



# Introducing Model Inversion Attacks on Automatic Speaker Recognition

Karla Pizzi<sup>\*,1,2</sup>, Franziska Boenisch<sup>\*,1</sup>, Ugur Sahin<sup>\*,1,2</sup>, Konstantin Böttinger<sup>1</sup>

<sup>1</sup>Fraunhofer AISEC, Germany; <sup>2</sup>Technical University Munich, Germany

[firstname.lastname]@aisec.fraunhofer.de

## Abstract

Model inversion (MI) attacks allow to reconstruct average per-class representations of a machine learning (ML) model's training data. It has been shown that in scenarios where each class corresponds to a different individual, such as face classifiers, this represents a severe privacy risk. In this work, we explore a new application for MI: the extraction of speakers' voices from a speaker recognition system. We present an approach to (1) reconstruct audio samples from a trained ML model and (2) extract intermediate voice feature representations which provide valuable insights into the speakers' biometrics.

Therefore, we propose an extension of MI attacks which we call *sliding model inversion*. Our sliding MI extends standard MI by iteratively inverting overlapping chunks of the audio samples and thereby leveraging the sequential properties of audio data for enhanced inversion performance. We show that one can use the inverted audio data to generate spoofed audio samples to impersonate a speaker, and execute voice-protected commands for highly secured systems on their behalf. To the best of our knowledge, our work is the first one extending MI attacks to audio data, and our results highlight the security risks resulting from the extraction of the biometric data in that setup.

**Index Terms:** speaker recognition, model inversion, privacy

## 1. Introduction

Privacy analysis of audio data has shown that speech parameters, such as accent, rhythm, or acoustic properties of speech inherently carry biometric information about the speakers, such as their age, gender, physical health, and geographical origin [1]. Therefore, it is important for machine learning (ML) models in speaker recognition not to leak information about their training data. However, recent research [2, 3, 4] suggests that ML models are, in general, vulnerable to privacy attacks. One particular attack is *model inversion* (MI) [2] which allows an attacker to retrieve abstract representations for individual classes of the target model's training data. With speaker recognition systems treating each individual as their own class, MI attacks have the potential to cause severe privacy breaches [2]. So far, the feasibility of MI attacks on speaker recognition systems and audio data has never been tested, thus, the question if information on the speakers can be maliciously retrieved remained open.

We are the first to show how to adapt and apply MI attacks for audio data. We do so by targeting SincNet, a state-of-the-art neural network (NN)-based speaker recognition model [5, 6]. We show that MI attacks are able to infer both entire audio samples and d-vectors as intermediate representations of the speakers' voice characteristics from the trained target model. Further, we propose the *sliding model inversion*, a novel form of the standard MI attack that leverages sequential processing properties of the audio data to improve inversion success. While with standard MI, the target model successfully identifies up

to 54% of the inverted audio samples as their correct speaker class, with our novel attack, we achieve up to 90% accuracy. Also, our sliding MI manages to decrease the distance between original and inverted sample in the d-vector representation, hence, yielding higher fidelity inversions. For directly inverting d-vectors, our experiments show that even standard MI achieves 100% identification success. These results highlight the vulnerability of speaker recognition models to privacy attacks. As a proof-of-concept to showcase that our MI can be exploited as a departure point for further attacks against speaker recognition, we, furthermore, explore using inverted audio samples as inputs for deepfake generation. Such deepfakes could be used to fool voice identification with arbitrary speech samples or to execute any speech command on behalf of the speakers under attack. While our generated deepfakes do not perfectly fool a human listener, as an informal evaluation conducted by the authors shows, they illustrate that privacy attacks can not only be used to disclose sensitive information about the individuals the model was trained on; additionally, they can severely threaten the security of systems relying on voice biometrics. Our contributions can be summarized as follows:

- We successfully apply MI attacks on speaker recognition models to invert entire audio samples and d-vectors and experimentally evaluate what kind of random initialization works best as an input for MI attacks on audio data.
- We introduce a novel *sliding MI* which exploits properties of sequential and chunk-wise audio processing.
- We show the feasibility of generating deepfakes based on the inferred audio samples.

## 2. Background and Related Work

The following section provides background information on speaker recognition systems and attacks against their privacy.

**Speaker recognition.** In this paper, we use a SincNet-based [5] text-independent speaker recognition system. This system uses NNs to extract voice features into so-called d-vectors and adds a classification layer on top of these. The input to the system consists of raw audio waves and the outputs is a per-class probability score over all possible classes (i.e., speakers). Overall, the system is composed of three submodels (see Figure 1) 1) SincNet, a convolutional NN that resembles a band-pass filter; 2) a multi-layer perceptron (MLP) calculating the d-vectors [7]; and 3) a fully connected layer to calculate the probabilities per speaker. It achieves a reported classification error<sup>1</sup> of  $5.772 \cdot 10^{-3}$  on the TIMIT [8] test data set (measured at sentence level). In its current version, SincNet does not provide any dedicated privacy-preserving mechanisms.

\*Authors contributed equally.

<sup>1</sup>See <https://pythonlang.dev/repo/mravanelli-sincnet/>.




Figure 1: **Speaker Recognition Model.** The speaker recognition model and its three submodels: SincNet obtains ① raw audio input and generates ② features. These features are input to an MLP which generates the ③  $d$ -vectors. A single layer performs ④ classification on them. In experiment 1, we invert full audio samples, while in experiment 2, we invert the  $d$ -vectors.

---

**Algorithm 1** Standard Model Inversion Attack [2]

---

```

function MI(input vector  $x_0$ , target class  $t$ , iterations  $\alpha$ , patience  $\beta$ , minimum cost threshold  $\gamma$ , learning rate for gradient descent  $\lambda$ )
  for  $i \leftarrow 1, \dots, \alpha$  do
     $x_i \leftarrow x_{i-1} - \lambda \cdot \nabla c(x_{i-1})$ 
    if  $c(x_i) \geq \max(c(x_{i-1}), \dots, c(x_{i-\beta}))$  then
      break
    end if
    if  $c(x_i) \leq \gamma$  then
      break
    end if
  end for
  return  $[\operatorname{argmin}_{x_i} c(x_i), \min_{x_i} c(x_i)]$ 

```

---

**Privacy in speaker recognition.** The ISO/IEC norm 24745:2011 proposes three general requirements to ensure individual privacy: irreversibility, renewability, and unlinkability of the protected data [9]. To reach this goal in speaker recognition systems, several solutions have been proposed and discussed [9]. However, aiming for any *anonymization* to protect the privacy of speakers in a speaker recognition model goes directly against the purpose of the system, which is to identify speakers based on their individual characteristics.

**MI attacks.** In MI attacks [2], the attacker exploits an ML model’s prediction confidence for inverting individual training classes (see Algorithm 1). More formally speaking, a MI attack can be expressed as follows: let  $f$  be the target model under attack. It is trained to map from an  $n$ -dimensional input data point  $x$  to an  $m$ -dimensional vector  $p$  indicating the probability per class, such that  $f : x \mapsto p$ , with  $\mathbb{R}^n \rightarrow [0, 1]^m$ . To invert the model, we define an objective function in order to use gradient descent. This function is called cost function  $c(x)$  and basically defines how close we are to the information we would like to reconstruct. We set  $c(x) = 1 - p_t$ , where  $t$  denotes the target class we would like to gain information about. Starting from a randomly initialized input sample  $x_0$ , we calculate its cost  $c(x_0)$ . With this at hand, we apply the gradient descent algorithm for  $\alpha$  iterations with a learning rate  $\lambda$  to alter the original input. The aim is to minimize the costs for a specific class [2], such that the resulting data sample is a representation of that class. In speaker recognition, every speaker denotes their own class. Hence, MI can reveal representations of the data from every single individual in the training data set. Particularly, this data can encode biometric as well as paralinguistic features.

---

**Algorithm 2** Sliding Model Inversion Attack

---

```

function SMI(target class  $t$ , length  $l$ , stride  $s$ , window size  $w$ ,  $\alpha, \beta, \gamma, \lambda$  as in Algorithm 1)
   $\text{inverted}[0, \dots, l] \leftarrow [\mathcal{N}(\mu, \sigma^2)]^l$ 
  for  $k \leftarrow 0, \dots, (l - s)$  according to stride  $s$  do
     $x \leftarrow \text{inverted}[k : (k + w)]$ 
    for  $i \leftarrow 1, \dots, \alpha$  do
       $x_i \leftarrow x_{i-1} - \lambda \cdot \nabla c(x_{i-1})$ 
      if  $c(x_i) \geq \max(c(x_{i-1}), \dots, c(x_{i-\beta}))$  then
        break
      end if
      if  $c(x_i) \leq \gamma$  then
        break
      end if
    end for
     $\text{inverted}[k : (k + w)] \leftarrow \operatorname{argmin}_x c(x)$ 
  end for
  return  $\text{inverted}[\frac{w}{2} : (l - \frac{w}{2})]$ 

```

---

### 3. Sliding MI Attack

In this section, we present our novel sliding MI attack. It extends standard MI (see Algorithm 1) to sequentially and chunkwise processed data, e.g., audio data. Instead of using MI to invert every chunk of data separately, our sliding ML iteratively inverts overlapping chunks. This way, some of the input to the MI is already inverted, and hence, in wave-form similar to actual speech data. Thereby, MI can more successfully invert it into representatives of the original speech data.

Our sliding MI consists of the following steps: (1) We invert the first window of a randomly initialized input vector. Note that different random initializations yield inversions of different quality. We experiment with several different types of initialization settings as our first main experiment, described in more detail in Section 4.1. (2) Then, we replace the first window of the input vector by the resulting inverted data. (3) Next, we iteratively calculate the inverted data for the subsequent input window. This input’s first part consists of the previously inverted data, its second part stems from the random input vector. Note that the amount of overlap for the sliding window determines the proportion of the previously inverted data and the randomly initialized input vector that are used for calculating the inversion during the MI attack. Its value depends on the *stride* of our inversion. For our experiments, we use a stride of 500 samples (roughly 30ms). Since in this new method, updates rely on the output of the previous inversion, we cannot use parallelization as a speed-up. Instead, the stride determines the computational overhead. By increasing the stride value, computational time can be decreased.

For a visualisation of our novel approach, see Figure 2. Note that in addition to the hyperparameters of the standard MI, we need to specify the length of our input data  $l$ , the stride  $s$ , and the window size  $w$ . While  $s$  specifies the overlap between subsequent inversions,  $w$  determines the lengths of the data chunks that are inverted. For each chunk, inversion is performed as an iterative process as in the standard MI. Since the beginning and the end of the inverted vector are iterated less, we cut the returned vector to half the window size. See Algorithm 2 for a formal introduction of our novel sliding ML.




Figure 2: **Sliding MI.** ① We initialize a random input vector. ② Starting from the beginning, we invert the first window based on this vector and replace the vector’s first part by our inverted data. ③ For the all subsequent windows, we use parts of the previously inverted vector and fill the remainder with the input vector to apply MI. ④ We then iteratively replace the input vector with our inverted data.

## 4. Experiments

We conduct three experiments: The *first experiment* is similar to MI in other domains, i.e., it inverts random input vectors back to the original input data domain. In the *second experiment* we do not invert the whole NN, but only the layers up to the d-vectors, which provide unique voice features of an individual (see Figure 1). Inverting these vectors instead of full audio samples reduces the computational costs of the attack. Our *third experiment* shows that in speaker recognition, our MI attack enables us to impersonate individual speakers, and to synthesize speech samples for them. The spoofing is performed based on the inverted audio samples from the experiment one.

We attack the NN-based speaker recognition system using SincNet [5], trained on the TIMIT dataset [8], and we use the pretrained model provided by Ravanelli et al. [5]. To perform our MI attack, we assume the attacker to have white-box access to the target model. This is the case, for example, when the speaker recognition model is deployed to a user-device, e.g. for biometric identification. Further, the attacker needs a unique identifier of their target individual under attack to know which class of the training data to invert. This can be, for example, the name or some pseudonymized combination of characters, in the case of the TIMIT dataset, e.g. “FGMB0”.

We quantify the success of our experiments as follows:

- *Percentage of correctly classified inverted samples.* We quantify the classification accuracy of the original target speaker model on both the inverted audio data and the inverted d-vectors. An inverted sample is “correctly classified” if it is classified as the correct original speaker.
- *Euclidean distance between original and reconstructed d-vector.* We measure the Euclidean distance between both d-vectors to specify the similarity between the respective samples.

The first metric allows us to analyse if the MI may be considered successful with respect to the target model, i.e., it answers the question of how successfully this model can be fooled. The second metric, in contrast, focuses on the inverted samples’ similarity to the original samples. Hence, it quantifies the similarity from the perspective of a human listener.

Since MI generates average representations of training classes, we use the target model’s classification accuracy on averaged per-speaker samples as a baseline (97.84% and 75.97% of correct classification on train and test data, respectively). In the following, we present our overall experimental setup to then describe every experiment in more detail.

### 4.1. Experiment 1: Invert Audio Samples

In the first experiment, we use MI to calculate full inverted audio samples which could be used to trick the speaker recognition model under attack without human listeners present. The experiment is designed to answer the following three questions: (1) Is it possible to successfully generate inverted audio samples for speaker recognition? (2) Which kind of randomly initialized input vector to the MI attack produces the most successful inverted audio samples (with respect to the classification as the original speaker)? (3) How does our new sliding MI approach improve the results in comparison to standard MI?

**Experiment.** Over all experiments, the audio data chunks are 3200 samples (or 200ms) long, and our sliding MI uses a stride of 500 samples (roughly 30ms). We evaluate the following (random) initializations:

- *Plain inputs:* all zeros, all ones, and all minus ones;
- *Noises:* white, pink, brown, violet, and blue noise. We generated them with methods pre-implemented in the `python-acoustics` library and applied the tanh-function to transform them to the interval range  $[-1, 1]$  with zero mean in a non-linear manner;
- *Samplings from distributions,* such as uniform (ranges  $[0, 1]$  or  $[-1, 1]$ ), Gaussian (with  $\mu = 0$  or  $\sigma = 0.2$ ), Laplace (with  $\mu = 0$  or  $b = 0.07$ ), Gumbel (with  $\mu = 0$  or  $\beta = 0.1$ ), and von Mises (with  $\mu = 0$  or  $\kappa = 0.1$ );
- *Samples from another dataset:* Librispeech [10], used as a plain input or averaged over 50, 100, 150 input vectors or with white noise ( $0.85 \cdot \text{input vector} + 0.15 \cdot \text{noise}$ ).

For the optimization process within the MI algorithm (see Section 2), we optimize two parameters, namely the maximum number of iterations  $\alpha$  and the learning rate  $l$ . For all experiments, setting  $\alpha = 1000$  showed to be sufficient. The optimal  $l$  depends on the experiment and is reported in the results. In our evaluation, for every input with its settings and optimal  $l$ , we report the following metrics:

- *MI Accuracy:* the percentage of correctly classified inverted samples;
- *# Correct Speakers:* the number of correctly classified speakers;
- *Avg. Eucl. Distance with Std. Deviation:* the average Euclidean distance of the inverted sample to original samples of the same speaker in the d-vector space, calculated on the successfully inverted audio samples.

For the Euclidean distances within the d-vector space, the average within-speaker distance in the original training samples can serve as a baseline ( $1.923 \cdot 10^{-1}$ ).

**Results.** (1) *Is it possible to successfully generate inverted audio samples for speaker recognition?* By looking at the distribution of the inverted samples, we observe that it does not fully match the distribution of the original data. Since the inverted samples also sound differently from the original ones— they do not necessary sound like speech—hence, they do not allow to fool a *human* listener. However, the results suggest that standard MI is good enough to fool the classification of the *automatic* speaker recognition system with an accuracy of up to 54.76%. The classification accuracy can be significantly improved to 90.48% through our new sliding MI. For speaker

recognition models in charge of identity control for a highly secured system, this accuracy on inverted data would be beyond acceptable. Our novel sliding MI, also reduces the Euclidean distance between inverted and original samples for some input vectors, in comparison to a standard MI. However, despite this decrease in distance, the reconstructed speech samples still do not sound very close to the original speaker.

(2) *Which kind of randomly initialized input vector to the MI attack produces the most successful inverted audio samples (with respect to the classification as the original speaker)?* We can also conclude that not all input vectors to the MI are equally suited to create inverted audio samples which successfully fool the speaker model: plain input vectors achieve the lowest quality in inversion with respect to classification accuracy. We assume that this is due to the difficulty of transforming constant vectors into speech-like wave forms through optimization. It seems that random initialization or data that is already in wave form are more suited inputs to MI for audio data: The best classification accuracy can be achieved with white noise and tanh activation. Brown noise exhibits the poorest performance, yet it exhibits a relatively small mean Euclidean distance. The Laplace distributions achieves the overall highest results. See Table 1 for an overview of results.

(3) *How does our new sliding MI approach improve the results in comparison to standard MI?* We observe from the results in Table 1 that sliding MI exhibits a higher performance than standard MI. While with standard MI, the accuracy of the target model on the inverted data is 54%, depending on the random initialization, our sliding MI yields above 90% accuracy.

#### 4.2. Experiment 2: Partial MI to Invert d-Vectors

While in previous applications on other data types, only a complete MI back to the original input domain is valuable, this is different for speaker recognition: the d-vectors, which are feature representations of the voice samples, already carry important paralinguistic information that can, for example, be used to generate spoofed audio samples [11]. Inverting simple d-vectors instead of full audio samples reduces the computational costs of the attack (since it does not need to be performed sequentially), and can be performed with standard MI attacks. Therefore, in our second experiment we set out to invert a model on the intermediate layers with the aim to answer the following two questions: (1) Is it possible to successfully invert d-vector? (2) Which input vector produces the most successful inverted d-vector (with respect to the classification as the original speaker)?

**Experiment.** We apply partial inversion by removing the SincNet and MLP part of the network and only focusing on the submodel 3 for the d-vector inversion (see Figure 1).

**Results.** (1) *Is it possible to successfully invert d-vector?* Our findings show that we can reconstruct d-vectors that are successfully classified as the original speaker (classification accuracy of the target model on them reaches 100%). This outperforms the baseline where we measure accuracy of the target model on averaged per-speaker samples (97.84%). However, even for the best-performing input (zeros), the Euclidean distance of  $0.84 \pm 0.028$  is clearly above the baseline average within speaker distance ( $1.923 \cdot 10^{-1}$ ). To evaluate whether the inverted d-vectors still leak the individual speakers’ privacy, we perform a principal component analysis (PCA) and train a binary classifier to predict the individuals’ gender. We fit the PCA to the TIMIT




Figure 3: *PCA on d-Vectors.* PCA fitted on the TIMIT test dataset (blue: female; green: male) and used to transform the inverted d-vectors (purple: female; red: male). Results indicate that inverted d-vectors reveals the individuals’ gender.

test dataset and transform the inverted d-vectors. Our results are visualized in Figure 3. They suggest that the inverted d-vector leak gender privacy.

(2) *Which input vector produces the most successful inverted d-vector (with respect to the classification as the original speaker)?* Since d-vectors and audio data have different properties, they require different input vectors for successful inversion. Sound noises are good inputs for audio samples. However, our observations suggest that initializing the input vector with sound noises does not yield high-quality inverted d-vector representations. Instead, we found plain zeros perform best when inverting d-vectors.

#### 4.3. Experiment 3: Create Deepfakes

Even though the inverted samples are classified correctly by the target model, they do not necessarily carry useful information for human listeners. With the following experiments, we focus on the question: (1) Based on the inverted audio samples, is it possible to generate audio data that resembles the original speaker for a human listener? The experiment can be considered as a proof-of-concept to demonstrate further security risks in speaker recognition systems made possible by our attack.

**Experiment.** To generate the deepfakes, we use the work from [11]. Their architecture consist of three parts: (1) speaker encoder, (2) speech synthesizer, and (3) vocoder. The speaker encoder is used to create a d-vector out of an audio file, which characterizes the audio sample in vector space. With this information, the speech synthesizer creates the mel-spectrogram by using the d-vectors. Finally, vocoder grabs the mel-spectrogram to perform frequency to time domain conversion of. The underlying speech synthesizer is Tacotron 2 [12] and vocoder is Wavenet [13]. We use the inverted audio samples from our sliding MI, and the inverted d-vectors as input for the method. In principle, inverted audio samples can be fed directly into the speaker encoder for the deepfake generation. However, to use our inverted d-vectors (2048 dimensions), we have to transform them to match the deepfake model’s speaker encoding, as it expects d-vectors with 256 dimensions. To do so, we train an MLP to map one vector space to another. The MLP has two hidden

	Sample Type	Learning R.	MI Accuracy	# Correct Speakers	Avg. E. Dist. $\pm$ Std. Dev.
plain inputs	ones	$1 \cdot 10^{-05}$	0.43%	2	$0.800 \pm 0.0416$
		0.2	5.41%	25	$0.696 \pm 0.0663$
	zeros	$1 \cdot 10^{-08}$	0.65%	3	$0.792 \pm 0.0377$
0.5		8.87%	41	$0.720 \pm 0.0682$	
dists	Gumbel	0.01	54.76%	253	$0.797 \pm 0.0562$
		0.01	89.83%	415	$0.748 \pm 0.0561$
	Laplace	0.005	54.33%	251	$0.793 \pm 0.0555$
		0.005	90.48%	418	$0.752 \pm 0.0550$
noise	white-tanh	0.2	54.11%	250	$0.798 \pm 0.0549$
		0.2	88.31%	408	$0.757 \pm 0.0561$
	brown	0.01	5.84%	27	$0.676 \pm 0.0938$
		0.05	31.82%	147	$0.681 \pm 0.0679$
Librispeech samples	Librispeech sample	0.001	6.277%	29	$0.751 \pm 0.0768$
		0.2	22.73%	105	$0.696 \pm 0.0558$
	Sample noise + Librispeech sample	0.005	52.6%	243	$0.779 \pm 0.0673$
		0.01	90.26%	417	$0.757 \pm 0.0595$
	Librispeech mean	0.001	33.55%	155	$0.750 \pm 0.0665$
		0.01	56.06%	259	$0.707 \pm 0.0713$
	Sample noise + Libspeech mean	0.005	45.89%	212	$0.776 \pm 0.0604$
		0.01	80.30%	371	$0.757 \pm 0.0600$

Table 1: **Results for standard MI (gray) and sliding MI (black)** calculated for inverted audio samples of the 462 speakers in the TIMIT dataset (326 men, 136 women). Sliding MI improves standard MI (higher accuracy and lower average Euclidean distance.)

layers with 1024 and 512 neurons and an output layer with 256 neurons. We use tanh activation in the first two layers. We create the training set for this transformation by feeding sound files to our and the deepfake’s speaker encoder. Our encodings are used as input while their encodings are treated as the outputs to be learned.

**Results.** (1) *Based on the inverted audio samples, is it possible to generate audio data that resembles the original speaker for a human listener?* As reported in experiment 2, our d-vectors, though correctly classified as the original speakers, were far away from the speaker’s original d-vectors in the Euclidean space. The question if a generated sample sounds similar to an original one is a semantic question and depends on the sensitivity of the context. From the authors’ perspective on inspection case-by-case, the d-vector-based deepfakes did not allow individual speakers characteristics to be recognized. However, based on the inverted audio samples from our novel sliding MI, we were able to generate a few good quality spoofed audio samples that resembled the original speaker.<sup>2</sup> With such samples at hand, an attacker could, hence, spoof someone’s identity solely based the inverted data from the pre-trained NN. We expect this to become even much more prevalent with more sophisticated deepfake generation systems in the future.

## 5. Countermeasures and Discussion

Speaker recognition systems heavily rely on learning individual per-speaker characteristics in order to fulfill the task they are designed for. Therefore, these systems always and necessarily contain information about the speaker data that they were trained on. Noising out individual speaker characteristics will result in drastically decreased performance of the systems. In particular, pseudonomization [14] and privacy methods that are

<sup>2</sup>Examples for the original audio data, averaged and inverted samples, and the spoofed audio data generated based on the inverted audio samples, are available at <https://www.dropbox.com/sh/ge6xx901aqmru9b/\AABUs3p4EwaN0n7g4Rgq8rwa>.

used in general speech systems (e.g., [15, 16, 17, 18, 19, 20]) render speaker recognition unusable for their original purpose.

**Protecting against MI attacks.** As an alternative, one can consider privacy protection methods that aim at impeding MI attacks. Most existing defenses from other domains focus on suppressing the model confidence score or reducing their utility. This can be done by injecting uniform noise to them [21], reducing their precision [2] or their dispersion [22]. The latter one leads to a decrease in the correlation between the input data and the scores, which renders MI attacks more inaccurate. With a similar aim, the use of regularization in the training loss function has been reported as a defense [23]. Additionally, hardware-oriented solutions to prevent an attacker from accessing the model parameters to decrease MI success, or at least preventing the extraction of intermediate features (see our experiment 2) can be applied [24].

**Differential privacy.** Initial work empirically showed that Differential Privacy (DP) [25] can reduce the success of MI attack’s [2] when using a very large amount of noise, which, in return, drastically degrades the model’s performance. Later work suggests that DP training for ML models cannot at all prevent MI attacks [26] because its aim is to dissimulate the presence of a data point in a specific data set and not to protect privacy over classes of data.

**Limitations.** So far, MI attacks require the availability of an NN’s confidence scores, and the attack’s success depends largely on the random initializations. Especially, many speaker recognition tools depend on the cosine similarity [27, 28] for which the algorithm would need to be updated. Also, the quality of the spoofed audio samples is limited by the deepfake creation methods. As a consequence, the practical impact of our attack might currently still be limited. However, with new and ever more powerful privacy attacks and deepfake methods being proposed, the threat space of exploiting privacy attacks to

violate security of speaker recognition systems will gain importance. It is, hence, important to create awareness and to consider and protect privacy and security jointly, rather than separately.

## 6. Conclusion

In this work, for the first time, we successfully perform MI attacks on audio data. Therefore, we introduce a novel sliding MI method which leverages the sequential properties of the audio data for improved inversion. We experimentally evaluate the attack's success on a state-of-the-art speaker recognition system. Our results indicate that our inverted audio samples can be used as a departure point for further attacks against the security of the target system. Thereby, we highlight the importance of implementing adequate privacy protection in such systems.

## 7. Acknowledgements

This research was supported by the Bavarian Ministry of Economic Affairs, Regional Development and Energy. The authors acknowledge J. Williams for comments on the manuscript.

## 8. References

- [1] J. L. Kröger, O. H.-M. Lutz, and P. Raschke, "Privacy implications of voice and speech analysis—information disclosure by inference," in *IFIP International Summer School on Privacy and Identity Management*. Springer, 2019, pp. 242–258.
- [2] M. Fredrikson, S. Jha, and T. Ristenpart, "Model inversion attacks that exploit confidence information and basic countermeasures," in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, 2015, pp. 1322–1333.
- [3] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, "Membership inference attacks against machine learning models," in *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2017, pp. 3–18.
- [4] K. Ganju, Q. Wang, W. Yang, C. A. Gunter, and N. Borisov, "Property inference attacks on fully connected neural networks using permutation invariant representations," in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, 2018, pp. 619–633.
- [5] M. Ravanelli and Y. Bengio, "Speaker recognition from raw waveform with sinnet," in *2018 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2018, pp. 1021–1028.
- [6] M. Ravanelli, T. Parcollet, A. Rouhe, P. Plantinga, E. Rastorgueva, L. Lugosch, N. Dawalatabad, C. Ju-Chieh, A. Heba, F. Grondin, W. Aris, C.-F. Liao, S. Cornell, S.-L. Yeh, H. Na, Y. Gao, S.-W. Fu, C. Subakan, R. De Mori, and Y. Bengio, "Speechbrain," <https://github.com/speechbrain/speechbrain>, 2021.
- [7] E. Variani, X. Lei, E. McDermott, I. L. Moreno, and J. Gonzalez-Dominguez, "Deep neural networks for small footprint text-dependent speaker verification," in *2014 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2014, pp. 4052–4056.
- [8] J. S. Garofolo, "Timit acoustic phonetic continuous speech corpus," *Linguistic Data Consortium*, 1993, 1993.
- [9] A. Nautsch, A. Jiménez, A. Treiber, J. Kolberg, C. Jasserand, E. Kindt, H. Delgado, M. Todisco, M. A. Hmani, A. Mtibaa *et al.*, "Preserving privacy in speaker and speech characterisation," *Computer Speech & Language*, vol. 58, pp. 441–480, 2019.
- [10] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [11] Y. Jia, Y. Zhang, R. J. Weiss, Q. Wang, J. Shen, F. Ren, Z. Chen, P. Nguyen, R. Pang, I. L. Moreno *et al.*, "Transfer learning from speaker verification to multispeaker text-to-speech synthesis," *arXiv preprint arXiv:1806.04558*, 2018.
- [12] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerrv-Ryan *et al.*, "Natural tts synthesis by conditioning wavenet on mel spectrogram predictions," in *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2018, pp. 4779–4783.
- [13] A. v. d. Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, "Wavenet: A generative model for raw audio," *arXiv preprint arXiv:1609.03499*, 2016.
- [14] P.-G. Noé, A. Nautsch, N. Evans, J. Patino, J.-F. Bonastre, N. Tomashenko, and D. Matrouf, "Towards a unified assessment framework of speech pseudonymisation," *Computer Speech & Language*, vol. 72, p. 101299, 2022.
- [15] R. Aloufi, H. Haddadi, and D. Boyle, "A tandem framework balancing privacy and security for voice user interfaces," *arXiv preprint arXiv:2107.10045*, 2021.
- [16] Q. Jin, A. R. Toth, T. Schultz, and A. W. Black, "Voice convergin: Speaker de-identification by voice transformation," in *2009 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2009, pp. 3909–3912.
- [17] B. M. L. Srivastava, N. Vauquier, M. Sahidullah, A. Bellet, M. Tommasi, and E. Vincent, "Evaluating voice conversion-based privacy protection against informed attackers," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 2802–2806.
- [18] B. M. L. Srivastava, N. Tomashenko, X. Wang, E. Vincent, J. Yamagishi, M. Maouche, A. Bellet, and M. Tommasi, "Design choices for x-vector based speaker anonymization," *arXiv preprint arXiv:2005.08601*, 2020.
- [19] J. Qian, H. Du, J. Hou, L. Chen, T. Jung, and X.-Y. Li, "Hide-behind: Enjoy voice input with voiceprint unclonability and anonymity," in *Proceedings of the 16th ACM Conference on Embedded Networked Sensor Systems*, 2018, pp. 82–94.
- [20] B. M. L. Srivastava, A. Bellet, M. Tommasi, and E. Vincent, "Privacy-preserving adversarial representation learning in asr: Reality or illusion?" *arXiv preprint arXiv:1911.04913*, 2019.
- [21] A. Salem, A. Bhattacharya, M. Backes, M. Fritz, and Y. Zhang, "Updates-leak: Data set inference and reconstruction attacks in online learning," in *29th {USENIX} Security Symposium ({USENIX} Security 20)*, 2020, pp. 1291–1308.
- [22] Z. Yang, B. Shao, B. Xuan, E.-C. Chang, and F. Zhang, "Defending model inversion and membership inference attacks via prediction purification," *arXiv preprint arXiv:2005.03915*, 2020.
- [23] T. Wang, Y. Zhang, and R. Jia, "Improving robustness to model inversion attacks via mutual information regularization," *arXiv preprint arXiv:2009.05241*, 2020.
- [24] Q. Xu, M. T. Arafain, and G. Qu, "Midas: Model inversion defenses using an approximate memory system," in *2020 Asian Hardware Oriented Security and Trust Symposium (AsianHOST)*. IEEE, 2020, pp. 1–4.
- [25] C. Dwork, "Differential privacy," in *International Colloquium on Automata, Languages, and Programming*. Springer, 2006, pp. 1–12.
- [26] Y. Zhang, R. Jia, H. Pei, W. Wang, B. Li, and D. Song, "The secret revealer: Generative model-inversion attacks against deep neural networks," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 253–261.
- [27] Y. Zhang, Z. Lv, H. Wu, S. Zhang, P. Hu, Z. Wu, H.-y. Lee, and H. Meng, "Mfa-conformer: Multi-scale feature aggregation conformer for automatic speaker verification," *arXiv preprint arXiv:2203.15249*, 2022.
- [28] B. Desplanques, J. Thienpondt, and K. Demuynck, "Ecapatdnn: Emphasized channel attention, propagation and aggregation in tdnn based speaker verification," *arXiv preprint arXiv:2005.07143*, 2020.