

Point to the Hidden: Exposing Speech Audio Splicing via Signal Pointer Nets

Denise Moussa^{1,2}, Germans Hirsch², Sebastian Wankerl³, Christian Riess²

¹Federal Criminal Police Office (BKA), Germany

²IT Security Infrastructures Lab, Friedrich-Alexander University of Erlangen-Nürnberg, Germany

³Data Science Chair, Julius-Maximilians University of Würzburg, Germany

{denise.moussa, christian.riess}@fau.de

Abstract

Verifying the integrity of voice recording evidence for criminal investigations is an integral part of an audio forensic analyst's work. Here, one focus is on detecting deletion or insertion operations, so called audio splicing. While this is a rather easy approach to alter spoken statements, careful editing can yield quite convincing results. For difficult cases or big amounts of data, automated tools can support in detecting potential editing locations. To this end, several analytical and deep learning methods have been proposed by now. Still, few address unconstrained splicing scenarios as expected in practice. With SigPointer, we propose a pointer network framework for continuous input that uncovers splice locations naturally and more efficiently than existing works. Extensive experiments on forensically challenging data like strongly compressed and noisy signals quantify the benefit of the pointer mechanism with performance increases between about 6 to 10 percentage points.¹

Index Terms: Audio Splicing Localisation, Audio Forensics, Pointer Networks

1. Introduction

In today's digital era, more and more speech recordings like voice messages, recorded phone calls or audio tracks of videos are produced and possibly post-processed and shared via the internet. Consequently, they often contain important cues for criminal investigations, too. With powerful tools, either commercial or free, as for example Audacity [1], the hurdles for editing operations have become low. Forensic audio analysts are thus often assigned to verify the integrity of material relevant to court cases. Audio splicing (which subsumes deletion, copying and insertion of speech segments) is an effective and easy-to-perform manipulation that violates integrity. For example, the sentence *I do not agree* is easily inverted by deleting the signal segment containing *not* and merging the remaining parts. Simple post-processing steps, such as saving forgeries using lossy compression, *e.g.* in MP3 format, can further weaken or obscure editing cues. Furthermore, the workload for analysts strongly increases with forgery quality and amount of data. Up to now, several methods have been proposed to assist with localising splices in speech material. However, they are mostly inapplicable to unconstrained signal characteristics. With this work, we address the current limitations and propose a novel and natural approach to audio splicing localisation.

1.1. Existing Approaches to Audio Splicing Localisation

Audio splicing localisation is mostly targeted with analytical and deep learning (DL) based methods that focus on specific

features to detect signal inconsistencies. By example, some previous works examine the consistency of specific audio formats [2–4], rely on splices of recordings from different devices [5, 6], or detect changes of the recording environment's noise levels [7, 8], acoustic impulse [9] or both [10]. Others also search for atypical changes in the subtle (and fragile) electric network frequency (ENF) [11–13]. Due to the rise of convincing audio synthesis techniques, several works specifically aim at detecting artificial segments amidst original speech [14, 15].

In practice, forensic analysts are confronted with audio samples from unconstrained sources, which implies, *e.g.*, arbitrary recording parameters, quality, formats, or post-processing operations. So, methods relying on the presence of very specific features might not be applicable if those are not present in certain audio signals. Indeed, recently, several DL approaches with unconstrained feature extraction have been proposed [16–21], however many of these methods are still preliminary. Some either target audio splicing detection but omit localisation [16–18], another work examines only two fixed splicing patterns, and its generalization remains unclear [19]. In addition, the small and non-diverse Free Spoken Digit Dataset [22] is often used to construct spliced samples [16–18], and frequently, material from different speakers instead of one is merged [16–19], which excludes the highly relevant and more difficult-to-detect case of forged statements of one person. Zeng *et al.* [20] consider more miscellaneous spliced forgeries from one speaker and employ a ResNet-18 [23] method for chunks of audio spectrogram frames. However, this is only fit for coarse splicing localisation within windows of 32 to 64 frames. For frame-level granularity, a sequence-to-sequence (seq2seq) Transformer [24] model has been proposed [21]. It outperforms several convolutional neural network (CNN) classifiers on challenging data, but still leaves room for improvement concerning well-made forgeries.

1.2. Audio Splicing Localisation via Pointer Mechanisms

We propose a major improvement over existing methods for audio splicing localisation by regarding this task as a pure pointing problem. Pointers predict a conditional probability distribution over elements of a sequence and were originally designed for approximating combinatorial optimization problems [25]. They can thus directly locate parts of the input series, in contrast to traditional seq2seq networks, where a mapping to a fixed set of target tokens is learned. Pointer mechanisms were recently also integrated into the Transformer [24] architecture, where mixtures of pointer and token generation components solve natural language processing tasks [26, 27]. In our case, we want the neural network (NN) to indicate splice locations by pointing to the respective input signal positions. This appears more nat-

¹code: <https://www.cs1.tf.fau.de/research/multimedia-security>

ural and efficient than step-wise classifying segments of fixed size [20] or learning a mapping to a fixed vocabulary [21]. Existing pointer methods however operate on categorical input. We thus design SigPointer, a Transformer [24] based pointer network for continuous input signals. We benchmark against existing works on audio splicing localisation and analyse the influence of our network’s components. SigPointer proves to perform best on forensically challenging data, both under seen and unseen quality degradations.

2. Proposed Method

In this section, we describe our proposed pointing method for signals, as well as training strategies and datasets.

2.1. SigPointer for Continuous Input Signals

We define audio splicing localisation as a pointing task. The input to our encoder-decoder network (Fig. 1) is a time series of signal representation vectors with length N , formally $\mathbf{S} = [\mathbf{s}_0, \mathbf{s}_1, \dots, \mathbf{s}_N]$. Additionally, \mathbf{s}_{-1} denotes $\langle eos \rangle$, a special token to which the pointer mechanism can point to denote that decoding is finished [27]. It is omitted in the following definitions for simplicity. The encoder part of our network (left) mostly aligns with the original Transformer [24] model. It is made of a stack of layers, each implementing multi-head self-attention followed by a fully-connected (FC) layer and layer-normalisation. However, unlike existing pointer methods on categorical data (cf. Sec. 1.2), we skip learning input embeddings and feed the raw data into the network, since we already operate on continuous, dense data. The encoded representation $\mathbf{H} = [\mathbf{h}_0, \mathbf{h}_1, \dots, \mathbf{h}_N]$ of \mathbf{S} , where $\mathbf{h}_n \in \mathbb{R}^l$ with latent size l , then serves as input to the decoder (middle). The decoder solely consists of a stack of multi-head attention layers [24]. Per time step t , it takes \mathbf{H} and one position vector $\mathbf{z}_{t^*} \in \mathbb{R}^l$ per previously predicted $\hat{y}_{t^*} \in \mathcal{I} = \{0, 1, \dots, N\}$ with $t^* < t$ to indicate the next splice location \hat{y}_t . More details about this decoding process are given in the following two paragraphs.

Pointer Mechanism To predict the full output index sequence $\hat{\mathbf{y}}$, we need the conditional probability distribution over all input positions $P(\hat{\mathbf{y}}|\mathbf{S}) = \prod_{t=0}^T P(\hat{y}_t|\mathbf{H}, \mathbf{Z})$. Thus, for each time step, we extract the cross-attention scores between \mathbf{H} and \mathbf{Z} from the last decoder layer for all M attention heads, yielding $\mathbf{A}_t^{M \times |\mathbf{S}|}$, with $a_{m,s} \in \mathbb{R}$ (Fig. 1, right). Unlike related Transformer pointer methods, we do not mix pointer and token generation tasks [26, 27], so all the model’s attention heads can be reserved for computing one final pointer result. We thus average and normalize all attention heads’ values to yield the discrete index probability distribution as $\hat{\mathbf{p}}_t = \text{softmax}(\overline{\mathbf{A}_t})_1^M$. During inference, $\hat{y}_t = \text{argmax}(\hat{\mathbf{p}}_t)$ is computed to yield the final splice point or $\langle eos \rangle$ if no (further) splice point is detected.

Slim Decoding Existing seq2seq (pointer) approaches on categorical data traditionally use the semantic of the previously decoded elements \hat{y}_{t^*} to decode the next \hat{y}_t from \mathbf{H} . Hereby, \hat{y}_{t^*} is projected to latent size l with an embedding layer E_m and positional encodings are added to yield the final representation. Contrary, we only use sinusoidal encoding vectors [24] $\mathbf{z}_{t^*} = \mathbf{e}_{t^*}$ for each \hat{y}_{t^*} . The vectors \mathbf{z}_{t^*} thus only preserve the relative ordering and number of already decoded output items. In several tests, we observed that this slim information is sufficient, since training some E_m and reusing y_{t^*} as decoder target input did not show any advantage. In fact, the performance even slightly degraded. We reason that this is due to sparser available context between sequence elements compared to approaches on

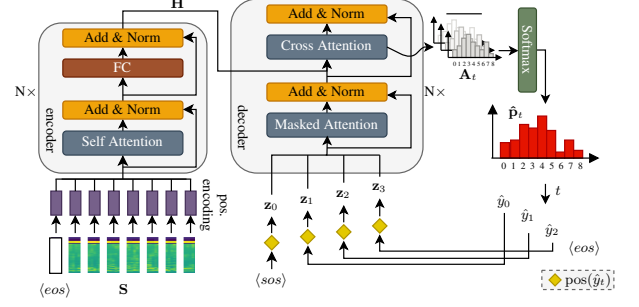


Figure 1: SigPointer model for locating splices in audio signals

discrete data, e.g., word tokens from a natural language [26, 27]. Differently from the latter, splicing points carry few syntactical information (locations) and are generally not expected to have rich dependencies among each other. So, previously found splice locations do not provide clues about future ones.

2.2. Choosing Model and Training Parameters

We conduct a parameter search of 200 trials with optuna [28] (v.2.10) and PyTorch [29] (v.1.10.2) on 4 consumer GPUs. The search space encompasses $n_e, n_d \in [1, 8]$ encoder and decoder layers, $h \in \{1, 3, 9\}$ heads (divisors of latent size $l = 279$), $f_d \in [2^7, 2^{11}]$ for the FC layer dimension, dropout $d \in [0.01, 0.05, \dots, 0.3]$ within the attention layers, learning rate $l_r \in [1e^{-5}, 5e^{-5}, \dots, 5e^{-3}]$ for the Adam optimizer and batch size $b \in [32, 64, 128, 256, 350]$. We choose the best configuration $\mathcal{C} : (n_e, n_d, h, f_d, d, l_r, b) = (7, 1, 9, 2^7, 0.1, 5e^{-4}, 350)$ for our network and use Glorot weight initialisation [30]. SigPointer converges significantly better with regression as opposed to classification losses. The cosine distance loss $d_c = \frac{\|\hat{\mathbf{p}}_t - \mathbf{p}_t\|}{\|\hat{\mathbf{p}}_t\| \|\mathbf{p}_t\|} \in [0, 1]$ proves to be most stable, where \mathbf{p}_t are the one-hot encoded target splice indices. Note that a more exact metric for our task is the mean absolute error between the predicted positions $\hat{\mathbf{y}}$ and targets \mathbf{y} . Since it guided training better but is non-differential, we only use it as validation loss.

2.3. Datasets and Training Strategy

All datasets are generated using our previously introduced pipeline [21] and cover forensically challenging scenarios. We use the anechoic ACE dataset [31] as source set, such that the model does not adapt to unintended, acoustic side channels in forged samples. Samples between 3 s and 45 s are created from n_s source segments of the same speaker from the same or different simulated environments. Post-processing that may obscure splicing is applied in the form of additive Gaussian noise and either single AMR-NB or MP3 compression, all in randomly sampled strength. Adding noise to a signal can serve as an easy method to mask tampering, just as compression which however may also be unintentionally introduced by re-saving the edited version or sharing via social media services. The final feature representation with a time resolution of 500 ms is a concatenation of the Mel spectrogram, MFCCs and spectral centroids, yielding the feature dimension $l = 279$. For more details about the generation pipeline we refer to the respective work [21].

For training, we employ curriculum learning [32] in a three-stage process with a train/validation split of 500k/30k for each stage. Training is conducted upon model convergence with 100 epochs and a patience of 20 epochs. The first dataset covers

samples with $n \in [0, 1]$ splices and no post-processing, the second extends to post-processing and the third includes both multi-splicing with $n \in [0, 5]$ and post-processing. As in existing work [21], we employ cross-dataset testing and generate our test sets from the Hi-Fi TTS set [33] as described in Sec. 3.

3. Experiments

We evaluate our proposed method on forensically challenging data and test against four existing methods and two custom baselines to analyse the benefits of the pointer framework.

3.1. Baseline Models

We train each NN as described in Sec. 2.3, where all but the pointer approaches employ the BCE instead of the cosine loss.

CNNs Most existing works rely on CNNs and natively have too strong limitations to be directly used for our task (cf. Sec. 1.1). We reimplement three approaches [16, 18, 20]. The first two frameworks [16, 18] only cover binary classification of spliced vs. non-spliced signals with custom CNN models, while Zeng *et al.* [20] propose localisation with a ResNet-18 [23] classifier. They employ a sliding window approach over chunks of frames, where the window stride $s = 1$ accounts for exact (frame-level) detection. The classifier’s decision per frame is inferred from averaged probabilities of multiple windows. This approximative method is however only fit for larger s [20]. Our reimplementation confirms long training times and poor detection on frame level, so we use a custom splice localisation strategy for CNNs. It is based on our previously introduced framework [21], where baseline CNNs are extended to classify at maximum n splicing positions with $n_o = n$ output layers. This however shows poor performance for $n > 1$. Instead, we set n_o to the maximum expected number of input frames, 90 in our case, and perform a binary classification per frame which proves to be more stable w.r.t. higher n . Splitting in smaller segments accounts for signals with lengths exceeding n_o . We exclude two methods because of no details on the model specifications [17] and very restrictive splicing assumptions where a relaxation to our more unconstrained task is not straight forward [19].

Seq2Seq Transformer We also re-train our seq2seq model from previous work on multi-splicing localisation [21]. Given a signal, the encoder outputs its latent representation from which a sequence of splicing points is decoded step-by-step using a fixed vocabulary set.

SigPointer_{C_M} For a direct comparison of the pointer against the seq2seq framework, we instantiate SigPointer with the Transformer configuration C_M in [21].

Transformer encoder SigPointer employs autoregressive decoding (cf. Sec. 2.1). To quantify its influence we test against the capacity of a plain, non-autoregressive Transformer encoder. Thereby, we project the encoder memory $\mathbf{H}^{l \times N}$ (Fig. 1) to size $2 \times N$ to perform a per frame classification as for the CNN baselines. For a fair comparison we again conduct a hyper-parameter search (Sec. 2.2), but double the search space for the number of encoder layers to $n_e \in [1, 16]$ to account for the missing decoder capacity. The best model configuration yields $\mathcal{C} : (n_e, h, f_d, d, l_r, b) = (12, 9, 2^{11}, 0.1, 1e^{-4}, 64)$.

3.2. Performance Metrics

We report the average Jaccard index J expressing the similarity of prediction $\hat{\mathbf{y}}$ and ground truth \mathbf{y} as intersection over union $J = \frac{|\hat{\mathbf{y}} \cap \mathbf{y}|}{|\hat{\mathbf{y}} \cup \mathbf{y}|}$, as well as average recall $R = \frac{|\hat{\mathbf{y}} \cap \mathbf{y}|}{|\hat{\mathbf{y}}|}$ of 5 training runs with different seeds. Note that the order of the predicted

points is irrelevant for both metrics. We evaluate both exact localisation and coarser granularity by binning the input signal by $f \in [1, 2, 3, 4]$ frames, denoted as $\text{Bin} = f$.

3.3. Evaluation of the Pointer Mechanism

For this experiment, we generate a dataset of 30k samples with uniformly sampled $n \in [0, 5]$ splicing positions from the Hi-Fi TTS test pool [33], including single compression and noise post-processing as described in Sec. 2.3. The size and performance results of all models are listed in Tab. 1. The CNN methods (rows 1 to 3) are inferior to the Transformer-based approaches (rows 4 to 7). Jadhav *et al.*’s [16] large but very shallow network performs worst, followed by Chuchra *et al.*’s [18] small but deeper model with 12 layers and the best and deepest ResNet-18 CNN baseline. Solving the same classification task (Sec. 3.1) with the Transformer encoder greatly improves splicing localisation. It also slightly surpasses the best proposed seq2seq model from related work [21]. However, it also uses about 2.3 as much trainable parameters. Training the seq2seq model in our pointer framework (SigPointer_{C_M}) demonstrates the benefit of our proposed approach. The missing seq2seq vocabulary mapping component slightly reduces the network size, still the Jaccard index and recall increase by approximately 5.0 percentage points (pp) and 4.8 pp and outperform both the Transformer encoder and the original seq2seq model. We achieve the best performance with optimized hyperparameters (cf. Sec. 2.2), yielding the even smaller SigPointer*. Notably, the decoder is reduced to $1 < 5$ layers compared to SigPointer_{C_M} which suffices for our slim decoding strategy (cf. Sec. 2.1). The advantage of about 5.1 pp and 5.2 pp to SigPointer_{C_M} with $J = 0.5184 > 0.4670$ and $R = 0.5719 > 0.5202$ for $\text{Bin} = 1$ steadily decreases, reaching 3.8 pp (J) and 3.0 pp (R) for $\text{Bin} = 4$. We thus assume that SigPointer_{C_M} is only slightly less sensitive to splicing points than the optimized SigPointer* but notably less exact in localisation.

3.4. Influence of Splices per Input

In Fig. 2 we report the Jaccard coefficients on our Hi-Fi TTS test set w.r.t. the number of splice positions per sample. Evidently, all models recognise non-spliced inputs relatively well, while localising actual splice positions correctly proves to be more difficult on this challenging dataset (Fig. 2a-2d). The best performing SigPointer* (red) achieves $J = 0.5187$ for single splices and still $J = 0.3918$ for $n = 5$ splices. When allowing coarser signal binning (Fig. 2b-2d), the performance increases considerably up to $J = 0.7268$ and $J = 0.6104$ for $n = 1, 5$ and $\text{Bin} = 4$ (Fig. 2d). As stated in Sec. 3.3, the less accurate SigPointer_{C_M} (orange) benefits from coarser bins, but cannot outperform SigPointer*. The Transformer encoder (grey) performs well for $n = 1, 2$, but drops below the Transformer seq2seq model’s [21] (green) performance for $n \geq 3$. The CNNs (blue) are weakest, where especially the custom models [16, 18] exhibit low sensitivity to splicing and thus cannot profit from coarser signal binning (Fig. 2b-2d). In summary, SigPointer* outperforms all baselines in terms of sensitivity to splicing and exactness of localisation, despite its small model size of 3.4 M parameters.

3.5. Robustness to Complex Processing Chains

We test the generalization ability of the models trained only on single compression and additive Gaussian noise post-processing to even stronger obscured splicing points. We thus run the ro-

Table 1: *Model size and performance (mean \pm SD of 5 training runs) on the Hi-Fi TTS test set with $n \in [0, 5]$ splices per sample. Jaccard Index J and recall R are evaluated for exact localisation (Bin = 1) and w.r.t $f \in [2, 3, 4]$ frame binning of the input signal.*

Model	Params	Bin = 1		Bin = 2		Bin = 3		Bin = 4	
		J	R	J	R	J	R	J	R
Jadhav [16]	205.15 M	0.2189 \pm 0.004	0.2369 \pm 0.004	0.2547 \pm 0.004	0.2768 \pm 0.004	0.2745 \pm 0.004	0.2982 \pm 0.004	0.2952 \pm 0.004	0.3191 \pm 0.004
Chuchra [18]	134.77 K	0.2964 \pm 0.004	0.3089 \pm 0.004	0.3280 \pm 0.004	0.3373 \pm 0.005	0.3445 \pm 0.005	0.3509 \pm 0.005	0.3533 \pm 0.005	0.3594 \pm 0.005
Zeng [20]	11.17 M	0.3274 \pm 0.003	0.3867 \pm 0.002	0.4249 \pm 0.003	0.4921 \pm 0.004	0.4734 \pm 0.004	0.5413 \pm 0.005	0.5162 \pm 0.004	0.5831 \pm 0.005
Transf. enc.	17.51 M	0.4377 \pm 0.005	0.4855 \pm 0.007	0.5386 \pm 0.005	0.5825 \pm 0.008	0.5839 \pm 0.006	0.6231 \pm 0.009	0.6231 \pm 0.006	0.6574 \pm 0.010
Moussa [21]	7.62 M	0.4198 \pm 0.005	0.4747 \pm 0.008	0.5263 \pm 0.008	0.5794 \pm 0.011	0.5679 \pm 0.010	0.6187 \pm 0.014	0.6123 \pm 0.011	0.6595 \pm 0.015
SigPointer _{C_M}	7.57 M	0.4670 \pm 0.003	0.5202 \pm 0.003	0.5767 \pm 0.003	0.6265 \pm 0.006	0.6160 \pm 0.003	0.6620 \pm 0.006	0.6648 \pm 0.004	0.7061 \pm 0.008
SigPointer*	3.40 M	0.5184 \pm 0.006	0.5719 \pm 0.011	0.6228 \pm 0.012	0.6675 \pm 0.018	0.6607 \pm 0.016	0.7002 \pm 0.022	0.6977 \pm 0.019	0.7322 \pm 0.025

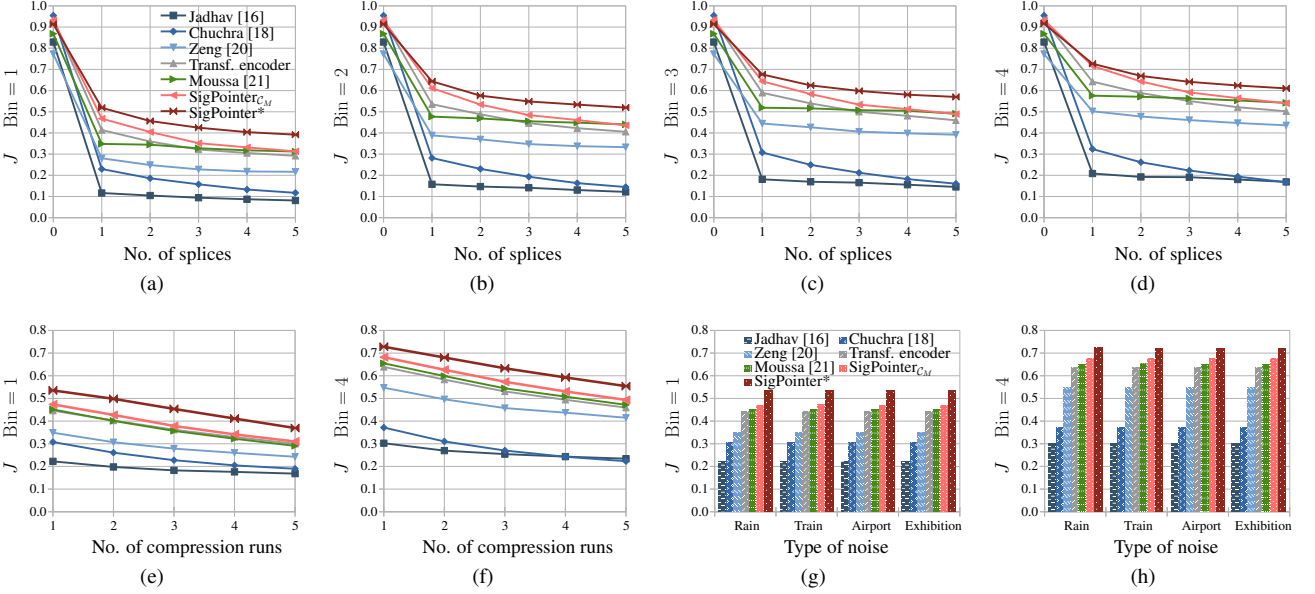


Figure 2: *Jaccard index J (mean of 5 training runs) for $n \in [0, 5]$ splices (Figure 2a-2d) and robustness towards out-of-distribution multi-compression and real noise post-processing (Fig. 2e-2h). The pointer framework (red tones) is clearly superior in all tests.*

business experiments from our previous work on the available test sets à 10k samples from Hi-Fi TTS [21]. They consist of 0 to 5 times spliced samples subjected to $n_c \in [1, 5]$ AMR-NB or MP3 compression runs (5 test sets) and additive real noise post-processing (4 test sets). Both compression and noise strength are randomly sampled [21]. Figures 2e-2h show the results for Bin = 1, 4. Most previously described trends in model performance (Sec. 3.3 and Sec. 3.4) also show in this robustness experiment. However, the seq2seq Transformer [21] this time surpasses the simpler Transformer encoder and shows slightly better robustness in all experiments. SigPointer* again outperforms all models. Compared to the best existing model [21] (green), for Bin = 1 it increases localisation ability by (on average) 8.4 pp for multi compression and 9.1 pp for additive real noise (Fig. 2e, Fig. 2g). Surprisingly, despite differing complexity, the type of noise has little influence on the performance.

3.6. Limitations

In our tests, SigPointer models are more sensitive to weight initialisation compared to the other considered NNs. The number of epochs until the model converges can thus vary strongly, so early pruning of weak runs is recommended in practice.

Also note that one design advantage of our model can be a pitfall in practice. SigPointer can process signals of arbitrary length, contrary to classifiers that are constrained by their number of output layers or seq2seq methods that are indirectly limited by their learned output vocabulary mapping (cf. Sec. 3.1). However, we empirically observed that the pointer adapts to signal lengths seen in training and does barely search for splices out of known ranges. The behaviour of inherent adaptation to problem sizes is already known from literature [34, 35]. To mitigate this issue, we thus strongly recommend cutting test samples into multiple separate segments that fit into the training distribution, as it was also done for the comparison methods.

4. Conclusions

With SigPointer, we present a novel and more natural approach to the task of audio splicing localisation with the help of pointer mechanisms. Our focus is on aiding with difficult-to-detect splice positions that pose by example a problem in forensic analysts' daily work. In several tests on in- and out-of-distribution data, we quantified the advantage of our pointer framework for continuous signals and outperform existing approaches by a large margin, even with a much smaller model size.

5. References

- [1] Audacity, “Audacity ® — Free, open source, cross-platform audio software for multi-track recording and editing,” accessed: 2023-02-27, <https://www.audacityteam.org/>.
- [2] R. Yang, Z. Qu, and J. Huang, “Detecting Digital Audio Forgeries by Checking Frame Offsets,” in *Proceedings of the 10th ACM Workshop on Multimedia and Security*, 2008, pp. 21–26.
- [3] A. J. Cooper, “Detecting Butt-Spliced Edits in Forensic Digital Audio Recordings,” in *Audio Engineering Society Conference: 39th International Conference: Audio Forensics: Practices and Challenges*, 2010.
- [4] Z. Xiang, P. Bestagini, S. Tubaro, and E. J. Delp, “Forensic Analysis and Localization of Multiply Compressed MP3 Audio Using Transformers,” in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2022, pp. 2929–2933.
- [5] L. Cuccovillo, S. Mann, M. Tagliasacchi, and P. Aichroth, “Audio Tampering Detection via Microphone Classification,” in *2013 IEEE 15th International Workshop on Multimedia Signal Processing*, 2013, pp. 177–182.
- [6] G. Baldini and I. Amerini, “Microphone Identification based on Spectral Entropy with Convolutional Neural Network,” in *IEEE International Workshop on Information Forensics and Security*, 2022, pp. 1–6.
- [7] X. Pan, X. Zhang, and S. Lyu, “Detecting Splicing in Digital Audios Using Local Noise Level Estimation,” in *2012 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2012, pp. 1841–1844.
- [8] D. Yan, M. Dong, and J. Gao, “Exposing Speech Transsplicing Forgery with Noise Level Inconsistency,” *Security and Communication Networks*, vol. 2021, 2021.
- [9] D. Capoferri, C. Borrelli, P. Bestagini, F. Antonacci, A. Sarti, and S. Tubaro, “Speech Audio Splicing Detection and Localization Exploiting Reverberation Cues,” in *2020 IEEE International Workshop on Information Forensics and Security*, 2020, pp. 1–6.
- [10] H. Zhao, Y. Chen, R. Wang, and H. Malik, “Audio Splicing Detection and Localization Using Environmental Signature,” *Multimedia Tools and Applications*, vol. 76, no. 12, pp. 13 897–13 927, 2017.
- [11] P. A. Esquef, J. A. Apolinário, and L. W. Biscainho, “Improved Edit Detection in Speech via ENF Patterns,” in *2015 IEEE International Workshop on Information Forensics and Security*, 2015, pp. 1–6.
- [12] X. Lin and X. Kang, “Supervised Audio Tampering Detection Using an Autoregressive Model,” in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2017, pp. 2142–2146.
- [13] M. Mao, Z. Xiao, X. Kang, X. Li, and L. Xiao, “Electric Network Frequency Based Audio Forensics Using Convolutional Neural Networks,” in *IFIP International Conference on Digital Forensics*, 2020, pp. 253–270.
- [14] B. Zhang and T. Sim, “Localizing Fake Segments in Speech,” in *26th International Conference on Pattern Recognition*, 2022, pp. 3224–3230.
- [15] K. Zhang, S. Liang, S. Nie, S. He, J. Pan, X. Zhang, H. Ma, and J. Yi, “A Robust Deep Audio Splicing Detection Method via Singularity Detection Feature,” in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2022, pp. 2919–2923.
- [16] S. Jadhav, R. Patole, and P. Rege, “Audio Splicing Detection Using Convolutional Neural Network,” in *IEEE 10th International Conference on Computing, Communication and Networking Technologies*, 2019, pp. 1–5.
- [17] R. Patole and P. P. Rege, “Machine Learning–Based Splicing Detection in Digital Audio Recordings for Audio Forensics,” *Journal of the Audio Engineering Society*, vol. 69, no. 11, pp. 793–804, 2021.
- [18] A. Chuchra, M. Kaur, and S. Gupta, “A Deep Learning Approach for Splicing Detection in Digital Audios,” in *Congress on Intelligent Systems*, 2022, pp. 543–558.
- [19] Z. Zhang, X. Zhao, and X. Yi, “ASLNet: An Encoder-Decoder Architecture for Audio Splicing Detection and Localization,” *Security and Communication Networks*, vol. 2022, 2022.
- [20] Z. Zeng and Z. Wu, “Audio Splicing Localization: Can We Accurately Locate the Splicing Tampering?” in *The 13th International Symposium on Chinese Spoken Language Processing*, 2022.
- [21] D. Moussa, G. Hirsch, and C. Riess, “Towards Unconstrained Audio Splicing Detection and Localization with Neural Networks,” in *ICPR International Workshops and Challenges*, 2022.
- [22] Z. Jackson, *Free Spoken Digit Dataset (FSDD)*, <https://github.com/Jakobovski/free-spoken-digit-dataset>, 2018.
- [23] K. He, X. Zhang, S. Ren, and J. Sun, “Deep Residual Learning for Image Recognition,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.
- [24] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention Is All You Need,” in *Advances in Neural Information Processing Systems*, 2017, pp. 5998–6008.
- [25] O. Vinyals, M. Fortunato, and N. Jaitly, “Pointer Networks,” *Advances in Neural Information Processing Systems*, vol. 28, 2015.
- [26] S. Enarvi, M. Amoia, M. D.-A. Teba, B. Delaney, F. Diehl, S. Hahn, K. Harris, L. McGrath, Y. Pan, J. Pinto *et al.*, “Generating Medical Reports from Patient-Doctor Conversations Using Sequence-to-Sequence Models,” in *Proceedings of the First Workshop on Natural Language Processing for Medical Conversations*, 2020, pp. 22–30.
- [27] H. Yan, J. Dai, T. Ji, X. Qiu, and Z. Zhang, “A Unified Generative Framework for Aspect-based Sentiment Analysis,” in *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 2021, pp. 2416–2429.
- [28] T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama, “Optuna: A Next-generation Hyperparameter Optimization Framework,” in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2019.
- [29] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, “PyTorch: An Imperative Style, High-Performance Deep Learning Library,” in *Advances in Neural Information Processing Systems*, 2019, pp. 8024–8035.
- [30] X. Glorot and Y. Bengio, “Understanding the Difficulty of Training Deep Feedforward Neural Networks,” in *Proceedings of the 13th International Conference on Artificial Intelligence and Statistics*, 2010, pp. 249–256.
- [31] J. Eaton, N. D. Gaubitch, A. H. Moore, and P. A. Naylor, “Estimation of Room Acoustic Parameters: The ACE Challenge,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 24, no. 10, pp. 1681–1693, 2016.
- [32] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, “Curriculum learning,” in *Proceedings of the 26th Annual International Conference on Machine Learning*, 2009, pp. 41–48.
- [33] E. Bakhturina, V. Lavrukhin, B. Ginsburg, and Y. Zhang, “Hi-Fi Multi-Speaker English TTS Dataset,” in *INTERSPEECH 2021*, 2021, pp. 2776–2780.
- [34] W. Kool, H. van Hoof, and M. Welling, “Attention, Learn to Solve Routing Problems!” in *International Conference on Learning Representations*, 2018.
- [35] C. Anil, Y. Wu, A. J. Andreassen, A. Lewkowycz, V. Misra, V. V. Ramasesh, A. Slone, G. Gur-Ari, E. Dyer, and B. Neyshabur, “Exploring Length Generalization in Large Language Models,” in *Advances in Neural Information Processing Systems*, A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho, Eds., 2022.