



Generating gender-ambiguous voices for privacy-preserving speech recognition

Dimitrios Stoidis, Andrea Cavallaro

Centre for Intelligent Sensing, Queen Mary University of London, UK

dimitrios.stoidis@qmul.ac.uk, a.cavallaro@qmul.ac.uk

Abstract

Our voice encodes a uniquely identifiable pattern which can be used to infer private attributes, such as gender or identity, that an individual might wish not to reveal when using a speech recognition service. To prevent attribute inference attacks alongside speech recognition tasks, we present a generative adversarial network, GenGAN, that synthesises voices that conceal the gender or identity of a speaker. The proposed network includes a generator with a U-Net architecture that learns to fool a discriminator. We condition the generator only on gender information and use an adversarial loss between signal distortion and privacy preservation. We show that GenGAN improves the trade-off between privacy and utility compared to privacy-preserving representation learning methods that consider gender information as a sensitive attribute to protect.

Index Terms: audio privacy, speech recognition, speaker verification, gender recognition, generative adversarial networks

1. Introduction

The human voice is shaped by the physical characteristics of the speaker and the spoken language. The voiceprint, which uniquely defines each individual [1], contains attributes of the speaker that can be inferred by voice-based services [2, 3]. To mitigate this risk and protect the identity of a speaker in Automatic Speech Recognition (ASR) tasks, adversarial training on privacy and utility objectives can be used [4, 5]. Adversarial representation learning determines the network weights that minimise the likelihood of finding the correct labels corresponding to identity (or gender) and removes information related to these attributes from the encoded representation [4]. Training with a speaker-adversarial branch acting as a gradient reversal layer has been used to remove speaker identity from the learned representation [6]. Adversarial feature extraction is used to improve the privacy-utility trade-off, when gender classification is considered as the utility objective and speaker identification accuracy as the privacy objective [5].

Gender information is typically used to condition models preserving the identity of a speaker. However, only a handful of methods explicitly consider gender as a sensitive attribute to protect [7, 8, 9, 10, 11]. A hybrid model combining Variational Autoencoders and Generative Adversarial Networks (GANs) can be used to protect gender information through voice conversion with a disentanglement approach targeted for the speech recognition task [7]. Two encoders are trained to independently encode content and speaker identity information that is then used to hide (or mask) gender information. Privacy methods that operate at feature-level have been used to disentangle gender information from x-vectors [12] with adversarial learning and an encoder-decoder based architecture [8]. Because this adversarial method removes the unwanted information at the level of the feature representation instead of the speech waveform,

it is not useful for tasks such as speech recognition. Vector Quantized Variational Autoencoders (VQ-VAE) [13] are used to create disentangled representations of linguistic content and speaker identity [9], and both identity and gender [11]. Linguistic content is quantised into a discrete latent space using a learned codebook. The decoding stage reconstructs the speech waveform by combining learned embedding spaces encoding the selected attributes. These methods have limited reconstruction capabilities and induce distortion at the decoding stage, when quantised content information is reconstructed as speech. PCMelGAN [10] synthesises speech using a generative adversarial approach that considers gender as an attribute to protect and reconstructs mel-scale spectrograms using the MelGAN vocoder [14] to maintain the utility (accuracy) in a digit recognition task. However, the dataset used by this work is composed of utterances of spoken digits, which is limited in vocabulary and size. Finally, PCMelGAN is based on PCGAN [15], which uses a filtering module to replace the sensitive information in speech with generated synthetic information. However, we will show that results can be improved *without* this additional process.

We aim to produce gender-ambiguous [16] voices. To this end we first determine the extent to which modifying gender information impacts the identity information of a speaker from a privacy-preservation perspective. Next, we produce a lightweight generative method that protects against the inference of gender and identity of synthesised speech signals. We achieve this without considering any identity information and maintain the utility of speech without explicitly optimising for the ASR task.

We propose GenGAN, a generative privacy transformation method that conceals gender and much of the identity information of speakers by synthesising voices with gender-ambiguous characteristics. To the best of our knowledge GenGAN is the first attempt to create gender-ambiguous voices in a privacy-preserving setting. The generator samples from a designed distribution that models a gender-ambiguous voice and learns to smooth spectral differences between genders. GenGAN is only conditioned on gender information, without considering any information on speaker identity during training. Furthermore, content information is preserved in the transformed signal independently of the ASR used for the speech transcription task.

2. Proposed approach

2.1. Attack scenario

Let the identity and gender of a speaker be the personal information to protect in an ASR scenario. An attacker attempts to infer, from the privacy-transformed audio, the gender of the speaker by classification and their identity by verification.

We assume that the attacker has access to the data (anonymised utterances) produced by the privacy-preserving




Figure 1: Evaluation of our privacy transformation method on utility and privacy objectives. The transformed audio is used to perform privacy-preserving speech recognition. The privacy objective assesses the ability of an attacker to recognise the gender and verify the identity of a speaker given the transformed gender-ambiguous speech signal. KEY – \mathcal{X} : original waveform, \mathcal{X}' : transformed waveform, \mathcal{Y} : ground-truth gender, \mathcal{Y}' : transformed gender, ASR: automatic speech recognition.

transformation and shared with the speech recognition service (see Fig. 1). We also assume the attacker has no knowledge of the applied privacy-transformation.

2.2. GenGAN

By assuming an attack scenario where voice signals are accessed prior to the downstream task (speech recognition), we are interested in reproducing voice and hence operate on the input waveforms, which are converted into 80-band mel-spectrograms. We consider a Generator, \mathcal{G} , and a Discriminator, \mathcal{D} , to be trained in an adversarial game between privacy and utility objectives. \mathcal{G} has a U-Net architecture [17] with a contracting and an expanding path, while \mathcal{D} consists of a simple AlexNet [18] architecture. \mathcal{G} produces audio data with the aim to fool \mathcal{D} , whose task is to discriminate between original and synthetic audio data. We maximise utility by minimising the distortion in the generated audio and minimise the risk of privacy inference by training the model to learn gender-ambiguous information. \mathcal{D} learns to discriminate between true and generated gender data, conditioned only on gender information. By conditioning only on gender information, we aim to distort the generated voice to protect both gender and identity.

The Generator, \mathcal{G} , takes as input mel-spectrograms \mathcal{M} , a noise vector \mathcal{Z} and the labels of the sensitive attribute to protect \mathcal{Y} , and synthesises the transformed spectrograms \mathcal{M}' (see Fig. 2). During training, a batch of n audio signals, \mathcal{X} , and their corresponding labels, \mathcal{Y} , representing the sensitive attribute (gender) are sampled uniformly at random from the dataset:

$$(x_1, y_1), \dots, (x_n, y_n) \sim \mathcal{X}_{train}. \quad (1)$$

The audio signals \mathcal{X} are converted to mel-spectrograms $m_i \in \mathcal{M} \forall i = 1, \dots, n$, and normalized such that amplitudes are bounded in $[0, 1]$ with

$$m_i = \mathcal{FFT}(x_i), \quad (2)$$

where \mathcal{FFT} is the Short-Time Fourier Transform.

As we consider a binary encoding for gender, we propose to sample from a synthetic distribution

$$\{\hat{y}_1, \dots, \hat{y}_n\} \in \mathcal{Y}_N \sim \mathcal{N}(0.5, 0.05). \quad (3)$$

As $\mathcal{N}(0.5, 0.05)$ is centred around 0.5 (i.e. equidistant from the ground-truth labels), \mathcal{G} learns to smooth feature-level differences between the spectral representations of the two gender labels and synthesises a new voice that is gender-ambiguous.

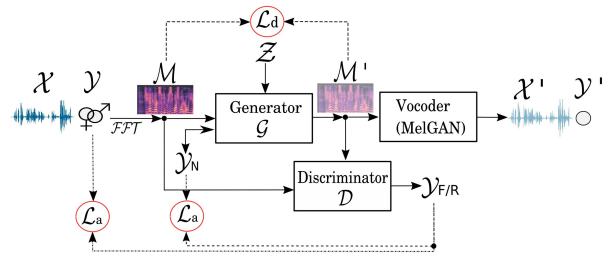


Figure 2: Privacy transformation and training procedure with the losses we use for GenGAN. The time-signal waveform \mathcal{X} is first converted into mel-spectrogram and then fed as input to the Generator along with the noise vector $\mathcal{Z} \sim \mathcal{N}(0, 1)$ and the synthetic non-binary gender label \mathcal{Y}_N . The generated mel-spectrogram \mathcal{M}' is fed to the Discriminator to predict if the sample corresponds to the ground-truth binary gender or non-binary gender. We use the MelGAN pre-trained vocoder to perform the mel-spectrogram inversion. KEY – \mathcal{X} : original waveform, \mathcal{X}' : transformed waveform, \mathcal{M} : original mel-spectrogram, \mathcal{M}' : transformed mel-spectrogram, $\mathcal{Y}(\varnothing)$: ground-truth gender label, \mathcal{Y}_N : non-binary gender vector, $\mathcal{Y}'(\circ)$: synthesised gender, $\mathcal{Y}_{F/R}$: predicted gender with Discriminator from real (original) and fake (generated) spectrograms, \mathcal{L}_d : distortion loss, \mathcal{L}_a : adversarial loss, \mathcal{FFT} : Fast Fourier Transform.

We select a small distribution variance ($\sigma^2 = 0.05$) for the synthetic distribution to minimise the overlap with the gender distributions. The noise vector $\mathcal{Z} \sim \mathcal{N}(0, 1)$ is inserted at the bottleneck of \mathcal{G} , at the transition of the contracting to the expansive path to ensure synthesised voices are different from the original ones, increasing reconstruction variability. \mathcal{Z} is reshaped and concatenated with the last convolutional layer ending the contracting path in UNet before being expanded back to a mel-spectrogram \mathcal{M}' .

The generator loss \mathcal{L}_G is computed as an adversarial game between reducing distortion on utility and maximising privacy by learning a non-binary gender. We take the Mean Squared Error between \mathcal{M} and \mathcal{M}' as distortion loss \mathcal{L}_d over the L_1 distance used in PCMelGAN to produce a smoother output. The adversarial loss \mathcal{L}_a is a cross-entropy loss between ground-truth gender \mathcal{Y} and predicted gender \mathcal{Y}_F , maximising the log probability of generating samples drawn from the synthetic distribution. The generator loss is the sum

$$\mathcal{L}_G = \mathcal{L}_d(\mathcal{M}, \mathcal{M}') + \epsilon \mathcal{L}_a(\mathcal{Y}, \mathcal{Y}_F), \quad (4)$$

where $\epsilon \in [0, 1]$ represents the dis-utility budget [19] in the privacy-utility trade-off.

The discriminator loss function is composed of two losses with respect to real \mathcal{M} or generated spectrogram \mathcal{M}' . The real loss is taken between the prediction of \mathcal{D} on real data \mathcal{Y}_R , and ground-truth gender \mathcal{Y} . The fake loss is taken between the prediction of \mathcal{D} on generated data \mathcal{Y}_F , and non-binary gender \mathcal{Y}_N . Both losses are computed as cross-entropies:

$$\mathcal{L}_D = \mathcal{L}_a(\mathcal{Y}, \mathcal{Y}_R) + \mathcal{L}_a(\mathcal{Y}_N, \mathcal{Y}_F). \quad (5)$$

Finally, to transform the generated spectrogram \mathcal{M}' back to raw waveform \mathcal{X}' , we use the MelGAN vocoder [14], a non-autoregressive conditional waveform synthesis model.

Figure 3 shows sample spectrograms of different utterances spoken by a male and by a female speaker. The bottom

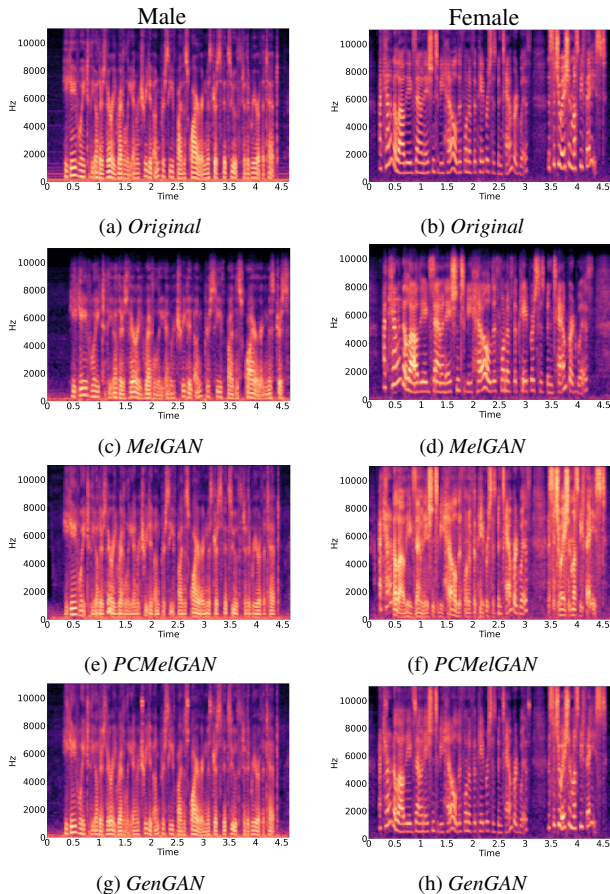


Figure 3: Spectrograms of two different utterances spoken by a male (left column) and a female speaker (right column). The privacy-preserving transformation affects the high formant frequencies associated with gender (see (e), (f), (g), (h)). (c) and (d) show the distortion induced on high frequencies by the MelGAN vocoder with respect to the original utterances.

row shows the privacy-transformed spectrograms produced by GenGAN. GenGAN synthesises higher resolution spectrograms than PCMelGAN [10], while both methods affect high formant frequencies associated with gender identification [20], smoothing out differences between male and female voices.

3. Validation

In this section we perform the validation of GenGAN in comparison with existing attribute-protecting privacy-preserving methods. We use the LibriSpeech dataset [21], which contains over 100 hours of transcribed clean speech and over 250 different speakers with balanced gender distribution.

3.1. Performance measures

We now describe the measures we use to compare the performance of different models in the scenario presented in Section 2.1.

We view speaker *identity* in the context of a speaker verification task, which consists in determining whether a trial utterance corresponds to the utterance of a certain, enrolled speaker. We chose speaker verification instead of speaker identification

as in our scenario the attacker uses a speaker verification system to recover the identity of a speaker from anonymised speech. We consider the equal error rate (EER) to assess the speaker verification performance. In biometrics verification systems, a high accuracy corresponds to a low EER, as a higher rejection over false acceptance rate is desired. To assess privacy, the target randomness in verifying the speaker identity [22] corresponds to an EER around 50%.

We measure success in *gender* inference as the standard binary classification accuracy, where the sum of correct predictions is taken over the total number of predictions. As in speaker verification, the target accuracy is 50%, which represents randomness in gender prediction. We compute the discrepancy from the 50% randomness value both on gender recognition and speaker verification and introduce a new measure. The measure normalises the EER and the gender recognition accuracy (GR) such that the absolute difference from the 50% accuracy is retained. The normalised Gender Recognition (gr) and normalised Equal Error Rate (eer) perform conversions given by $gr = 100 - 2 \times |GR - 50|$, and $eer = 100 - 2 \times |EER - 50|$. A value of 100 for gr (or eer) denotes the highest level of privacy.

We use the transcription accuracy of the speech recognition results as the utility measure. The Word Error Rate (WER) represents the utility of the spoken signal and is computed by taking the Levenshtein [23] distance between the words of the correct, expected transcription and the transcription provided by the speech recognition service. To facilitate the comparison between privacy and utility, rather than the error rate we consider the correct recognition rate with the use of the Word Accuracy metric (A_w) [11, 24], where $A_w = 100 - WER$. $A_w = 100$ denotes no transcription errors.

3.2. Classifiers

For *speech recognition*, we use Quartznet [25], an end-to-end neural acoustic model trained with Connectionist Temporal Classification (CTC) [26] loss, based on the Jasper architecture [27]. Our testing reported an initial performance on the LibriSpeech test-clean set of A_w of 95.64% (or 4.36% WER).

For *speaker verification*, we extract speaker-identifying features with a modified version of ResNet-34 [28] with reduced computational cost, which we refer to as SpeakerNet [29], pre-trained on VoxCeleb2 [30] dataset and reported an EER of 5.73% in our experiments when tested on the LibriSpeech test-clean set.

For *gender classification*, we use a deep convolutional binary classifier trained on spectrograms with 5 stacked one-dimensional convolutional layers followed by batch normalisation and max pooling [11], which we refer to as GenderNet. The output layer is a fully connected layer that outputs a pair of predictions for each binary class, passing through a sigmoid function. We tested the gender classifier on the LibriSpeech clean test set and reported an accuracy of 91.37%.

3.3. Methods under comparison

We compare methods that consider gender as an attribute to protect in an attribute inference scenario. VQ-VAE-Gen [11] considers gender information as private and assesses the impact of the privacy-transformation on speaker identity on LibriSpeech. The Client-Hybrid [7] model protects gender information from inference by using speaker identity information and provides results on the same dataset. We also compare our method with PCMelGAN [10] anonymisation method which uses a filtering

Table 1: Comparison of privacy and utility results for various models on the Librispeech [21] test clean set. KEY – A_w : Word Accuracy computed with QuartzNet [25], EER: Equal Error Rate computed with SpeakerNet [29], eer : normalised EER, GR: Gender Recognition computed with GenderNet [11], gr : normalised Gender Recognition, †: RG (Random Gender) setting [11] is reported. \star : a random accuracy close to 50% is desired for high privacy. Results for GenGAN are obtained with $\epsilon = 0.001$ and $\mathcal{Y}_N \sim (0.5, 0.05)$. The MelGAN model is reported along with the Original signal to assess the impact of the spectrogram inversion without privacy-preserving transformation.

Model	Utility $A_w \uparrow$	Privacy			
		EER \star	$eer \uparrow$	GR \star	$gr \uparrow$
Original	95.64	5.73	11.46	91.37	17.26
MelGAN [14]	93.22	17.42	34.84	89.04	21.92
Client-Hybrid [7]	71.54	–	–	53.90	92.20
VQ-VAE-Gen [11]†	26.84	51.88	96.24	50.01	99.98
PCMelGAN [10]	66.81	41.95	83.90	48.39	96.78
GenGAN (ours)	76.64	38.37	76.74	53.63	92.74

process to remove gender information and conditions the model with identity information during training. We did not modify the architecture of PCMelGAN and trained it from scratch on the LibriSpeech clean train set. We improve and address the limitations in PCMelGAN’s pipeline with GenGAN’s implementation. We train PCMelGAN with the additional Filter network prior to the generation and keep the original loss function used in [10] for the same number of epochs and identical hyper-parameter setting. We use the same pre-trained MelGAN vocoder model in our experiments with PCMelGAN and GenGAN. We also use the same pre-trained models to evaluate the privacy tasks for GenGAN, PCMelGAN and VQ-VAE-Gen. For utility, we use the same ASR for GenGAN and PCMelGAN models for a fair comparison. We run all our experiments on a single Tesla V100 GPU with 32GB memory. The models were trained for 100 epochs by shuffling batches, a learning rate of 0.001 and Adam optimiser [31].

3.4. Discussion

In Tab. 1 “Original” refers to the original audio prior to any transformation. We also include the MelGAN [14] vocoder model to assess the distortion induced by the mel-spectrogram inversion and compare the effect of this operation on utility and privacy of the generated signal.

The values of privacy and utility for Client-Hybrid [7] and VQ-VAE-Gen [11] are taken from the literature. The values corresponding to the ‘Hybrid’ model with fine-tuning on two speakers were chosen for Client-Hybrid, and the ‘Random Gender’ (RG) setting, which considers only gender information during training, for VQ-VAE-Gen. Although VQ-VAE-Gen provides the best privacy guarantees both on gender and identity inference, it is also the worst performing in terms of utility with an A_w value of 26.84%. The Client-Hybrid model performs equally well in terms of utility and privacy, however the selected model for comparison was fine-tuned on two speakers from the train set while no results on the speaker verification task are provided. Results for PCMelGAN show that the privacy objectives are achieved but utility remains low ($A_w = 66.81$). GenGAN improves upon the utility with respect to existing related models without being trained explicitly to optimise for the speech recognition task. GenGAN is a simplified version of PCMel-




Figure 4: Comparisons of the privacy-utility trade-off for the methods compared in Table 1. Values close to the top-right corner denote higher utility and privacy. GenGAN provides the best utility and comparable privacy performance. KEY – A_w : Word Accuracy, gr : normalised Gender Recognition, eer : normalised Equal Error Rate.

GAN with modified cost functions that improve the privacy-utility trade-off with respect to methods considering gender information. In terms of privacy, the 76.74 eer value suggests that some residual identity information is still present in the generated speech, since identity information was not explicitly manipulated in the GenGAN pipeline.

Privacy-utility results are reported in Fig. 4. The larger eer , gr and A_w , the higher the privacy and the higher the utility. GenGAN reaches the best trade-off with improved reconstruction capabilities and comparable privacy performance.

4. Conclusion

We proposed GenGAN, a generative adversarial network that synthesises gender-ambiguous voices that can be used for privacy-preserving speech recognition scenarios. The generator and discriminator are adversarially trained to limit the distortion in the resulting gender-ambiguous signal. Our model improves the privacy-utility trade-off with respect to existing methods.

Future work includes improving the signal reconstruction capabilities of the network without compromising privacy and assessing the naturalness of the generated voices with subjective evaluations.

5. References

- [1] L. G. Kersta, “Voiceprint Identification,” in *Nature*, vol. 196, Dec 1962, pp. 1253–1257.
- [2] D. Toney, D. Feinberg, and K. Richmond, “Acoustic Features for Profiling Mobile Users of Conversational Interfaces,” in *International Conference on Mobile Human-Computer Interaction (MobileHCI)*. Berlin, Heidelberg: Springer, 2004, pp. 394–398.
- [3] C. Y. Huang, Y. Y. Lin, H. Y. Lee, and L. S. Lee, “Defending Your Voice: Adversarial Attack on Voice Conversion,” in *IEEE Spoken Language Technology Workshop (SLT)*, 2021, pp. 552–559.
- [4] B. M. L. Srivastava, A. Bellet, M. Tommasi, and E. Vincent, “Privacy-Preserving Adversarial Representation Learning in ASR: Reality or Illusion?” in *Proceedings INTERSPEECH 2019-20th Annual Conference of the International Speech Communication Association*, 2019.
- [5] A. Nelus and R. Martin, “Gender Discrimination Versus Speaker Identification Through Privacy-Aware Adversarial Feature Ex-

- traction,” in *Speech Communication; 13th ITG-Symposium*, 2018, pp. 1–5.
- [6] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky, “Domain-adversarial Training of Neural Networks,” in *The Journal of Machine Learning Research (JMLR)*, vol. 17, no. 1, 2016, pp. 2096–2030.
 - [7] P. Wu, P. Liang, R. Salakhutdinov, and L.-P. Morency, “Understanding the Tradeoffs in Client-Side Privacy for Speech Recognition,” *arXiv preprint arXiv:2101.08919*, Jan. 2021.
 - [8] P.-G. No , M. Mohammadamini, D. Matrouf, T. Parcollet, A. Nautsch, and J.-F. Bonastre, “Adversarial Disentanglement of Speaker Representation for Attribute-Driven Privacy Preservation,” in *Proceedings INTERSPEECH 2021*, Brno, Czech Republic, Aug 2021. [Online]. Available: <https://hal.archives-ouvertes.fr/hal-03046920>
 - [9] R. Aloufi, H. Haddadi, and D. Boyle, “Privacy-preserving Voice Analysis via Disentangled Representations,” *Proceedings of the 2020 ACM SIGSAC Conference on Cloud Computing Security Workshop*, Nov 2020. [Online]. Available: <http://dx.doi.org/10.1145/3411495.3421355>
 - [10] D. Ericsson, A.  stberg, E. L. Zec, J. Martinsson, and O. Mogren, “Adversarial Representation Learning for Private Speech Generation,” in *International Conference on Machine Learning, Workshop on Self-supervision in Audio and Speech (SAS)*, 2020.
 - [11] D. Stoidis and A. Cavallaro, “Protecting Gender and Identity with Disentangled Speech Representations,” in *Proceedings INTERSPEECH 21-22nd Annual Conference of the International Speech Communication Association*, 2021, pp. 1699–1703.
 - [12] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, “X-Vectors: Robust DNN Embeddings for Speaker Recognition,” in *Proceedings International Conference on Acoustics, Speech and Signal Processing, IEEE*, 2018, pp. 5329–5333.
 - [13] A. Van Den Oord, O. Vinyals, and K. Kavukcuoglu, “Neural Discrete Representation Learning,” in *Advances in Neural Information Processing Systems*, 2017, pp. 6306–6315.
 - [14] K. Kumar, R. Kumar, T. de Boissiere, L. Gestin, W. Z. Teoh, J. Sotelo, A. de Br bisson, Y. Bengio, and A. C. Courville, “MelGAN: Generative Adversarial Networks for Conditional Waveform Synthesis,” in *Advances in Neural Information Processing Systems*, vol. 32, 2019, p. 14910–14921.
 - [15] J. Martinsson, E. L. Zec, D. Gillblad, and O. Mogren, “Adversarial Representation Learning for Synthetic Replacement of Private Attributes,” in *IEEE International Conference on Big Data*, Dec. 2021.
 - [16] S. J. Sutton, “Gender Ambiguous, Not Genderless: Designing Gender in Voice User Interfaces (VUIs) with Sensitivity,” in *Proceedings of the 2nd Conference on Conversational User Interfaces*, ser. CUI ’20. New York, NY, USA: Association for Computing Machinery, 2020. [Online]. Available: <https://doi.org/10.1145/3405755.3406123>
 - [17] O. Ronneberger, P. Fischer, and T. Brox, “U-Net: Convolutional Networks for Biomedical Image Segmentation,” in *Proceedings Medical Image Computing and Computer Assisted Intervention*, 2015.
 - [18] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “ImageNet Classification with Deep Convolutional Neural Networks,” in *Advances in Neural Information Processing Systems*, vol. 25, Dec. 2012, p. 1097–1105.
 - [19] A. Tripathy, Y. Wang, and P. Ishwar, “Privacy-preserving Adversarial Networks,” in *57th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, 2019, pp. 495–505.
 - [20] M. S. F. Poon and M. L. Ng, “The Role of Fundamental Frequency and Formants in Voice Gender identification,” *Speech, Language and Hearing*, vol. 18, no. 3, pp. 161–165, 2015.
 - [21] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: an ASR Corpus Based on Public Domain Audio Books,” in *International Conference on Acoustics, Speech and Signal Processing, IEEE*, South Brisbane, Queensland, Australia, April 2015, pp. 5206–5210.
 - [22] N. Tomashenko, B. M. L. Srivastava, X. Wang, E. Vincent, A. Nautsch, J. Yamagishi, N. Evans, J. Patino, J.-F. Bonastre, P.-G. No , and M. Todisco, “Introducing the VoicePrivacy Initiative,” in *Proceedings INTERSPEECH 2020-21st Annual Conference of the International Speech Communication Association, IEEE*, 2020, pp. 1693–1697.
 - [23] V. I. Levenshtein, “Binary Codes Capable of Correcting Deletions, Insertions, and Reversals,” in *Soviet physics doklady*, vol. 10, 1966, pp. 707–710.
 - [24] J. Qian, H. Du, J. Hou, L. Chen, T. Jung, and X.-Y. Li, “Hide-behind: Enjoy Voice Input with Voiceprint Unclonability and Anonymity,” in *Proceedings International Conference on Embedded Networked Sensor Systems*, 2018, pp. 82–94.
 - [25] S. Kriman, S. Beliaev, B. Ginsburg, J. Huang, O. Kuchaiev, V. Lavrukhin, R. Leary, J. Li, and Y. Zhang, “Quartznet: Deep Automatic Speech Recognition with 1D Time-Channel Separable Convolutions,” in *Proceedings International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE*, May 2020, pp. 6124–6128.
 - [26] A. Graves, S. Fern andez, F. Gomez, and J. Schmidhuber, “Connectionist Temporal Classification: Labelling Unsegmented Sequence Data with Recurrent Neural Networks,” in *Proceedings of the 23rd International Conference on Machine Learning*, 2006, pp. 369–376.
 - [27] J. Li, V. Lavrukhin, B. Ginsburg, R. Leary, O. Kuchaiev, J. M. Cohen, H. Nguyen, and R. Gadde, “Jasper: An End-to-End Convolutional Neural Acoustic Model,” in *Proceedings INTERSPEECH 2019-20th Annual Conference of the International Speech Communication Association, IEEE*, 2019, pp. 71–75.
 - [28] K. He, X. Zhang, S. Ren, and J. Sun, “Deep Residual Learning for Image Recognition,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Dec. 2016, pp. 770–778.
 - [29] J. S. Chung, J. Huh, S. Mun, M. Lee, H.-S. Heo, S. Choe, C. Ham, S. Jung, B.-J. Lee, and I. Han, “In Defence of Metric Learning for Speaker Recognition,” in *Proceedings INTERSPEECH 2020-21st Annual Conference of the International Speech Communication Association*, Sep. 2020, pp. 2977–2981.
 - [30] J. S. Chung, A. Nagrani, and A. Zisserman, “Voxceleb2: Deep speaker recognition,” in *Proceedings INTERSPEECH 2018-19th Annual Conference of the International Speech Communication Association*, 2018, pp. 1086–1090.
 - [31] D. Kingma and J. Ba, “Adam: A Method for Stochastic Optimization,” in *3rd International Conference on Learning Representations (ICLR)*, San Diego, CA, USA, Dec. 2015.