The First VoicePrivacy Attacker Challenge Evaluation Plan

Version 1.0

Natalia Tomashenko¹, Xiaoxiao Miao², Emmanuel Vincent¹, and Junichi Yamagishi³

¹Inria, France ²Singapore Institute of Technology, Singapore ³National Institute of Informatics, Tokyo, Japan

https://www.voiceprivacychallenge.org/attacker/ attacker.challenge@inria.fr

Abstract

The First VoicePrivacy Attacker Challenge is a new kind of challenge organized as part of the VoicePrivacy initiative [1] and supported by **ICASSP 2025** as the **SP Grand Challenge**.¹ It focuses on developing **attacker systems against voice anonymization**, which will be evaluated against a set of anonymization systems submitted to the VoicePrivacy 2024 Challenge². Training, development, and evaluation datasets are provided along with a baseline attacker system. Participants shall develop their attacker systems in the form of automatic speaker verification systems and submit their scores on the development and evaluation data to the organizers. To do so, they can use any additional training data and models, provided that they are openly available and declared before the specified deadline. The metric for evaluation is equal error rate (EER). Results will be presented at the ICASSP 2025 special session to which 5 selected top-ranked participants will be invited to submit and present their challenge systems.

1 Context

Speech encapsulates a wealth of personal, private data, e.g., age and gender, health and emotional state, racial or ethnic origin, geographical background, social identity, and socio-economic status [2]. Formed in 2020, the VoicePrivacy initiative [1] is promoting the development of privacy preservation solutions for speech technology via a series of competitive benchmarking challenges, with common datasets, protocols and metrics. In this context, privacy preservation is classically formulated as a game between *users* who process their utterances (referred to as *trial* utterances) with a privacy preservation system prior to sharing with others, and *attackers* who access these processed utterances or data derived from them and wish to infer information about the users. The level of privacy offered by a given solution is measured as the lowest error rate among all attackers.

The first three VoicePrivacy Challenge editions [3–5] focused on the development of voice anonymization systems. In particular, the systems submitted to the VoicePrivacy 2024 Challenge had to meet the following requirements: (a) output a speech waveform; (b) conceal the speaker identity at the *utterance level*; (c) not distort the linguistic and emotional content. The processed utterances sound as if they were uttered by another speaker, which we refer to as a *pseudo-speaker*. The pseudo-speaker is selected independently for every utterance, and it can be an artificial voice not corresponding to any real speaker. In practice, many voice anonymization systems select the pseudo-speaker or modify prosody in a random or semi-random way using a random number generator. A *semi-informed attack model* [6] was assumed, whereby attackers have access to the voice anonymization system (but not to the random numbers drawn by that system for each utterance, if any), and they seek to re-identify the original speaker behind each anonymized trial utterance. Specifically, an ECAPA-TDNN [7] automatic speaker verification (ASV) system was provided by the organizers and trained by the participants on data anonymized using their anonymization system. While this attack model is undeniably the most realistic to date, the provided attacker system is not its strongest possible implementation as it does not exploit spoken content similarities, specific pseudo-speaker selection strategies [8], or stronger ASV architectures [9], among others.

¹https://2025.ieeeicassp.org/sp-grand-challenges/#gc7

 $^{^{2}}$ The VoicePrivacy Challenge focuses on strengthening voice anonymization systems from the user's perspective, often assuming fixed attack scenarios, which may not fully reflect practical use cases, as real-world attacks can exploit any available clues and resources. In contrast, the Attacker Challenge aims to develop more robust and practical attacker systems capable of challenging various advanced anonymization systems from the attacker's point of view.

To ensure a fair and reliable privacy assessment, it is essential to find the strongest possible attacker against every anonymization system. Hence, the current challenge edition takes the attacker's perspective and focuses on the development of attacker systems against voice anonymization systems.

2 Task

Participants are required to develop one or more attacker systems against one or more voice anonymization systems selected among three VoicePrivacy 2024 Challenge baselines [4] and four systems developed by the VoicePrivacy 2024 Challenge participants. For each speaker of interest, the attacker is assumed to have access to one or more utterances spoken by that speaker, which are referred to as *enrollment* utterances. The attacker system shall output an ASV score for every given pair of trial utterance and enrollment speaker, where higher (resp., lower) scores correspond to same-speaker (resp., different-speaker) pairs.

To develop and evaluate their attacker system against a given voice anonymization system, in line with the assumed semi-informed attack model, participants have access to:

- (a) anonymized trial utterances;
- (b) original and anonymized enrollment utterances;
- (c) original and anonymized training data (as well as other publicly available training resources that will be specified in Section 3) for the ASV system;
- (d) a written description of the voice anonymization system;
- (e) the software implementation of that voice anonymization system when available.

3 Data

For each voice anonymization system, participants are provided with training, development and evaluation data anonymized using that system. A detailed description of these datasets is presented below and in Table 1.

	\mathbf{Subset}		Female	Male	Total	#Utterances
Training	LibriSpeech:	train-clean-360	439	482	921	104,014
Development	LibriSpeech	Enrollment	15	14	29	343
	dev-clean	Trial	20	20	40	1,978
Evolution	LibriSpeech	Enrollment	16	13	29	438
Evaluation	test-clean	Trial	20	20	40	1,496

Table 1: Number of speakers and utterances in the attacker training, development, and evaluation sets.

Training resources. The training set is the *train-clean-360* subset of the *LibriSpeech* [10] corpus. Besides the provided anonymized training data, participants are allowed to use the original *train-clean-360* data. In addition, participants are allowed to use other resources such as speech corpora and pretrained models to develop their attacker systems.

Requirements for training data and models

- All the proposed training data and pretrained models (e.g., wav2vec, WavLM, HuBERT, etc.) should be openly available to everyone at no cost.
- Each registered participant can submit a list of proposed data and models (with the corresponding URLs) to the organizers at attacker.challenge@inria.fr by **13th October**.
- The organizers will verify these requests and publish the list of training data and pretrained models allowed for training attacker models in an updated version of the evaluation plan to be shared with the participants on **15th October**. Any other data or models not included the list will not be allowed for training attacker models.

Development and evaluation data. The development and evaluation sets comprise *LibriSpeech dev-clean* and *LibriSpeech test-clean*. Besides the provided anonymized enrollment data, the participants are allowed to use the original enrollment data.

Voice anonymization systems. The voice anonymization systems to be attacked include three baseline systems (B3, B4, and B5) [4] and four selected systems developed by the VoicePrivacy 2024 Challenge participants (T8-5, T10-2, T12-5, and T25-1) [11]. The participants' systems were chosen according to their anonymization

performance in the highest privacy category (EER $\geq 40\%$), excluding cascaded anonymization systems based on automatic speech recognition (ASR) followed by text-to-speech (TTS). Thus, a total of seven systems are to be attacked:

- **B3** based on phonetic transcription, pitch and energy modification, and artificial pseudo-speaker embedding generation [4, 12].
- B4 based on neural audio codec language modeling [4,13].
- **B5** based on vector quantized bottleneck (VQ-BN) features extracted from an ASR model and on original pitch [4, 14].
- **T8-5** (team *JHU-CLSP*, system "Admixture (p = 0.4)" [15]) random selection of one of two methods for each utterance (with probability p for the second method): (1) a cascaded ASR-TTS system with Whisper [16] for ASR and VITS [17] for TTS and (2) a k-nearest neighbor (kNN) voice conversion (VC) system operating on WavLM [18] features.
- **T10-2** (team *NPU-NTU*, system "*C*4" [19]) neural audio codec, with a specific disentanglement strategy for linguistic content, speaker identity and emotional state.
- T12-5 (team NTU-NPU, system "3" [20]) based on B5, with additional pitch smoothing.
- **T25-1** (team *USTC-PolyU*, system "*large: ESD+LibriTTS*" [21]) disentanglement of content (VQ-BN as in **B5**) and style (global style token (GST) [22]) features and emotion transfer from target speaker utterances.

The code of **B3**, **B4**, and **B5** is available on GitHub³ and can be used to develop attacker systems by, e.g., generating different or additional training data to train those systems.

4 Evaluation metric

We use the equal error rate (EER) metric to evaluate the attacker's performance. This metric has been used in all VoicePrivacy Challenge editions. For every given pair of trial utterance and enrollment speaker, the attacker system outputs an ASV score from which a same-speaker vs. different-speaker decision is made by thresholding. Denoting by $P_{\rm fa}(\theta)$ and $P_{\rm miss}(\theta)$ the false alarm and miss rates at threshold θ , the EER metric corresponds to the threshold $\theta_{\rm EER}$ at which the two detection error rates are equal, i.e., EER = $P_{\rm fa}(\theta_{\rm EER}) = P_{\rm miss}(\theta_{\rm EER})$. The lower this metric, the stronger the attacker. The number of same-speaker and different-speaker trials in the development and evaluation datasets is given in Table 2. The attackers will be ranked separately for each voice anonymization system.

Subs	set	Trials	Female	Male	Total
Dovelopment	LibriSpeech	Same-speaker	704	644	1,348
Development	dev-clean	Different-speaker	14,566	12,796	$27,\!362$
Evaluation	LibriSpeech	Same-speaker	548	449	997
Divaluation	test-clean	t-clean Different-speaker 1		$9,\!457$	$20,\!653$

Table 2: Number of speaker verification trials.

5 Baseline attacker system

As a baseline, we consider the attacker system used in the VoicePrivacy 2024 Challenge [4] (see Figure 1). The ASV system (denoted ASV_{eval}^{anon}) is an ECAPA-TDNN [7] with 512 channels in the convolution frame layers, implemented by adapting the *SpeechBrain* [23] *VoxCeleb* recipe to *LibriSpeech*, and it is trained on anonymized training data. For a given trial utterance and enrollment speaker, the attacker computes the average speaker embedding of all anonymized enrollment utterances from that speaker and compares it to the speaker embedding of the anonymized trial utterance. Results for this baseline attacker system are reported in Table 3. The code to train the baseline attacker systems for given anonymized data is available in the GitHub VoicePrivacy 2024 Challenge repository: https://github.com/Voice-Privacy-Challenge/Voice-Privacy-Challenge-2024.⁴

6 Evaluation rules

• Participants are free to develop their own attacker systems, using components of the provided baseline or not. They are encouraged (but not required) to submit results for each anonymization system and to design attacker systems that target the specific weaknesses of each anonymization system.

³https://github.com/Voice-Privacy-Challenge/Voice-Privacy-Challenge-2024 ⁴See Step 2: Evaluation, run_evaluation.py



Figure 1: Baseline attacker: training ASV_{eval}^{anon} on anonymized training data and using it to compare anonymized trial and enrollment data.

Table 3: EER achieved by the baseline attacker system ASV_{eval}^{anon} on data processed by different anonymization systems vs. EER achieved on the original (Orig.) unprocessed data by an ASV model trained on original data.

Dataset	Gender	EER $(\%)$							
Dataset		Orig.	B3	B4	$\mathbf{B5}$	T8-5	T10-2	T12-5	T25-1
LibriSpeech-dev	female	10.51	28.43	34.37	35.82	39.63	43.63	43.32	42.65
	male	0.93	22.04	31.06	32.92	40.84	40.04	44.10	40.06
Average dev		5.72	25.24	32.71	34.37	40.24	41.83	43.71	41.36
LibriSpeech-test	female	8.76	27.92	29.37	33.95	42.50	41.97	43.61	42.34
	male	0.42	26.72	31.16	34.73	40.05	38.75	41.88	41.92
Average eval		4.59	27.32	30.26	34.34	41.28	40.36	42.75	42.13

- Participants can use the training resources and development datasets specified in Section 3 in order to train their system and tune hyperparameters. The use of additional speech data and models is allowed provided that they are publicly and freely available and declared by email to attacker.challenge@inria.fr before the data declaration deadline (see Table 4).
- To compute the score for the pair {set of enrollment utterances, trial utterance} only the utterances included in this pair can be used from the evaluation data.
- Anonymization system authors and members of their team are welcome to participate. Their results will be considered as official only when attacking other systems than their own. If the authors release the code for the corresponding anonymization system before the deadline for "Declaration of additional training/development data and models" (see Table 4), they can also attack their own anonymization system and participate in official ranking for this system.

7 Registration and submission of results

Registration. Participants/teams are requested to register for the evaluation. Registration should be performed **once only** for each participating entity using the registration form. Participants will receive a confirmation email within 48 hours after successful registration, otherwise or in case of any questions they should contact the organizers: attacker.challenge@inria.fr.

Submission of results. Each participant may submit scores/results for one or more attacker systems, each targeting all anonymization systems or only some of them. Each single submission should include the EER and corresponding ASV scores (for the development and evaluation data) obtained with the proposed attacker system for 4 trial lists (in the same format as generated by the baseline attacker system⁵):

- data/libri_dev_trials_f/trials (example: libri_dev_enrolls-libri_dev_trials_f scores)
- data/libri_dev_trials_m/trials (example: libri_dev_enrolls-libri_dev_trials_m scores)

⁵Example: link

- data/libri test trials f/trials (example: libri test enrolls-libri test trials f scores)
- data/libri_test_trials_m/trials (example: libri_test_enrolls-libri_test_trials_m scores)

All data should be submitted in the form of a single compressed archive.

Each participant should also submit a single, detailed system description. All submissions should be made according to the schedule below. Submissions received after the deadline will be marked as 'late' submissions, without exception. System descriptions will be made publicly available on the Challenge website. Further details concerning the submission procedure will be published via the participants mailing list and via the VoicePrivacy Attacker Challenge website.

Special session at ICASSP 2025. Results will be presented at the ICASSP 2025 special session to which 5 selected top-ranked participants will be invited to submit and present their challenge systems. All participants will be invited to submit the extended versions of their papers to the SPSC 2025 Symposium. Accepted papers will be published in the ICASSP proceedings. According to https://2025.ieeeicassp.org/call-for-gc-proposals/: "The review process is coordinated by the challenge organizers and the SPGC chairs. All 2-page papers should be covered by an ICASSP registration and should be presented in person at the conference."

8 Schedule

The result submission deadline is **5th December 2024**.

Release of the training, development and evaluation data, baselines and evaluation software	September 2024
Declaration of additional training/development data and models	13th October 2024
Publication of the full final list of training data and models	15th October 2024
Deadline for participants to submit scores, evaluation results and system descriptions	5th December 2024
Deadline for participants to submit 2-page papers to ICASSP-2015 (by invitation only)	9th December 2024
Paper Acceptance Notification	30th December 2024

Table 4:	Important	dates
----------	-----------	-------

9 Acknowledgement

This work was supported by the French National Research Agency under project Speech Privacy and project IPoP of the Cybersecurity PEPR.

References

- N. Tomashenko, B. M. L. Srivastava, X. Wang, E. Vincent, A. Nautsch, J. Yamagishi, N. Evans, J. Patino, J.-F. Bonastre, P.-G. Noé, and M. Todisco, "Introducing the VoicePrivacy Initiative," in *Proc. Interspeech* 2020, 2020, pp. 1693–1697. [Online]. Available: http://dx.doi.org/10.21437/Interspeech.2020-1333
- [2] A. Nautsch, A. Jimenez, A. Treiber, J. Kolberg, C. Jasserand, E. Kindt, H. Delgado et al., "Preserving privacy in speaker and speech characterisation," Computer Speech and Language, vol. 58, pp. 441–480, 2019.
- [3] N. Tomashenko, X. Wang, E. Vincent, J. Patino, B. M. L. Srivastava, P.-G. Noé, A. Nautsch, N. Evans, J. Yamagishi, B. O'Brien, A. Chanclu, J.-F. Bonastre, M. Todisco, and M. Maouche, "The VoicePrivacy 2020 Challenge: Results and findings," *Computer Speech and Language*, vol. 74, 2022, https://arxiv.org/pdf/2109.00648.pdf. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0885230822000080
- [4] N. Tomashenko, X. Miao, P. Champion, S. Meyer, X. Wang, E. Vincent, M. Panariello, N. Evans, J. Yamagishi, and M. Todisco, "The VoicePrivacy 2024 challenge evaluation plan," arXiv preprint arXiv:2404.02677, 2024.
- [5] M. Panariello, N. Tomashenko, X. Wang, X. Miao, P. Champion, H. Nourtel, M. Todisco, N. Evans, E. Vincent, and J. Yamagishi, "The voiceprivacy 2022 challenge: Progress and perspectives in voice anonymisation," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2024.
- [6] 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 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2020, pp. 2802–2806.

- [7] B. Desplanques, J. Thienpondt, and K. Demuynck, "ECAPA-TDNN: Emphasized channel attention, propagation and aggregation in TDNN based speaker verification," in *Interspeech*, 2020, pp. 3830–3834.
- [8] P. Champion, T. Thebaud, G. Le Lan, A. Larcher, and D. Jouvet, "On the invertibility of a voice privacy system using embedding alignment," in 2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU), 2021, pp. 191–197.
- [9] C. Zeng, X. Wang, E. Cooper, X. Miao, and J. Yamagishi, "Attention back-end for automatic speaker verification with multiple enrollment utterances," in 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2022, pp. 6717–6721.
- [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), 2015, pp. 5206–5210.
- [11] N. Tomashenko, X. Miao, P. Champion, S. Meyer, X. Wang, E. Vincent, M. Panariello, N. Evans, J. Yamagishi, and M. Todisco, "The VoicePrivacy 2024 challenge," 2024. [Online]. Available: https://www.voiceprivacychallenge.org/vp2024/docs/VPC-2024-.pdf
- [12] S. Meyer, F. Lux, J. Koch, P. Denisov, P. Tilli, and N. T. Vu, "Prosody is not identity: A speaker anonymization approach using prosody cloning," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [13] M. Panariello, F. Nespoli, M. Todisco, and N. Evans, "Speaker anonymization using neural audio codec language models," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024, pp. 4725–4729.
- [14] P. Champion, "Anonymizing speech: Evaluating and designing speaker anonymization techniques," Ph.D. dissertation, Université de Lorraine, 2023.
- [15] H. L. Xinyuan, Z. Cai, A. Garg, K. Duh, L. P. García-Perera, S. Khudanpur, N. Andrews, and M. Wiesner, "HLTCOE JHU submission to the Voice Privacy challenge 2024," arXiv preprint arXiv:2409.08913, 2024.
- [16] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," in *International Conference on Machine Learning*. PMLR, 2023, pp. 28492–28518.
- [17] J. Kim, J. Kong, and J. Son, "Conditional variational autoencoder with adversarial learning for end-to-end text-to-speech," in *International Conference on Machine Learning (ICML)*, 2021, pp. 5530–5540.
- [18] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao, J. Wu, L. Zhou, S. Ren, Y. Qian, Y. Qian, J. Wu, M. Zeng, and F. Wei, "WavLM: Large-scale self-supervised pre-training for full stack speech processing," *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1505–1518, 2022.
- [19] J. Yao, N. Kuzmin, Q. Wang, P. Guo, Z. Ning, D. Guo, K. A. Lee, E.-S. Chng, and L. Xie, "NPU-NTU System for Voice Privacy 2024 Challenge," arXiv preprint arXiv:2409.04173, 2024.
- [20] N. Kuzmin, H.-T. Luong, J. Yao, L. Xie, and K. A. Lee, "NTU-NPU System for Voice Privacy 2024 Challenge," SPSC 2024, 2024. [Online]. Available: https://www.voiceprivacychallenge.org/vp2024/docs/ T12 NTU-NPU System for Voice Privacy 2024 Challenge.pdf
- [21] W. Gu, Z. Liu, L. Chen, R. Wang, C. Guo, W. Guo, K. A. Lee, and Z.-H. Ling, "USTC-PolyU system for the VoicePrivacy 2024 Challenge," SPSC 2024, 2024. [Online]. Available: https://www.voiceprivacychallenge. org/vp2024/docs/T25____USTC-PolyU_system_for_the_VoicePrivacy_2024_Challenge.pdf
- [22] Y. Wang, D. Stanton, Y. Zhang, R.-S. Ryan, E. Battenberg, J. Shor, Y. Xiao, F. Ren, Y. Jia, and R. A. Saurous, "Style tokens: Unsupervised style modeling, control and transfer in end-to-end speech synthesis," in *International Conference on Machine Learning (ICML)*, 2018, pp. 5180–5189.
- [23] M. Ravanelli, T. Parcollet, P. Plantinga, A. Rouhe, S. Cornell, L. Lugosch, C. Subakan, N. Dawalatabad, A. Heba, J. Zhong, J.-C. Chou, S.-L. Yeh, S.-W. Fu, C.-F. Liao, E. Rastorgueva, F. Grondin, W. Aris, H. Na, Y. Gao, R. D. Mori, and Y. Bengio, "SpeechBrain: A general-purpose speech toolkit," arXiv preprint arXiv:2106.04624, 2021.