Understanding the Tradeoffs in Client-side Privacy for Downstream Speech Tasks

Peter Wu, Paul Pu Liang, Jiatong Shi, Ruslan Salakhutdinov, Shinji Watanabe, Louis-Philippe Morency Carnegie Mellon University, PA, USA E-mail: peterw1@cs.cmu.edu

Abstract—As users increasingly rely on cloud-based computing services, it is important to ensure that uploaded speech data remains private. Existing solutions rely either on server-side methods or focus on hiding speaker identity. While these approaches reduce certain security concerns, they do not give users clientside control over whether their biometric information is sent to the server. In this paper, we formally define client-side privacy and discuss its unique technical challenges requiring 1) direct manipulation of raw data on client devices, 2) adaptability with a broad range of server-side processing models, and 3) low time and space complexity for compatibility with limited-bandwidth devices. These unique challenges require a new class of models that achieve fidelity in reconstruction, privacy preservation of sensitive personal attributes, and efficiency during training and inference. As a step towards client-side privacy for speech recognition, we investigate three techniques spanning signal processing, disentangled representation learning, and adversarial training. Through a series gender and accent masking tasks, we observe that each method has its unique strengths, but none manage to effectively balance the trade-offs between performance, privacy, and complexity. These insights call for more research in clientside privacy to ensure a safer deployment of cloud-based speech processing services.

I. INTRODUCTION

Users increasingly rely on cloud-based machine learning models to process their personal data [1–7]. For example, in cloud-based automatic speech recognition (ASR) systems, audio data recorded on client-side mobile devices are typically uploaded to centralized servers for server-side processing [5–8], which enables general server-side ASR models to improve over time and enjoy economies of scale. However, there is growing concern that sending raw speech data to the cloud leaves users vulnerable to giving away sensitive personal biometric information such as gender, age, race, and other social constructs [9].

In order to have full control over their privacy, users should be able to encrypt their data themselves before uploading to downstream applications. We refer to this type of privacy as **client-side privacy**, which requires removing sensitive information on the client device while keeping the resulting data compatible with cloud-based services. For example, for a cloud-based ASR service, a client-side privacy algorithm needs to remove biometric information from raw speech data while keeping the resulting audio signal useful for training the ASR model on the cloud.

As a step towards achieving client-side privacy for speech data, we contribute the following in this work:

- 1) We formally define client-side privacy and describe its unique technical challenges requiring 1) direct manipulation and regeneration of raw data on client devices, 2) adaptability with a broad range of general server-side processing models, and 3) low time and space complexity for compatibility with limited-bandwidth client devices.
- 2) We study three different client-side privacy approaches for speech: signal processing, disentangled representation learning, and adversarial training.
- 3) We conduct experiments on protecting gender and accent information for downstream ASR systems and provide an empirical comparison of current approaches.

We find that each of our three approaches performs well on a subset of metrics, and quantify remaining areas for improvement using multiple privacy metrics. Based on these insights, we propose several extensions for future work and call for more research in client-side privacy to ensure safe cloud-based speech processing.

We proceed by discussing related privacy algorithms in Section II. In Section III, we formalize client-side privacy and describe its unique technical challenges. Then, we describe our client-side privacy approaches for downstream ASR in Section IV and detail the experiments in Section V. Finally, we summarize our results and propose future directions in Section VII. All our code and other supplementary material can be found at https://github.com/peter-yh-wu/speech-privacy.

II. RELATED WORK

A. Client- and Server-side Privacy

One way to view privacy algorithms is by whether they preserve privacy on the client-side, the server-side, or both. Client-side algorithms execute operations on the user's local device, and server-side algorithms run on a remote server [10–12]. For example, one way for cloud services to strengthen data privacy is by encrypting on the client-side and decrypting on the server-side [13]. For a privacy algorithm to be client-side only, all operations must run on-device without any additional work needed on the server.

Based on these definitions, we can categorize existing ways to preserve privacy in machine learning models. Currently, utilizing cryptography algorithms like secure multi-party computation (SMC) and fully homomorphic encryption (FHE) requires both client and server-side components [14, 15]. Other popular approaches like federated learning and Private Aggregation of Teacher Ensembles (PATE) require both client and server-side modifications as well [11, 12, 16]. Serverside only approaches also exist, including differentially private SGD (DP-SGD) and global differential privacy [17, 18]. We identify two types of approaches that are client-side only: (1) local differential privacy and (2) methods we refer to as client-side transforms. We define client-side transforms as algorithms that anonymize sensitive information on-device while preserving content needed for downstream tasks. In other words, client-side privacy can be obtained using performant client-side transforms. We note that being entirely ondevice, client-side transforms can be used in conjunction with the aforementioned privacy algorithms by simply applying the client-side transforms first. All three approaches that we study in this paper are client-side transforms, as detailed in Section III-A. We observe that client-side transforms are not constrained by trade-offs inherent in local differential privacy. Figure 1 summarizes the aforementioned categorization of privacy algorithms.

We note that many client-side transforms exist for downstream tasks simpler than ASR. For example, for the downstream task of storage, client-side encryption is sufficient [19]. In contrast, complex downstream tasks like ASR require that the output of the client-side transforms preserve complex content like transcribable audio. Thus, a challenge arises from ensuring that this resulting content also lacks sensitive information like biometrics. Other vulnerabilities like lexicalbased ones and sensitive contextual information can also be protected using client-side transforms [20–22]. Since methods like changing language usage can mitigate these vulnerabilities much more easily than biometrics in speech, we focus on client-side privacy for downstream speech tasks here [9, 23].



Fig. 1: Privacy approaches. Unlike other algorithms, our clientside transforms are housed entirely on the client side and circumvent trade offs inherent in differential privacy.

B. Privacy-preserving Speech Processing

ASR models are getting larger and more powerful, making the case for putting them on the server side stronger [24]. Since the best ASR models are currently on the server side, measures must be taken to ensure user data sent to the server remains private. Early research on privacy for speech data has focused on voice encryption, which aims to make the original audio hard to recover from the encrypted data [25–27]. We note that these methods cannot be only on the client side since they require the receiver to decrypt the signal. Another more recent direction focuses on hiding speaker identity [28–34]. A range of methods relying on server-side operations have been proposed, including rearranging audio segments on the server [31] or leveraging server communication protocols [32]. Similar to research in voice encryption, these studies, to our knowledge, do not address the privacy of sensitive information beyond speaker identity like race, gender, or accent.

Speaker anonymization generally refers to approaches that hide speaker identity on the client side [29, 34, 35]. Namely, these works leverage voice conversion techniques to transform raw speech into that of another speaker [28, 29, 33-35]. Current approaches are predominantly neural, utilizing adversarial, disentanglement, or other encoder-decoder-related architectures [29, 30, 33-39]. Since voice conversion has already shown success in anonymizing speaker identity, our client side transforms in this work extend these ideas. Among related work, several address biometric information like gender [30, 38-42], but to our knowledge only two evaluate on complex downstream tasks like ASR [30, 42]. Since both of these works report high ASR error rates, namely word error rates (WER) above 60% on LibriSpeech [43], they are unable to maintain downstream performance while preserving privacy. In our paper, we study three distinct approaches that achieve lower WER while preserving privacy.

We focus on complex downstream tasks like ASR in this work since differential privacy or on-device approaches may be preferable for simpler tasks like classification [44–47]. In Section V, we also show that differential privacy is not suitable for ensuring privacy in downstream ASR tasks. Additionally, it is much easier to anonymize speaker identity than biometrics like gender, since the former generally requires a much smaller user data distribution shift than the latter [48]. Thus, in this paper, we study how well our client-side transforms can anonymize gender and accent, being the first to our knowledge to explore the latter. We note that speaker anonymization is a subtask of our client-side privacy task defined in Section I and detailed below.

III. CLIENT-SIDE PRIVACY

As defined in Section I, client-side privacy refers to privacy obtained only via on-device operations, which remove sensitive information while preserving content needed for downstream tasks. We proceed to formalize the problem statement in Section III-A and discuss the technical challenges in Section III-B.

A. Problem Statement

We start with a set of users \mathcal{U} each of which has access to a client-side device $m_u, u \in \mathcal{U}$. On each client-side device, data x_u is collected which potentially contains information about their private attributes y_u such as gender, age, race, or accent.

While it is ideal to leverage shared data collected at a large scale across users, it is also imperative to prevent leakage of private attributes y_u outside of the client's device. Therefore, the goal in client-side privacy is to learn an *encrypted* signal x'_u from x_u using a privacy-preserving function $f_{\theta} : x_u \to x'_u$ with parameters θ . f_{θ} should learn an encrypted signal x'_u that balances both fidelity and privacy:

1. *Fidelity*: Given a downstream model trained for a certain task (e.g., ASR) defined on the server, the performance of model on the encrypted signal x'_u should be as close as possible to that of the original signal x_u .

2. *Privacy*: One should not be able to decode the private attributes y_u from encrypted signals x'_u regardless of the function used to predict private attributes.

3. *Efficiency*: $|\theta|$ should be small and applying the encoding function $f_{\theta}: x_u \to x'_u$ should be fast for cheap inference and storage on resource-constrained mobile devices.

To avoid confounding factors, both evaluation models (ASR and private attribute classifier) are trained on data completely separate from those used to train encryption approaches. In this paper, we measure fidelity using ASR performance, namely character error rate (CER) and word error rate (WER), and privacy using gender and accent classification accuracy. In other words, low ASR error rate and low classification accuracy would indicate high fidelity and high privacy, respectively. Section IV contains the efficiency of each of our three approaches, and further details are described in Section V.

B. Technical Challenges

Client-side privacy for downstream speech tasks essentially requires one to *re-generate* raw audio with data-level private attributes masked out. This poses several compelling challenges: (1) it requires directly manipulating the user's audio [34, 35] and re-generating high-dimensional raw audio, (2) the encrypted audio must still be compatible with downstream server tasks without any modification on the server (i.e., for speech tasks, we require that the encrypted data is still comprehensible for downstream tasks), which requires preserving information at the high-dimensional data level rather than the feature level, and (3) methods for client-side privacy must have low time and space complexity to be compatible with limited-bandwidth client devices.

As a result, client-side privacy presents novel challenges over commonly studied server-side methods, particularly on fidelity and efficiency perspectives. Furthermore, it is much more challenging to preserve privacy for cluster-level attributes such as race, gender, and accent as compared to individuallevel attributes such as speaker identity. This is because transforming data across clusters requires a larger distribution shift than transformations to a new speaker, who could be in the same cluster [48].

IV. CLIENT-SIDE TRANSFORMS FOR DOWNSTREAM ASR

Given our definition in Section II, client-side transforms are one approach to obtaining client-side privacy as defined in Section I. Namely, client-side transforms anonymize sensitive information on-device while preserving content needed for downstream tasks. In this paper, we study three client-side transforms adapted from existing voice conversion literature and analyze their pros and cons.

A. Pitch Standardization

For our first client-side transform approach, we perform pitch standardization using signal processing [49, 50]. Specifically, we shift the average pitch of each utterance to a predefined value while preserving formants. For each utterance, we calculate its fundamental frequency (F_0) value and then perform a pitch shift from that value to a reference F_0 value. We calculate the F_0 sequence of an utterance using REAPER¹, and define the utterance F_0 value as the average of the nonnegative F_0 values. We then use the Rubber Band Library to perform pitch shifting with formant preservation using a phase vocoder.² For utterance u with F_0 value of f_u , we shift its pitch by $12 \log_2(f_r/f_u)$ semitones, where f_r is the reference F_0 value. This approach easily has the highest efficiency out of our three since it does not depend on a neural model.

B. Disentangled Representation Learning

Our second proposed approach uses variational autoencoders (VAE) to disentangle private attributes from non-private speech features [30, 42, 51]. VAEs allow us to learn a set of latent representations that best reconstruct a given input audio signal, while enforcing disentanglement into a set of speaker-dependent and speaker-independent factors [52]. We define a private attribute encoder $e(x; \theta_p)$ that encodes the input audio signal x into speaker-dependent private factors z_p and a content encoder $e(x; \theta_c)$ into speaker-independent content factors z_c . Since the private factors should capture the private attributes y using a classifier, our goal is to them exclude the private factor when decoding the encrypted signal. We optimize the following loss function:

$$\mathcal{L}_{\text{dis}} = \|d(e(x;\theta_c);\theta_d) - x\|_1 - \lambda_p \log P(y|e(x;\theta_c)) \quad (1) + \lambda_{\text{dis}} \text{KL}\left([e(x;\theta_c), e(x;\theta_p)] \parallel \mathcal{N}(0, I_d)\right), \quad (2)$$

where $d(z; \theta_d)$ is a decoder from latent space z back into audio space. The first term measures reconstruction of the input signal, the second term measures how well z_p captures the private attributes y, and the third term measures disentanglement of z_c and z_p by ensuring minimal correlated entries. λ_p and λ_{dis} are tunable hyperparameters controlling the tradeoff between privacy disentanglement and performance.

We train a convolutional VAE using the hyperparameters described in Chou et al. [51]. The model has log-magnitude spectrograms as its input and output acoustic features, and we use the Griffin-Lim algorithm to convert the model output into waveforms [53]. Instance normalization is also added to the content encoder $e(x; \theta_c)$ in order to remove speaker information and an adaptive instance normalization layer is added to the decoder in order to add the desired speaker

¹https://github.com/google/REAPER

²https://github.com/breakfastquay/rubberband

information [54, 55], which allows reconstruction of content while transforming speaker information. As far as we are aware, this architecture is considered fairly recent among the voice conversion literature [35, 36, 51]. Given the success of related architectures in anonymizing speaker identity [42, 56], we study this model's efficacy in anonymizing biometrics like gender and accent.

C. Adversarial Training

Reconstruction of high-dimensional signals is difficult and has been shown to cause poor generation quality [57–59]. Our final approach attempts to fix this by using adversarial training to ensure high-fidelity generation of encrypted audio [56, 60, 61]. In the first stage, we still assume disentanglement into of input audio signal x into speaker-dependent private factors z_p and speaker-independent content factors z_c learned using a VAE in Section IV-B. The second stages train a generative adversarial network (GAN) [62] to generate realistic audio. The generator is conditioned on the content factor z_c and a new speaker label. The discriminator predicts whether an audio sample was from the true dataset or from the generator, and also classifies the speaker. The loss function for the generator is

$$\log c_2(x) + \log(1 - c_2(g(x, y))) - \log P_{c'_2}(y|g(x)), \quad (3)$$

where g is the generator, and the loss function for the discriminator is

$$-\log c_2(x) - \log(1 - c_2(g(x, y))) - \log P_{c'_2}(y|x).$$
(4)

This model is suitable for our work because it explicitly separates speaker identity from content twice. We additionally modify this architecture by substituting the speaker label with either the gender label or the accent label. We refer to these three approaches as the speaker, gender, and accent adversarial approaches. While this architecture predates our disentangled one, we observe that our modified methods outperform the disentangled approach on multiple metrics, as detailed in Section V. We train a convolutional autoencoder using the hyperparameters described in Chou et al. [60]. The model has log-magnitude spectrograms as its input and output acoustic features, and we use the Griffin-Lim algorithm to convert the model output into waveforms [53]. Thus, we note that our reported results in Section V can be further improved using more complex vocoders [63].

V. EXPERIMENTS

Our experiments test whether our proposed approaches are able to balance the trade-offs in fidelity, privacy, and complexity required for client-side privacy. We test these approaches on masking gender and accents in speech recognition³.

A. Setup

Datasets: We train all of our encryption models on the VCTK corpus, and the ASR model on LibriSpeech [43, 64]. For both our VCTK and LibriSpeech experiments, we test on speakers unseen during training. We train our privacy attribute classifer on the respective dataset used during testing. In our LibriSpeech experiments, the ASR model and the gender classifier are both evaluated on the test-clean subset. In our VCTK experiments, we evaluate on a hold-out set of 20 speakers comprised of 10 males and 10 females.

Classifier: We use VGGVox model, a modified version of the VGG-16 CNN, as our privacy attribute classifier [65, 66], slightly modifying the network by adding a ReLU activation followed by a fully-connected layer with size-2 output. We approximate the data available to an adversary by training in two stages: 1. we pre-train the classifier on 100 hours of labeled, unmodified speech from the train data, and 2. we fine-tune the classifier on the encrypted speech of a handful of speakers from the same subset. We measure the privacy attribute masking ability of each encryption approach by calculating the VGGVox model's privacy attribute classification accuracy on an encrypted versions of the test subset after being finetuned on data encrypted using the respective approach.

ASR model: Unless mentioned otherwise, we use a pretrained ESPNet Transformer ASR model to evaluate downstream ASR performance [67]. This model was trained on 960 hours of LibriSpeech data.

B. Privacy-Fidelity Tradeoff

Table I compares the gender classification accuracy with the CER and WER for different levels of Gaussian noise added to the VCTK test data at the waveform level. We increment the standard deviation of the noise by 0.01 for each subsequent experiment. As expected, we observe a negative correlation between classification accuracy and ASR performance. In other words, these results reveal a tradeoff between privacy and fidelity.

TABLE I: Tradeoff between privacy and fidelity on the VCTK test set for different levels of added Gaussian noise. We observe a negative correlation between classification accuracy and ASR performance, as expected.

Noise	Classification	CER	WER
0.00	0.99	4.5	9.5
0.01	0.91	11.2	26.7
0.02	0.79	19.3	31.9
0.03	0.68	28.5	41.0
0.04	0.61	33.9	48.1

C. Gender Classification

Table II describes the gender classification accuracy on VCTK and LibriSpeech using gender-masked audio. All proposed approaches can successfully mask gender when no encrypted training examples are available. However, given encrypted training data from a male and female speaker, the

³Our code and models are publicly available at https://github.com/peter-yhwu/speech-privacy.

signal processing samples become much easier to classify than those generated from the other approaches. This makes sense as pitch shifting would retain underlying speaker qualities like that could be readily identified by a neural classifier. Also, the adversarial approach using the gender-based loss outperforms the speaker-based loss approach here, which reflects how the former explicitly learns to hide gender information.

TABLE II: Gender classification accuracy on two datasets using gender-masked audio. The integer n in each column denotes the number of speakers whose encrypted audio was used to finetune the classifier, where n/2 are male and n/2 are female. All our voice conversion approaches perform better than the baseline without any masking for n = 0. For higher n values, the disentanglement approach performs the best, followed by the adversarial approached with the modified gender loss.

VCTK	0	2	4	20
No Masking	0.991	0.991	0.991	0.991
Signal Processing	0.590	0.987	0.990	0.997
Disentanglement	0.590	0.590	0.598	0.757
Adversarial (Speaker)	0.591	0.824	0.840	0.935
Adversarial (Gender)	0.590	0.707	0.740	0.877
LibriSpeech	0	2	4	20
No Masking	0.972	0.972	0.972	0.972
Signal Processing	0.422	0.869	0.887	0.933
Disentanglement	0.590	0.627	0.702	0.781
Adversarial (Speaker)	0.580	0.712	0.714	0.838
Adversarial (Gender)	0.570	0.621	0.628	0.833

D. ASR after Gender Encryption

Table III describes the ASR results when transcribing gender-encrypted data, measured using mean character and word error rates. For each approach, we provide results on both VCTK and LibriSpeech [43, 64]. All ASR models are pretrained on LibriSpeech as aforementioned. The VCTK model is further finetuned on the data from speakers outside our test set. All finetuned models are tuned on the respective converted data of 20 train speakers. The signal processing approach performs the best, potentially since Griffin-Lim and output distribution priors inherent in neural network architectures introduce artifacts. The adversarial approach using the genderbased loss again outperforms the speaker-based loss approach. This reflects how the former can model less style information than the latter and thus can model more content. For the LibriSpeech dataset, our adversarial method with the modified gender loss performs similarly to the disentangled method in the finetuned scenario. Moreover, our disentanglement and adversarial approaches do not improve when finetuned for the VCTK experiments. This suggests that the ASR model may not be robust enough or our neural converted samples may be acting like adversarial samples during the training procedure [68]. We also note that, compared to Table I, our voice conversion approaches generally achieve lower WER for fixed privacy performances. This suggests that our client-side

transforms described in Section IV are more suitable for clientside privacy than other approaches like differential privacy.

TABLE III: ASR performance on two datasets using gendermasked audio. Among our voice conversion approaches, the signal processing method performs the best. For the LibriSpeech dataset, our adversarial method with the modified gender loss performs similarly to the disentangled method in the finetuned scenario. Moreover, our disentenglement and adversarial approaches do not improve when finetuned for the VCTK experiments. This suggests that the ASR model may not be robust enough or our neural converted samples may be acting like adversarial samples during the training procedure [68].

	VCTK			
	CER		v	VER
Method	0-Shot	Finetune	0-Shot	Finetune
No Masking	4.5	3.4	9.5	4.8
Signal Processing	15.0	7.7	24.7	9.8
Disentanglement	21.1	21.1	35.0	35.0
Adversarial (Speaker)	31.1	31.1	48.5	48.5
Adversarial (Gender)	25.0	25.0	40.1	40.1
		Libris	Speech	
	CER		WER	
Method	0-Shot	Finetune	0-Shot	Finetune
No Masking	2.4	2.4	4.6	4.6
Signal Processing	5.0	5.0	8.8	8.8
Disentanglement	15.5	15.5	25.0	25.0
-				
Adversarial (Speaker)	29.7	17.5	47.5	28.0

E. Gender Listening Tests

In addition to our automatic metrics, we perform mean opinion score (MOS) preference tests using eight human listeners. Namely, we compare the signal processing, the disentanglement, and the adversarial gender-based loss approaches for the VCTK gender encryption task. For the MOS test, we ask listeners to rate audio samples on a naturalness scale of 1 to 5. We use 40 utterances for each test, where 2 are randomly chosen from each test speaker. In other words, each listener listens to 120 unique audio clips. Table IV summarizes these results. The signal processing approach performs the best for female speakers and the worst for male speakers. This is likely due to the reference F0 being from a female speaker and the relative absence of artifacts. Also, while the disentanglement approach outperforms the adversarial one in both the classification and ASR metrics, listeners consistently rated the latter higher. This suggests that the disentanglement approach may be standardizing the audio to waveforms that are unnatural to people but suitable for downstream ASR systems. We perceive such utterances as robotic but understandable.

F. Accent Classification

Table V describes the accent classification accuracy on VCTK using accent-masked audio. Given that the largest class

TABLE IV: Listening test results on VCTK gender encryption approaches. Our signal processing and adversarial approaches perform the best. While our signal processing approach performs the best for female speakers, our adversarial approach does the best for males, suggesting that the latter is more robust to different speaker attributes.

MOS	М	F	Both
Signal Processing Disentanglement Adversarial (Gender)	$\begin{array}{c} 1.8 \pm 0.2 \\ 2.4 \pm 0.5 \\ 2.7 \pm 0.6 \end{array}$	$\begin{array}{c} 4.3 \pm 0.2 \\ 2.4 \pm 0.5 \\ 3.7 \pm 0.3 \end{array}$	$3.3 \pm 0.2 \\ 2.4 \pm 0.5 \\ 3.3 \pm 0.4$

in the test set contains 31% of the samples, we observe that both the disentanglement and adversarial approaches are able to successfully fool the accent classifier. Moreover, our adversarial approach using the modified accent loss performs the best, indicating the usefulness of our modified loss function.

TABLE V: Automatic metric results for accent encryption approaches. Given that the largest class in the test set has 31% of the samples, our results indicate that both our disentanglement and adversarial approaches successfully fooled the accent classifier. Moreover, our adversarial approach using the modified accent loss performs the best.

Classification	0	20
Largest Class	0.31	0.31
No Masking	0.36	0.36
Disentanglement	0.29	0.29
Adversarial (Speaker)	0.25	0.25
Adversarial (Gender)	0.25	0.25
Adversarial (Accent)	0.23	0.23

G. ASR after Accent Encryption

Table VI describes the ASR results when transcribing gender-encrypted data, measured using mean character and word error rates. Our experimental setup here follows that of the gender ASR experiment. We observe trends similar to Section V-D. Additionally, ASR results here are consistently better than those in the gender encryption task. This reflects how transforming gender requires a larger data distribution shift than transforming accent.

TABLE VI: Automatic metric results for accent encryption approaches. As with our other ASR experiments, we observe that our disentanglement approach outperforms our adversarial one. We also note that these ASR results are consistently better than those for the gender experiment, reflecting the larger data distribution shift required for transforming gender.

Speech Recognition	CER	WER
Disentanglement	17.5	29.9
Adversarial (Speaker)	26.5	42.4
Adversarial (Gender)	19.5	32.4
Adversarial (Accent)	23.1	37.4

VI. KEY TAKEAWAYS

In this section, we outline several key takeaways from our experimental results which we hope will help practitioners apply privacy-preserving methods for client-side privacy.

1. Pitch standardization approaches unfortunately are not very effective in keeping gender private: 78.6% gender classification accuracy for n = 0, and 95.1% using pitch shift, both of which are much higher than those of the neural methods. They yield artifacts that were readily recognizable by the adversary gender classifier, which implies poorer performance in maintaining privacy. Namely, when observing the attention map of the classifier, we noticed that the classifier learned to identify specific patterns that resulted from the pitch shift. Thus, the gender classification accuracy of the neural methods was much lower than those of the signal processing methods.

2. VAEs and GANs: VAEs, through use of an encoder, are suitable to learn latent disentangled representations [69, 70] which are useful in our task of disentangling content from private attributes. GANs are also suitable for learning latent representations. While they have been used less in the disentanglement literature, adding attribute-specific loss functions can disentangle sensitive information from content well. While we found GANs to be harder to train than VAEs, GANs that converge appear to perform better.

3. *Memory*: The large differences in memory consumption are a consequence of the large memory costs of using neural models compared to signal processing approaches. Overall, our conclusions point out a ripe opportunity for future work to reconcile the privacy benefits of neural methods with the performance and memory advantages of signal processing approaches.

VII. CONCLUSION AND FUTURE DIRECTIONS

In this work, we setup the problem of achieving clientside privacy for speech representation learning that does not rely on any server-side privacy guarantees. We formalized several desirable properties regarding performance, privacy, and computation and performed a large-scale empirical study of existing approaches. We find that while GAN-based approaches currently have the best tradeoff between gender masking, downstream performance, and memory usage, all existing approaches still fall short of ideal performance. Our initial empirical analysis opens the door towards more reliable evaluations of the tradeoffs underlying privacy-preserving approaches on the client-side, a property crucial for safe realworld deployment of speech systems at scale across mobile devices. In addition to developing privacy-preserving algorithms that satisfy the various desiderata as outlined in this paper, future work should also analyze other downstream speech tasks besides speech recognition, such as speech translation.

REFERENCES

 Y. Chen *et al.*, "Fedhealth: A federated transfer learning framework for wearable healthcare," *IEEE Intelligent Systems*, 2020.

- [2] R. C. Geyer, T. Klein, and M. Nabi, "Differentially private federated learning: A client level perspective," *arXiv preprint arXiv*:1712.07557, 2017.
- [3] P. P. Liang *et al.*, "Think locally, act globally: Federated learning with local and global representations," 2020.
- [4] J. Xu and F. Wang, "Federated learning for healthcare informatics," *Journal of Healthcare Informatics Research*, pp. 1 – 19, 2021.
- [5] I. McGraw *et al.*, "Personalized speech recognition on mobile devices," *ICASSP*, pp. 5955–5959, 2016.
- [6] X. Gong *et al.*, "Model extraction attacks and defenses on cloud-based machine learning models," *IEEE Communications Magazine*, vol. 58, no. 12, pp. 83–89, 2020.
- [7] N. Anggraini, A. Kurniawan, L. K. Wardhani, and N. Hakiem, "Speech recognition application for the speech impaired using the android-based google cloud speech api," *Telkomnika*, vol. 16, no. 6, pp. 2733–2739, 2018.
- [8] D. Leroy *et al.*, "Federated learning for keyword spotting," *ICASSP*, pp. 6341–6345, 2019.
- [9] R. Singh, Profiling Humans from their Voice, 01 2019.
- [10] J. Jing, A. S. Helal, and A. Elmagarmid, "Client-server computing in mobile environments," *CSUR*, vol. 31, no. 2, pp. 117–157, 1999.
- [11] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated learning: Challenges, methods, and future directions," *IEEE Signal Processing Magazine*, vol. 37, no. 3, pp. 50–60, 2020.
- [12] S. Caldas *et al.*, "Leaf: A benchmark for federated settings," *Workshop on Federated Learning for Data Privacy and Confidentiality*, 2018.
- [13] J.-J. Hwang, H.-K. Chuang, Y.-C. Hsu, and C.-H. Wu, "A business model for cloud computing based on a separate encryption and decryption service," in 2011 International Conference on Information Science and Applications, 2011, pp. 1–7.
- [14] C. Zhao *et al.*, "Secure multi-party computation: Theory, practice and applications," *IS*, vol. 476, pp. 357–372, 2019.
- [15] X. Sun *et al.*, "Private machine learning classification based on fully homomorphic encryption," *IEEE Transactions on Emerging Topics in Computing*, vol. 8, no. 2, pp. 352–364, 2018.
- [16] N. Papernot *et al.*, "Scalable private learning with PATE," *ICLR*, 2018.
- [17] M. Abadi et al., "Deep learning with differential privacy," in Proceedings of the 2016 ACM SIGSAC conference on computer and communications security, 2016, pp. 308– 318.
- [18] E. De Cristofaro, "A critical overview of privacy in machine learning," *IEEE Security Privacy*, vol. 19, no. 4, pp. 19–27, 2021.
- [19] D. C. Wilson and G. Ateniese, ""to share or not to share" in client-side encrypted clouds," in *ICIS*. Springer, 2014, pp. 401–412.
- [20] S. Silessi, C. Varol, and M. Karabatak, "Identifying

gender from sms text messages," in ICMLA, 2016, pp. 488-491.

- [21] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, "Membership inference attacks against machine learning models," in 2017 IEEE Symposium on Security and Privacy (SP), 2017, pp. 3–18.
- [22] A. Caliskan Islam, J. Walsh, and R. Greenstadt, "Privacy detective: Detecting private information and collective privacy behavior in a large social network," in *Proceedings of the 13th Workshop on Privacy in the Electronic Society*, 2014, pp. 35–46.
- [23] T. Sun *et al.*, "Mitigating gender bias in natural language processing: Literature review," in ACL, 2019, pp. 1630– 1640.
- [24] A. Baevski, H. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," *NeurIPS*, 2020.
- [25] S. C. Kak and N. Jayant, "On speech encryption using waveform scrambling," *Bell System Technical Journal*, vol. 56, no. 5, pp. 781–808, 1977.
- [26] S. Sridharan, E. Dawson, and B. Goldburg, "Fast fourier transform based speech encryption system," *IEE Proceedings I-Communications, Speech and Vision*, vol. 138, no. 3, pp. 215–223, 1991.
- [27] P. Smaragdis and M. Shashanka, "A framework for secure speech recognition," *TASLP*, vol. 15, no. 4, pp. 1404– 1413, 2007.
- [28] Q. Jin, A. R. Toth, T. Schultz, and A. W. Black, "Speaker de-identification via voice transformation," in ASRU. IEEE, 2009, pp. 529–533.
- [29] N. Tomashenko *et al.*, "Introducing the voiceprivacy initiative," *Interspeech*, 2020.
- [30] R. Aloufi, H. Haddadi, and D. Boyle, "Privacy-preserving voice analysis via disentangled representations," ACM SIGSAC, 2020.
- [31] S. Ahmed, A. R. Chowdhury, K. Fawaz, and P. Ramanathan, "Preech: A system for privacy-preserving speech transcription," in USENIX. USENIX Association, Aug. 2020, pp. 2703–2720.
- [32] Z. Ma, Y. Liu, X. Liu, J. Ma, and F. Li, "Privacypreserving outsourced speech recognition for smart iot devices," *IEEE Internet of Things Journal*, vol. 6, no. 5, pp. 8406–8420, 2019.
- [33] A. Nautsch *et al.*, "Preserving privacy in speaker and speech characterisation," *CSL*, vol. 58, pp. 441–480, 2019.
- [34] B. M. L. Srivastava *et al.*, "Design choices for x-vector based speaker anonymization," *Interspeech*, 2020.
- [35] W.-C. Huang, T. Hayashi, S. Watanabe, and T. Toda, "The sequence-to-sequence baseline for the voice conversion challenge 2020: Cascading asr and tts," *arXiv* preprint arXiv:2010.02434, 2020.
- [36] B. Sisman, J. Yamagishi, S. King, and H. Li, "An overview of voice conversion and its challenges: From statistical modeling to deep learning," *TASLP*, vol. 29, pp. 132–157, 2021.

- [37] L. Zheng *et al.*, "When automatic voice disguise meets automatic speaker verification," *IEEE Transactions on Information Forensics and Security*, vol. 16, pp. 824– 837, 2020.
- [38] P.-G. Noé *et al.*, "Adversarial disentanglement of speaker representation for attribute-driven privacy preservation," *arXiv preprint arXiv:2012.04454*, 2020.
- [39] D. Ericsson *et al.*, "Adversarial representation learning for private speech generation," *arXiv preprint arXiv:2006.09114*, 2020.
- [40] J. Chen *et al.*, "Using keyword spotting and replacement for speech anonymization," in *ICME*. IEEE, 2007, pp. 548–551.
- [41] K. Kondo and H. Sakurai, "Gender-dependent babble maskers created from multi-speaker speech for speech privacy protection," in 2014 Tenth International Conference on Intelligent Information Hiding and Multimedia Signal Processing, 2014, pp. 251–254.
- [42] D. Stoidis and A. Cavallaro, "Protecting gender and identity with disentangled speech representations," *arXiv preprint arXiv:2104.11051*, 2021.
- [43] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *ICASSP*. IEEE, 2015, pp. 5206–5210.
- [44] Z. Ji, Z. C. Lipton, and C. Elkan, "Differential privacy and machine learning: a survey and review," arXiv preprint arXiv:1412.7584, 2014.
- [45] P. Wu, S. K. Rallabandi, A. W. Black, and E. Nyberg, "Ordinal triplet loss: Investigating sleepiness detection from speech." in *Interspeech*, 2019, pp. 2403–2407.
- [46] P. Wu, Y. Zhong, and A. W. Black, "Automatically identifying language family from acoustic examples in low resource scenarios," *arXiv preprint arXiv:2012.00876*, 2020.
- [47] P. P. Liang, P. Wu, L. Ziyin, L.-P. Morency, and R. Salakhutdinov, "Cross-modal generalization: Learning in low resource modalities via meta-alignment," *arXiv* preprint arXiv:2012.02813, 2020.
- [48] Y. Wang, J. Du, L.-R. Dai, and C.-H. Lee, "A gender mixture detection approach to unsupervised singlechannel speech separation based on deep neural networks," *TASLP*, vol. 25, no. 7, pp. 1535–1546, 2017.
- [49] A. Mousa, "Voice conversion using pitch shifting algorithm by time stretching with psola and re-sampling," *Journal of electrical engineering*, vol. 61, no. 1, p. 57, 2010.
- [50] R. H. Laskar, K. Banerjee, F. A. Talukdar, and K. S. Rao, "A pitch synchronous approach to design voice conversion system using source-filter correlation," *International Journal of Speech Technology*, vol. 15, no. 3, pp. 419– 431, 2012.
- [51] J.-c. Chou and H.-y. Lee, "One-shot voice conversion by separating speaker and content representations with instance normalization," in *Interspeech*, 2019.
- [52] C. P. Burgess *et al.*, "Understanding disentangling in βvae," *arXiv preprint arXiv:1804.03599*, 2018.

- [53] D. Griffin and Jae Lim, "Signal estimation from modified short-time fourier transform," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 32, no. 2, pp. 236–243, 1984.
- [54] D. Ulyanov, A. Vedaldi, and V. Lempitsky, "Instance normalization: The missing ingredient for fast stylization," *arXiv preprint arXiv:1607.08022*, 2016.
- [55] X. Huang and S. Belongie, "Arbitrary style transfer in real-time with adaptive instance normalization," in *ICCV*, 2017, pp. 1510–1519.
- [56] I.-C. Yoo *et al.*, "Speaker anonymization for personal information protection using voice conversion techniques," *IEEE Access*, vol. 8, pp. 198 637–198 645, 2020.
- [57] P. Dhariwal *et al.*, "Jukebox: A generative model for music," *arXiv preprint arXiv:2005.00341*, 2020.
- [58] S. Rallabandi, P. Wu, and A. W. Black, "Submission from cmu for blizzard challenge 2019," *Proceedings of Blizzard Challenge*, vol. 2019, 2019.
- [59] M. Ranzato *et al.*, "Video (language) modeling: a baseline for generative models of natural videos," *arXiv preprint arXiv*:1412.6604, 2014.
- [60] J.-c. Chou, C.-c. Yeh, H.-y. Lee, and L.-S. Lee, "Multitarget voice conversion without parallel data by adversarially learning disentangled audio representations," in *Interspeech*, 2018.
- [61] K. Zhou, B. Sisman, R. Liu, and H. Li, "Seen and unseen emotional style transfer for voice conversion with a new emotional speech dataset," in *ICASSP*, 2021, pp. 920– 924.
- [62] I. Goodfellow *et al.*, "Generative adversarial nets," *NIPS*, vol. 27, 2014.
- [63] A. v. d. Oord *et al.*, "Wavenet: A generative model for raw audio," in *SSW*, 2016.
- [64] C. Veaux, J. Yamagishi, K. MacDonald *et al.*, "CSTR VCTK corpus: English multi-speaker corpus for cstr voice cloning toolkit," 2016.
- [65] A. Nagrani, J. S. Chung, and A. Zisserman, "Voxceleb: A large-scale speaker identification dataset," in *Interspeech*, 2017.
- [66] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *ICLR*, 2015.
- [67] S. Watanabe *et al.*, "ESPnet: End-to-end speech processing toolkit," in *Interspeech*, 2018, pp. 2207–2211.
- [68] I. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," in *ICLR*, 2015.
- [69] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner, "betavae: Learning basic visual concepts with a constrained variational framework," *ICLR*, 2016.
- [70] F. Locatello, S. Bauer, M. Lucic, G. Raetsch, S. Gelly, B. Schölkopf, and O. Bachem, "Challenging common assumptions in the unsupervised learning of disentangled representations," in *ICML*. PMLR, 2019, pp. 4114– 4124.