Why Voice Privacy Researchers Should Worry About Attribute Inference?

Mehtab Ur Rahman[1], *Eulalie Thiombiano*[2], *Martha Larson*[1,2]

[1]Centre for Language Studies, Radboud University, Netherlands
[2]Institute for Computing and Information Sciences, Radboud University, Netherlands
(mehtab.rahman, eulalie.thiombiano, martha.larson)@ru.nl

1. The Importance of Attribute Inference for Voice Privacy

Speech carries more than linguistic content. In addition to the spoken message, it also conveys information about the speaker, including identity, emotion, and other demographic or personal characteristics [1]. Such information can be privacy sensitive and should be protected. Most voice privacy research has focused on anonymization, where the main goal is to hide the speaker's identity while preserving linguistic content or other attributes [2, 3]. However, protecting identity alone may not be enough. Even when identity is concealed, other speaker attributes may still remain inferable from speech and can continue to reveal sensitive information about the speaker.

Attribute inference raises important privacy concerns for two reasons. First, even when a speaker's identity is known, inferring additional attributes reveals unintended personal information. Second, if the speaker's identity is unknown, inferred attributes can serve as quasi identifiers. When several attributes are combined into a speaker profile, they can narrow the set of possible speakers, reduce the anonymity set, increase the uniqueness risk, and make singling out possible. As a result, attribute leakage may still undermine privacy even when identity protection appears successful. The resulting information may then be used for profiling, discrimination, exclusion, surveillance, or unfair decision making.

2. Key Concerns

A central concern is not only whether identity is protected, but which personal attributes remain inferable from protected speech. From a privacy perspective, the key question is what information remains accessible to an attacker after anonymization. A speech signal may appear anonymized, yet still allow an attacker to infer attributes such as age, gender, or accent. Such residual attribute leakage weakens the privacy guarantees of speech protection systems.

Leakage is not limited to correct inference. Even incorrect but consistent predictions can enable speaker linking or tracking across utterances. From a privacy perspective, what matters is not only whether the inferred attributes are semantically correct but also whether they remain stable enough to distinguish or track speakers.

Another concern is that preserving utility may unintentionally retain exploitable attributes. In many settings, anonymization systems are expected to preserve useful information. For example, the VPC 2024 emphasizes preserving emotion-related utility [3]. However, if anonymized speech preserves too many attributes, then the resulting speaker profile may still enable indirect identification, even when identity cues are weakened.

3. Research Directions

These concerns suggest that speech privacy research should move beyond identity-centred evaluation and adopt a broader view of privacy. Privacy should be assessed not only in terms of re-identification risk but also in terms of the mechanisms that enable it, such as combinations of inferred attributes (i.e., speaker profile). If multiple attributes remain inferable after protection, they may still form a stable speaker profile that reduces anonymity and increases the risk of singling out.

Future speech protection methods should be designed to suppress not only identity cues, but also attribute leakage, depending on the intended threat model and application scenario. This means that anonymization systems should be evaluated with respect to what information remains available to an attacker, rather than only whether speaker verification performance is reduced. More attention should be given to which attributes should be protected first.

Future work should evaluate the stability and consistency of inferred attributes under anonymization. Privacy risk is not determined only by whether inferred labels match ground truth, but also by whether the same speaker is assigned a stable inferred profile across utterances. Finally, speech privacy should more carefully examine the tension between privacy and utility. Preserving useful information may be desirable for downstream tasks, but retaining too many speaker attributes can weaken privacy protection. A key direction for future research is to first define realistic threat models in consultation with stakeholders and in alignment with the legal framework. These models should specify the attacker's capabilities and intended use of inferred attributes, providing a principled basis for balancing utility and protection against attribute inference.

Taken together, these directions argue for a broader privacy framework in speech privacy research, one that considers identity, attribute leakage, utility and indirect identification jointly rather than in isolation.

4. References

[1] T. Bäckström, "Privacy in speech technology," *Proceedings of the IEEE*, vol. 113, no. 7, pp. 668–692, 2025.

[2] N. Tomashenko, X. Wang, E. Vincent, J. Patino, B. M. L. Srivastava, P.-G. Noé, A. Nautsch, N. Evans, J. Yamagishi, B. O'Brien *et al.*, "The VoicePrivacy 2020 Challenge: Results and findings," *Computer Speech & Language*, vol. 74, p. 101362, 2022.

[3] 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.