



# Use Cases for Voice Anonymization

Sarina Meyer, Ngoc Thang Vu

Institute for Natural Language Processing, University of Stuttgart, Germany

sarina.meyer@ims.uni-stuttgart.de

## Abstract

The performance of a voice anonymization system is typically measured according to its ability to hide the speaker's identity and keep the data's utility for downstream tasks. This means that the requirements the anonymization should fulfill depend on the context in which it is used and may differ greatly between use cases. However, these use cases are rarely specified in research papers. In this paper, we study the implications of use case-specific requirements on the design of voice anonymization methods. We perform an extensive literature analysis and user study to collect possible use cases and to understand the expectations of the general public towards such tools. Based on these studies, we propose the first taxonomy of use cases for voice anonymization, and derive a set of requirements and design criteria for method development and evaluation. Using this scheme, we propose to focus more on use case-oriented research and development of voice anonymization systems.

**Index Terms:** voice anonymization, voice privacy

## 1. Introduction

The goal of voice anonymization is to modify a speech recording in such a way that it does not contain any information that allow to recognize the speakers in it. At the same time, it should still be possible to use the audio for other purposes. These could be the interaction with a voice assistant [1–3], a therapy session [4] or the training of a text-to-speech (TTS) system [5]. As these downstream applications have very different requirements to an audio, the exact setup of a voice anonymization system depends on its use case.

Most approaches follow the definitions of the Voice Privacy Challenge (VPC) which has been held in biennial rhythm since 2020 [6]. In this challenge, the downstream tasks are specified by models used for objective evaluation and by subjective annotations. The exact evaluation setup and thus the utility focus changes each VPC edition which also affects the focus in the research community. For example, in VPC 2020, utility was evaluated via speech recognition (ASR) and subjective naturalness and intelligibility, whereas in 2024 [7], it was ASR and speech emotion recognition (SER) without subjective evaluation.

While it is known in the research community that the downstream task affects the structure and requirements of the voice anonymization system, in most publications, the downstream task remains underspecified and needs to be deduced from the evaluation metrics. The task is often framed as a downstream machine learning model that is applied to the anonymized data for example for ASR [6–8], SER [7, 9, 10], or health state detection [11–14]. However, voice anonymization is not only relevant if the audio is used for an automatic prediction method but also in situations where no such model is applied, such as




Figure 1: *Methodology of this paper. We perform a use case collection consisting of a literature analysis and user study to develop a use case taxonomy for voice anonymization. Using this categorization, we define use case-dependent requirements and design criteria for anonymization systems.*

online conversations between strangers. In such use cases, the requirements for voice anonymization systems differ from the one for downstream model applications.

Although the need for a taxonomy of use cases has been expressed before [15], to our knowledge, there has not been an extensive discussion of possible use cases for voice anonymization and their implications for the design of such systems so far. In this paper, we thus close this gap by creating a first taxonomy of voice anonymization use cases. As shown in Figure 1, we base this taxonomy on an extensive use case collection consisting of a literature analysis and a user study. We analyze how different use cases affect the requirements an anonymization method needs to fulfill, and define system design criteria based on these requirements.

Our contributions are the following:

- We propose a taxonomy of use cases for voice anonymization based on expert and non-expert perspectives and expectations. We collect these opinions using an extensive literature review and a large-scale international user study.
- We define requirements that an anonymization needs to fulfill in different use cases and that influence the design of voice anonymization approaches.
- Based on these requirements, we propose a use case-based development scheme consisting of a list of design criteria that help to decide how to develop and evaluate an anonymization method for a specific use case.

## 2. Background and Related Work

### 2.1. Voice Anonymization as defined in the VPC

In the speech research community, the term *Voice Anonymization* (speaker anonymization, de-identification) refers to the task of manipulating a speech audio such that the identity of

a speaker cannot be detected anymore while the audio is still usable for a specific purpose, generally called the *downstream task*. In 2020, after a few individual methods have been proposed previously [16–25], the first common definitions and standards for the task were defined by the first VPC [6, 26]. Since then, most publications follow these definitions or the versions of the VPC 2022 [8] or VPC 2024 [7].

In the VPC 2020 [6, 26], the task is described as a game between users publishing data and attackers aiming to extract private information from this data. To prevent attackers from achieving this, users seek to remove personal, identifying information from the data while ensuring that all downstream goals can be achieved. For this, a voice anonymization method is designed that should (a) return a speech waveform, (b) hide the identity of the speaker, (c) not change other speech characteristics, (d) keep the same pseudo-voice for all utterances of a speaker, with different pseudo-voices for different speakers. Condition (c) was changed to linguistic content and paralinguistic attributes in VPC 2022 [8], and to linguistic and emotional content in VPC 2024 [7]. The latter further dropped condition (d). While the exact nature of the downstream goals is not specified in the VPC descriptions, (c) is evaluated using a range of objective and subjective utility metrics, such as ASR, pitch correlation, SER, or human annotations of naturalness and intelligibility. The privacy condition (b) has been further evaluated with different automatic speaker verification (ASV) methods.

These definitions lack a certain specificity, especially regarding condition (c). It is not clear if utility should only be preserved for the specific models used during evaluation, or if it is expected to e.g., keep all paralinguistic attributes that do not correspond to speaker identity. In the case of the latter, it is not clear how to separate identifying attributes from non-identifying ones. These vague conditions lead to a mismatch in expectations and assumptions between groups. For example, it has been argued that a cascading system of ASR and TTS is not suitable for voice anonymization [27, 28], yet, this is not reflected in the task definitions and thus such systems have been proposed by other researchers [29, 30]. Without specifying a use case for the anonymization, it is difficult to come to a consensus about what approaches are suitable or even acceptable.

## 2.2. Use Cases in Voice Anonymization

The topic of use cases has been addressed briefly in previous voice privacy research. For instance, [31] states that the main applications for speech technology are telecommunication and human-computer interfaces, and discusses privacy for these cases. However, the work stays only on a high level in summarizing what kind of sensitive information is processed in a specific application type (e.g., ASR), and how its processing influences privacy. By examining the legal and technical situation of privacy in speech data, [15] describe how and when different data sensitivity categories apply to speech data. They conclude that there is a lack of common understanding between legal and technical communities and thus a need for taxonomies across several dimensions, including use cases. To our knowledge, such a taxonomy has not yet been proposed. [32], however, identified a lack of detail and consistency in papers presenting voice anonymization approaches. They therefore propose a scenario of use scheme to specify the attack and protection models an approach has been designed for. While this scheme would certainly facilitate the comparison of different approaches, it does not support researchers in their decision on how to design a system. On the other hand, [33] examine techniques and chal-

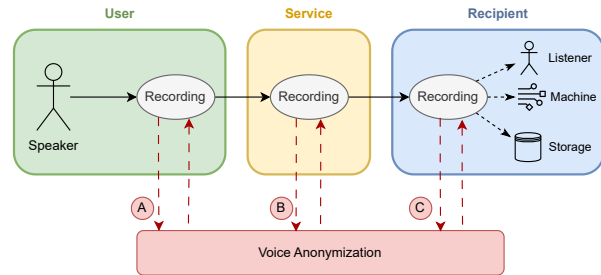


Figure 2: Process of a recording in a general voice anonymization application and actors involved in it. There are three options where to perform voice anonymization: (a) on the user side, (b) on the service side, (c) on the recipient side.

lenges of content privacy in audio recordings. They identify a need in specifying the downstream application when choosing a masking technique because different factors influence the applicability of and requirement for the privacy-preserving method. While focusing on private information in speech content rather than voice characteristics, their categorization of privacy threats according to the way how data is transmitted between different kinds of human or machine actors and how it can be intercepted resembles our use case taxonomy in Section 5.

As addressed in Section 2.1 and further discussed in Section 4.1, most voice anonymization publications do not specify what kind of use case they target. A few works, however, approach a specific scenario and evaluate their anonymization for that, such as psychotherapy sessions [4], civic dialogue networks [34], various health diagnostic settings [11–14], and the anonymization of training data for machine learning models [5, 35]. These approaches are still preliminary work for a few specific use cases but they already show that there is a wide variety of different recording settings, target groups, and general objectives. In this paper, we will address these issues in more detail.

## 3. General Application Scheme of Voice Anonymization

In a general sense, a *use case* or *application* for voice anonymization is a situation involving audio containing human speech (*speech recording*), a reason for why this recording was created and/or is shared (*purpose*), and one or several reasons for why it should be anonymized (*threats*). Within such a situation, there are different actors interacting with the recording, with different interests regarding its contents and modifications. A simplified scheme of these actors is shown in Figure 2.

The first actor is the *speaker* whose speech is being recorded and processed. We will use the terms *speaker* and *user* interchangeably for this role, even if the speaker is not aware of being recorded. Regardless of whether the speaker knows or has access to the recording, the first version of it exists on the user side before then being given to some *service*. This service might just transmit the recording to its final destination, e.g., as an online meeting platform. It might also perform some simple processing steps or store the audio for a while until it can be further processed by the last actor, the *recipient*. The recipient is closely linked to the purpose of the recording, and can be one or several humans (*listener*), algorithms and models processing the audio (*machine*), or a long-term *storage* of the recording. Some actors are combined as the same entity, e.g., the transmission service might be the same as the recipient of the recording.

Each actor has their own version of the recording. Voice anonymization could be performed at any of those stages (A-C in Fig. 2), depending on the use case and factors like trust into each actor, security of the transmission, and device limitations. It is commonly assumed that the anonymization should either be performed on the user (A) or service side (B), due to a lack of trust in the recipient or the requirement to not have to change the downstream application. However, there could be a scenario in which the anonymization on the side of the recipient (C) is preferred, for example when creating a speech dataset.

When developing a voice anonymization method, it is necessary to consider at which point the anonymization should be performed. This affects technical restrictions but also which party can influence the anonymization because each actor has different interests in the recording. The user might want the recording to be only usable for one specific purpose while the recipient might have an interest in using it for several different tasks. Thus, in the one case, anonymization should remove as much information as possible and keep only what is necessary for the downstream task, while in the other case, as much information as possible should be kept and only identifying information should be removed. Which of these strategies to follow, must be decided at an early stage of the system development.

## 4. Use Case Collection

Since no extensive examination of use cases for voice anonymization has been performed yet, we first collect statements about applications of such systems from voice privacy literature. This literature analysis gives insights into what kind of scenarios and objectives research has been targeting so far, and how the success of voice anonymization in these scenarios has been evaluated. As the literature only mirrors the perspective of experts on this task, we extend this analysis by performing a large-scale user study. By including answers of members of the general public across the globe, we aim to understand when and why non-expert users of speech technology would want to have voice anonymization being performed. We analyze and categorize the different statements about anonymization use cases, motivations, and requirements from literature and user study using techniques of qualitative content analysis [36]. We start by collecting all statements for a topic from a subset of publications or responses, group them into categories according to their similarity, and assign the remaining statements to these categories. New categories are formed if necessary and the overall categorization is checked and revised at the end.

### 4.1. Literature Analysis

In order to understand the research perspective on voice anonymization, we analyze a collection of 110 publications [1–14, 16–30, 32, 34, 35, 37–114]. Most papers present anonymization approaches [1, 2, 11, 12, 16–25, 27, 29, 30, 37–97] but we also include work focusing on evaluation techniques for voice anonymization [3–5, 9, 10, 13, 14, 28, 98–112], papers describing the VPC [6–8, 26] and other papers about voice anonymization [32, 34, 35, 113, 114]. We analyze these papers according to three aspects: (a) the anonymization situation (use case) and reason (privacy threat), (b) the requirements for the anonymization (i.e., which information to keep and which to remove during the process), and (c) the evaluation that is either performed or proposed for an anonymization system. We give the percentage of papers that mention a certain aspect to estimate its relevance in the research community.

#### 4.1.1. Anonymization Situation

Most papers (72%) mention one or several general use cases for voice anonymization. Half of the papers (51%) highlight the automatic processing of speech data, e.g., to preserve user privacy during their interactions with voice assistants and IoT, tools analyzing the medical conditions of users, or for using speech as training data in machine learning. 35% of papers see a benefit in anonymizing data that is intended for human listeners, such as in social media, medical consultations, customer service, or court recordings. Another type of use cases includes the general anonymization in data collections, mentioned in 22% of papers.

Although most papers state such use cases as motivations, only 16% explicitly express which use case they are aiming for, whereas the majority leaves this open. If it is specified explicitly, the approaches are mainly designed for health applications, anonymizing model training data, or data sharing.

53% of papers also give one or several reasons for why voice anonymization should be performed. These reasons are generally divided into three attack categories. The main one is protection against spoofing attacks (35%), usually in the form of voice cloning to get access to personal information or to generate harmful fake content about a speaker. Another threat are profiling attacks (26%) in which personal information about a user (e.g., their age) is extracted, often from different sources, in order to create a profile of that user, e.g., for targeted advertisement. Finally, linkage attacks (20%) describe a scenario in which data is linked to a specific person, usually by using speaker recognition.

#### 4.1.2. Anonymization Requirements

The requirements of an anonymization system are usually specified along two dimensions: what information the system should remove (i.e., anonymize) and what it should keep (i.e., preserve) from the input audio. Naturally, the main type of information that is aimed to be removed is anything related to speaker identity (94%). However, also other personal attributes could be aimed to be removed, such as gender (4%), emotion (1%), or sensitive speech content (4%).

There is less consistency when it comes to what information needs to be kept. The preservation of linguistic information is often seen as a natural requirement of voice anonymization [28] but is only named in 65% of papers. Several papers state that the intelligibility should be preserved (18%). Although similar, it is important to note that intelligibility and the preservation of linguistic content are distinct concepts. An anonymized audio might be intelligible but of different linguistic information than the original, or it might degrade the rate of intelligibility while keeping the exact words. Other information that is regularly targeted to be preserved are naturalness (24%), emotions (14%), voice or audio quality (11%), voice distinctiveness (10%), and prosody (10%). 4 papers [46, 49, 59, 92] aim for *asynchronous anonymization* in which identity information is only removed for machine inference but unchanged for human perception, e.g., by adding adversarial noise to the audio.

#### 4.1.3. Evaluation

Almost all papers present some form of objective evaluation (96%) while a subjective evaluation is performed in 33% of cases. The objective evaluation is often split into assessing the privacy requirements (what is removed) and utility requirements (what is kept). For privacy evaluation, speaker verification (76%) and identification (18%) are most often performed.

Utility evaluation often includes the use of ASR (75%), SER (15%), or metrics to assess voice distinctiveness (16%), pitch correlation (13%), and audio or voice quality (15%). Subjective evaluation is often measured as naturalness (17%), intelligibility (12%), speaker similarity (13%) or audio quality (10%).

We observed that not all papers are consistent in evaluating all requirements that they mention in a paper. For example, 12 papers (11%) claim to keep naturalness during anonymization but do not evaluate this.

## 4.2. User Study

The user study was performed on Prolific<sup>1</sup> from September 2024 to May 2025. Besides demographic information, participants were asked to answer in free text form in what scenarios and why they could imagine voice anonymization should be used, both in *general* and specifically for their *personal* voice. Participation in the study took around 5-7 minutes and was compensated according to minimum wage in Germany (12.80 Euro per hour). In total, 170 participants from around the world who were fluent in English were recruited. 25 had to be excluded due to off-topic answers, resulting in a total of 145 participants. The pool of participants was diverse in terms of gender (52% female, 46% male, 2% other or unknown), and age (48% below 30 years, 28% between 30 and 39, 13% between 40 and 49, 10% above 50 years). The participants came from a set of 24 countries (40% from Europe, 26% from Africa, 29% from North America, 6% other), with the majority coming from the US (28%), South Africa (24%), and United Kingdom (23%). 74% of participants stated that English was (one of) their native language(s), with 15% of participants giving multiple native languages. Most participants had a higher education (44% Bachelor's, 24% Master's, 6% PhD as highest degree), and all participants stated to have at least some basic knowledge about artificial intelligence.

### 4.2.1. Anonymization Use Cases

Almost all participants named at least one use case for *general* and *personal* scenarios. A few participants (9%) stated that they would not want their voice to be anonymized, either because they do not trust the technology or because they would find it confusing. Otherwise, there were no clear differences in answers for *general* and *personal* situations, so we merge the answers to both questions together and give the percentages of participants naming a use case for either question.

The use case that has been mentioned most involves *legal situations*, e.g., when a victim or witness gives a statement in court (33%). Other frequently named situations are *online interactions with strangers* (20%) and *calls to customer service* (18%). Given that the participants are likely to frequently take part in other studies on Prolific, 19% also mentioned *research situations* involving speech.

Overall, a large variety of use cases were named which will be discussed in more detail in Section 5. Most participants (83%) listed at least one use case concerning the interaction between humans. Human-computer communication and data storage scenarios were mentioned by 17% and 32%, respectively.

### 4.2.2. Reasons for Anonymization

80% of participants gave not only the use cases but also reasons for why the voice should be anonymized in these scenarios.

The main reason is the wish to simply *hide the identity or stay anonymous* (51%), without further explanation of why this should be necessary. Furthermore, *safety* (27%), like protection from retaliation, and protection from *harassment* (6%) were given by several participants. Other important reasons are to ensure the *confidentiality* of shared information (12%), especially when talking about medical or financial topics, and to *protect the voice from being cloned* or used in *model training* (9%). 8% of participants wish to increase the *security* of their data through anonymization, e.g., in case of data leaks, and 6% want to avoid a *misuse* of their data without consent.

Besides protecting their identity, several participants also mentioned reasons that involved hiding certain attributes in their voice, for example because they *disliked their voice* (3%), want to generally *hide how it sounds like* (3%), or do not want to *reveal certain aspects* about them like their gender or accent (7%). 10% of participants further stated that voice anonymization would give them *confidence to speak freely* and give candid responses. This shows that voice anonymization should not only be seen as a technique for privacy protection, but also to reduce bias, fear, and lack of confidence.

## 4.3. Comparison of Literature and User Study

When comparing the summaries of use cases described in literature and user study in Sections 4.1.1 and 4.2.1, it is clear that there are certain differences in what researchers focus on when developing a voice anonymization system and what users expect from it. While use cases involving the interaction between humans and computers are given in 51% of papers and most papers rely only on objective instead of subjective evaluation, such use cases are only mentioned by 17% of participants in the user study. On the other hand, 83% of participants list use cases for human-human interactions, which are only named in 35% of papers. This suggests that there are needs for voice anonymization from the perspective of the general public that are not addressed in research. We also observed a larger diversity of use cases and reasons for anonymization given by participants than in publications. While the literature analysis might not fully reflect all assumptions and motivations of researchers but only the ones explicitly stated in papers, it indicates that researchers might need to broaden their perspective on possible applications for voice anonymization.

## 5. Use Case Taxonomy

Based on the use cases named in the literature and user study, we create 25 use case categories by grouping the single use cases based on the purpose of the recording, the nature of its contents, or the voluntariness of its creation. Following the distinction between recipients described in Section 3, we divide the categories into three main groups (**human-human interactions, human-computer interactions, data storage**) and 9 subgroups. The overall taxonomy is shown in Figure 3.

Note that some recordings can be used for several use cases, e.g., the interaction with a voice assistant and the usage of that data for model training. Since the categories are purely based on the outcome of our use case collection, there might be use cases that are not included here. Furthermore, some participants in the user study stated that anonymization should always be possible for any kind of recordings, either to have it as a personal choice, or because certain user groups like children or people in witness protection programs should always be protected.

<sup>1</sup><https://www.prolific.com>

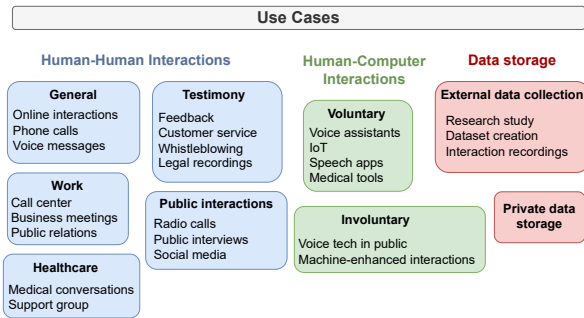


Figure 3: Taxonomy of use cases for voice anonymization.

### 5.1. A: Human-Human Interactions

This category comprises situations in which speech data is recorded and shared for a purpose in which people listen to this data, mostly as a means of communication. The recipients might or might not be known to the user, and could be one person or a group of people.

**A1: General.** These are use cases without a very specific purpose but just with a general communicative aim. This includes **online interactions** where users want to keep the anonymity of their online presence also in voice conversations. Similarly, in **phone calls**, especially when the number of the caller is suppressed or the caller is a stranger, users would like to have the option to keep their voice anonymous in case the other person is a scammer trying to record the user for future voice cloning scams, or in other way untrusted. A third use case in this group is **voice messages**. Some users are aware that these voice messages would be saved on servers that might not be secure enough and therefore want to keep their voice identity hidden. Others stated that they might not know or trust the receiver of the voice message who could for example forward the message to other people without the user’s consent.

**A2: Testimony.** This subcategory contains use cases in which the user is giving some kind of testimony. The use cases differ in their gravity and impact, and range from general feedback situations to crime revealings. As such, the lowest stakes are in situations where users give **feedback** in an environment in which they are known as a person. This could be at the work place or in any other situation in which the user is known to the receivers of the feedback, and wishes to stay anonymous in order to be able to freely express their concerns or complaints. A second scenario is the classic **customer service** situation. Here, it is unlikely that the user is known personally to the customer service employee but needs to be aware that other people might listen to the call or might even record it. Depending on the kind of service, serious complaints might result in disadvantages for the user. Anonymity is crucial if the testimony reveals crimes or ethical misconduct, and the user needs to be protected from retaliation. This is the case for **whistleblowing** in which an insider reveals secret information about an organization. Another situation are **legal recordings** which could be statements by witnesses or victims about a crime. The main difference between whistleblowing and legal recordings is that whistleblowers generally aim to reach the general public with their information, while legal recordings are usually intended for a smaller group of people such as the police or in court.

**A3: Privacy at work.** This category consists of scenarios in which the user is at work and therefore might need to do or say something that they do not want to be connected to their per-

sonal identity. This could be in a **call center** situation in which the call center employee should be anonymous such that the caller is not aware of their identity in case they want to express their dissatisfaction with a company in the form of personal retaliation. It could also be in **business meetings**, with company internal or external members, in which candid comments might be discouraged unless voices are anonymized. It could also be a work-related recording that is directed to the general public in the form of a **public relations** task. Especially if the job is controversial or even dangerous, it is important to give the option to hide the speaker’s identity. Even if the user themselves agrees with the values of their own work, they might live in an environment where it would be dangerous to admit to this.

**A4: Public interactions.** Related to the previous use case, any recording that is intended for a large group of unknown people can pose a certain privacy risk. A user cannot be certain that the recording would not be listened to by someone they know personally, or that the data will not be misused. Especially if the contents of the recording are sensitive, for example in case of **public interviews** or talks about specific topics that are broadcasted to TV, radio or online platforms, anonymization might be important. This could also apply to other public interactions, such as **radio calls** in which the user calls a radio station (e.g., to participate in a quiz) and this call gets broadcasted. In modern times, online presence in **social media** is more prevalent. Subjects stated that they would like to have their voice anonymized in social media such that their audio cannot be used for threats like voice cloning, that their personal contacts cannot identify them, or that their voice cannot be used as background track in other social media posts.

**A5: Healthcare.** Any situation that deals with sensitive information about a person making them vulnerable requires special protection. This is especially true in case of healthcare topics. One scenario are **medical conversations** such as online medical consultations or therapy sessions. Another use case are situations in which a person seeks for help, for example in a **support group** or helpline setting. Offering the option to anonymize the voice of such people or patients in general might encourage more people to seek help without having to fear personal shame or legal consequences.

### 5.2. B: Human-Computer Interactions

In human-computer interactions, speech recordings are intended to be processed by machines, for example to perform ASR or analysis tasks. We distinguish between voluntary and involuntary recordings in this category.

**B1: Voluntary Interactions.** Recordings in which the user chooses to interact with a machine via voice count as voluntary interactions. This includes **voice assistants**, **IoT**, and smaller **speech apps** that are designed for a specific purpose (e.g., transcription or language learning). These applications might process the data on device or in the cloud, and differ in the kind of information they process. Another use case in this category are **medical tools** that are intended to support human practitioners in diagnosing and monitoring patients’ health states either within or outside of medical sessions.

**B2: Involuntary Interactions.** This category consists of use cases in which a user does not get to choose whether they want to interact with a machine, and might not even be aware that their speech is being recorded and processed. These are generally situations of **voice tech in public** places. This comprises surveillance technology but also situations where there was never an intention of processing the data of uninformed




Figure 4: Requirements and design criteria for voice anonymization systems that are influenced by its use case.

people, such as voice assistants in cafes or doctor’s offices. In such cases, it would be necessary to detect and filter out which interaction is intentional and should be kept, and which should be discarded or at least anonymized. Another case are situations that we refer to as **machine-enhanced interactions**. These are typically interactions that are intended for a human recipient but automatically enhanced with voice technology, even if the user does not want this, e.g., phone calls that include an automatic processing of the request before transferring them to a human. Often, the user is forced to interact with the machine if they want to continue the call and might not know how their speech data is processed or stored.

### 5.3. C: Data storage

The third group consists of use cases of general data collection and storage. There are no immediate goals that the user tries to achieve with their recording but rather long-term storage or future processing purposes. The data might later be used in a use case from **A** or **B**, but this future application might not be defined at the time of recording and storage. Thus, data in this group generally require a broader utility than in the other cases.

**C1: External data collection.** These are use cases in which the speech data is recorded or shared with an external party, for example to create a voice dataset. This includes **research study** situations such as voice-based surveys and interviews, or **dataset creation** scenarios in which the data is intended for further analysis or model training. Another type are **interaction recordings** which are intended for a different use case but then stored by a company for quality improvement or training purposes, such as recorded customer service calls.

**C2: Private data storage.** There are also cases in which users might want to let external entities store their data but not process or access it. This would for example occur in case of voice messages or private videos stored in the cloud. Though this could also be solved by encryption, anonymization might be a viable option if the users do not care or even wish that the data will be altered.

## 6. Implications for Voice Anonymization Systems

The large number of diverse use cases shows that voice anonymization is needed in a variety of situations and environments. These range from daily life scenarios that most people experience on a regular basis to less common but particularly sensitive circumstances such as whistleblowing. We also found that hiding one’s identity is not always the only purpose of voice anonymization. Sometimes, the feeling of being anonymous or someone else, a protection against voice cloning, secure data storage, or the reduction of attribute-based biases can be more important than complete anonymity. Thus, depending on the

use case, there are different conditions that a voice anonymization system needs to fulfill. As shown in Figure 4, these are first reflected in different general requirements to the anonymized data, and in turn influence how the model and evaluation need to be designed. As a consequence, the choice of a target use case leads to strong implications for the development of an anonymization system.

### 6.1. Requirements for Voice Anonymization

Each anonymization method needs to fulfill certain requirements for its data to be considered as anonymous and suitable for further processing. These requirements affect the general applicability of a method in a specific scenario and are therefore defined by its use case.

#### 6.1.1. Privacy requirement

The first requirement addresses the question of what information to remove from the recording in order to achieve the required level of privacy. Generally, this defaults to all features that could be used by a detection model or a human listener to recognize the identity of any individual in an audio (*identifying information*). However, depending on the use case, it might be necessary to extend this to other personal attributes such as age, gender, accent, or health state. For example, when anonymizing a recording where potential listeners might know the original speaker personally and where the pool of possible speakers is limited (e.g., in feedback scenarios), the user’s accent might reveal their identity even if the voice attributes are changed. It is therefore necessary to evaluate the exact privacy requirements for a specific use case, and possibly even for a specific user.

#### 6.1.2. Utility requirement

The question of what information should be kept during anonymization is closely related to the purpose of the recording and thus its utility. In some cases, only information strictly necessary for a use case needs to be kept. For example, if an audio is only intended for ASR (e.g., as part of a dictation app), the anonymization might not need to preserve the naturalness or voice diversity of the audio. On the other hand, especially in use cases involving human-human interaction, it might be preferable to keep more than what is strictly necessary. For example, in a phone call, not only the linguistic content should be preserved and intelligible, but the voice should also be rather pleasant to listen to. For some use cases, it might be beneficial to not keep the respective property as displayed in the original audio but to modify them in a way that they improve utility. For example, if the speaker has a strong accent that the recipient might struggle to understand, the anonymization could change the accent to a more familiar one. In a multi-speaker environment, increasing the perceptual differences between voices might make it easier for listeners to distinguish between different speakers.

#### 6.1.3. Output requirements

One basic definition of voice anonymization is that the output should be a waveform. However, the exact specifications for the output format are not given. For example, it is rarely specified if the audio should have the same length and speed as the original, or if this could or should be changed. These decisions can affect the privacy protection [112] but also depend on the purpose of the recording or on other processing steps that are being performed in parallel to voice anonymization. For instance, the anonymized audio might need to be synced with a video that

might itself have been anonymized [4]. It could also be that the linguistic content needs to be anonymized [22, 115] which might need to happen before, during, or after voice anonymization.

## 6.2. System Design Criteria

Combined with the requirements, the use case further influences criteria of anonymization model design and its evaluation.

### 6.2.1. Model Design Criteria

The use case and requirements affect different criteria for model design, ranging from the design of the anonymization process to inference limitations.

**Target voice selection.** Since anonymization typically changes the appearance of the speaker's voice, the selection of the target voice is a big topic in the field. In some approaches, a reference speaker from a dataset is used as target speaker [47, 88], other works focus on creating voices that do not correspond to real speakers, either as a combination of a pool of reference speakers [25, 39, 56, 58], a modification of the original voice [55, 70, 94], or by using a generation model [62, 63, 68, 95]. Besides impacting the performance of the anonymization, there are several aspects that should be considered when implementing a selection method:

- **Voice Origin:** The first question is where the target voice should come from. Depending on the purpose, recipient and ethical factors, it might be acceptable or even beneficial to reuse or modify the voice of an existing speaker, or it might be necessary to generate an artificial voice. Selecting an existing speaker as the target voice is generally simpler but might be harmful for that speaker or confusing, especially if the voice is known to the recipient.
- **Voice Diversity:** Another aspect is the diversity of different voices created by one anonymization system. In some cases, it might be acceptable if several users are anonymized with the same target voice, for example if the recordings are intended for different recipients. However, especially in situations involving several users, it might be necessary to select a new voice for each user or even each utterance. In those cases, it has to be decided how different these voices should be, and if this voice diversity should exactly match the diversity of the original voices (e.g., in a conversation).
- **Durability:** A more technical question is how long the assignment of original voice to target voice should be kept. This depends on if the target voice for one speaker should be consistent for several utterances, sessions or even all times the speaker uses the respective application. From a security and performance point of view, choosing a new voice for each utterance of a speaker is preferable because then speaker mappings do not need to be tracked or stored, and in multi-speaker recordings, speaker diarization only needs to separate utterances but not speakers. Furthermore, it is more stable against privacy attacks because the attacker cannot learn this mapping. However, in most use cases it would be distracting if the user's voice would change with each new utterance, for example in therapy sessions. Especially in recordings containing the voices of several speakers, the information about how many speakers and who speaks what could be lost.
- **Voice Attributes:** Since a voice is connected to several attributes of a speaker, changing a voice would also affect how the user is perceived. In some cases, this is preferred, e.g., to reduce attribute-specific biases, but it might also be relevant to keep certain aspects, for example if they are given

by the context or content. For instance, in a legal or medical recording, the gender of a speaker should generally be kept, whereas it might be preferable to choose features of an opposite or neutral gender for recordings at the work place.

**Multi-speaker support.** The handling of multiple voices by different speakers in one recording has not been addressed much in voice anonymization research yet [96]. However, several use cases consist of multi-speaker interactions, especially in cases of human-human conversations (e.g., online interactions, business meetings, support groups) and general data collections. If multiple voices are present in a recording, the first question is if all of them need to be anonymized or if some could or should stay unmodified (e.g., the therapist or interviewer). Another question is if all voices are combined in one audio channel, or if each voice comes from a different source (e.g., in online meetings). Depending on this, speaker diarization might need to be performed. Multi-speaker anonymization also affects the question of whether the anonymization or target voice selection should be performed on user- or service-side. If everything is performed on the user-side, it might be difficult to ensure sufficient voice diversity in the anonymized result.

**Language support.** While voice anonymization research often focuses on English-only data, several approaches for language-independent or multilingual anonymization have been proposed [58, 86, 97]. For a use case, it needs to be assessed which languages might be present in the audios and whether all of them need to be supported by the anonymization. Linguistic phenomena like code-switching and the use of foreign words should also be taken into account for this decision.

**Run-time.** The run-time of different voice anonymization systems has been discussed in the context of designing methods that could be run in real-time or even streaming [59, 77, 90]. Such real-time anonymization systems are necessary for several use cases, such as phone calls, business meetings or support groups. In other use cases, a run-time that is a bit slower than real-time might be acceptable, such as the use of medical tools or social media. There are even some use cases in which anonymization could be slower and run in the background, e.g., the anonymization of datasets or legal recordings. One question related to the run-time is if the anonymized audio should be presented to the user for a check before it is sent to the recipient. Such a check could give the user the possibility to make sure that the anonymization worked correctly and produced audio of sufficient privacy, utility and quality. This is obviously not possible if the interaction is in real-time but could be useful for example for social media or feedback scenarios.

**Device limitations.** The device that the anonymization should run on restricts the space and resources that are available during the process. Thus, several lightweight approaches have been proposed [41, 77, 89, 90]. The device limitations depend on whether the anonymization should be performed on the user-side or on a server. User-side often means that no GPU is available and that memory is limited. There is more freedom when performing the operation on a server, but this would require to transfer the non-anonymized audio to the server, posing additional privacy and security risks. A related question is whether an internet connection needs to be available during anonymization. These questions depend on the sensitivity of the speech contents and the severity of a privacy leak.

**Audio characteristics.** A recording can include more than just voices but information related to the environment in which it has been created (e.g., background sounds, microphone type, room reverberation). If this is the case, it should be discussed

how to process this. The information might give indications about the identity of the speaker and then needs to be changed during anonymization. In some situations or use cases, however, it might be preferred to keep certain aspects of the original audio characteristics. For example, if a person gives a statement to an interviewer during a protest, they might not want to remove or modify the background sounds as long as nobody could be identified from them.

**Reversibility.** In the earlier works of voice anonymization, the ability to reverse an anonymization process to retrieve the original audio from it was seen as an requirement [17, 19, 21]. Nowadays, there is a consensus that anonymization should be irreversible instead [113] to fully preserve speaker privacy. In fact, from a legal perspective, anonymization is defined as an irreversible process, while the reversible pendant is called pseudonymization [116]. Especially in the case of private data storage, a reversible modification has the benefit that the user would be able to retrieve the data back in its near-original form. While reversible pseudonymization shares certain similarities to encryption, it might have some advantages over existing methods, and thus might be a preferred solution for a use case.

### 6.2.2. Evaluation Design Criteria

Another set of design criteria that depends both on the requirements and the model design concerns the evaluation of the anonymization system. Generally, everything that is defined by the requirements needs to be evaluated. Model design choices need to be included in the evaluation if they are important for a use case. For example, if the run-time is crucial for the suitability of an anonymization method in a certain scenario, it must be measured and reported.

Depending on the utility requirements, the utility evaluation can be either performed objectively, subjectively or both. Not every use case requires an objective and subjective utility evaluation. For example, if the anonymized recording is only intended for human recipients but not for machines, an objective utility evaluation such as ASR is not necessary and can only be used as an approximation of the human perception. On the other hand, if an audio should only be processed by a machine and not by a human, the naturalness of the audio likely does not matter and thus a subjective evaluation is not necessary.

Privacy evaluation should always be performed objectively to test the worst case scenarios for privacy. In most cases, privacy evaluation should measure the ability of the anonymization to hide speaker identity (e.g., using ASV) but it could be extended to other privacy metrics such that different privacy threats are considered, e.g., voice cloning. Privacy evaluation should include a subjective component if it is relevant that the anonymization can be perceived accordingly by human listeners. Besides empirical metrics, theoretical measures can be used to give mathematical proofs for privacy levels [31].

## 7. Future Perspectives

We are aware that the development and evaluation of voice anonymization systems is already complicated and time-consuming, and that even in simplified settings, the task is far from being considered as solved. By proposing to target specific use cases when developing a method, we do not aim to make the task more complicated, but instead to increase openness about assumptions and limitations underlying a technique. We do not expect researchers to fulfill all criteria for a use case or to perform the perfect evaluation. Instead, it is important to address

the limitations of an approach or evaluation to raise the awareness that this method might not perform in all situations as presented. Furthermore, just because an approach does not achieve the perfect scores in privacy or utility evaluation, it is not automatically unsuitable for an actual application. Other factors like run-time, language support or perceptual quality might be more important for a use case than perfect anonymity, and even some imperfect anonymization might still be better than no protection at all. However, the (proven) performance and limitations should always be communicated to the users.

By developing methods regardless of a concrete use case, researchers tend to follow the evaluation framework of the VPC in order to achieve comparability to other methods. In this way, however, ideas that do not fit to this specific evaluation setup are being discouraged even though they might be beneficial in other anonymization scenarios. We therefore need to encourage researchers to explore different directions with different objectives within the realm of voice anonymization. To facilitate this, we propose to split this task into several subtasks, and to include separate tracks for each subtask in the VPC, e.g.: (a) Voice anonymization for human-computer interactions in which utility is defined by the performance of downstream models, (b) voice anonymization for either only human recipients or human and machine processing, in which factors regarding perceptual qualities of the audio need to be considered, (c) asynchronous anonymization [92] in which the identifiability of speakers should differ between human perception and automatic recognition, and (d) (reversible) pseudonymization in which it is possible to reconstruct the original audio, in part or in full.

During our user study, we observed that there is a need to further educate the general public about the harm of generating false audio content (e.g., in the form of audio deepfakes) and about ethical uses of voice anonymization tools. Almost 10% of participants stated that they would like to use voice anonymization technology to either improve the voice in their recordings or to prank other people. It is important to be aware that voice anonymization might be misused for illegal or unethical purposes, but it is also our responsibility as researchers to inform about what are and what are not the objectives of voice anonymization. We should ensure that our technology is unsuitable for generating deepfakes such that anonymized audios should always be recognizable as manipulated audios.

## 8. Conclusion

In this paper, we presented the to our knowledge first analysis of use cases for voice anonymization. Through an extensive literature and user study, we examined the differences between what users expect or want from an anonymization and what the current research focuses on. We find that the people from the general public have more interest in anonymization in scenarios that involve human-human interactions such as online communication and legal testimonies, while the research focuses more on situations involving human-computer interfaces. From the outcome of this use case collection, we derive a taxonomy of 25 use cases in voice anonymization, divided into applications of human-human and human-computer interaction, as well as data storage. We identify different requirements to the anonymization depending on the use case, and define criteria that need to be considered during system development and evaluation in order to meet these requirements. We hope that this scheme will help researchers in making decisions for developing new methods and to support open communication about assumptions and limitations of voice anonymization systems.

## 9. Acknowledgements

This work is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project: Multilingual Controllable Voice Privacy (VoiPy) - Project number 533241795.

## 10. References

- [1] P. Champion, D. Jouvét, and A. Larcher, “A Study of F0 Modification for X-Vector Based Speech Pseudo-Anonymization Across Gender,” in *The Second AAI Workshop on Privacy-Preserving Artificial Intelligence (PPAI)*, online, United States, Nov. 2020. [Online]. Available: <https://hal.archives-ouvertes.fr/hal-02995862>
- [2] C. O. Mawalim and M. Unoki, “Improving security in mcadams coefficient-based speaker anonymization by watermarking method,” in *2021 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (AP-SIPA ASC)*, 2021, pp. 1627–1633.
- [3] S. Zhang, Z. Li, and A. Das, “Privacy measurement of physical attributes on voice anonymity,” in *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking*, ser. ACM MobiCom '24. New York, NY, USA: Association for Computing Machinery, 2024, p. 1650–1652. [Online]. Available: <https://doi.org/10.1145/3636534.3697448>
- [4] C. Franzreb, A. Das, H. Gieseler, E. C. Jahn, T. Polzehl, and S. Möller, “Towards audiovisual anonymization for remote psychotherapy: a subjective evaluation,” in *4th Symposium on Security and Privacy in Speech Communication*, 2024, pp. 102–110.
- [5] W.-C. Huang, Y.-C. Wu, and T. Toda, “Multi-speaker text-to-speech training with speaker anonymized data,” *IEEE Signal Processing Letters*, vol. 31, pp. 2995–2999, 2024.
- [6] 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, “The voiceprivacy 2020 challenge evaluation plan,” 2022. [Online]. Available: <https://arxiv.org/abs/2205.07123>
- [7] 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,” 2024. [Online]. Available: <https://arxiv.org/abs/2404.02677>
- [8] N. Tomashenko, X. Wang, X. Miao, H. Nourtel, P. Champion, M. Todisco, E. Vincent, N. Evans, J. Yamagishi, and J.-F. Bonastre, “The voiceprivacy 2022 challenge evaluation plan,” 2022. [Online]. Available: <https://arxiv.org/abs/2203.12468>
- [9] H. Nourtel, P. Champion, D. Jouvét, A. Larcher, and M. Tahon, “Evaluation of speaker anonymization on emotional speech,” in *2021 ISCA Symposium on Security and Privacy in Speech Communication*, 2021, pp. 62–66.
- [10] S. Zhang, Z. Li, and A. Das, “Voicepm: A robust privacy measurement on voice anonymity,” in *Proceedings of the 16th ACM Conference on Security and Privacy in Wireless and Mobile Networks*, ser. WiSec '23. New York, NY, USA: Association for Computing Machinery, 2023, p. 215–226. [Online]. Available: <https://doi.org/10.1145/3558482.3590175>
- [11] S. Ghosh, M. Jouaiti, A. Das, Y. Sinha, T. Polzehl, I. Siegert, and S. Stober, “Anonymising elderly and pathological speech: Voice conversion using dds and query-by-example,” in *Interspeech 2024*, 2024, pp. 4438–4442.
- [12] V. Ravuri, R. Gutierrez-Osuna, and T. Chaspari, “Preserving mental health information in speech anonymization,” in *2022 10th International Conference on Affective Computing and Intelligent Interaction Workshops and Demos (ACIIW)*, 2022, pp. 1–8.
- [13] S. Tayebi Arasteh, T. Arias-Vergara, P. A. Perez Toro, T. Weise, K. Packhäuser, M. Schuster, E. Nöth, A. Maier, and S. H. Yang, “Addressing challenges in speaker anonymization to maintain utility while ensuring privacy of pathological speech,” *Communications Medicine*, vol. 4, 2024, cRIS-Team Scopus Importer:2024-10-04.
- [14] Y. Zhu, M. Imoussaïne-Aïkous, C. Côté-Lussier, and T. H. Falk, “On the impact of voice anonymization on speech diagnostic applications: A case study on covid-19 detection,” *IEEE Transactions on Information Forensics and Security*, vol. 19, pp. 5151–5165, 2024.
- [15] A. Nautsch, C. Jasserand, E. Kindt, M. Todisco, I. Trancoso, and N. Evans, “The gdpr speech data: Reflections of legal and technology communities, first steps towards a common understanding,” in *Interspeech 2019*, 2019, pp. 3695–3699.
- [16] Q. Jin, A. R. Toth, T. Schultz, and A. W. Black, “Speaker de-identification via voice transformation,” in *2009 IEEE Workshop on Automatic Speech Recognition Understanding*, 2009, pp. 529–533.
- [17] —, “Voice convergin: Speaker de-identification by voice transformation,” in *Proceedings of the 2009 IEEE International Conference on Acoustics, Speech and Signal Processing*, ser. ICASSP '09. USA: IEEE Computer Society, 2009, p. 3909–3912. [Online]. Available: <https://doi.org/10.1109/ICASSP.2009.4960482>
- [18] M. Pobar and I. Ipšić, “Online speaker de-identification using voice transformation,” in *2014 37th International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO)*, 2014, pp. 1264–1267.
- [19] M. Abou-Zleikha, Z.-H. Tan, M. G. Christensen, and S. H. Jensen, “A discriminative approach for speaker selection in speaker de-identification systems,” in *2015 23rd European Signal Processing Conference (EUSIPCO)*, 2015, pp. 2102–2106.
- [20] T. Justin, V. Štruc, S. Dobrišek, B. Vesnicer, I. Ipšić, and F. Mihelič, “Speaker de-identification using diphone recognition and speech synthesis,” in *2015 11th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG)*, vol. 04, 2015, pp. 1–7.
- [21] C. Magariños, P. Lopez-Otero, L. Docio-Fernandez, E. Rodríguez-Banga, D. Erro, and C. Garcia-Mateo, “Reversible speaker de-identification using pre-trained transformation functions,” *Computer Speech Language*, vol. 46, pp. 36–52, 2017. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0885230816302959>
- [22] J. Qian, H. Du, J. Hou, L. Chen, T. Jung, X. Li, Y. Wang, and Y. Deng, “Voicemask: Anonymize and sanitize voice input on mobile devices,” *ArXiv*, vol. abs/1711.11460, 2017.
- [23] F. Bahmaninezhad, C. Zhang, and J. Hansen, “Convolutional neural network based speaker de-identification,” in *The Speaker and Language Recognition Workshop (Odyssey 2018)*, 2018, pp. 255–260.
- [24] J. Qian, F. Han, J. Hou, C. Zhang, Y. Wang, and X.-Y. Li, “Towards privacy-preserving speech data publishing,” in *IEEE INFOCOM 2018 - IEEE Conference on Computer Communications*, 2018, pp. 1079–1087.
- [25] F. Fang, X. Wang, J. Yamagishi, I. Echizen, M. Todisco, N. Evans, and J.-F. Bonastre, “Speaker Anonymization Using X-vector and Neural Waveform Models,” in *Proc. 10th ISCA Workshop on Speech Synthesis (SSW 10)*, 2019, pp. 155–160.
- [26] 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 *Interspeech 2020*, 2020, pp. 1693–1697.
- [27] A. S. Shamsabadi, B. M. L. Srivastava, A. Bellet, N. Vauquier, E. Vincent, M. Maouche, M. Tommasi, and N. Papernot, “Differentially private speaker anonymization,” *Proc. Priv. Enhancing Technol.*, vol. 2023, no. 1, pp. 98–114, 2023. [Online]. Available: <https://doi.org/10.56553/popets-2023-0007>
- [28] 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 Trans. Audio, Speech and Lang. Proc.*, vol. 32, p. 3477–3491, Jul. 2024. [Online]. Available: <https://doi.org/10.1109/TASLP.2024.3430530>

- [29] S. Meyer, P. Tilli, F. Lux, P. Denisov, J. Koch, and N. T. Vu, "Cascade of phonetic speech recognition, speaker embeddings gan and multispeaker speech synthesis for the voiceprivacy 2022 challenge," in *2nd Symposium on Security and Privacy in Speech Communication*, 2022.
- [30] J. Lee, T. Park, and Y. You, "Voice anonymization using emotion-enriched feature integration with stt and tts models," in *4th Symposium on Security and Privacy in Speech Communication*, 2024, pp. 50–54.
- [31] T. Bäckström, "Privacy in speech technology," 2024. [Online]. Available: <https://arxiv.org/abs/2305.05227>
- [32] M. U. Rahman, M. Larson, L. ten Bosch, and C. Tejedor-García, "Scenario of use scheme: Threat modelling for speaker privacy protection in the medical domain," in *4th Symposium on Security and Privacy in Speech Communication*, 2024, pp. 21–25.
- [33] J. Williams, K. Pizzi, S. Das, and P.-G. Noé, "New challenges for content privacy in speech and audio," in *2nd Symposium on Security and Privacy in Speech Communication*, 2022, pp. 1–6.
- [34] W. Kang, M. A. Hughes, and D. Roy, "Anonymization of voices in spaces for civic dialogue: Measuring impact on empathy, trust, and feeling heard," *Proc. ACM Hum.-Comput. Interact.*, vol. 8, no. CSCW2, Nov. 2024. [Online]. Available: <https://doi.org/10.1145/3687021>
- [35] X. Miao, X. Wang, E. Cooper, J. Yamagishi, N. Evans, M. Todisco, J.-F. Bonastre, and M. Rouvier, "Synvox2: Towards a privacy-friendly voxceleb2 dataset," in *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024, pp. 11 421–11 425.
- [36] H.-F. Hsieh and S. E. Shannon, "Three approaches to qualitative content analysis," *Qualitative Health Research*, vol. 15, no. 9, pp. 1277–1288, 2005.
- [37] R. Aloufi, H. Haddadi, and D. Boyle, "Privacy-preserving voice analysis via disentangled representations," in *Proceedings of the 2020 ACM SIGSAC Conference on Cloud Computing Security Workshop*, ser. CCSW'20. New York, NY, USA: Association for Computing Machinery, 2020, p. 1–14. [Online]. Available: <https://doi.org/10.1145/3411495.3421355>
- [38] Y. Han, Y. Cao, S. Li, Q. Ma, and M. Yoshikawa, "Voice-indistinguishability – protecting voiceprint with differential privacy under an untrusted server," in *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 2125–2127. [Online]. Available: <https://doi.org/10.1145/3372297.3420025>
- [39] 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," in *Proc. Interspeech 2020*, 2020, pp. 1713–1717.
- [40] I.-C. Yoo, K. Lee, S. Leem, H. Oh, B. Ko, and D. Yook, "Speaker anonymization for personal information protection using voice conversion techniques," *IEEE Access*, vol. 8, pp. 198 637–198 645, 2020.
- [41] H. Kai, S. Takamichi, S. Shiota, and H. Kiya, "Lightweight voice anonymization based on data-driven optimization of cascaded voice modification modules," in *2021 IEEE Spoken Language Technology Workshop (SLT)*, 2021, pp. 560–566.
- [42] J. Patino, N. Tomashenko, M. Todisco, A. Nautsch, and N. Evans, "Speaker anonymisation using the mcadams coefficient," in *Interspeech 2021*, 2021, pp. 1099–1103.
- [43] G. P. Prajapati, D. K. Singh, P. P. Amin, and H. A. Patil, "Voice Privacy Through x-Vector and CycleGAN-Based Anonymization," in *Proc. Interspeech 2021*, 2021, pp. 1684–1688.
- [44] J. Qian, H. Du, J. Hou, L. Chen, T. Jung, and X.-Y. Li, "Speech sanitizer: Speech content desensitization and voice anonymization," *IEEE Transactions on Dependable and Secure Computing*, vol. 18, no. 6, pp. 2631–2642, 2021.
- [45] D. Stoidis and A. Cavallaro, "Protecting Gender and Identity with Disentangled Speech Representations," in *Proc. Interspeech 2021*, 2021, pp. 1699–1703.
- [46] A. Agarwal, A. Swain, and S. R. Mahadeva Prasanna, "Speaker anonymization for machines using sinusoidal model," in *2022 IEEE International Conference on Signal Processing and Communications (SPCOM)*, 2022, pp. 1–5.
- [47] C. Pierre, A. Larcher, and D. Jouvét, "Are disentangled representations all you need to build speaker anonymization systems?" in *Proc. Interspeech 2022*, 2022, pp. 2793–2797.
- [48] H.-P. Chang, I.-C. Yoo, C. Jeong, and D. Yook, "Zero-shot unseen speaker anonymization via voice conversion," *IEEE Access*, vol. 10, pp. 130 190–130 199, 2022.
- [49] M. Chen, L. Lu, J. Yu, Y. Chen, Z. Ba, F. Lin, and K. Ren, "Privacy-utility balanced voice de-identification using adversarial examples," 2022. [Online]. Available: <https://arxiv.org/abs/2211.05446>
- [50] M. Costante, M. Matassoni, and A. Brutti, "Using seq2seq voice conversion with pre-trained representations for audio anonymization: experimental insights," in *2022 IEEE International Smart Cities Conference (ISC2)*, 2022, pp. 1–7.
- [51] S. P. Dubagunta, R. J. van Son, and M. Magimai.-Doss, "Adjustable deterministic pseudonymization of speech," *Computer Speech Language*, vol. 72, p. 101284, 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0885230821000851>
- [52] A. Hernandez, P. A. Pérez-Toro, J. C. Vázquez-Correa, J. R. Orozco-Arroyave, A. Maier, and S. H. Yang, "Self-supervised speech representations preserve speech characteristics while anonymizing voices," 2022. [Online]. Available: <https://arxiv.org/abs/2204.01677>
- [53] Y. Hu, R. Li, S. Wang, F. Tao, and Z. Sun, "Speechhide: A hybrid privacy-preserving mechanism for speech content and voiceprint in speech data sharing," in *2022 7th IEEE International Conference on Data Science in Cyberspace (DSC)*, 2022, pp. 345–352.
- [54] M. Maouche, B. M. L. Srivastava, N. Vauquier, A. Bellet, M. Tommasi, and E. Vincent, "Enhancing speech privacy with slicing," in *Interspeech 2022*, 2022, pp. 5025–5029.
- [55] C. O. Mawalim, K. Galajit, J. Karnjana, S. Kidani, and M. Unoki, "Speaker anonymization by modifying fundamental frequency and x-vector singular value," *Computer Speech Language*, vol. 73, p. 101326, 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0885230821001194>
- [56] S. Meyer, F. Lux, P. Denisov, J. Koch, P. Tilli, and N. T. Vu, "Speaker anonymization with phonetic intermediate representations," in *Interspeech 2022*, 2022, pp. 4925–4929.
- [57] X. Miao, X. Wang, E. Cooper, J. Yamagishi, and N. Tomashenko, "Analyzing Language-Independent Speaker Anonymization Framework under Unseen Conditions," in *Proc. Interspeech 2022*, 2022, pp. 4426–4430.
- [58] —, "Language-Independent Speaker Anonymization Approach Using Self-Supervised Pre-Trained Models," in *Proc. The Speaker and Language Recognition Workshop (Odyssey 2022)*, 2022, pp. 279–286.
- [59] P. O' Reilly, A. Bugler, K. Bhandari, M. Morrison, and B. Pardo, "Voiceblock: Privacy through real-time adversarial attacks with audio-to-audio models," in *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, Eds., vol. 35. Curran Associates, Inc., 2022, pp. 30 058–30 070. [Online]. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2022/file/c204d12afa0175285e5aac65188808b4-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2022/file/c204d12afa0175285e5aac65188808b4-Paper-Conference.pdf)
- [60] J. M. Perero-Codosero, F. M. Espinoza-Cuadros, and L. A. Hernández-Gómez, "X-vector anonymization using autoencoders and adversarial training for preserving speech privacy," *Computer Speech Language*, vol. 74, p. 101351, 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S088523082200002X>

- [61] L. Tavi, T. Kinnunen, and R. González Hautamäki, “Improving speaker de-identification with functional data analysis of f0 trajectories,” *Speech Communication*, vol. 140, pp. 1–10, 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0167639322000498>
- [62] H. Turner, G. Lovisotto, and I. Martinovic, “Generating identities with mixture models for speaker anonymization,” *Computer Speech Language*, vol. 72, p. 101318, 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0885230821001133>
- [63] R. Yuan, Y. Wu, J. Li, and J. Kim, “Deid-vc: Speaker de-identification via zero-shot pseudo voice conversion,” in *Interspeech 2022*, 2022, pp. 2593–2597.
- [64] M. Chen, L. Lu, J. Wang, J. Yu, Y. Chen, Z. Wang, Z. Ba, F. Lin, and K. Ren, “Voicecloak: Adversarial example enabled voice de-identification with balanced privacy and utility,” *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, vol. 7, no. 2, Jun. 2023. [Online]. Available: <https://doi.org/10.1145/3596266>
- [65] J. Deng, F. Teng, Y. Chen, X. Chen, Z. Wang, and W. Xu, “V-cloak: intelligibility-, naturalness- & timbre-preserving real-time voice anonymization,” in *Proceedings of the 32nd USENIX Conference on Security Symposium*, ser. SEC ’23. USA: USENIX Association, 2023.
- [66] P. Gupta, S. Singh, G. P. Prajapati, and H. A. Patil, *Voice Privacy in Biometrics*. Cham: Springer International Publishing, 2023, pp. 1–29. [Online]. Available: [https://doi.org/10.1007/978-3-031-15816-2\\_1](https://doi.org/10.1007/978-3-031-15816-2_1)
- [67] Y. Lv, J. Yao, P. Chen, H. Zhou, H. Lu, and L. Xie, “Salt: Distinguishable speaker anonymization through latent space transformation,” in *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2023, pp. 1–8.
- [68] S. Meyer, P. Tilli, P. Denisov, F. Lux, J. Koch, and N. T. Vu, “Anonymizing speech with generative adversarial networks to preserve speaker privacy,” in *2022 IEEE Spoken Language Technology Workshop (SLT)*, 2023, pp. 912–919.
- [69] 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 *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [70] X. Miao, X. Wang, E. Cooper, J. Yamagishi, and N. Tomashenko, “Speaker anonymization using orthogonal householder neural network,” *IEEE/ACM Trans. Audio, Speech and Lang. Proc.*, vol. 31, p. 3681–3695, Sep. 2023. [Online]. Available: <https://doi.org/10.1109/TASLP.2023.3313429>
- [71] F. Nespoli, D. Barreda, J. Bitzer, and P. A. Naylor, “Two-stage voice anonymization for enhanced privacy,” in *Interspeech 2023*, 2023, pp. 3854–3858.
- [72] M. Tran and M. Soleymani, “A speech representation anonymization framework via selective noise perturbation,” in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [73] J. Yao, Q. Wang, Y. Lei, P. Guo, L. Xie, N. Wang, and J. Liu, “Distinguishable speaker anonymization based on formant and fundamental frequency scaling,” in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [74] S. Akti, T. N. Nguyen, Y. Liu, and A. Waibel, “Voice privacy - investigating voice conversion architecture with different bottleneck features,” in *4th Symposium on Security and Privacy in Speech Communication*, 2024, pp. 44–49.
- [75] O. L. Blouch, R. BAKARI, and N. Gengembre, “Tuning dissc for voice privacy challenge 2024,” in *4th Symposium on Security and Privacy in Speech Communication*, 2024, pp. 111–115.
- [76] Z. Cai, H. L. Xinyuan, A. Garg, L. P. García-Perera, K. Duh, S. Khudanpur, N. Andrews, and M. Wiesner, “Privacy versus emotion preservation trade-offs in emotion-preserving speaker anonymization,” in *2024 IEEE Spoken Language Technology Workshop (SLT)*, 2024, pp. 409–414.
- [77] W. Chen, W. Tang, Y. Meng, and Y. Zhang, “An hasm-assisted voice disguise scheme for emotion recognition of iot-enabled voice interface,” *IEEE Internet of Things Journal*, vol. 11, no. 22, pp. 36 397–36 409, 2024.
- [78] L. Chen, K. A. Lee, W. Guo, and Z.-H. Ling, “Modeling pseudo-speaker uncertainty in voice anonymization,” in *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024, pp. 11 601–11 605.
- [79] X. Chen, S. Li, J. Li, H. Huang, Y. Cao, and L. He, “Reprogramming self-supervised learning-based speech representations for speaker anonymization,” in *Proceedings of the 5th ACM International Conference on Multimedia in Asia*, ser. MMAsia ’23. New York, NY, USA: Association for Computing Machinery, 2024. [Online]. Available: <https://doi.org/10.1145/3595916.3626366>
- [80] M. Cheng, X. Diao, S. Cheng, and W. Liu, *SAIC: Integration of Speech Anonymization and Identity Classification*. Cham: Springer Nature Switzerland, 2024, pp. 295–306. [Online]. Available: [https://doi.org/10.1007/978-3-031-63592-2\\_22](https://doi.org/10.1007/978-3-031-63592-2_22)
- [81] A. Das, C. Franzreb, T. Herzog, P. Pirlet, and T. Polzehl, “Comparing speech anonymization efficacy by voice conversion using knn and disentangled speaker feature representations,” in *4th Symposium on Security and Privacy in Speech Communication*, 2024, pp. 121–126.
- [82] A. Das, C. Franzreb, S. Ghosh, T. Polzehl, and S. Möller, “Speaker: Towards privacy ensuring decoder only speech reconstruction through disentanglement for german speech anonymization using any-to-many voice conversion,” in *4th Symposium on Security and Privacy in Speech Communication*, 2024, pp. 86–91.
- [83] W. Gu, Z. Liu, L. Chen, R. Wang, C. Guo, W. Guo, K. A. Lee, and Z.-H. Ling, “A voice anonymization method based on content and non-content disentanglement for emotion preservation,” in *4th Symposium on Security and Privacy in Speech Communication*, 2024, pp. 116–120.
- [84] F. Huang, K. Zeng, and W. Zhu, “Diffvc+: Improving diffusion-based voice conversion for speaker anonymization,” in *Interspeech 2024*, 2024, pp. 4453–4457.
- [85] M. Matassoni, S. Fong, and A. Brutti, “Speaker anonymization: Disentangling speaker features from pre-trained speech embeddings for voice conversion,” *Applied Sciences*, vol. 14, no. 9, 2024. [Online]. Available: <https://www.mdpi.com/2076-3417/14/9/3876>
- [86] S. Meyer, F. Lux, and N. T. Vu, “Probing the feasibility of multi-lingual speaker anonymization,” in *Interspeech 2024*, 2024, pp. 4448–4452.
- [87] M. Panariello, M. Todisco, and N. Evans, “Preserving spoken content in voice anonymisation with character-level vocoder conditioning,” in *4th Symposium on Security and Privacy in Speech Communication*, 2024, pp. 12–16.
- [88] M. Panariello, F. Nespoli, M. Todisco, and N. Evans, “Speaker anonymization using neural audio codec language models,” in *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024, pp. 4725–4729.
- [89] J. Pohlhausen, F. Nespoli, and J. Bitzer, “Enhancing speech privacy with lpc modifications,” in *4th Symposium on Security and Privacy in Speech Communication*, 2024, pp. 80–85.
- [90] W. Quamer and R. Gutierrez-Osuna, “End-to-end streaming model for low-latency speech anonymization,” in *2024 IEEE Spoken Language Technology Workshop (SLT)*, 2024, pp. 727–734.
- [91] D. K. Singh, G. P. Prajapati, and H. A. Patil, “Voice privacy using time-scale and pitch modification,” *SN Comput. Sci.*, vol. 5, no. 2, Jan. 2024. [Online]. Available: <https://doi.org/10.1007/s42979-023-02549-8>
- [92] R. Wang, L. Chen, K. A. Lee, and Z.-H. Ling, “Asynchronous voice anonymization using adversarial perturbation on speaker embedding,” in *Interspeech 2024*, 2024, pp. 4443–4447.

- [93] J. J. Webber, O. Watts, G. E. Henter, J. Williams, and S. King, "Voice conversion-based privacy through adversarial information hiding," in *4th Symposium on Security and Privacy in Speech Communication*, 2024, pp. 39–43.
- [94] J. Yao, Q. Wang, P. Guo, Z. Ning, and L. Xie, "Distinctive and natural speaker anonymization via singular value transformation-assisted matrix," *IEEE/ACM Trans. Audio, Speech and Lang. Proc.*, vol. 32, p. 2944–2956, Jun. 2024. [Online]. Available: <https://doi.org/10.1109/TASLP.2024.3407600>
- [95] L. Chen, W. Gu, K. A. Lee, W. Guo, and Z.-H. Ling, "Pseudo-speaker distribution learning in voice anonymization," *IEEE Transactions on Audio, Speech and Language Processing*, vol. 33, pp. 272–285, 2025.
- [96] X. Miao, R. Tao, C. Zeng, and X. Wang, "A benchmark for multi-speaker anonymization," *IEEE Transactions on Information Forensics and Security*, vol. 20, pp. 3819–3833, 2025.
- [97] J. Yao, Q. Wang, P. Guo, Z. Ning, Y. Yang, Y. Pan, and L. Xie, "Musa: Multi-lingual speaker anonymization via serial disentanglement," *IEEE Transactions on Audio, Speech and Language Processing*, vol. 33, pp. 1664–1674, 2025.
- [98] M. Maouche, B. M. L. Srivastava, N. Vauquier, A. Bellet, M. Tommasi, and E. Vincent, "A Comparative Study of Speech Anonymization Metrics," in *Proc. Interspeech 2020*, 2020, pp. 1708–1712.
- [99] A. Nautsch, J. Patino, N. Tomashenko, J. Yamagishi, P.-G. Noé, J.-F. Bonastre, M. Todisco, and N. Evans, "The Privacy ZEBRA: Zero Evidence Biometric Recognition Assessment," in *Proc. Interspeech 2020*, 2020, pp. 1698–1702.
- [100] P.-G. Noé, J.-F. Bonastre, D. Matrouf, N. Tomashenko, A. Nautsch, and N. Evans, "Speech Pseudonymisation Assessment Using Voice Similarity Matrices," in *Proc. Interspeech 2020*, 2020, pp. 1718–1722.
- [101] B. M. Lal 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)*, 2020, pp. 2802–2806.
- [102] 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," *Comput. Speech Lang.*, vol. 74, no. C, Jul. 2022. [Online]. Available: <https://doi.org/10.1016/j.csl.2022.101362>
- [103] P. Champion, D. Jouvét, and A. Larcher, "Evaluating x-vector-based speaker anonymization under white-box assessment," in *Speech and Computer*, A. Karpov and R. Potapova, Eds. Cham: Springer International Publishing, 2021, pp. 100–111.
- [104] P. Champion, T. Thebaud, G. Le Lan, A. Larcher, and D. Jouvét, "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.
- [105] K. Baeg, Y. Han, and B.-K. Jeon, "Dnn based speaker meta information estimation using privacy-preserving speech data," in *2022 IEEE International Conference on Big Data and Smart Computing (BigComp)*, 2022, pp. 88–91.
- [106] W. Liu, J. Li, C. Wei, M. Sun, X. Zhang, and Y. Li, "A novel method to evaluate the privacy protection in speaker anonymization," in *Artificial Intelligence and Security*, X. Sun, X. Zhang, Z. Xia, and E. Bertino, Eds. Cham: Springer International Publishing, 2022, pp. 627–636.
- [107] C. Franzreb, T. Polzehl, and S. Möller, "A comprehensive evaluation framework for speaker anonymization systems," in *3rd Symposium on Security and Privacy in Speech Communication*, 2023, pp. 65–72.
- [108] M. Panariello, M. Todisco, and N. Evans, "Vocoder drift in x-vector-based speaker anonymization," in *Interspeech 2023*, 2023, pp. 2863–2867.
- [109] A. Leschanowsky, E. Gaznepoglu, and N. Peters, "Voice anonymization for all-bias evaluation of the voice privacy challenge baseline systems," in *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024, pp. 4785–4789.
- [110] S. Meyer, X. Miao, and N. T. Vu, "Voicepat: An efficient open-source evaluation toolkit for voice privacy research," *IEEE Open Journal of Signal Processing*, vol. 5, pp. 257–265, 2024.
- [111] J. Williams, K. Pizzi, N. Tomashenko, and S. Das, "Anonymizing speaker voices: Easy to imitate, difficult to recognize?" in *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024, pp. 12 491–12 495.
- [112] N. Tomashenko, E. Vincent, and M. Tommasi, "Analysis of speech temporal dynamics in the context of speaker verification and voice anonymization," in *ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2025, pp. 1–5.
- [113] A. Nautsch, A. Jiménez, A. Treiber, J. Kolberg, C. Jasserand, E. Kindt, H. Delgado, M. Todisco, M. A. Hmani, A. Mtibaa, M. A. Abdelraheem, A. Abad, F. Teixeira, D. Matrouf, M. Gomez-Barrero, D. Petrovska-Delacrétaz, G. Chollet, N. Evans, T. Schneider, J.-F. Bonastre, B. Raj, I. Trancoso, and C. Busch, "Preserving privacy in speaker and speech characterisation," *Computer Speech Language*, vol. 58, pp. 441–480, 2019. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0885230818303875>
- [114] B. O'Brien, N. Tomashenko, A. Chanclu, and J.-F. Bonastre, "Anonymous speaker clusters: Making distinctions between anonymised speech recordings with clustering interface," in *Interspeech 2021*, 2021, pp. 3580–3584.
- [115] J. Williams, K. Pizzi, P.-G. Noe, and S. Das, "Exploratory evaluation of speech content masking," in *Speech Communication; 15th ITG Conference*, 2023, pp. 215–219.
- [116] P. Kamocki and I. Siegert, "Pseudonymisation of speech data as an alternative approach to GDPR compliance," in *Proceedings of the Workshop on Ethical and Legal Issues in Human Language Technologies and Multilingual De-Identification of Sensitive Data In Language Resources within the 13th Language Resources and Evaluation Conference*, I. Siegert, M. Rigault, and V. Arranz, Eds. Marseille, France: European Language Resources Association, Jun. 2022, pp. 17–21. [Online]. Available: <https://aclanthology.org/2022.legal-1.4/>