from the *bona fide* utterances were labeled as *spooft* and *bona fide*, respectively. Each segment or utterance that contains one or more generated frames was labeled as *spooft*, otherwise *bona fide*.

Step 5. Post-processing: To balance the subsets, intra-speech generated segment ratios of the concatenated audio were quantized into ten levels, and then random sampling was done for each level. As a result, audio files with a small intra-speech generated segment ratio and those with a large intra-speech generated segment ratio were equally included in the database and in each subset.³ The number of files that we randomly sampled was also the same as the *spooft* class of the ASVspoof 2019 LA database.

C. Database Statistics

The total number of samples for each temporal resolution is shown in Table II. The sample size increases as resolution increases from utterance (utt.) to 20 ms. The percentage of samples belonging to the *spooft* class in each temporal resolution is shown in Table III. If any part of the generated waveform is

³The result of this process is that there is less mismatch between the subsets regarding the intra-speech generated segment ratio. It is also possible to design and adopt mismatched conditions; however, in this paper, we consider the matched condition to simplify our analysis.

TABLE II

NUMBER (IN THOUSANDS) OF SAMPLES IN EACH TEMPORAL RESOLUTION

Subset	Temporal resolution (ms)						utt.
	20	40	80	160	320	640	
Train	4,347	2,167	1,077	532	260	124	25
Dev.	4,364	2,176	1,082	535	261	125	25
Eval.	12,129	6,047	3,006	1,485	725	346	71

TABLE III

PERCENTAGES (%) OF SPOOF CLASS IN EACH TEMPORAL RESOLUTION

Subset	Temporal resolution (ms)						utt.
	20	40	80	160	320	640	
Train	43.79	45.10	47.76	53.00	61.51	73.77	89.83
Dev.	42.96	44.30	47.02	52.31	60.86	73.41	89.74
Eval.	38.46	39.81	42.60	48.03	57.52	71.19	89.68

contained in the target segment to be annotated, it is labeled as *spooft*. We can see that the finer the temporal resolution of the segment, the less likely it is to contain generated audio signals; thus, the number of samples in the *bona fide* class increases at finer-grained resolutions. As can be seen from the table, the data have a bias toward the *spooft* class on a per utterance basis, but the extreme bias is eliminated on finer segment bases. In multi-task learning using labels at multiple temporal resolutions, such differences in bias may lead to more robust model learning.

D. Limitations

In the PartialSpoof database, the variable-length speech segments found by the VADs are replaced without considering the meaning of sentences and words as well as the phonemes before and after the segments. Therefore, the speech in the database is not partially-spoofed speech that is intended to deceive listeners and deliver wrong linguistic messages to them but an approximation of it. On the other hand, since partially-spoofed segments of the PartialSpoof database are of variable length, CMs built on this database can execute segment-level detection at various temporal resolutions, and their accuracy can be evaluated.

IV. OVERVIEW OF COUNTERMEASURES FOR PS SCENARIO

In this section we switch to the defender's side and explain the tasks of CMs and existing CM architectures for the PS scenario. We use the notations listed in Table IV to facilitate the explanation.

A. Task Definition

A CM for the PS scenario is required to conduct utterance- and segment-level detection on an input waveform $x_{1:T}$. The two tasks can be defined more formally as follows.

TABLE IV
LIST OF MATH NOTATIONS USED IN THIS PAPER

$x_t \in \mathbb{R}$	a waveform sampling point at time t
$\mathbf{x}_{1:T} = (x_1, x_2, \dots, x_T)$	a waveform with T sampling points
$\mathbf{a}_n \in \mathbb{R}^{D_a}$	an acoustic feature vector at the n -th frame
$\mathbf{a}_{1:N} = (\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_N)$	an acoustic feature sequence with N frames
$\mathbf{h}_m \in \mathbb{R}^{D_h}$	a hidden feature vector at the m -th segment
$\mathbf{h}_{1:M} = (\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_M)$	a hidden feature sequence with M segments
$y^u \in \{\textit{bona fide}, \textit{spoof}\}$	an utterance-level label
$\mathbf{e}^u \in \mathbb{R}^{D_e}$	an utterance-level embedding vector
$s^u \in \mathbb{R}$	an utterance-level CM score
$\mathbf{y}_{1:M} = (y_1, y_2, \dots, y_M)$	segment-level labels where y_m is the label for the m -th segment
$\mathbf{e}_{1:M} = (\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_M)$	segment-level embeddings where \mathbf{e}_m is for the m -th segment
$\mathbf{s}_{1:M} = (s_1, s_2, \dots, s_M)$	segment-level CM scores where s_m is the score for the m -th segment

- *Utterance-level detection*: learn a function f_θ to convert $\mathbf{x}_{1:T}$ to an utterance-level score s^u

$$f_\theta : \mathbb{R}^{T \times 1} \rightarrow \mathbb{R}; \quad \mathbf{x}_{1:T} \mapsto s^u, \quad (1)$$

- *Segment-level detection*: learn a function f_ρ to convert $\mathbf{x}_{1:T}$ to segment-level scores $\mathbf{s}_{1:M}$

$$f_\rho : \mathbb{R}^{T \times 1} \rightarrow \mathbb{R}^{M \times 1}; \quad \mathbf{x}_{1:T} \mapsto \mathbf{s}_{1:M}. \quad (2)$$

Depending on the CM, f_θ and f_ρ may share common architecture components and parameters since both tasks may be tackled by networks which operate upon similar hidden features.

In real applications, a CM needs to assign a *bona fide* or *spoof* label to the trial by comparing the s^u with an application-dependent threshold. It may also classify each segment on the basis of s_m , $\forall m \in \{1, \dots, M\}$. In this study, we follow the conventions in the research community and do not require the CM to produce hard decisions. Given s^u and $\mathbf{s}_{1:M}$ computed from test set trials, we calculate EERs to measure the utterance- and segment-level detection performance of the CM.

B. Existing CM Architectures for PS Scenario

The definition of the utterance-level detection task in the PS scenario is identical to that in the LA (also PA and DF) scenario. Segment-level detection is different and more challenging because the required output is a sequence of scores, where each score must be derived from segments of comparatively shorter duration. Accordingly, CMs that conduct only utterance-level detection for the LA or other scenarios cannot be applied directly to the PS scenario.

Previous studies modified the utterance-level CMs for other conventional scenarios so that the CMs can do segment-level detection as well [5], [6], [16]. One approach involves the addition of a sub-network for segment-level detection to a conventional utterance-level CM. Another approach involves the application of a non-trainable re-scoring step to derive segment-level scores $\mathbf{s}_{1:M}$. We describe these modifications after first introducing the utterance-level CM.

1) *Conventional Utterance-Level CM*: A conventional utterance-level CM consists of a front-end and back-end,

as shown in Fig. 2(a). The front-end extracts an acoustic feature sequence $\mathbf{a}_{1:N}$ from the input $\mathbf{x}_{1:T}$, after which the back-end converts $\mathbf{a}_{1:N}$ into an utterance-level CM score s^u . The CM usually uses DSP algorithms in the front-end, and the conversion from $\mathbf{x}_{1:T}$ to $\mathbf{a}_{1:N}$ is hardwired by the DSP designer. In contrast, the back-end can be implemented using various machine learning models that can transform $\mathbf{a}_{1:N}$ into a scalar s^u .

We now explain a back-end based on a light convolutional neural network (LCNN) since it is related to this study and used in many CMs. It computes the utterance-level score by $\mathbf{a}_{1:N} \xrightarrow{\text{LCNN}} \mathbf{h}_{1:M} \xrightarrow{\text{Pooling}} \mathbf{e}^u \xrightarrow{\text{Linear}} s^u$, where $\mathbf{h}_{1:M}$ is a hidden feature sequence with length $M \leq N$ determined by the LCNN stride, and \mathbf{e}^u is an embedding vector of fixed dimension produced by global average pooling. Finally, s^u is computed from \mathbf{e}^u using linear transformation. The back-end is typically trained by minimizing the standard cross-entropy (CE) or other advanced criteria given an utterance-level score s^u and ground-truth label y^u for a sufficiently large dataset.

2) *Adapting Utterance-Level CM for PS Scenario*: The conventional utterance-level CM has to be modified so that it can produce segment-level scores for the PS scenario. Two studies proposed to revise the LCNN-based back-end [5], [16]. As shown in Fig. 2(c), [5] and [16] remove pooling layer(s) from the utterance-level CM to derive $\mathbf{s}_{1:M}$ from $\mathbf{h}_{1:M}$. While there are differences in the implementations, the common idea is to add a sub-network to transform hidden features into segment-level scores $\mathbf{h}_{1:M} \xrightarrow{\text{Linear}} \mathbf{s}_{1:M}$. The revised CM may have a bipartite structure as shown in Fig. 2(b). Note that the linear transformation is applied independently to each segment, i.e., $s_m = \text{Linear}(\mathbf{h}_m)$, $\forall m \in \{1, \dots, M\}$.

Segment scores $\mathbf{s}_{1:M}$ are deterministically aligned with the input $\mathbf{a}_{1:N}$. Specifically, if the LCNN uses a stride equal to 1 for all layers, M will be equal to N , and each s_m is aligned with $\mathbf{a}_{n=m}$ ⁴. For strides greater than 1, the score-sequence length M becomes shorter than N , and the s_m is aligned with the $(\frac{N}{M}m)$ -th input frame⁵. No matter how the two sequences are aligned, s_m reflects the degree to which the m -th segment is *bona fide*, and thus can be used for segment-level detection. Training can be conducted by minimizing the CE loss between $\mathbf{s}_{1:M}$ and segment labels $\mathbf{y}_{1:M}$ and the utterance-level CE loss between the s^u and y^u . This method is referred to as multi-task learning in our previous study [6].

It is also possible to carry out both utterance-level and segment-level detection using only the segment-level detection branch, as shown in Fig. 2(c). Here, the utterance-level score s^u is derived from the segment-level scores $\mathbf{s}_{1:M}$, for example, in accordance with $s^u = \min_m s_m$ [6]. An advantage of this approach is the need to compute only the segment-level CE loss during training.

⁴Strictly speaking, s_m is computed given the acoustic features within the receptive field of the CNN, and the receptive field is centered around $\mathbf{a}_{n=m}$ and covers adjacent frames. The size of the receptive field is determined by the convolution kernel and dilation size.

⁵For example, if the CNN consists of k layers each with a stride equal to 2, M will be equal to $N/2^k$, and s_m will be aligned with \mathbf{a}_n , where $n = 2^k m$.

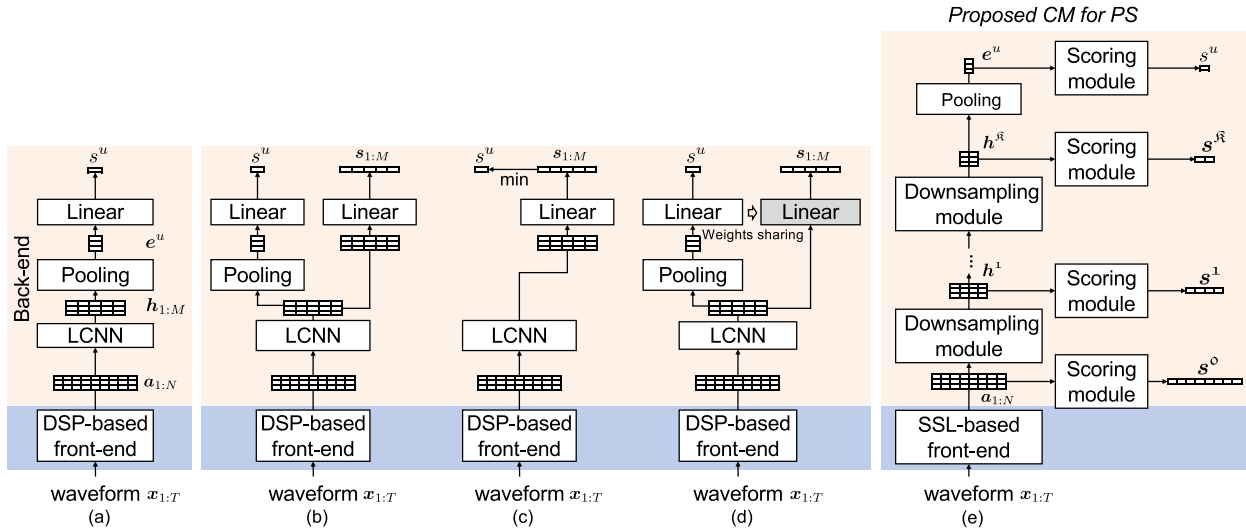


Fig. 2. Conventional utterance-level CM (a), existing CMs for PS scenario (b-d), and proposed CM for PS (e). Grey box in (d) indicates that linear layer is copied from the branch on its left side. Conventional studies [5], [16], [17] and [23] used (a), [6] and [24] used (b), [5], [6] and [16] used (c), and [6] used (d). Note that this figure illustrates the CMs during scoring.

3) *Attaching Re-Scoring Step for PS Scenario*: The second approach involves not the adaptation of the utterance-level CM but the addition of a non-trainable re-scoring step to derive segment-level scores $s_{1:M}$ [5]. For the LCNN-based utterance-level CM, the transformation $h_{1:M} \xrightarrow{\text{Pooling}} e^u \xrightarrow{\text{Linear}} s^u$ can be written as $s^u = w e^u = \frac{1}{M} \sum_{m=1}^M w h_m$, where w is the parameter of the linear layer. We can interpret $w h_m$ as a score for the m -th segment. Therefore, given the transformation vector w for utterance-level scoring, the re-scoring step to compute the $s_{1:M}$ can be defined as

$$h_{1:M} \xrightarrow{\text{Linear}} s_{1:M} : s_m = w h_m, \forall m \in [1, M]. \quad (3)$$

This approach is illustrated in Fig. 2(d).

An advantage of this approach is that the utterance-level CM can be used directly for the PS scenario without further training. It does not require segment-level labels. However, its potential is limited because without re-training, there is no adaptation to the spoofed segments in the PS scenario [5].

C. CMs Built for ADD Challenge 2022

One track of the ADD challenge 2022 is utterance-level detection of partially-spoofed audio [17]. Most participants used a conventional utterance-level CM, as shown in Fig. 2(a) (e.g. [23]) whereas one team used a similar multi-task structure, as shown in Fig. 2(b) [24].⁶

V. PROPOSED CM FOR PS SCENARIO

A. Motivation

The above CMs have achieved encouraging results in existing studies [5], [6], [16], but there is still room for improvement. For example, compared with a deterministic DSP-based front-end, a

trainable data-driven front-end may extract more discriminative features. As for the back-end, the commonly used LCNN only learns to conduct segment-level detection at a fixed temporal resolution. For example, using the LCNN configuration in [5], the back-end transforms the input $a_{1:N}$ into hidden features $h_{1:M}$, where $M = N/16$, and segment scores s_m are computed once for every 16 frames. So we can hypothesize that it would be better to use a more flexible DNN architecture that can leverage segment labels at different temporal resolutions during training and conduct segment-level detection accordingly. The hypothesis motivated us to propose a new CM that consists of a self-supervised learning (SSL)-based front-end and a back-end that supports multiple temporal resolutions. The new front-end is expected to extract more discriminative acoustic features in a data-driven manner, while the new back-end enables the CM to better identify spoofed segments with varied length. Fig. 2(e) illustrates the architecture of our CM.⁷

B. Front-End

An SSL speech model is a DNN that processes a waveform using trainable non-linear transformations. Because they are trained using task-agnostic self-supervised criteria on speech data from various domains, pre-trained SSL models can extract robust and informative acoustic features for many down-stream tasks [25]. SSL-based front-ends have also been shown to improve CM performance for the LA scenario [25], [26], [27].

Inspired by the above studies, we use an SSL-based front-end to extract acoustic features $a_{1:N}$ from an input $x_{1:T}$. While there are many types of SSL models, we investigated wav2vec 2.0 [7] and HuBERT [28]. These SSL models consist of a CNN-based encoder and cascade of Transformer blocks [29]. Following previous studies [25], we extract output features from all the Transformer blocks and use their trainable weighted sum as

⁶Interested readers are encouraged to see the challenge website and its details at <http://addchallenge.cn/add2022>.

⁷<https://github.com/nii-yamagishilab/PartialSpoof>

$\mathbf{a}_{1:N}$. During CM training, we fine-tune the pre-trained SSL front-end in conjunction with the back-end.

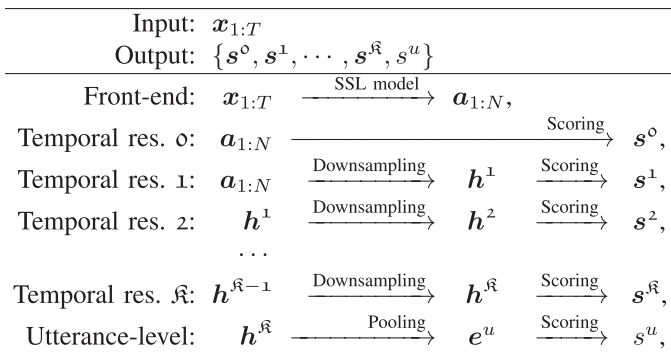
Note that the number of acoustic feature frames N is determined by the waveform length T and CNN encoder configuration. For the pre-trained SSL models we tested, the relationship is $N = T/320$. This means that, given an input waveform with a sampling rate of 16 kHz, the extracted acoustic features $\mathbf{a}_{1:N}$ have a ‘frame shift’ of 20 ms.

C. Back-End

Given $\mathbf{a}_{1:N}$, our back-end computes $\mathbf{s}_{1:M}$ at multiple temporal resolutions. This enables the proposed CM to handle the challenging segment-level detection task in a more flexible manner: by learning multiple functions $\{f_\rho^0, f_\rho^1, \dots, f_\rho^{\mathfrak{K}}\}$, where the function f_ρ^ℓ converts $\mathbf{x}_{1:T}$ to segment-level scores $\mathbf{s}_{1:M}^\ell$ at the ℓ -th temporal resolution, $\forall \ell \in \{0, \dots, \mathfrak{K}\}$. Without loss of generality, we assume there are \mathfrak{K} temporal resolutions in total. Lower indices indicate finer temporal resolutions. To simplify the notation, we drop the subscript in data sequences (e.g., $\mathbf{s}_{1:M}^\ell$ becomes \mathbf{s}^ℓ) for the rest of the paper.

As illustrated in Fig. 2(e), the back-end computes $\mathbf{s}_{1:M}$ at different temporal resolutions in a fine-to-coarse order. First, the acoustic features $\mathbf{a}_{1:N}$ are used to compute the score sequence $\mathbf{s}^0 \in \mathbb{R}^N$ at the frame level, which is the finest temporal resolution supported by the proposed CM. The n -th score s_n^0 indicates the likelihood that the n -th frame is *bona fide*. Next, $\mathbf{a}_{1:N}$ is down-sampled to \mathbf{h}^1 and used to compute \mathbf{s}^1 . This process is repeated to compute segment-level scores for each temporal resolution.

For utterance-level detection, the hidden feature sequences $\mathbf{h}^{\mathfrak{K}}$ at the lowest temporal resolution are pooled into the utterance-level embedding vector \mathbf{e}^u which is then transformed into the utterance-level score s^u . This scoring module has the same architecture as those for segment-level detection. Combined with the front-end, the proposed CM computes the scores for the PS scenario in the following order:



The scoring modules at different temporal resolutions are independent but use the same architecture, and we compare a few candidate architectures in Section VI-B. The down-sampling modules are also independent but use the same architecture. They contain a max-pooling operator with a stride equal to 2, followed by a 1D convolution layer with a kernel size of 1 and equal number of output and input channels. Therefore, if the input to the back-end $\mathbf{a}_{1:N}$ is of dimension $\mathbb{R}^{N \times D}$, \mathbf{h}^ℓ at the ℓ -th

temporal resolution will be in $\mathbb{R}^{N/2^\ell \times D}$, and the corresponding score sequence \mathbf{s}^ℓ is of dimension $\mathbb{R}^{N/2^\ell \times 1}$. This means that the temporal resolution of the $\mathbf{s}_{1:M}$ is reduced by a factor of 2 after each down-sampling module.

Since acoustic features are extracted every 20 ms (i.e., the ‘frame shift’ of the SSL-based front-end), a frame-level score s_n^0 is produced every 20 ms. This is the finest supported temporal resolution. For the ℓ -th temporal resolution, one segment score s_m is produced every $20 \times 2^\ell$ ms. We set $\mathfrak{K} = 5$, and the indices $\ell = 0$ to $\ell = 5$ correspond to temporal resolutions of 20, 40, 80, 160, 320, and 640 ms, respectively. This setting was determined empirically, taking into account the ground-truth labels provided in the extended PartialSpoof database. Although the temporal resolutions do not have explicit linguistic meanings, we hope that each time resolution may capture the following content:

- 40 ms: consonants and parts of vowels,
- 80 and 160 ms: consonants and vowels, and
- 320 ms: monosyllabic words and parts of longer words.

D. Training Strategies

Supposing that ground-truth labels are available for all temporal resolutions, the proposed CM can be trained using two types of strategies:

1) *Multiple Temporal Resolutions*: For the ℓ -th temporal resolution, a loss \mathcal{L}^ℓ is computed given the score \mathbf{s}^ℓ and label \mathbf{y}^ℓ . Similar to the CMs described in Section IV, the loss function can be CE or other advanced metrics. The gross loss function can thus be written as $\mathcal{L} = \sum_{\ell=0}^{\mathfrak{K}-1} \mathcal{L}^\ell$ ⁸, where we use $\mathcal{L}^{\mathfrak{K}+1}$ to denote the utterance-level loss.

2) *Single Temporal Resolution*: As described earlier, CM training using multiple temporal resolutions requires ground-truth labels for all temporal resolutions. While such labels are provided with the PartialSpoof database, they might not be available for other databases. In such a case, the CM can still be learned without application of the scoring modules and loss functions at temporal resolutions that have no segment-level labels. A special case is the learning of a CM for a single temporal resolution. For example, if we only have segment-level labels at the ℓ -th temporal resolution, training can be carried out using only the back-end layers required to compute \mathbf{s}^ℓ , in which case the training loss is computed over \mathbf{s}^ℓ and \mathbf{y}^ℓ . After training, the CM can only produce scores \mathbf{s}^ℓ , but not scores for other segment or utterance levels.

VI. EXPERIMENTS

We conducted experiments to test the proposed CM on the PartialSpoof database. We first conducted a pilot study to select the most suitable SSL-based front-end and back-end scoring module. We then conducted a comparative study to compare the

⁸In practice, the CM can be trained initially for the finest temporal resolution before more coarse resolutions, or the other way around. Furthermore, we may either fix or continue to fine-tune the scoring modules of a previous temporal resolution when moving to the next one. However, no obvious significant difference was observed among the aforementioned methods in our preliminary experiments. Hence, in this study we trained the CM at all temporal resolutions simultaneously.




Fig. 3. Comparison of proposed CMs trained at single temporal resolution and multiple temporal resolutions training strategies. Two figures on the left side show EERs of segment-level detection, and the right-side figure shows EERs of utterance-level detection. Each CM was trained three times with different random seeds.

two CM training strategies. After that, we analyzed differences in the tendencies of the proposed CM’s performance on the development and evaluation sets. Furthermore, we conducted a cross-scenario experiment using the ASVspoof 2019 LA database [18], the goal of which was to show that the proposed CM optimized for the PS scenario can also be applied to the LA scenario. After explaining the experimental setup and model configurations, we describe each experiment in detail.

A. Experimental Setup and Model Configurations

All CMs were trained by minimizing the so-called P2SGrad-based mean square error [30]. This criterion was used because it was found to be more stable and slightly superior to the conventional CE used in our previous study related to the LA scenario [30]. During training, we used the Adam optimizer with a default configuration ($\beta_1 = 0.9, \beta_2 = 0.999, \epsilon = 10^{-8}$). The learning rate was initialized with 1×10^{-5} and halved every 10 epochs.

We did not use any data augmentation, voice activity detection, or feature normalization during training, nor did we trim the input trials. All experiments were repeated three times with different random seeds for CM initialization, except for the pre-trained SSL front-end. The averaged results of the three runs are reported. Each round of training was conducted using a single Nvidia Tesla A100 card.

B. Comparing Front- and Back-Ends for Proposed CM

As explained in Section V, we used a pre-trained SSL model for the front-end and designed a powerful scoring module for the back-end. In this experiment, we compared several candidates for both. Performance was measured using the utterance-level detection EER on the development set of the PartialSpoof database. The CM was trained following the same conventional utterance-level training as in the LA scenario, that is, using the single temporal resolution training strategy with utterance-level labels only.

For the front-end, we tested the four pre-trained SSL models listed in Table V; *w2v2-base* is based on the Wav2vec 2.0 Base model, while *w2v2-large* and *w2v2-xlsr* are based on the Wav2vec 2.0 Large model. The difference is that the

TABLE V
SSL MODELS MENTIONED IN SECTION VI-B

ID	Model type	Data for pre-training	# of paras	Feat. dim.
<i>w2v2-base</i>	Wav2vec 2.0 Base	Librispeech	95.04 m	768
<i>w2v2-large</i>	Wav2vec 2.0 Large	CommonVoice, Switchboard [31], Libri-Light [32], Fisher [33]	317.38 m	1024
<i>w2v2-xlsr</i>	Wav2vec 2.0 Large	LibriSpeech [34], CommonVoice [35], BA-BEL [36]	317.38 m	1024
<i>wavlm-large</i>	HuBERT Large	Libri-Light [32], Gi-gaSpeech [37], Vox-Populi [38]	316.62 m	1024

Wav2vec 2.0 Base model has only 12 Transformer [29] blocks while *Wav2vec 2.0 Large* has 24. There are further differences in the dimension of the output features of the Transformer block, as shown in Table V. The *w2v2-large* and *w2v2-xlsr* models, both based on *Wav2vec 2.0 Large*, differ in the data used for pre-training. These three SSL models were included in the experiment because they have been shown to give reliable performance for the LA scenario [26]. The last candidate *wavlm-large* is based on HuBERT [28]. It was included because it was the top entry on the leaderboard of SUPERB⁹ when we started this study.

For the scoring module in the back-end, we compared the following architectures:

- a single fully-connected (FC) layer;
- an FC layer after a bidirectional long short-term memory (BLSTM) [39] layer;
- an FC layer after two BLSTM layers;
- a single gated multilayer perceptron (gMLP) block [40];
- five gMLP blocks.

A gMLP block is similar to basic Multilayer Perceptrons (MLPs) with a gating unit. It was found to be simple and powerful compared with other alternatives.

We compared each of the above components in terms of the EER on the development set of the PartialSpoof database. To reduce the time cost, we used *w2v2-base* as the front-end when comparing the back-end scoring modules. The EERs listed in Table VI show that using five gMLP blocks gives the best

⁹SUPERB: <https://superbenchmark.org/>

TABLE VI
COMPARISON OF DIFFERENT FRONT-ENDS AND BACK-ENDS

Front-end	Back-end	EER (%) on dev. set
w2v2-base	1 FC	2.89
	1 BLSTM + 1 FC	2.84
	2 BLSTM + 1 FC	2.49
	1 gMLP	4.28
	5 gMLP	2.22
w2v2-base		2.22
w2v2-large	5 gMLP	0.35
w2v2-xlsr		0.65
wavlm-large		0.64

performance. We then compared the SSL models while using the architecture of five gMLP blocks in the back-end scoring module. The results indicate that the three Large SSL models outperformed w2v2-base, and that w2v2-large performed best. Given these results, all following experiments were conducted using w2v2-large and five gMLP layers.

C. Comparing Training Strategies on the Proposed CM

In this section, we compared the two CM training strategies, that is, training at a single temporal resolution or multiple resolutions. Performance was measured on the development set and evaluation set of the PartialSpoof database.

We first explored the two training strategies on segment-level detection. The results are shown in Fig. 3(a). For the proposed CM trained at a single temporal resolution (black circle), since the segment-level ground-truth labels are available at the six segment-level temporal resolutions, we trained six versions of the proposed CM, one for each segment-level temporal resolution. After training, each CM version produces scores at the corresponding temporal resolution. We also trained the CM at multiple temporal resolutions (red triangle) that can derive scores for all resolutions at once.

Results suggest that segment-level detection is more difficult at a higher temporal resolution, but it is not impossible. For the highest temporal resolution of 20 ms, the proposed CM achieved an EER of around 10% on the evaluation set. This EER is reasonably good considering the fact that each segment consists of only one frame and is extremely short. Besides, although the multiple-resolution CM is slightly better than the single-resolution counterpart at the resolution of 640 ms, it performed worse than the single-resolution CMs at fine-grained segment levels (20 ~ 320 ms). There is thus room for improvement of the fine-grained segment-level detection.

We then compared performance on utterance-level detection. The results are shown in Fig. 3(b). No matter which training strategy was used, the proposed CM achieved an EER below 1% on the evaluation set for utterance level detection. Furthermore, the CM trained at multiple temporal resolutions outperformed the CMs trained at a single temporal resolution for utterance-level detection. This indicates that the coarse-grained detection tasks can benefit from the training strategies using multiple temporal resolutions.

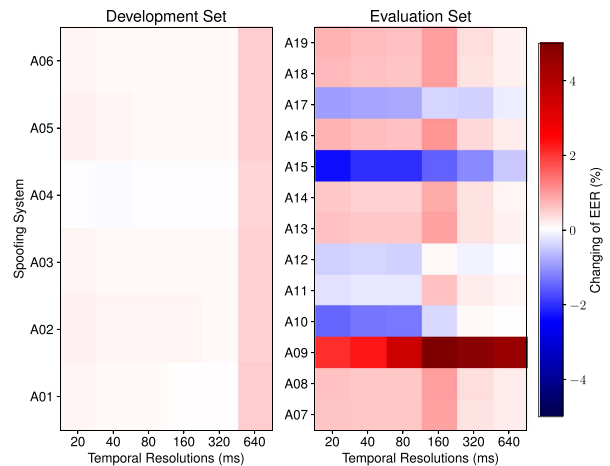


Fig. 4. Results of leave-one-out experiments on development and evaluation sets. Color represents changing of EER (%).

Results of the above experiments suggest that we need to select a suitable temporal resolution and training strategy depending on our goal: (1) For the detection at the segment level, training should be applied at the specific target temporal resolution; (2) in terms of utterance-level detection, the use of more fine-grained information is expected to improve CM performance.

Besides, the differences between the EERs on the evaluation and development sets demonstrate the difficulty of segment-level detection and show that it is difficult to generalize to the evaluation set. We thus explore the differences between them in the next subsection.

D. EER Gap Between Development and Evaluation Sets

We have two hypotheses to explain the significant differences between the EERs on the development and evaluation sets, and we examined both using the CM trained at multiple temporal resolutions.

1) *Hypothesis 1: More difficult spoofing systems exist in the evaluation set.* For the LA scenario and ASVspoof 2019 LA database, existing studies have measured the EER over fully spoofed trials from each spoofing system and found that a voice-conversion-based spoofing system called A17 significantly increased the CMs' EERs on the evaluation set [3], [41]. Since the PartialSpoof database used spoofed trials from the ASVspoof 2019 LA database, we hypothesize that the strong attacks in the evaluation set may have led to the higher EER. Because the segments in a partially-spoofed audio sample can be from different spoofing methods, we did not use the same investigation method as previous studies on the LA scenario and instead used a leave-one-out evaluation approach. For each spoofing method, we excluded the test trials that contain at least one segment produced by that spoofing method. We then measured the EER on the remaining trials and compared it with the original EER computed on all the test trials. We conducted the analysis on both the development and evaluation sets.

The results are shown in Fig. 4. Each column in the figure corresponds to the results on segment-level detection at one

TABLE VII

EERs (%) OF DIFFERENT CMS ON PARTIALSPOOF EVALUATION SET. COLUMN TYPES LISTS EACH CM IN ACCORDANCE WITH CATEGORIES IN FIG. 2. ALL CMS WERE TRAINED ON PARTIALSPOOF TRAINING SET, ALTHOUGH THEY MAY HAVE USED LABELS AT DIFFERENT TEMPORAL RESOLUTIONS. PREVIOUS STUDIES [5], [6], [16] USED ONLY ONE OR TWO TEMPORAL RESOLUTIONS

Ref.	Model ID ¹ (in ref.)	Types (in Fig. 2)	Training Resolutions	Front- end	Back- end	Temporal resolutions (ms)							
						20	40	80	160	320	640	utt.	
[16]	Utterance Segment ²	a c	utt. 10 ms	CQCC	LCNN	27.17							
[5]	Utterance Segment	a c	utt. 160 ms	LFCC ³	LCNN-BLSTM				40.20 16.21			6.19 8.61	
[6]	Utterance Segment	d c	utt. 160 ms						44.00 15.93			6.33 7.69	
	UttU SegU	d c	160 ms, utt. 160 ms, utt.	LFCC	SELCNN-BLSTM				20.04 17.75			9.96 7.04	
	MulBS UttBW	b b	160 ms, utt. 160 ms, utt.						17.55 17.77			5.90 5.66	
	SegBW	b	160 ms, utt.						16.60			6.07	
	Proposed	Multi reso.	e	20~640, utt.	w2v2-large	5gmlp	12.84	11.94	10.92	9.24	6.34	5.19	0.49

¹ The Model ID in the second column is defined by the corresponding references in the first column Ref.

² The original resolution in [16] was 8 ms, a value determined by the Constant Q Cepstral Coefficients (CQCC) extraction configuration [42]. We used 10 ms so that the results from this system are comparable to those of the other systems. Results for 20 ms were calculated using the average method [16].

³ LFCC stands for linear frequency cepstral coefficients.

temporal resolution, while each row corresponds to the results of analysis on a specific spoofing method. The IDs of the spoofing systems such as A07 were inherited from those in the ASVspoof 2019 LA database. It was observed that, if the spoofed trials that contain spoofed segments from A15¹⁰ were removed, the EER on the rest of the evaluation set decreased significantly. For our CM, A15 was the strongest spoofing system in the PS scenario. Leaving out other spoofing systems such as A10, A11, A12 and A17 also decreased the EER but to a lesser extent. In contrast, the EERs on the development set did not decrease as much as the change in the EERs on the evaluation set no matter which spoofing method was left out.

As hypothesized, partially-spoofed segments produced by some of the spoofing systems in the evaluation set were more difficult to detect. Particularly, A15 was the strongest attack. This may be one reason that EERs in the evaluation set are higher than those in the development set. However, the EER on the evaluation set was still higher than on the development set even when we removed A15. We thus have another hypothesis to explain the EER gap between the development and evaluation sets.

2) *Hypothesis 2: Detecting spoofed segments with fewer concatenated boundaries on the evaluation set is more difficult than on the development set.* As described in Section III-B, the partially-spoofed trials in PartialSpoof were created on the basis of substitution and concatenation. The overlap-add-based concatenation may bring in artifacts around the concatenated boundaries, which are expected to be useful for the CM. However, a segment to be scored by the CM may not contain any concatenated boundaries. Furthermore, if such a segment is from an unseen spoofing attack, the CM is more likely to make a mistake. Hence, we hypothesize that our CM performed worse on the evaluation set because it was not perfectly designed to

¹⁰ A15 is a hybrid spoofing system that uses a TTS voice as a source speaker and WaveNet as a vocoder for voice conversion.

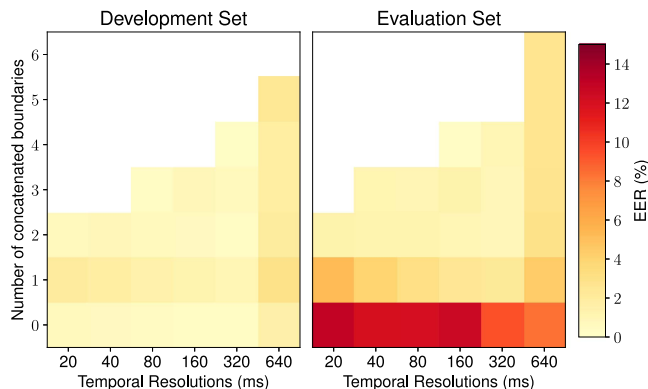


Fig. 5. Breakdown of EER on number of concatenated boundaries.

distinguish the spoofed segments without any concatenation point from those with one or more concatenated boundaries. The performance on the development set may receive less impact since all the spoofing methods were seen during the training of the CM.

To verify the hypothesis, we computed the breakdown EERs on the basis of the number of concatenated boundaries in the segments. The results are listed in Fig. 5. Each row in the figure lists the EERs evaluated on segments that have a corresponding number of concatenated boundaries, and each column is for EERs at one temporal resolution. We observe that the performance of our CM on the evaluation set degraded when the number of concatenated boundaries decreased. Especially, when the number of concatenated boundaries was zero, the CM's EER significantly increased at temporal resolutions ranging from 20 to 160 ms. In contrast, the EERs on the development set varied (but to a much lesser extent) regardless of the change in the number of concatenated boundaries. As hypothesized, our CM made more errors when facing unseen spoofed segments with few concatenated boundaries in the evaluation set.

TABLE VIII
CROSS-SCENARIO STUDY IN UTTERANCE LEVEL (EER %). (ASVspoof 2019
LA AND PARTIALSPOOF DATABASES WERE APPLIED FOR LA AND PS
SCENARIOS SEPARATELY.)

	Train	LA		PS	
		Dev.	Eval.	Dev.	Eval.
Single reso. at utt. level	LA	0.04	0.83	12.22	14.19
Single reso. at utt. level	PS	0.57	0.77	0.35	0.64
Multiple reso. at all levels		0.24	0.90	0.20	0.49

E. Comparing Proposed and Conventional CMs

We next compared the proposed CM with the conventional CMs [5], [6], [16] explained in Section IV-B. Comparisons are made to a CM trained with the multiple temporal resolutions. EERs for the PartialSpoof evaluation set are summarized in Table VII.

The first noteworthy observation was that the conventional CMs could only directly detect segments at a single temporal resolution. This is due to the inflexible back-end structure, as discussed in Section V-A. In contrast, our proposed CM can detect spoofed segments at different temporal resolutions simultaneously. It also significantly outperformed the conventional CMs at the corresponding segment and utterance levels. These results indicate that the proposed CM is more suitable for the PS scenario.

F. Cross-Scenario Study

Since a CM for the PS scenario can execute utterance-level detection, it can also be applied to the LA scenario. To determine whether the improvements brought by multi-resolution training for the PS scenario also translate to the LA scenario, we investigated the performance of the proposed CM on utterance-level detection for both LA and PS scenarios. The data for these scenarios were from the ASVspoof 2019 LA and PartialSpoof databases, respectively.

These experiments involved two training settings, one using the training set of the ASVspoof 2019 LA database and the other using that of the PartialSpoof database. For the ASVspoof 2019 LA database, since the training data only contains utterance-level labels, the proposed CM was trained using the single temporal resolution training strategy for utterance-level detection. For CMs trained using the PartialSpoof data, we used the two CMs trained using either the single temporal resolution strategy or the multiple resolution one as described in Section V-D. The first used only utterance-level labels and was trained using the same strategy as the CM trained on the ASVspoof 2019 LA database. The latter used the full set of utterance and segment labels from the PartialSpoof database and trained using the multiple temporal resolution strategy.

We used each CM to produce scores for the ASVspoof 2019 LA and PartialSpoof development and evaluation sets for the utterance-level detection. The results are shown in Table VIII. Focusing first on single resolution and utterance-level training, the performance for the ASVspoof 2019 LA database was shown to be competitive. However, the EERs increased significantly

beyond 10% in the PartialSpoof dataset. This was expected because the CMs trained using the LA data are not exposed to partial spoofs, thus perform poorly. When trained using the PartialSpoof dataset; however, performance in the case of PS data improved significantly, while performance when tested using the LA data remained competitive.

Further improvements were achieved when the proposed CM was trained at multiple temporal resolutions. In summary, when trained on the PartialSpoof dataset, the proposed CM demonstrated better utterance-level detection performance for both LA and PS scenarios. Using PartialSpoof training data is beneficial even when the CM is tested on LA data.¹¹

VII. CONCLUSION

We reported a new spoofing scenario, PS, in which only a fraction of speech segments are spoofed, with the remaining segments containing bona fide speech. Successful approaches to spoofing detection in this scenario can be applied at either the utterance level or segment level. The latter requires the assessment of spoofing classifiers using a database of bona fide and spoofed speech, with the latter labeled at the segment level. We described the new PartialSpoof database which is labeled at multiple temporal resolutions from 20 to 640 ms.

After formulating CM tasks required to tackle the PS scenario, we introduced SSL models as an enhanced front-end and proposed new neural architectures and training strategies that exploit segment-level labels for simultaneous, multi-resolution training. Experimental results suggest that CMs and training strategies should be adapted to a specific goal. Utterance-level detection can benefit from the use of more fine-grained information during training, whereas the comparatively more challenging task of segment-level detection calls for matched resolution.

There were significant differences between the results of segment-level detection on the development and evaluation sets. We thus investigated two hypotheses to explain the differences: (1) More difficult spoofing systems exist in the evaluation set. This hypothesis is supported by a series of leave-one-out experiments for each resolution, and significant changes in the EER were observed after removing certain unknown spoofing systems. Among those unknown spoofing systems, A15 was found to be the strongest attack in the PartialSpoof dataset. (2) Detecting spoofed segments with fewer concatenated boundaries on the evaluation set is more difficult than on the development set. This is supported by an analysis of the breakdown EERs on the spoofed data with different numbers of concatenated boundaries. The performance of the CM on the evaluation set was much worse when there were fewer concatenated boundaries within the segments. In contrast, the performance on the development set was less affected. How to overcome this issue for the segment-level detection is worth exploring in the future.

¹¹In addition to the cross-scenario study, it is also possible to combine the two databases and see if they are complementary to each other. This investigation requires significant changes in the proposed CM structure so that model training can be effectively conducted even in situations where parts of audio files in the training database do not have segment-level labels, which is beyond the scope of this paper and is a future analysis topic.

We also conducted a cross-scenario study on the LA and PS scenarios. The proposed CM was shown to achieve the best reported utterance-level detection result for the ASVspoof 2019 LA database (an EER of 0.77% for the evaluation set). The PS scenario is a realistic, important, timely, and challenging spoofing scenario, which warrants greater attention in the future. The more conventional, utterance-level CMs considered in the LA scenario largely fail in the face of only partially-spoofed utterances, indicating their vulnerability to manipulation through such attacks. Results nonetheless show scope for improvement.

Our future work will include explicit use of linguistic information for the detection of short spoofed segments embedded in otherwise bona fide utterances and robust training strategies to handle imbalanced classes. An important challenge will be to increase the number of TTS and VC methods for creating partially-spoofed speech and to use more up-to-date methods.

ACKNOWLEDGMENT

We would like to thank Prof. Kiyoshi Honda and Dr. Jose Patino for their comments.

REFERENCES

- [1] Z. Wu et al., "ASVspoof 2015: The first automatic speaker verification spoofing and countermeasures challenge," in *Proc. Interspeech*, 2015, pp. 2037–2041.
- [2] T. Kinnunen et al., "The ASVspoof 2017 challenge: Assessing the limits of replay spoofing attack detection," in *Proc. Interspeech*, 2017, pp. 2–6.
- [3] A. Nautsch et al., "ASVspoof 2019: Spoofing countermeasures for the detection of synthesized, converted and replayed speech," *IEEE Trans. Biometrics, Behav., Identity. Sci.*, vol. 3, no. 2, pp. 252–265, Apr. 2021.
- [4] J. Yamagishi et al., "ASVspoof 2021: Accelerating progress in spoofed and deepfake speech detection," in *Proc. ASVspoof Workshop*, 2021, pp. 47–54.
- [5] L. Zhang, X. Wang, E. Cooper, J. Yamagishi, J. Patino, and N. Evans, "An initial investigation for detecting partially spoofed audio," in *Proc. Interspeech*, 2021, pp. 4264–4268.
- [6] L. Zhang, X. Wang, E. Cooper, and J. Yamagishi, "Multi-task learning in utterance-level and segmental-level spoof detection," in *Proc. ASVspoof Workshop*, 2021 pp. 9–15.
- [7] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," in *Proc. Adv. Neural Inf. Process. Syst.*, 2020, vol. 33, pp. 12449–12460.
- [8] S. Chen et al., "WavLM: Large-scale self-supervised pre-training for full stack speech processing," *IEEE J. Sel. Topics Signal Process.*, vol. 16, no. 6, pp. 1505–1518, Oct. 2022.
- [9] M. Zakariah, M. K. Khan, and H. Malik, "Digital multimedia audio forensics: Past, present and future," *Multimedia Tools Appl.*, vol. 77, pp. 1009–1040, 2018.
- [10] P. R. Bevinamarad and M. S. Shirdonkar, "Audio forgery detection techniques: Present and past review," in *Proc. IEEE 4th Int. Conf. Trends Electron. Inform.*, 2020, pp. 613–618.
- [11] N. Campbell, "CHATR: A high-definition speech re-sequencing system," in *Proc. 3rd ASA/ASJ Joint Meeting*, 1996, pp. 1223–1228.
- [12] P. Taylor, *Text-to-Speech Synthesis*. Cambridge, U.K.: Cambridge Univ. Press, 2009.
- [13] D. Tan, L. Deng, Y. T. Yeung, X. Jiang, X. Chen, and T. Lee, "EditSpeech: A text based speech editing system using partial inference and bidirectional fusion," in *Proc. IEEE Autom. Speech Recognit. Understanding Workshop*, 2021, pp. 626–633.
- [14] M. Morrison, L. Rencker, Z. Jin, N. J. Bryan, J.-P. Caceres, and B. Pardo, "Context-aware prosody correction for text-based speech editing," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2021, pp. 7038–7042.
- [15] "Descript," 2022. [Online]. Available: <https://www.descript.com/>
- [16] J. Yi et al., "Half-truth: A partially fake audio detection dataset," in *Proc. Interspeech*, 2021, pp. 1654–1658.
- [17] J. Yi et al., "Add 2022: The first audio deep synthesis detection challenge," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2022, pp. 9216–9220.
- [18] X. Wang et al., "ASVspoof 2019: A large-scale public database of synthesized, converted and replayed speech," *Comput. Speech Lang.*, vol. 64, 2020, Art. no. 101114.
- [19] International Telecommunication Union, Recommendation G.191: Software Tools and Audio Coding Standardization, Nov.11 2005.
- [20] D. Povey et al., "The Kaldi speech recognition toolkit," in *Proc. Autom. Speech Recognit. Understanding*, 2011.
- [21] T. Kinnunen and H. Li, "An overview of text-independent speaker recognition: From features to supervectors," *Speech Commun.*, vol. 52, no. 1, pp. 12–40, 2010.
- [22] M. Lavechin, M.-P. Gill, R. Bousbib, H. Bredin, and L. P. Garcia-Perera, "End-to-end domain-adversarial voice activity detection," in *Proc. Interspeech*, 2020, pp. 3685–3689.
- [23] J. M. Martín-Doñas and A. Álvarez, "The vicomtech audio deepfake detection system based on wav2vec2 for the 2022 ADD challenge," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2022, pp. 9241–9245.
- [24] H. Wu et al., "Partially fake audio detection by self-attention-based fake span discovery," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2022, pp. 9236–9240.
- [25] S. W. Yang et al., "SUPERB: Speech processing universal Performance benchmark," in *Proc. Interspeech*, 2021, pp. 1194–1198.
- [26] X. Wang and J. Yamagishi, "Investigating self-supervised front ends for speech spoofing countermeasures," in *Proc. Odyssey*, 2022, pp. 100–106.
- [27] J. Shor, A. Jansen, W. Han, D. Park, and Y. Zhang, "Universal paralinguistic speech representations using self-supervised conformers," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2022, pp. 3169–3173.
- [28] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, "HuBERT: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 29, pp. 3451–3460, 2021.
- [29] A. Vaswani et al., "Attention is all you need," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 5998–6008.
- [30] X. Wang and J. Yamagishi, "A comparative study on recent neural spoofing countermeasures for synthetic speech detection," in *Proc. Interspeech*, 2021, pp. 4259–4263.
- [31] J. J. Godfrey, E. C. Holliman, and J. McDaniel, "SWITCHBOARD: Telephone speech corpus for research and development," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 1992, pp. 517–520.
- [32] J. Kahn et al., "Libri-light: A benchmark for ASR with limited or no supervision," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2020, pp. 7669–7673.
- [33] C. Cieri, D. Miller, and K. Walker, "The fisher corpus: A resource for the next generations of speech-to-text," in *Proc. Int. Conf. Lang. Resour. Eval.*, 2004, pp. 69–71.
- [34] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An ASR corpus based on public domain audio books," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2015, pp. 5206–5210.
- [35] R. Ardila et al., "Common voice: A massively-multilingual speech corpus," in *Proc. 12th Lang. Resour. Eval. Conf.*, 2020, pp. 4218–4222.
- [36] IARPA, "Babel program," 2011. [Online]. Available: <https://www.iarpa.gov/research-programs/babel>
- [37] G. Chen et al., "GigaSpeech: An evolving, multi-domain ASR corpus with 10,000 hours of transcribed audio," in *Proc. Interspeech*, 2021, pp. 3670–3674.
- [38] C. Wang et al., "VoxPopuli: A large-scale multilingual speech corpus for representation learning, semi-supervised learning and interpretation," in *Proc. ACL Int. Joint Conf. Natural Lang. Process.*, 2021, pp. 993–1003.
- [39] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [40] H. Liu, Z. Dai, D. So, and Q. V. Le, "Pay attention to MLPS," in *Advances in Neural Information Processing Systems*, vol. 34. Red Hook, NY, USA: Curran Associates, Inc., 2021, pp. 9204–9215.
- [41] H. Tak, J. Patino, M. Todisco, A. Nautsch, N. Evans, and A. Larcher, "End-to-end anti-spoofing with rawnet2," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2021, pp. 6369–6373.
- [42] M. Todisco, H. Delgado, and N. Evans, "Constant Q cepstral coefficients: A spoofing countermeasure for automatic speaker verification," *Comput. Speech Lang.*, vol. 45, pp. 516–535, 2017.




Lin Zhang (Student Member, IEEE) received the M.S. degree from Tianjin University, Tianjin, China, in 2020. Then, she was a Research Assistant with the SMIIP Lab, Duke Kunshan University, Suzhou, China, in 2020. She is currently working toward the Ph.D. degree with the SOKENDAI/National Institute of Informatics, Tokyo, Japan. Her research interests include anti-spoofing, speaker recognition, and machine learning.



Nicholas Evans (Member, IEEE) is currently a Professor with EURECOM, France, where he heads research in audio security and privacy. He is a Co-founder of the community led, ASVspoof, SASV, and VoicePrivacy challenge series. He participated in the EU FP7 Tabula Rasa and EU H2020 OCTAVE projects, both involving antispoofing. Today, his team is leading the EU H2020 TReSPAsS-ETN Project, a training initiative in security and privacy for multiple biometric characteristics. He co-edited the second and third editions of the *Handbook of Biometric Anti-Spoofing*. Previously, he served on the IEEE Speech and Language Technical Committee. He is an Associate Editor for the IEEE TRANSACTION ON BIOMETRICS, BEHAVIOR, AND IDENTITY SCIENCE.

Spoofing. Previously, he served on the IEEE Speech and Language Technical Committee. He is an Associate Editor for the IEEE TRANSACTION ON BIOMETRICS, BEHAVIOR, AND IDENTITY SCIENCE.



Xin Wang (Member, IEEE) received the B.E. degree from the University of Electronic Science and Technology of China, Chengdu, China, in 2012, the M.S. degree from the University of Science and Technology of China, Hefei, China, in 2015, and the Ph.D. degree from SOKENDAI/National Institute of Informatics, Tokyo, Japan, in 2018. He is currently a Project Assistant Professor with the National Institute of Informatics. His research interests include statistical speech synthesis, speech security, and machine learning. He is also a Co-organizer of the latest ASVspoof and VoicePrivacy challenges.

ASVspoof and VoicePrivacy challenges.



Junichi Yamagishi (Senior Member, IEEE) received the Ph.D. degree from the Tokyo Institute of Technology (Tokyo Tech), Tokyo, Japan, in 2006. From 2007 to 2013, he was a Research Fellow with the Centre for Speech Technology Research, The University of Edinburgh, Edinburgh, U.K. He was an Associate Professor with the National Institute of Informatics (NII), Tokyo, Japan, in 2013. He is currently a Professor with NII. His research interests include speech processing, machine learning, signal processing, biometrics, digital media cloning, and media forensics.

He was a Co-organizer for the bi-annual ASVspoof Challenge and the bi-annual Voice Conversion Challenge. He also was a Member of the IEEE Speech and Language Technical Committee during 2013–2019, an Associate Editor for the IEEE/ACM TRANSACTIONS ON AUDIO SPEECH AND LANGUAGE PROCESSING during 2014–2017, and a Chairperson of ISCA SynSIG during 2017–2021. He is currently a PI of JST-CREST and ANR supported VoicePersonae Project and the Senior Area Editor of the IEEE/ACM TRANSACTIONS ON AUDIO, SPEECH, AND LANGUAGE PROCESSING.



Erica Cooper (Member, IEEE) received the B.Sc. degree and M.Eng. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 2009 and 2010, respectively, and the Ph.D. degree in computer science from Columbia University, New York, NY, USA, in 2019. Since 2019, she has been a Project Researcher with the National Institute of Informatics, Tokyo, Japan. Her research interests include statistical machine learning and speech synthesis. Dr. Cooper was the recipient of the 3rd Prize in the CSAW Voice Biometrics and Speech Synthesis Competition, Computer Science Service Award from Columbia University, and Best Poster Award in the Speech Processing Courses in Crete.

CSAW Voice Biometrics and Speech Synthesis Competition, Computer Science Service Award from Columbia University, and Best Poster Award in the Speech Processing Courses in Crete.