



End-to-End Spectro-Temporal Graph Attention Networks for Speaker Verification Anti-Spoofing and Speech Deepfake Detection

Hemlata Tak¹, Jee-weon Jung², Jose Patino¹, Madhu Kamble¹,
Massimiliano Todisco¹ and Nicholas Evans¹

¹EURECOM, Sophia Antipolis, France, ²Naver Corporation, South Korea

lastname@eurecom.fr, jeeweon.jung@navercorp.com

Abstract

Artefacts that serve to distinguish bona fide speech from spoofed or deepfake speech are known to reside in specific sub-bands and temporal segments. Various approaches can be used to capture and model such artefacts, however, none works well across a spectrum of diverse spoofing attacks. Reliable detection then often depends upon the fusion of multiple detection systems, each tuned to detect different forms of attack. In this paper we show that better performance can be achieved when the fusion is performed within the model itself and when the representation is learned automatically from raw waveform inputs. The principal contribution is a spectro-temporal graph attention network (GAT) which learns the relationship between cues spanning different sub-bands and temporal intervals. Using a model-level graph fusion of spectral (S) and temporal (T) sub-graphs and a graph pooling strategy to improve discrimination, the proposed RawGAT-ST model achieves an equal error rate of 1.06% for the ASVspoof 2019 logical access database. This is one of the best results reported to date and is reproducible using an open source implementation.

1. Introduction

It is well known that cues indicative of spoofing attacks reside in specific sub-bands or temporal segments [1–13]. Prior work [9–12, 14] showed that these can be learned automatically using a model with either spectral or temporal attention. Our most recent work [13] showed the merit of graph attention networks (GATs) to learn the relationships between cues located in different sub-bands or different temporal intervals. We observed that different attention mechanisms spanning either the spectral or the temporal domain, work more or less well for different sets of spoofing attack and that the benefit of both can be exploited through their fusion at the score level. We also found that neither model works as well on its own for the full set of diverse spoofing attacks in the ASVspoof 2019 logical access database.

Motivated by psychoacoustics studies [15–18] which show the power of the human auditory system to select simultaneously the most discriminative spectral bands and temporal segments for a given task, and inspired by the power of GATs to model complex relationships embedded within graph representations, we have explored fusion within the model itself (earlier fusion). The idea is to extend the modelling of relationships between *either* different sub-bands *or* different temporal segments using separate models and attention mechanisms to the modelling of relationships spanning different spectro-temporal intervals using a GAT with combined spectro-temporal attention (GAT-ST). The approach facilitates the aggregation of com-

plementary, discriminative information simultaneously in both domains.

In a further extension to our past work, the GAT-ST reported in this paper operates directly upon the raw waveform. Such a fully end-to-end approach is designed to maximise the potential of capturing discriminative cues in both spectral and temporal domains. Inspired by work in speaker verification [19–22] and in building upon our end-to-end anti-spoofing solution reported in [12], the proposed RawGAT-ST model uses a one-dimensional sinc convolution layer to ingest raw audio. The principal contributions of this paper hence include:

- a fully end-to-end architecture comprising feature representation learning and a GAT;
- a novel spectro-temporal GAT which learns the relationships between cues at different sub-band and temporal intervals;
- a new graph pooling strategy to reduce the graph dimension and to improve discrimination;
- an exploration of different model-level, graph fusion strategies.

The remainder of this paper is organized as follows. Section 2 describes related works. Section 3 provides an introduction to GATs and describes how they can be applied to anti-spoofing. The proposed RawGAT-ST model is described in Section 4. Experiments and results are presented in Sections 5 and 6. Finally, the paper is concluded in Section 7.

2. Related works

In recent years, graph neural networks (GNNs) [23–27] have attracted growing attention, especially variants such as graph convolution networks (GCNs) [28] or GATs [29]. A number of studies have shown the utility and appeal of graph modeling [30–33] for various speech processing tasks. Zhang et al. [30] applied GCNs to a few-shot audio classification task to derive an attention vector which helps to improve the discrimination between different audio examples. Jung et al. [32] demonstrated the use of GATs as a back-end classifier to model the relationships between enrollment and test utterances for a speaker verification task. Panagiotis et al. [33] used GCNs to exploit the spatial correlations between the different channels (nodes) for a multi-channel speech enhancement problem.

Our previous work [13] demonstrated how GATs can be used to model spoofing attack artefacts. This is achieved using a self attention mechanism to emphasize the most informative sub-bands or temporal intervals and the relationships between them. We applied GATs separately to model the relationships in either spectral (GAT-S) or temporal (GAT-T) domains and

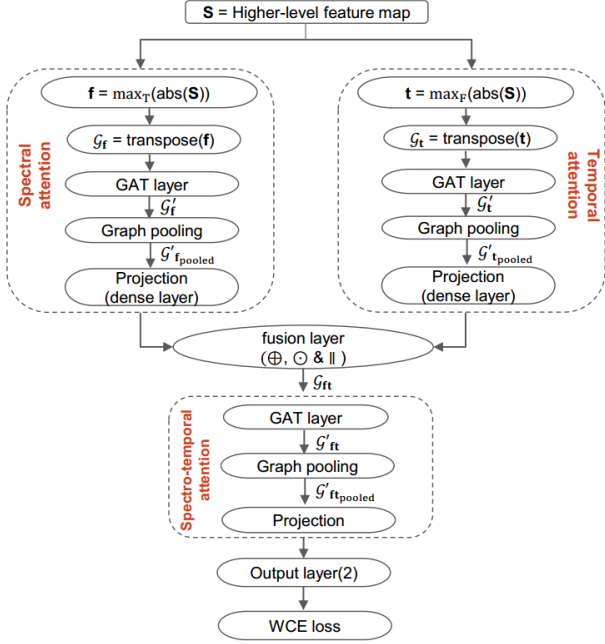


Figure 1: The proposed RawGAT-ST model architecture. The top-left spectral attention block captures discriminative spectral information. The top-right temporal attention block captures discriminative temporal information. Fusion is performed at the model level (middle ellipse). The bottom block comprises the spectro-temporal graph attention model.

demonstrated their complementarity through a score-level fusion (i.e., *late* fusion). Our hypothesis is that the integration of these two models (i.e., *early* fusion) has better potential to leverage complementary information (spectral and temporal) and further boost performance while using a *single* model. We present a summary of our original GAT-S/T approach in the next section. The new contribution is described in Section 4.

3. Graph attention networks for anti-spoofing

The work in [13] demonstrated the use of GATs to model the relationships between artefacts in different sub-bands or temporal intervals. The GAT learning process is illustrated in Algorithm 1 and is described in the following. A graph is defined as:

$$\mathcal{G}(N, \mathcal{E}, \mathbf{h}), \quad (1)$$

where N is the set of nodes, and \mathcal{E} represents the edges between all possible node connections, including self-connections. \mathcal{G} is formed from a higher-level feature representation \mathbf{h} (e.g. the output feature map of a residual network), where $\mathbf{h} \in \mathbb{R}^{N \times d}$. The number of nodes in the set N is the number of spectral sub-bands remaining after temporal averaging applied to each channel or, conversely, the number of temporal frames after spectral averaging. d is the feature dimensionality of each node and is equal to the number of channels in the feature map.

The GAT layer operates upon an input graph \mathcal{G} to produce an output graph \mathcal{G}' . It aggregates neighboring nodes using learnable weights via a self-attention mechanism. Node features are projected onto another representation learned from the minimisation of a training loss function using an affine transform (dense layer). Nodes are aggregated with attention weights

Algorithm 1 The GAT learning process

Input: $\mathcal{G}(N, \mathcal{E}, \mathbf{h})$

$\mathbf{h}_n \in \mathbb{R}^d, n \in N$

for $n \in N$ **do**

$\forall u \in \mathcal{M}(n) \cup \{n\}, \mathcal{M}(n)$ (set of nodes neighboring n)

$\alpha_{u,n} \leftarrow \text{softmax}(W_{\text{map}}(\mathbf{h}_n \odot \mathbf{h}_u))$

$m_n \leftarrow \sum_u \alpha_{u,n} \mathbf{h}_u$, node aggregation w.r.t Eq. 3

$\mathbf{o}_n \leftarrow \text{SeLU}(\text{BN}(W_{\text{att}}(m_n) + W_{\text{res}}(\mathbf{h}_n)))$, $\mathbf{o}_n \in \mathbb{R}^{d'}$

end

Output: $\mathcal{G}' \leftarrow \mathbf{o}$

which reflect the connective strength (relationship) between a given node pair.

The output graph \mathcal{G}' comprises a set of nodes \mathbf{o} , where each node \mathbf{o}_n is derived according to:

$$\mathbf{o}_n = \text{SeLU}(\text{BN}(W_{\text{att}}(m_n) + W_{\text{res}}(\mathbf{h}_n))), \quad (2)$$

where SeLU refers to a scaled exponential linear unit [34] activation function, BN refers to batch normalisation [35], m_n is the aggregated information for node n , and $\mathbf{h}_n \in \mathbb{R}^d$ represents the feature vector of node $n \in N$. Each output node $\mathbf{o}_n \in \mathbb{R}^{d'}$ has a target dimensionality $d' < d$. W_{att} projects the aggregated information for each node n to the target dimensionality d' . W_{res} projects the residual (skip) connection output to the same dimension.

The information from neighboring nodes is aggregated via self-attention according to:

$$m_n = \sum_{u \in \mathcal{M}(n) \cup \{n\}} \alpha_{u,n} \mathbf{h}_u, \quad (3)$$

where $\mathcal{M}(n)$ refers to the set of neighbouring nodes for node n , and $\alpha_{u,n}$ refers to the attention weight between nodes u and n . The GAT layer assigns learnable self-attention weights α to each edge. Weights reflect how informative one node is of another, where higher weights imply a higher connective strength.

The attention weights are derived according to:

$$\alpha_{u,n} = \frac{\exp(W_{\text{map}}(\mathbf{h}_n \odot \mathbf{h}_u))}{\sum_{w \in \mathcal{M}(n) \cup \{n\}} \exp(W_{\text{map}}(\mathbf{h}_n \odot \mathbf{h}_w))}, \quad (4)$$

where $W_{\text{map}} \in \mathbb{R}^d$ is the learnable map applied to the dot product and where \odot denotes element-wise multiplication.

4. RawGAT-ST model for anti-spoofing and speech deepfake detection

In this section, we introduce the proposed raw GAT with spectro-temporal attention (RawGAT-ST) model. It comprises four stages: i) learning higher-level semantic feature representations in truly end-to-end fashion by operating on the raw waveform; ii) a novel graph attention module with spectro-temporal attention; iii) a new graph pooling layer for discriminative node selection; iv) model-level fusion. The architecture is illustrated in Fig. 1 whereas a summary of the network configuration is shown in Table 1.

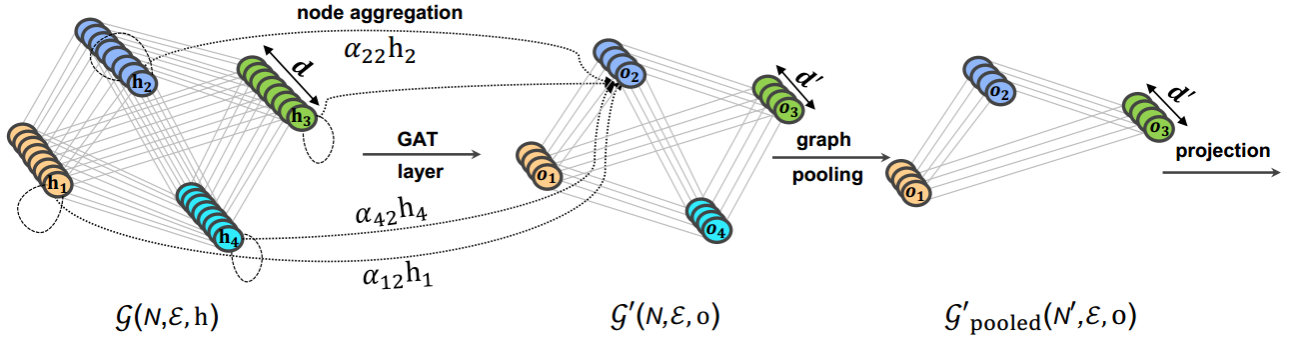


Figure 2: An illustration of the proposed graph-based process including a graph attention layer and a graph pooling layer for an input graph with 4 nodes. $h_n \in \mathbb{R}^d$ represents the node representation $n \in N$. The graph attention layer estimates attention weights α between each node pair, aggregating node information, and projecting to a new representation space $o_n \in \mathbb{R}^{d'}$ with lower output dimensionality d' . The graph pooling layer reduces the number of nodes to improve discrimination.

4.1. Front-end (higher-level) feature representation

In contrast to our prior work [13] which used hand-crafted features, the RawGAT-ST model operates directly upon the raw waveform [12, 36]. The literature shows that solutions based upon a bank of sinc filters are particularly effective in terms of both convergence stability and performance [12, 20, 37]. Accordingly, we use a sinc convolution layer for front-end feature learning similar to that reported in our previous work [12]. It performs time-domain convolution of the raw waveform with a set of parameterized sinc functions which correspond to rectangular band-pass filters [38, 39]. The centre frequencies of each filter in the filterbank are set according to a mel-scale. For all work reported in this paper, rather than learning cut-in and cut-off frequencies of each sinc filter, we used fixed cut-in and cut-off frequencies to alleviate over-fitting to training data.

The output of the sinc layer is transformed to a time-frequency representation by adding one additional channel dimension. The result is fed to a 2-dimensional (2D) residual network [40] to learn higher-level feature representations $\mathbf{S} \in \mathbb{R}^{C \times F \times T}$ where C , F and T refers to the number of channels, frequency bins and time samples respectively. We use a residual neural network identical to that reported in [12], except for the use of a 2D CNN instead of 1D CNN. A summary of the configuration is shown in Table 1. Each residual block layer consists of batch normalization (BN) with SeLU activation units, 2D convolution and a final max-pooling layer for data downsampling.

4.2. Spectro-temporal attention

The new approach to bring spectro-temporal graph attention to a single model is a core contribution of this work. An overview of the proposed RawGAT-ST model architecture is illustrated in Fig. 1. The input to the RawGAT-ST is a higher-level feature map (top of Fig. 1). The RawGAT-ST model itself comprises three principal blocks, each of which contains a single GAT layer: a spectral attention block (top-left of Fig. 1); a temporal attention block (top-right); a final spectro-temporal attention block (bottom). The spectral and temporal attention blocks have the goal of identifying spectral and temporal cues. The third, spectro-temporal attention block operates upon the pair of resulting graphs to model the relationships spanning both domains. All three blocks contain GAT and graph pooling layers. The operation of both processes is illustrated in Fig. 2. The GAT

Table 1: The details of RawGAT-ST model architecture. Numbers denoted in Sinc layer & Conv layer refers to (filter size, stride, number of filters). The output size are refers to (CNN channels, Freq, Time). Separate GAT layers are use for spectral and temporal attention blocks.

Layer	Input: 64600 samples	Output shape	
Sinc layer	Conv-1D(129,1,70)	(70,64472)	
	add channel (TF representation)	(1,70,64472)	
	Maxpool-2D(3)	(1,23,21490)	
BN & SeLU			
Res block	}	Conv-2D((2,3),1,32)	} × 2 (32,23,2387)
		BN & SeLU	
		Conv-2D((2,3),1,32)	
		Maxpool-2D((1,3))	
Res block	}	Conv-2D((2,3),1,64)	} × 4 (64,23,29)
		BN & SeLU	
		Conv-2D((2,3),1,64)	
		Maxpool-2D((1,3))	
Spectral-attention		Temporal-attention	
max _T (abs()) = (64, 23)		max _F (abs()) = (64, 29)	
GAT layer =(32,23)		GAT layer =(32,29)	
Graph pooling=(32,14)		Graph pooling=(32,23)	
Projection=(32,12)		Projection=(32,12)	
element-wise addition		(32,12)	
Fusion layer	element-wise multiplication	(32,12)	
	concatenation (along feature dim)	(64,12)	
Spectro-temporal attention	GAT layer	(16,12)	
	Graph pooling	(16,7)	
	Projection (along feature dim)	(1,7)	
Output	FC(2)	2	

layer operates upon an input graph $\mathcal{G}(N, \mathcal{E}, h)$ and produces an output graph $\mathcal{G}'(N, \mathcal{E}, o)$ by aggregating node information using self-attention weights between node pairs. The example in Fig. 2 shows the operation upon a 4-node input graph where the GAT layer reduces the node dimension from d to d' .

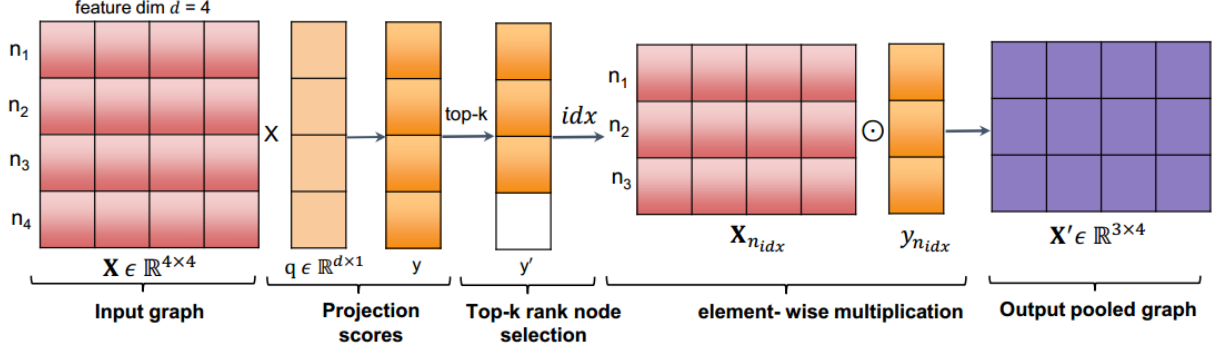


Figure 3: An illustration of the graph pooling layer [41]. The input graph $\mathbf{X} \in \mathbb{R}^{N \times d}$ has 4 nodes and feature dimensionality $d = 4$. $q \in \mathbb{R}^{d \times 1}$ is a learnable projection vector. Projection scores y are obtained from the dot product between q and the input feature vector X_n for each node. The indices idx corresponding to the nodes with the top- k highest projection scores are used to form a new pooled graph \mathbf{X}' from the element-wise multiplication of $\mathbf{X}_{n_{idx}}$ with $\text{sigmoid}(y_{n_{idx}})$ (Eq. 8).

This process is first applied separately with attention to either spectral or temporal domains to model the relationships between spectral or temporal artefacts present in the different sub-band or temporal intervals. Both spectral and temporal blocks operate upon the higher-level feature map $\mathbf{S} \in \mathbb{R}^{C \times F \times T}$. The spectral and temporal attention blocks first collapse temporal and spectral information respectively to a single dimension via max-pooling before the GAT layer. For the spectral attention block, max-pooling is applied to absolute values across the temporal dimension thereby giving a spectral feature map $\mathbf{f} \in \mathbb{R}^{C \times F}$:

$$\mathbf{f} = \max_T(\text{abs}(\mathbf{S})), \quad (5)$$

The temporal attention block operates instead across the spectral dimension giving a temporal feature map $\mathbf{t} \in \mathbb{R}^{C \times T}$:

$$\mathbf{t} = \max_F(\text{abs}(\mathbf{S})), \quad (6)$$

Since \mathbf{S} is derived from temporal data and hence contains both positive and negative values, use of absolute values in Eqs. 5 and 6 prevents meaningful negative-valued data from being discarded.

Graphs $\mathcal{G}_f \in \mathbb{R}^{N_f \times d}$ and $\mathcal{G}_t \in \mathbb{R}^{N_t \times d}$ are then constructed from the transpose of \mathbf{f} and \mathbf{t} feature maps, respectively. \mathcal{G}_f contains a set of 23 nodes (the number of spectral bins) whereas \mathcal{G}_t contains a set of 29 nodes (the number of temporal segments). Both graphs have a common dimensionality of $d = 64$. Separate GAT layers are then applied to both \mathcal{G}_f and \mathcal{G}_t to model the relationships between different sub-bands and temporal segments thereby producing a pair of new output graphs $\mathcal{G}'_f \in \mathbb{R}^{N'_f \times d'}$ and $\mathcal{G}'_t \in \mathbb{R}^{N'_t \times d'}$ with a common, reduced output feature dimensionality $d' = 32$.

4.3. Graph pooling

As illustrated in Fig. 1, a graph pooling layer is included in all three GAT blocks. Graph pooling generates more discriminative graphs by selecting a subset of the most informative nodes. In the example illustrated in Fig. 2 the graph pooling layer reduces the number of nodes from 4 to 3.

The approach is based upon the graph pooling layer in the recently proposed Graph U-Net architecture [41]. An example is illustrated in Fig. 3. The process is described in the following. Graph pooling uses a learnable projection vector $q \in \mathbb{R}^{d \times 1}$.

The dot-product between the input feature vector for each node $X_n \in \mathbb{R}^{1 \times d}$ where, $n \in N$, and q gives projection scores y_n :

$$y_n = X_n \cdot q, \quad (7)$$

Nodes in the input graph \mathbf{X} corresponding to the top- k -hot vector y' are then retained according to element-wise multiplication:

$$\mathbf{X}' = \mathbf{X}_{n_{idx}} \odot \text{sigmoid}(y_{n_{idx}}), \quad (8)$$

where the pooling ratio k is a hyperparameter, and where $X_{n_{idx}}$ and $y_{n_{idx}}$ are the node features and projection scores selected corresponding to the highest top- k indices, idx . All other nodes are purged. The complete process is illustrated in Fig. 3 for an input graph \mathbf{X} with 4 nodes, each of dimension $d = 4$. Finally, pooled graph representations $\mathcal{G}'_{f_{\text{pooled}}}$ and $\mathcal{G}'_{t_{\text{pooled}}}$ are generated from the original spectral \mathcal{G}'_f and temporal \mathcal{G}'_t output graphs, respectively. Since spectral and temporal pooled graphs $\mathcal{G}'_{f_{\text{pooled}}}$ and $\mathcal{G}'_{t_{\text{pooled}}}$ have different number of nodes, 14 for $\mathcal{G}'_{f_{\text{pooled}}}$ and 23 for $\mathcal{G}'_{t_{\text{pooled}}}$. Hence, both graph nodes are projected into the same dimensional space using an affine-transform to match the input node dimensionality of the fusion layer.

4.4. Model-level fusion

Model-level fusion (ellipse in Fig. 1) is used to exploit complementary information captured by the spectral and temporal attention graphs. We studied three different approaches to model-level fusion :

$$\mathcal{G}_{\mathbf{ft}} = \begin{cases} \mathcal{G}'_{f_{\text{pooled}}} \oplus \mathcal{G}'_{t_{\text{pooled}}} \\ \mathcal{G}'_{f_{\text{pooled}}} \odot \mathcal{G}'_{t_{\text{pooled}}} \\ \mathcal{G}'_{f_{\text{pooled}}} \parallel \mathcal{G}'_{t_{\text{pooled}}} \end{cases} \quad (9)$$

where the fused graph $\mathcal{G}_{\mathbf{ft}} \in \mathbb{R}^{N_{\mathbf{ft}} \times d_{\mathbf{ft}}}$ is generated from one of the fusion approaches in Eq. 9. It acts to combine the spectral pooled graph $\mathcal{G}'_{f_{\text{pooled}}}$ with the temporal pooled graph $\mathcal{G}'_{t_{\text{pooled}}}$. A fused graph $\mathcal{G}_{\mathbf{ft}}$ contains a set of 12 nodes and feature dimensionality $d_{\mathbf{ft}} = 32$. The three different operators in Eq. 9 are element-wise addition \oplus , multiplication \odot and concatenation \parallel . A third GAT layer is then applied to $\mathcal{G}_{\mathbf{ft}}$ to produce output graph $\mathcal{G}'_{\mathbf{ft}} \in \mathbb{R}^{N_{\mathbf{ft}} \times d'_{\mathbf{ft}}}$, where $d'_{\mathbf{ft}} = 16$ is the output feature dimensionality. Graph pooling is then applied one last time to generate a pooled graph $\mathcal{G}'_{\mathbf{ft}_{\text{pooled}}}$. The final two-class prediction (bona fide or spoofed) is then obtained using projection and output layers.

5. Experimental setup

Described in this section are the database, metrics and baseline systems used for our experiments, together with specific implementation details of the RawGAT-ST model.

5.1. Database and evaluation metric

Experiments were performed on the ASVspoof 2019 logical access (LA) database [42,43]. It has three independent partitions: train; development; evaluation. Spoofed speech is generated using a set of different speech synthesis and voice conversion algorithms [42]. The training and development partitions were created with a set of 6 different attacks (A01-A06), whereas the evaluation set was created with a set of 13 attacks (A07-A19). We used the minimum normalised tandem detection cost function (t-DCF) [44,45] as a primary metric but also report results in terms of the pooled equal error rate (EER) [46].

5.2. Baseline

The baseline is an end-to-end RawNet2 system [12]. It is among the best-performing, reproducible solutions. The first sinc layer is initialised with a bank of 128 mel-scaled filters. Each filter has an impulse response of 129 samples (8 ms duration) which is convolved with the raw waveform. The latter are truncated or concatenated to give segments of approximately 4 seconds duration (64,600 samples). It is followed by a residual network and a gated recurrent unit (GRU) to predict whether the input audio is bona fide or spoofed speech. Full details of the baseline system are available in [12].

Our temporal and spectral GAT systems introduced in [13] are not used as baselines in this work since they are not E2E systems; unlike the spectro-temporal GAT system introduced in this paper, they operate upon hand-crafted features. Nonetheless, we do report results for E2E temporal and spectral GAT solutions which are based upon the appropriate sub-blocks illustrated in Fig. 1. Use of these temporal and spectral GAT variants is necessary in order that differences in performance can be attributed solely to use of the proposed spectro-temporal GAT model and not also from differences in the first network layer. Details are described in Section 6.2.

5.3. RawGAT-ST implementation

In contrast to the baseline RawNet2 system and in order to reduce computational complexity, we reduced the number of filters in the first layer to 70. To improve generalisation, we added channel masking in similar fashion to the frequency masking [14,47,48] to mask (set to zero) the output of a random selection of contiguous sinc channels during training. The same channel mask is applied to all training data within each mini-batch. The number of masked channels is chosen from a uniform distribution between 0 and F_{mask} , where $F_{\text{mask}} = 14$ is the maximum number of masked channels selected based on minimum validation loss. While attention helps the model to focus on the most discriminative spectral sub-bands and temporal segments, channel masking improves generalization by ensuring that information at all sub-bands and segments is used at least to some extent. In contrast to usual practice, we also use fewer filters (32 and 64) in the first and second residual blocks to further protect generalisation to unseen attacks [49]. Graph pooling is applied with empirically selected pooling ratios of $k = 0.64, 0.81, \text{ and } 0.64$ for spectral, temporal and spectro-temporal attention blocks respectively.

The complete architecture is trained using the ASVspoof

2019 LA training partition to minimise a weighted cross entropy (WCE) loss function, where the ratio of weights assigned to bonafide and spoofed trials are 9:1 to manage the data imbalance in the training set. We used the standard Adam optimiser [50] with a mini-batch size of 10 and a fixed learning rate of 0.0001 and train for 300 epochs. The feature extractor and back-end classifier are jointly optimised using back-propagation [51]. The best model was selected based on the minimum validation loss (loss on development set). The spectro-temporal GAT model has only 0.44M parameters and is comparatively light weight compared to the baseline as well as other state of the art systems. All experiments were performed on a single GeForce RTX 3090 GPU and reproducible with the same random seed and GPU environment using an open source implementation¹.

6. Experimental Results

In this section, we present our results, an ablation study and a comparison of our results to those of existing, state-of-the-art single systems.

6.1. Results

Results are illustrated in Table 2, where the last two columns P1 and P2 indicate pooled min t-DCF and pooled EER results for the baseline system (RawNet2) and three different configurations of the new RawGAT-ST system. The variants involve the use of different spectro-temporal fusion strategies. Whereas all RawGAT-ST systems outperform the baseline by a substantial margin, the best result is obtained using the RawGAT-ST-mul system for which the t-DCF is 0.0335 (cf. 0.1547 for the baseline) and the EER is 1.06% (5.54%).

These results show that all RawGAT-ST systems are effective in exploiting spectro-temporal attention and to model the relationships between different spectro-temporal estimates, thereby improving the discrimination between spoofed and bona fide inputs. We now seek to demonstrate the merit of early fusion at the model level, rather than late fusion at the score level. This is done by way of ablation experiments.

6.2. Ablation study

Only through ablation experiments can we properly demonstrate the merit of the RawGAT-ST approach; we cannot use results from our previous work [13] since they were generated using hand-crafted features. Through ablation, we essentially remove one of the blocks or operations in the full RawGAT-ST architecture illustrated in Fig. 1. The remaining blocks and operations are then used as before. Results are illustrated in Table 3. The middle row highlighted in boldface is the RawGAT-ST-mul result that is also illustrated in Table 2.

Ablation of the spectral GAT attention block (top left in Fig. 1) leaves the system being capable of exploiting only temporal attention. Without spectral attention (first row of Table 3), performance degrades by 34% relative to the full system (0.0514 cf. 0.0335). The degradation without temporal attention (0.0385) is less severe (13% relative), indicating the greater importance of spectral attention versus temporal attention, even if both are beneficial. Last, we demonstrate the benefit of graph pooling by ablating the pooling layers in all three blocks. The relative degradation in performance of 58% (0.0788 cf. 0.0335) is even more substantial and shows the benefit of using graph pooling to concentrate on the most informative node features.

¹<https://github.com/eurecom-asp/RawGAT-ST-antispoofing>

Table 2: Results for the ASVspoof 2019 logical access (LA) database in terms of min t-DCF for each attack in the evaluation partition (A07-A19). Pooled min t-DCF (P1) and pooled EER (P2) are shown in the last two columns. Results shown for the end-to-end RawNet2 baseline system and the three RawGAT-ST based systems introduced in this paper.

System	A07	A08	A09	A10	A11	A12	A13	A14	A15	A16	A17	A18	A19	P1	P2
RawNet2-baseline [12, 13]	.098	.179	.073	.089	.042	.088	.020	.013	.073	.046	.240	.629	.058	0.1547	5.54
RawGAT-ST-add	.011	.030	.004	.016	.009	.019	.007	.007	.014	.027	.053	.104	.023	0.0373	1.15
RawGAT-ST-concat	.021	.027	.003	.027	.008	.029	.015	.008	.022	.029	.046	.120	.022	0.0388	1.23
RawGAT-ST-mul	.010	.016	.002	.012	.010	.010	.010	.009	.009	.023	.055	.080	.024	0.0335	1.06

Table 3: Results for ablation studies

System	min-tDCF	EER
w/o spectral attention	0.0514	1.87
w/o temporal attention	0.0385	1.13
w/ spectro-temporal attention	0.0335	1.06
w/o graph pooling	0.0788	2.47

6.3. Performance comparison

Illustrated in Table 4 is a comparison of performance for the proposed RawGAT-ST system and competing single systems reported in the literature [52]. The comparison shows that our system which uses GATs with self attention outperforms alternative attention approaches such as: Convolutional Block Attention Module (CBAM); Squeeze-and-Excitation (SE); Dual attention module with pooling and convolution operations. Furthermore, to the best of our knowledge, our approach is the best single model system reported in the literature to date. Our system is also one of only two that operates on the raw signal. With only 0.44M parameters, the proposed RawGAT-ST system is also among the least complex; only the Res-TSSDNet (0.35M) [36] and LCNN-LSTM-sum (0.27M) [53] have fewer parameters, while our system performs better. Despite the simplicity, the RawGAT-ST system outperforms our previous best temporal GAT system with late score fusion [13] (row 12 in Table 4) by 63% and 77% relative in terms of min t-DCF and EER respectively. This result also points towards the benefit of operating directly upon the raw signal in fully end-to-end fashion.

7. Conclusions

Inspired from previous findings which show that different attention mechanisms work more or less well for different forms of spoofing attack, we report in this paper a novel graph neural network approach to anti-spoofing and speech deepfake detection based on model-level spectro-temporal attention.

The RawGAT-ST solution utilises a self attention mechanism to learn the relationships between different spectro-temporal estimates and the most discriminative nodes within the resulting graph. The proposed solution also operates directly upon the raw waveform, without the use of hand-crafted features, and is among the least complex of all solutions reported thus far. Experimental results for the standard ASVspoof 2019 logical access database show the success of our approach and are the best reported to date by a substantial margin. Results for an ablation study show that spectral attention is more important than temporal attention, but that both are beneficial, and that graph pooling improves performance substantially. Last, a comparison of different systems, working either on the raw sig-

Table 4: Performance for the ASVspoof 2019 evaluation partition in terms of pooled min t-DCF and pooled EER for different state-of-the-art single systems and the best performing RawGAT-ST-mul system (boldface). Only the top two RawGAT-ST and Res-TSSDNet systems operate directly upon raw inputs; all others use some form of hand-crafted representations.

System	front-end	min-tDCF	EER
RawGAT-ST (mul)	Raw-audio	0.0335	1.06
Res-TSSDNet [36]	Raw-audio	0.0481	1.64
MCG-Res2Net50+CE [54]	CQT	0.0520	1.78
ResNet18-LMCL-FM [14]	LFB	0.0520	1.81
LCNN-LSTM-sum [53]	LFCC	0.0524	1.92
Capsule network [55]	LFCC	0.0538	1.97
Resnet18-OC-softmax [56]	LFCC	0.0590	2.19
MLCG-Res2Net50+CE [54]	CQT	0.0690	2.15
SE-Res2Net50 [57]	CQT	0.0743	2.50
LCNN-Dual attention [58]	LFCC	0.0777	2.76
Resnet18-AM-Softmax [56]	LFCC	0.0820	3.26
ResNet18-GAT-T [13]	LFB	0.0894	4.71
GMM [11]	LFCC	0.0904	3.50
ResNet18-GAT-S [13]	LFB	0.0914	4.48
PC-DARTS [59]	LFCC	0.0914	4.96
Siamese CNN [60]	LFCC	0.0930	3.79
LCNN-4CBAM [58]	LFCC	0.0939	3.67
LCNN+CE [61]	DASC-CQT	0.0940	3.13
LCNN [62]	LFCC	0.1000	5.06
LCNN+CE [63]	CQT	0.1020	4.07
ResNet [64]	CQT-MMPS	0.1190	3.72

nal or hand-crafted inputs show the benefit of a fully end-to-end approach.

We foresee a number of directions with potential to improve performance further. These ideas include the use of learnable pooling ratios to optimise the number of node features retained after graph pooling and the use of a learnable weighted selective fusion module to dynamically select the most discriminative features from spectral and temporal attention networks. Working with the latest ASVspoof 2021 challenge database [65], these ideas are the focus of our ongoing work.

8. Acknowledgements

This work is supported by the VoicePersonae project funded by the French Agence Nationale de la Recherche (ANR) and the Japan Science and Technology Agency (JST).

9. References

- [1] M. Sahidullah, T. Kinnunen, and C. Hanilci, "A comparison of features for synthetic speech detection," in *Proc. of the International Speech Communication Association (INTERSPEECH)*, 2015, pp. 2087–2091.
- [2] K. Sriskandaraja, V. Sethu, P. Ngoc Le, and E. Ambikairajah, "Investigation of sub-band discriminative information between spoofed and genuine speech," in *Proc. INTERSPEECH*, 2016, pp. 1710–1714.
- [3] M. Witkowski, S. Kacprzak, P. Zelasko, K. Kowalczyk, and J. Galka, "Audio replay attack detection using high-frequency features," in *Proc. INTERSPEECH*, 2017, pp. 27–31.
- [4] P. Nagarsheth, E. Khoury, K. Patil, and M. Garland, "Replay attack detection using DNN for channel discrimination," in *Proc. INTERSPEECH*, 2017, pp. 97–101.
- [5] L. Lin, R. Wang, and Y. Diqun, "A replay speech detection algorithm based on sub-band analysis," in *International Conference on Intelligent Information Processing (IIP)*, Nanning, China, 2018, pp. 337–345.
- [6] B. Chettri, D. Stoller, V. Morfi, M. A. M. Ramírez, E. Benetos, and B. L. Sturm, "Ensemble models for spoofing detection in automatic speaker verification," in *Proc. INTERSPEECH*, 2019, pp. 1118–1112.
- [7] J. Yang, R. K. Das, and H. Li, "Significance of subband features for synthetic speech detection," *IEEE Transactions on Information Forensics and Security*, 2019.
- [8] S. Garg, S. Bhilare, and V. Kanhangad, "Subband analysis for performance improvement of replay attack detection in speaker verification systems," in *International Conference on Identity, Security, and Behavior Analysis (ISBA)*, 2019.
- [9] B. Chettri, T. Kinnunen, and E. Benetos, "Subband modeling for spoofing detection in automatic speaker verification," in *Proc. Speaker Odyssey Workshop*, 2020.
- [10] H. Tak, J. Patino, A. Nautsch, N. Evans, and M. Todisco, "An explainability study of the constant Q cepstral coefficient spoofing countermeasure for automatic speaker verification," in *Proc. Speaker Odyssey Workshop*, 2020, pp. 333–340.
- [11] H. Tak, J. Patino, A. Nautsch, et al., "Spoofing Attack Detection using the Non-linear Fusion of Sub-band Classifiers," in *Proc. INTERSPEECH*, 2020, pp. 1106–1110.
- [12] H. Tak, J. Patino, M. Todisco, A. Nautsch, N. Evans, and A. Larcher, "End-to-end anti-spoofing with rawnet2," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021.
- [13] H. Tak, J.-w. Jung, J. Patino, M. Todisco, and N. Evans, "Graph attention networks for anti-spoofing," in *Proc. INTERSPEECH*, 2021.
- [14] T. Chen, A. Kumar, P. Nagarsheth, G. Sivaraman, and E. Khoury, "Generalization of audio deepfake detection," in *Proc. Speaker Odyssey Workshop*, 2020, pp. 132–137.
- [15] J. Youngberg and S. Boll, "Constant-q signal analysis and synthesis," in *Proc. ICASSP*, 1978, pp. 375–378.
- [16] J. Brown, "Calculation of a constant Q spectral transform," *Journal of the Acoustical Society of America (JASA)*, vol. 89, no. 1, pp. 425–434, 1991.
- [17] C. Schörkhuber, A. Klapuri, N. Holighaus, and M. Dörfler, "A matlab toolbox for efficient perfect reconstruction time-frequency transforms with log-frequency resolution," in *Proc. Audio Engineering Society International Conference on Semantic Audio*, London, UK, 2014.
- [18] M. Todisco, H. Delgado, and N. Evans, "A new feature for automatic speaker verification anti-spoofing: Constant Q cepstral coefficients," in *Proc. Speaker Odyssey Workshop*, 2016, pp. 249–252.
- [19] J. Jung, H. Heo, I. Yang, H. Shim, and H. Yu, "Avoiding speaker overfitting in end-to-end dnns using raw waveform for text-independent speaker verification," in *Proc. INTERSPEECH*, 2018, pp. 3583–3587.
- [20] M. Ravanelli and Y. Bengio, "Speaker recognition from raw waveform with sincnet," in *IEEE Proc. Spoken Language Technology Workshop (SLT)*, 2018, pp. 1021–1028.
- [21] M. R. Kamble, H. B. Sailor, H. A. Patil, and H. Li, "Advances in anti-spoofing: from the perspective of asvspoof challenges," *APSIPA Transactions on Signal and Information Processing*, vol. 9, 2020.
- [22] J.-w. Jung, S.-b. Kim, H.-j. Shim, J.-h. Kim, and H.-J. Yu, "Improved RawNet with Filter-wise Rescaling for Text-independent Speaker Verification using Raw Waveforms," in *Proc. INTERSPEECH*, 2020, pp. 1496–1500.
- [23] M. Gori, G. Monfardini, and F. Scarselli, "A new model for learning in graph domains," in *Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005. IEEE*, 2005, vol. 2, pp. 729–734.
- [24] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini, "The graph neural network model," *IEEE transactions on neural networks*, vol. 20, no. 1, pp. 61–80, 2009.
- [25] M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst, "Geometric deep learning: going beyond euclidean data," *IEEE Signal Processing Magazine*, vol. 34, no. 4, pp. 18–42, 2017.
- [26] W. L. Hamilton, R. Ying, and J. Leskovec, "Inductive representation learning on large graphs," in *Proc. Conference on Neural Information Processing Systems (NIPS)*, 2017, pp. 1025–1035.
- [27] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and S. Yu Philip, "A comprehensive survey on graph neural networks," *IEEE transactions on neural networks and learning systems*, vol. 32, no. 1, pp. 4–24, 2020.
- [28] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," in *International Conference on Learning Representations (ICLR)*, 2017.
- [29] P. Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio, "Graph attention networks," in *Proc. ICLR*, 2018.
- [30] S. Zhang, Y. Qin, K. Sun, and Y. Lin, "Few-shot audio classification with attentional graph neural networks," in *INTERSPEECH*, 2019, pp. 3649–3653.
- [31] R. Liu, B. Sisman, and H. Li, "Graphspeech: Syntax-aware graph attention network for neural speech synthesis," in *Proc. ICASSP*, 2021.
- [32] J.-w. Jung, H.-S. Heo, H.-J. Yu, and J. S. Chung, "Graph attention networks for speaker verification," in *Proc. ICASSP*, 2021.

- [33] P. Tzirakis, A. Kumar, and J. Donley, “Multi-channel speech enhancement using graph neural networks,” in *Proc. ICASSP*, 2021, pp. 3415–3419.
- [34] G. Klambauer, T. Unterthiner, A. Mayr, and S. Hochreiter, “Self-normalizing neural networks,” in *Proc. NIPS*, 2017, pp. 972–981.
- [35] S. Ioffe and C. Szegedy, “Batch normalization: Accelerating deep network training by reducing internal covariate shift,” in *International Conference on Machine Learning (ICML)*, 2015, pp. 448–456.
- [36] G. Hua, A. Beng jin teoh, and H. Zhang, “Towards end-to-end synthetic speech detection,” *IEEE Signal Processing Letters*, 2021.
- [37] J.-w. Jung, H.-S. Heo, J.-h. Kim, H.-j. Shim, and H.-J. Yu, “Rawnet: Advanced end-to-end deep neural network using raw waveforms for text-independent speaker verification,” in *Proc. INTERSPEECH*, 2019, pp. 1268–1272.
- [38] T. F. Quatieri, *Discrete-time speech signal processing: principles and practice*, Pearson Education, 2006.
- [39] J. R. Deller Jr, J. G. Proakis, and J. H. Hansen, *Discrete time processing of speech signals*, Prentice Hall PTR, 1993.
- [40] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proc. CVPR*, 2016, pp. 770–778.
- [41] H. Gao and S. Ji, “Graph u-nets,” in *international conference on machine learning (PMLR)*, 2019, pp. 2083–2092.
- [42] X. Wang, J. Yamagishi, M. Todisco, H. Delgado, A. Nautsch, N. Evans, M. Sahidullah, V. Vestman, T. Kinnunen, K. A. Lee, et al., “ASVspoof 2019: a large-scale public database of synthesized, converted and replayed speech,” *Computer Speech & Language (CSL)*, vol. 64, 2020, 101114.
- [43] M. Todisco, X. Wang, V. Vestman, M. Sahidullah, H. Delgado, A. Nautsch, et al., “ASVspoof 2019: Future horizons in spoofed and fake audio detection,” in *Proc. INTERSPEECH*, 2019, pp. 1008–1012.
- [44] T. Kinnunen, K. Lee, H. Delgado, N. Evans, M. Todisco, J. Sahidullah, M. and Yamagishi, and D. A. Reynolds, “t-DCF: A detection cost function for the tandem assessment of spoofing countermeasures and automatic speaker verification,” in *Proc. Speaker Odyssey Workshop*, 2018, pp. 312–319.
- [45] T. Kinnunen, H. Delgado, N. Evans, K. A. Lee, V. Vestman, A. Nautsch, M. Todisco, X. Wang, M. Sahidullah, J. Yamagishi, and D. A. Reynolds, “Tandem assessment of spoofing countermeasures and automatic speaker verification: Fundamentals,” *IEEE/ACM Transactions on Audio Speech and Language Processing (TASLP)*, vol. 28, pp. 2195 – 2210, 2020.
- [46] N. Brümmer and E. de Villiers, “The BOSARIS toolkit user guide: Theory, algorithms and code for binary classifier score processing,” 2011.
- [47] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, “SpecAugment: A simple data augmentation method for automatic speech recognition,” in *Proc. INTERSPEECH*, 2019, pp. 2613–2617.
- [48] H. Wang, Y. Zou, and W. Wang, “SpecAugment++: A hidden space data augmentation method for acoustic scene classification,” in *Proc. INTERSPEECH*, 2021.
- [49] P. Parasu, J. Epps, K. Sriskandaraja, and G. Suthokumar, “Investigating light-resnet architecture for spoofing detection under mismatched conditions,” in *Proc. INTERSPEECH*, 2020, pp. 1111–1115.
- [50] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [51] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*, MIT Press, 2016, <http://www.deeplearningbook.org>.
- [52] A. Nautsch, X. Wang, N. Evans, T. 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 (T-BIOM)*, vol. 3, 2021.
- [53] X. Wang and J. Yamagishi, “A comparative study on recent neural spoofing countermeasures for synthetic speech detection,” in *Proc. INTERSPEECH*, 2021.
- [54] X. Li, X. Wu, H. Lu, X. Liu, and H. Meng, “Channel-wise gated res2net: Towards robust detection of synthetic speech attacks,” in *Proc. INTERSPEECH*, 2021.
- [55] A. Luo, E. Li, Y. Liu, X. Kang, and Z. J. Wang, “A capsule network based approach for detection of audio spoofing attacks,” in *Proc. ICASSP*, 2021, pp. 6359–6363.
- [56] Y. Zhang, F. Jiang, and Z. Duan, “One-class learning towards synthetic voice spoofing detection,” *IEEE Signal Processing Letters*, vol. 28, pp. 937–941, 2021.
- [57] X. Li, N. Li, C. Weng, X. Liu, D. Su, D. Yu, and H. Meng, “Replay and synthetic speech detection with res2net architecture,” in *Proc. ICASSP*, 2021, pp. 6354–6358.
- [58] X. Ma, T. Liang, S. Zhang, S. Huang, and L. He, “Improved lightcnn with attention modules for asv spoofing detection,” in *2021 IEEE International Conference on Multimedia and Expo (ICME)*, 2021, pp. 1–6.
- [59] W. Ge, M. Panariello, J. Patino, M. Todisco, and N. Evans, “Partially-Connected Differentiable Architecture Search for Deepfake and Spoofing Detection,” in *Proc. INTERSPEECH*, 2021.
- [60] Z. Lei, Y. Yang, C. Liu, and J. Ye, “Siamese convolutional neural network using gaussian probability feature for spoofing speech detection,” in *Proc. INTERSPEECH*, 2020, pp. 1116–1120.
- [61] R. K. Das, J. Yang, and H. Li, “Data augmentation with signal companding for detection of logical access attacks,” in *Proc. ICASSP*, 2021, pp. 6349–6353.
- [62] G. Lavrentyeva, S. Novoselov, A. Tseren, et al., “STC antispooofing systems for the ASVspoof2019 challenge,” in *Proc. INTERSPEECH*, 2019, pp. 1033–1037.
- [63] Z. Wu, R. K. Das, J. Yang, and H. Li, “Light convolutional neural network with feature genuinization for detection of synthetic speech attacks,” in *Proc. INTERSPEECH*, 2020, pp. 1101–1105.
- [64] J. Yang, H. Wang, R. K. Das, and Y. Qian, “Modified magnitude-phase spectrum information for spoofing detection,” *IEEE/ACM TASLP*, vol. 29, pp. 1065–1078, 2021.
- [65] J. Yamagishi, X. Wang, M. Todisco, M. Sahidullah, J. Patino, A. Nautsch, X. Liu, K. A. Lee, T. Kinnunen, N. Evans, and H. Delgado, “ASVspoof2021: accelerating progress in spoofed and deep fake speech detection,” in *Proc. ASVspoof 2021 Workshop*, 2021.