Evaluating X-vector-based Speaker Anonymization under White-box Assessment

Pierre Champion^{1,2}, Denis Jovet¹, and Anthony Larcher²

¹ Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France
{`pierre.champion`, `denis.jovet`}@inria.fr

² Le Mans Université, LIUM, France
`anthony.larcher@univ-lemans.fr`

Abstract. In the scenario of the Voice Privacy challenge, anonymization is achieved by converting all utterances from a source speaker to match the same target identity; this identity being randomly selected. In this context, an attacker with maximum knowledge about the anonymization system can not infer the target identity. This article proposed to constrain the target selection to a specific identity, i.e., removing the random selection of identity, to evaluate the extreme threat under a white-box assessment (the attacker has complete knowledge about the system). Targeting a unique identity also allows us to investigate whether some target’s identities are better than others to anonymize a given speaker.

Keywords: Speaker anonymization · VoicePrivacy · Anonymization evaluation.

1 Introduction

In many applications, such as virtual assistants, speech signal is sent from the user device to the service provider’s servers in which data is collected, processed, and stored. Recent regulations, e.g., the General Data Protection Regulation (GDPR) [10] in the EU, emphasize on privacy preservation and protection of personal data. As speech data can reflect both biological and behavioral characteristics of the speaker, it is qualified as personal data [8]. The research reported in this article has been done using the Voice privacy Challenge framework [17], which is one of the first attempts of the speech community to evaluate research on this topic, by producing dedicated protocols, metrics, datasets and baselines.

Speaker anonymization is performed to suppress the personally identifiable paralinguistic information from a speech utterance while maintaining the linguistic content. This is also referred to as *speaker anonymization* [3] or *de-identification* [5]. Recently, Fang et al. [3] proposed an x-vector-based speaker anonymization system based on voice conversion where one of the hyper-parameters used to transform the voice is a target pseudo-speaker identity. In order to choose the anonymization system and hyper-parameters, for a given use case, we must evaluate and rank the performances of each anonymization method. In this first edition of the Voice privacy Challenge, the quality of anonymization is assessed

by using state-of-the-art speaker verification system together with an automatic speech recognition system that is used to evaluate the preservation of the linguistic content. In the context of privacy, a game-theoretic reasoning with two agents is defined as follows: one who wishes to anonymize a user’s speech (the service provider) and the other (the attacker) who will do everything possible to “break” this anonymity. Introduced in [15], the notion of attacker’s knowledge of the anonymization method was introduced, it defines three possible attack scenarios: black-box, white-box, and grey-box. In all approaches, privacy is evaluated by comparing original speech (accessible to the attacker) and anonymous speech published by a service provider. The goal of an effective anonymization system is to allow the publication of speech whose identity is difficult to link to a user’s identity, even if the attacker has some knowledge of the anonymization mechanism. In the black-box attacker scenario, privacy is measured by comparing clean, non-anonymized enrollment speech, and anonymized trial speech. In this scenario, the attacker is not aware that speech has been anonymized. In contrast, a white-box attacker is fully aware of the anonymization system and hyper-parameters. Grey-box attackers cover the whole range of possible scenario in-between black and white boxes, when the attacker has a partial knowledge of the anonymization system.

The target pseudo-speaker identity used to anonymize the voices can be generated following multiple strategies [15]. In the *permanent* strategy, used in the VoicePrivacy challenge, all utterances from a speaker are converted using the same target pseudo-speaker. The challenge defines a grey-box attacker that has access to the anonymization toolkit, and thus is able to transform any utterance from her enrollment dataset using the same system and target strategy. However, for a given source speaker, the target pseudo-speaker used by the attacker differs from the one used by the service provider due to the random pseudo-speaker selection process included in the *permanent* strategy. In case the attacker knows the pseudo-speaker, this grey-box scenario turns into a white-box scenario.

We modify the Voice Privacy Challenge scenario into a white-box assessment by targeting a same speaker identity for all utterances from all speakers. This target selection strategy allows us to evaluate the performance of the x-vector-based anonymization system on its own, and generate a report that does not assume that the attacker has knowledge deficiencies. Experiments are performed with a large group of target speaker identity in order to investigate the effect of this identity on the quality of the anonymization.

In the remaining of this article, we first describe the baseline system and voice conversion method in Section 2. We then introduce our experimental protocol in Section 3 and present our experimental results in Section 4. Eventually, we draw our conclusions and propose future avenues in Section 5.

2 Anonymization Technique

The anonymization of speaker’s identity can be performed with various methods [1, 3, 7, 12, 15]. In this article, our contributions are based on the *Baseline-1*

(referred to as the *baseline* in this article) of the VoicePrivacy challenge that anonymizes speech using x-vectors and neural waveform models [3].

2.1 The Voice Conversion System

Fig. 1. The speaker anonymization pipeline. Modules A, B, and C are parts of the baseline model. Module D is an enhancement used to coherently alter the F0 values with respect to the selected pseudo-speaker identity.

The baseline system introduced in [3] aims at separating speaker identity and linguistic content from an input speech utterance. Assuming that those features are disentangled, an anonymized speech waveform can be generated after altering only the features that encode the speaker’s identity. The anonymization system illustrated in Figure 1 can be decomposed into three groups of modules. Modules from the *group A* extract different features from the source signal: the fundamental frequency, the phonetic features encoding articulation of speech sounds (Phoneme Posterior-Grams (PPGs) [16]) and the speaker’s x-vector. *The module B* derives a new pseudo-speaker identity. The x-vector from each source input speaker is compared to a pool of external x-vectors to select the 200 furthest vectors; 100 of them are randomly selected and averaged to create an anonymized pseudo-speaker x-vector identity. Finally, *the module C* synthesizes a speech waveform from the pseudo-speaker x-vector together with the original PPGs features and F0.

As an enhancement to this baseline, [2] proposed to modify the F0 values (cf. module D in Figure 1) using the F0 mean and variance statistics associated with each of the x-vector of the pool.

2.2 Design Choices for Anonymization

In the evaluation plan of the VoicePrivacy Challenge, all utterances of a given speaker should be converted to match a single target pseudo-speaker. This strategy, described as *permanent* in [15], ensures that a one-to-one mapping exists between the source speaker identity and the anonymized speaker identity. This

requirement allows anonymized voices to be distinguishable from each other. This one-to-one mapping does not apply in-between speech anonymized by the service provider and speech anonymized by the attacker. Speech data anonymized by the service provider corresponds to the trial dataset, and speech available to the attacker corresponds to the enrollment dataset, see Section 3.1 for more detail. Figure 2a shows how the *permanent* strategy converts enrollment and trial speech to different anonymized pseudo-speakers.

To select the target pseudo-speaker identity, module B (from Figure 1) has many hyper-parameters that affect the selection mechanism. According to [2, 13, 14], the best anonymization results are achieved by picking the pseudo-speaker in a dense region of the x-vector space, randomly targeting male or female gender pseudo-speaker, and modifying the F0 values of the input speech so that it matches the F0 statistics of the real speakers used to generate the pseudo-identity.

The VoicePrivacy challenge’s protocol assumes that attackers have access to anonymized trial utterances and original enrollment utterances. During the challenge, two sets of tests are performed following black-box and grey-box attacker scenarios corresponding to situations where the enrollment utterances are original or have been anonymized. In the latter scenario, the attacker has partial knowledge of the system and is able to anonymize the enrollment utterances using the same anonymization system and *permanent* target selection strategy. The pseudo-speaker chosen for each of the enrollment speakers differs from the pseudo-speaker chosen for the trial speakers as the attacker does not have knowledge of the randomly selected speakers used to generate the pseudo-speaker identity. As the voices of the same speakers once anonymized by the

Fig. 2. Overview of target selection strategies. In the *permanent* strategy pseudo-speakers are randomly selected, first for the trial dataset and then for the enrollment dataset. While our *constant* strategy, always targets the same pseudo-speaker.

service provider (trial) and the attacker (enrollment) are different from each other, this leads to a rather good anonymization performance [17].

With the *permanent* target selection strategy used in the challenge, anonymized voices remain distinguishable and all utterances from the same original speaker are anonymized with the same pseudo-speaker. This process is referred to as pseudonymization. Providing attackers the target’s identity of each speaker in the *permanent* strategy evaluates the quality of the voice conversion system and the preservation of the voice distinctiveness [9]. In this paper, we evaluate the privacy of the anonymization technique with the best attacker, that is using the *constant* selection strategy.

3 Experimental Setup

In this evaluation, we want to provide a comprehensive assessment of the voice conversion toolkit under a white-box scenario where the attacker has full knowledge about the system. We change the game between the attacker and the service provider to use the *constant* identity selection strategy defined in [15]. In contrast to the anonymization performed in the VoicePrivacy challenge, the *constant* strategy defines a single pseudo-speaker identity for all speakers in a given dataset. We extend this property to all speakers of the trial and the enrollment datasets so that all of them should have the same anonymized voice identity. We expect good anonymization performance as the voices of all speakers should appear to be spoken by a single identity. This breaks the one-to-one requirement of the challenge, speakers will not be distinguishable from each other. Still, we believe this assessment is complementary to the realistic attacker-based evaluation of the challenge. Figure 2 shows the differences in the target selection strategies between the VoicePrivacy strategy (*permanent* strategy), and the strategy chosen for this study (our *constant* strategy). Experiments are performed with different target speaker identities to provide averaged global results about the voice conversion toolkit, and detailed analysis of the target identity effect on multiple source speakers.

3.1 Dataset

All evaluation tasks for the experiments follow the conditions presented on the publicly available VoicePrivacy challenge¹. The triphone extractor has been trained on the *train-clean-100* and *train-other-500* subsets of LibriSpeech. The x-vector extractor has been trained on VoxCeleb-1,2. The speech synthesis system has been trained on the *train-clean-100* subset of LibriTTS. Finally, the *train-other-500* subset of LibriTTS has been used to create a pool of x-vector and F0 statistics. The evaluation dataset is built from LibriSpeech *test-clean*. Details about the number of speakers and utterances in the enrollment and trial dataset are reported in Table 1.

¹ <https://github.com/Voice-Privacy-Challenge>

Table 1. Statistics of the evaluation dataset [17].

Subset		Female	Male	Total
Librispeech test-clean	Speakers in enrollment	16	13	29
	Speakers in trials	20	20	40
	Enrollment utterances	254	184	438
	Trial utterances	734	762	1496

3.2 Utility and Privacy Metrics

To evaluate the performance of the system in both privacy (*speaker’s concealing* capability) and utility (*content intelligibility*) two systems and metrics are used. To quantitatively evaluate the privacy, an x-vector-PLDA based Automatic Speaker Verification (ASV) architecture provided by the challenge organizers is used. The privacy protection is measured with the linkability metric: [4] introduced two different measures that are calculated based on mated and non-mated likelihood score distributions. In this case, mated scores are computed comparing anonymized speech of the same user. Whereas non-mated scores are computed comparing anonymized speech of different users. The local measure $D_{\leftrightarrow}$ is the local score-wise measure depending upon the likelihood ratio between the mated and non-mated sample’s score. The global linkability measure $D_{\leftrightarrow}^{\text{sys}}$ is the average value of $D_{\leftrightarrow}$ over all mated scores. To obtain the linkability score for a given speaker, the average is taken from all mated scores of this specific speaker. Work in [6] advocate the use of $D_{\leftrightarrow}^{\text{sys}}$ as a robust privacy metric. The lower the $D_{\leftrightarrow}^{\text{sys}}$, the better the speakers are anonymized. As the Equal Error Rate ($\text{EER}_{\%}$) measure is more often used in speaker verification, we present our result in terms of both $\text{EER}_{\%}$ and $D_{\leftrightarrow}^{\text{sys}}$. The higher the $\text{EER}_{\%}$, the better the speakers are anonymized.

For the utility, the pre-trained Automatic Speech Recognition (ASR) system provided by the challenge organizers is used to decode the anonymized speech and compute the Word Error Rate ($\text{WER}_{\%}$). In this evaluation, the $\text{WER}_{\%}$ measure is used to evaluate how the content is kept intelligible. The lower the $\text{WER}_{\%}$ is, the more intelligible the anonymized speech is. Both ASR and ASV systems are trained on LibriSpeech *train-clean-360* using Kaldi [11].

3.3 Evaluation Methodology

In this experiment, we run the anonymization and evaluation on 40 real target speaker’s identities that cover as best as possible the speaker space. Thus, to select the target speakers, we identify 20 female and 20 male clusters of x-vectors in the anonymization pool (*LibriTTS train-other-500*) using K-Means. We then pick the speaker x-vector that is the closest to the centroid of each cluster. The assessment of the anonymization performances will be done for each of those 40 target speaker identities. Real speaker identities are used instead of pseudo-speaker identities derived from multiples speakers.

To evaluate the performances of the anonymization system, we perform the following procedure for the 40 selected target speaker identities. First, we convert all utterances of *Librispeech test-clean* (data corresponding to enrollment and trial) and *Librispeech train-360* to the selected target speaker. Then, we train the ASV model on the anonymized *Librispeech train-360* dataset. Lastly, we evaluate the privacy performances for each of the speakers of *Librispeech test-clean* using the specially trained ASV model. As a result, we obtain one score for each trial speaker of *Librispeech test-clean* and target identities (that is a total of 29×40 scores). To evaluate the quality of the conversion process in terms of utility (linked to speech recognition performance), we use a pre-trained ASR model released for the VoicePrivacy Challenge. To study the effectiveness of the F0 anonymization (module D of Figure 1), we perform the anonymization process, and the associated evaluation, with and without the F0 transformation enabled.

4 Experimental Results

4.1 Global Results

Table 2 compares the anonymization performance on a global scale. The first line presents the linkability when no anonymization is performed (i.e., on original speech data), clean speech encapsulates the speaker’s information to a high degree ($D_{\leftrightarrow}^{\text{sys}}$ scores > 0.90). The second and third lines display the linkability when using the voice anonymization system without and with the F0 transformation.

Table 2. Linkability ($D_{\leftrightarrow}^{\text{sys}}$) and EER% scores for original and anonymized speech. For the anonymized data, the mean and standard deviation values are calculated over the 40 experiments (i.e., one for each target speaker). Results on anonymized data are given without and with F0 transformation.

	Female speakers		Male speakers	
	$D_{\leftrightarrow}^{\text{sys}}$	EER%	$D_{\leftrightarrow}^{\text{sys}}$	EER%
Original	0.90	7.66	0.96	1.11
Anon. without F0	0.72 ± 0.01	12.1 ± 0.6	0.77 ± 0.01	9.5 ± 0.6
Anon. with F0	0.74 ± 0.01	11.6 ± 0.6	0.75 ± 0.01	10.6 ± 0.8

In contrast with the original speech results, the anonymized results come from 40 ASV tests, each using a different speaker identity. The mean and standard deviation values are calculated from the 40 evaluations. From the linkability score difference between original and Anon. lines of table 2 we can conclude that speakers are less linkable to their true identity after applying the x-vector base anonymization system. From the original data to the anonymized speech, linkability scores drop by at least 17%, meaning the anonymization system has some

effectiveness. The comparison between the last two lines shows that modifying the F0 values does not help to remove speaker information from the speech signal, as shown by the standard deviation values. This means that the transformation applied was not strong enough to actually remove extra speaker information or that the ASV system is quite robust against such modification. The rather low standard deviation values across all scores show that there is no large variation when changing target speaker identity, this indicates that a given target speaker isn't more suited than another to anonymize the whole dataset. As there are not large scores differences between voice conversion without or with F0 values transformation, the following sections will only analyze the results of the anonymization system that includes the F0 linear transformation.

4.2 Detailed Analysis

Results on a single test speaker (speaker 5105)

Fig. 3. Linkability scores ($D_{\leftrightarrow}^{\text{sys}}$) obtained using the ASV retrained on anonymized data for each of the 40 target speakers (colored solid line) and on original speech (dotted black circle). Each axis corresponds to a target speaker. The anonymization system lowered the linkability scores, meaning better privacy is achieved.

We conducted a detailed analysis to check whether a specific target identity is more suited to anonymize one or more speakers of our test dataset. Figure 3 illustrate the visualization used for this study in the case of a single source speaker. The linkability $D_{\leftrightarrow}^{\text{sys}}$ scores are computed for speech anonymized with 40 different x-vectors and F0 statistics.

The black-dot circle indicates the linkability of the speaker on clean original speech signal (note that the original speech does not depend on the target speakers, hence the circle). And, for each of the 40 target speakers, the linkability is presented by the colored solid line. After transforming the speech with 40 target speakers, we can observe that none of the 40 targets are significantly better to anonymize this speaker’s voice. The variation between the anonymized linkability scores is more likely due to the difference between the ASV model training rather than a better target choice. This observation also applies to the 29 other speakers of our test dataset. It is also noteworthy that, out of the 40 target identities, 20 of them induce cross-gender voice conversion. Results for same-gender and cross-gender voice anonymization were found similar.

Fig. 4. Mean and standard deviation $D_{\leftrightarrow}^{\text{sys}}$ scores obtained for N speakers using re-trained ASVs on anonymized data for each of the 40 target speakers (mean and standard deviation corresponding to the colored solid line and light area) and on original speech (dotted black circle and grey area). Each axis corresponds to a target speaker.

Figure 4 shows the linkability $D_{\leftrightarrow}^{\text{sys}}$ scores of two groups to illustrate the common anonymization behavior: one for which the anonymization system did not remove speaker information, and the other for which the anonymization did remove some speaker information. For the poorly anonymized speakers, we observe that the distributions of linkability scores on anonymized speech (colored area) completely overlaps the distribution of linkability scores on original speech (grey area). The anonymization system did not remove any speaker information for half of our test speakers. On the other hand, for the well anonymized speakers, the anonymized speech and original speech scores distributions diverge. The difference is distinct, speaker information was removed by the anonymization system for the other half of our test speakers.

4.3 Utility Results

Across all experiments, we evaluated the utility for each 40 target identities. We performed the intelligibility test with the pre-trained ASR system of the VoiecPrivacy challenge. Figure 5 shows the utility scores after anonymization considering each of the 40 target identities. On original clean speech signal, the ASR systems score 4.15 WER%. When using the same model (trained on clean data) the overall WER% on the anonymized data reaches 7.30 WER% (average value over the 40 experiments). Retraining the ASR system on anonymized speech improves the WER% significantly [17]. The very high utility loss yielded when using speaker 2487 is due to a generalization issue of the anonymization system (with happens with and without the F0 modification), in this case, non-intelligible speech was generated at the beginning of some segments. We conducted an additional test using 100 randomly selected target speaker identities, and were able to find 3 target speaker x-vectors that have the same behavior, with the worse case having a score of 59.37 WER%. Informal listening test reported that the original speech that produced the faulty x-vector contained singing segments. Further analysis needs to be conducted.

Fig. 5. WER% scores obtained by the VoicePrivacy ASR evaluation systems for each of the 40 speakers and on original speech (dotted black line).

5 Conclusion

In this work, we evaluated and analyzed the x-vector-based speaker anonymization system proposed in [3] and F0 extension proposed in [2] under a white-box attack scenario approach and a *constant* target selection strategy. To assess the

long-term performance of the voice conversion system and the impact of each of the hyper-parameters (i.e. the x-vector target identity, the F0 linear transformation, and the random gender selection), the use of a *constant* target selection strategy is beneficial. This target selection strategy allows the attacker to have complete knowledge about the system, under those circumstances models and hyper-parameters design choices can be compared as the best attacker will be used to evaluate the anonymization performance.

The experiments done in [2, 14] showed that black and grey-box attackers can easily be fooled when enrollment speech and trial speech are anonymized in a different manner, i.e. by selecting the x-vector in different regions, applying the F0 linear transformation and random gender selection. On the contrary, we observed that neither the x-vector target, F0 transformation, or cross-gender conversion actually help to remove the speaker information from the speech signal. We showed that regardless of the hyper-parameters used, privacy protection stays the same. The detailed analysis showed that a given set of hyper-parameters does not help the anonymization system to better anonymize a given source speaker. Furthermore, we concluded that the anonymization system performance depends on the speaker to anonymize, half of our test speakers did not have their privacy improved.

We raise caution on the privacy evaluation procedure, as we’ve shown that system performance varies depending on the attacker’s knowledge and ASV system used. In future work, we plan to evaluate the source of the speaker information leakage that occurs through the phonetic features (PPGs).

Acknowledgments. This work was supported in part by the French National Research Agency under project DEEP-PRIVACY (ANR-18-CE23-0018) and Région Grand Est. Experiments were carried out using the Grid’5000 testbed, supported by a scientific interest group hosted by Inria and including CNRS, RENATER and several Universities as well as other organizations.

References

1. Bahmaninezhad, F., Zhang, C., Hansen, J.H.L.: Convolutional neural network based speaker de-identification. In: Odyssey (2018)
2. Champion, P., Jouvét, D., Larcher, A.: A Study of F0 Modification for X-Vector Based Speech Pseudonymization Across Gender. In: PPAI 2021 - The Second AAAI Workshop on Privacy-Preserving Artificial Intelligence (2021)
3. Fang, F., Wang, X., Yamagishi, J., Echizen, I., Todisco, M., Evans, N., Bonastre, J.F.: Speaker Anonymization Using X-vector and Neural Waveform Models. In: 10th ISCA Speech Synthesis Workshop (2019)
4. Gomez-Barrero, M., Galbally, J., Rathgeb, C., Busch, C.: General framework to evaluate unlinkability in biometric template protection systems. IEEE Transactions on Information Forensics and Security (2018)
5. Magariños, C., Lopez-Otero, P., Docío-Fernandez, L., Rodríguez-Banga, E., Erro, D., García-Mateo, C.: Reversible speaker de-identification using pre-trained transformation functions. Computer Speech & Language (2017)

6. Maouche, M., Srivastava, B.M.L., Vauquier, N., Bellet, A., Tommasi, M., Vincent, E.: A comparative study of speech anonymization metrics. In: Interspeech (2020)
7. McAdams, S.: Spectral fusion, spectral parsing and the formation of the auditory image. Ph. D. Thesis, Stanford (1984)
8. Nautsch, A., Jasserand, C., Kindt, E., Todisco, M., Trancoso, I., Evans, N.: The GDPR & Speech Data: Reflections of Legal and Technology Communities, First Steps Towards a Common Understanding. In: Interspeech (2019)
9. Noé, P.G., Bonastre, J.F., Matrouf, D., Tomashenko, N., Nautsch, A., Evans, N.: Speech Pseudonymisation Assessment Using Voice Similarity Matrices. In: Interspeech (2020)
10. Parliament, E., Council: Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC. General Data Protection Regulation (2016)
11. Povey, D., Ghoshal, A., Boulianne, G., Burget, L., Glembek, O., Goel, N., Hannemann, M., Motlíček, P., Qian, Y., Schwarz, P., Silovský, J., Stemmer, G., Vesel, K.: The Kaldi Speech Recognition Toolkit. IEEE Workshop on Automatic Speech Recognition and Understanding (2011)
12. Srivastava, B.M.L., Bellet, A., Tommasi, M., Vincent, E.: Privacy-Preserving Adversarial Representation Learning in ASR: Reality or Illusion? In: Interspeech (2019)
13. Srivastava, B.M.L., Maouche, M., Sahidullah, M., Vincent, E., Bellet, A.e.a.: Privacy and utility of x-vector based speaker anonymization. Transactions on Audio, Speech and Language Processing (2021)
14. Srivastava, B.M.L., Tomashenko, N., Wang, X., Vincent, E., Yamagishi, J., Maouche, M., Bellet, A., Tommasi, M.: Design choices for x-vector based speaker anonymization. Interspeech (2020)
15. Srivastava, B.M.L., Vauquier, N., Sahidullah, M., Bellet, A., Tommasi, M., Vincent, E.: Evaluating voice conversion-based privacy protection against informed attackers. In: IEEE International Conference on Acoustics, Speech, and Signal Processing (2020)
16. Sun, L., Li, K., Wang, H., Kang, S., Meng, H.: Phonetic posteriorgrams for many-to-one voice conversion without parallel data training. In: IEEE International Conference on Multimedia and Expo (2016)
17. Tomashenko, N., Srivastava, B.M.L., Wang, X., Vincent, E., Nautsch, A., Yamagishi, J., Evans, N., Patino, J., Bonastre, J.F., Noé, P.G., Todisco, M.: Introducing the VoicePrivacy Initiative. Interspeech (2020)