Text-to-Speech Pipeline for Swiss German - A comparison

Tobias Bollinger[1,*], Jan Deriu[2] and Manfred Vogel[1]

[1]*University of Applied Sciences and Arts Northwestern Switzerland, Windisch*

[2]*Zurich University of Applied Sciences, Winterthur*

Abstract
In this work, we studied the synthesis of Swiss German speech using different Text-to-Speech (TTS) models. We evaluated the TTS models on three corpora, and we found, that VITS models performed best, hence, using them for further testing. We also introduce a new method to evaluate TTS models by letting the discriminator of a trained vocoder GAN model predict whether a given waveform is human or synthesized. In summary, our best model delivers speech synthesis for different Swiss German dialects with previously unachieved quality.

Keywords
Speech Synthesis, Text to Speech, Swiss German, low resource

1. Introduction

Text-to-Speech (TTS) describes the task of synthesizing raw waveform of speech for a given text. While the field is making staggering progress for many large-resource languages, the work on Swiss German is still sparse. There are several reasons for this. First, Switzerland is a small country and around 60% speak Swiss German [1] making it a low resource setting. Second, there is no standardized text form, thus, the system has to handle standard German input to create Swiss German audio. Third, there is no single Swiss German language, on the contrary, there are many dialects that differ quite significantly. The dialects can also differ significantly in their grammar, vocabulary and phonetics. Thus, the setting for TTS for Swiss German is challenging.

On the other hand, there was a push in recent years to collect high-quality data for Swiss German [1, 2, 3]. This opens up for the possibility of training higher quality TTS systems, and also textural translation systems from standard German to different Swiss German dialects. For instance, [2] collected around 200 hours of Swiss German speech with the corresponding standard German translation, and [1] collected around 30 hours of high-quality Swiss German audio and provide Swiss German texts to specifically train TTS systems.

In this work, we make two contributions. First, we assess the quality of the variational inference with adversarial learning for end-to-end Text-to-Speech (VITS) [4] when trained on three different data sets, namely [1, 2, 3]. Second, we train a textual translation system for stan-

dard German to Swiss German text (although there is no formal grammar for it), which is used as prepossessing to input Swiss German text to the TTS system.

Our results show that using high-quality data and using a translation system for translating standard German text to Swiss German text, yields high quality synthesized speech. That is, it achieves an average of 4.1 on a 5-point scale (MOS), which is considered high quality.

2. Related Work

2.1. Speech synthesis in Swiss German

Text to Speech models are prominent for many languages like English, Japanese, Korean or German, but for Swiss German less work is done. In 2016, SlowSoft GmbH developed a speech synthesis system for Grisons and Romansh [3]. The online service Narakeet, a text-to-speech video maker, also supports one Swiss German speaker [5]. [1] published a Swiss German Corpus and also trained a Tacotron2 [6] model for Swiss German speech synthesis.

2.2. Deep learning-based synthesis

Current TTS methods can be categorized into cascaded and end-to-end methods. In cascaded TTS systems like Tacotron2 [6] or Fastspeech2 [7], the task of synthesis is divided into two steps. First, generating an intermediate representation of Mel spectrograms out of the text with a so-called acoustic model. Second, using the Mel spectrograms to generate the final wave form with a vocoder. End-to-end TTS systems try to jointly optimize the acoustic model and the vocoder. This system is used by the VITS model [4]. To get more control over the produced synthesis, zero-shot models like YourTTS [8] or Vall-E [9] provide a method to produce speech out of text and a

✉ tobias.bollinger@students.fhnw.ch (T. Bollinger);
jan.deriu@zhaw.ch (J. Deriu); manfred.vogel@fhnw.ch (M. Vogel)

[1]https://www.eda.admin.ch/aboutswitzerland/de/home/gesellschaft/sprachen/die-sprachen---fakten-und-zahlen.html

Table 1
Data set size used for training: SDS-200, SwissDial, SlowSoft.

SDS-200

Set	Fraction	Samples
Train	0.9	12'948
Valid	0.05	720
Test	0.05	720

SwissDial

Set	Fraction	Samples
Train	0.9	27'828
Valid	0.05	1'546
Test	0.05	1'547

Slowsoft

Set	Fraction	Samples
Train	0.9	2'061
Valid	0.05	115
Test	0.05	115

short speech input to imitate the voice of the speaker of the provided speech input.

3. Training Data

The focus of this work is to benchmark how well the end-to-end training of VITS performs on Swiss German data sets. For this, we used three different corpora, which differ on the quality of the audio and the number of samples per speaker.

3.1. SDS-200

The SDS-200 corpus contains 200 hours of Swiss German audio recordings from all dialects, along with their Standard German transcriptions. Note that in contrast to classical TTS corpora, the SDS-200 contains a large variety of different speakers as the primary purpose of the corpus is training Speech-to-Text models. The recordings represent one sentence each and have been crowd-sourced from the Swiss population during the "Schweizer Dialektsammlung" [2]. The audio files are compressed and there is a lot of background noise in the audios such as mouse clicking at the end of the audio. However, this is the largest corpus in our set, and it showed to be useful for training speech-to-text systems.

3.2. SwissDial

SwissDial is an annotated corpus of spoken Swiss German across 8 major dialects (AG, BE, BS, GR, LU, SG, VS, ZH). The data set includes around 3 hours of high quality audio per dialect, together with Swiss German and

Standard German transcripts [1]. The audio files have been recorded in a quiet room in front of a high-quality recording set-up to obtain clean recordings. We use the Swiss German and standard German transcripts to train the textual translation model. Note that in SwissDial the Grisons dialect has an extra 6 horus of audio samples compared to the other dialects.

3.3. Slowsoft

SlowSoft builds speech technology components with a focus on minority languages and dialects. The data, kindly provided for our research, contains phonetic transcriptions created by professionals, in the Grisons (GR) dialect [10].

Table 1 contains the overview of the data set sizes, which we used to train the TTS pipeline.

4. Pipeline

The pipeline consists of two main parts: first, a T5 [11] model fine-tuned on translating standard German text to Swiss German text, and second the VITS model trained on the data sets mentioned in the previous section.

4.1. T5 - German to Swiss German

The standard German to Swiss German text translation is based on a T5 model [11] fine-tuned on the SwissDial data set [1]. As this data set contains for each sample the German sentence alongside the Swiss German sentences for eight different dialects, we train the translation model to translate to the various dialects. For this, we used the Huggingface [12] implementation. To allow the model to handle the different dialects, we prepended the dialect tag to the input.

4.2. VITS - Swiss German Text to Speech

The Conditional Variational Autoencoder with Adversarial Learning for End-to-End Text-to-Speech is a parallel end-to-end TTS method. It adapts variational inference with normalizing flows and an adversarial training process [4]. For the experiment, the ESPNet Framework was used [13], which implements the VITS architecture.

The input text was split into a sequence by character or word. When a split by characters was used, it generalized way better than when the text was split by words and therefore produced better results.

We also tested a Grapheme-to-Phoneme (G2P) conversion for German text which uses eSpeak as backend [14]. But because the rule-based G2P algorithm only works for German and not Swiss German, it results in worse synthesized speech.

Different preprocessing steps were tested, for SDS-200 a noise reduction was applied and for all corpora silence at the start and end were trimmed. The reduction of noise resulted in a more unnatural result, therefore only the trim off the silence was used for preprocessing.

To allow the synthesis of different dialects, a speaker embedding with x-vectors [15] with an embedding dimension of 512 has been used.

Thus, we train three VITS models, one for each of the three different corpora described in 3.

5. Evaluation Metrics

The evaluation of Text-to-Speech is a challenging topic. Therefore, we perform an extensive automated evaluation using a variety of different metrics, as well as a human evaluation using a 5-point Mean Option Score (MOS). For the T5 model, we use the BLEU score [16] to evaluate and propose a qualitative analysis.

MCD Mel-Cepstral Distortion [17] is a measure which calculates how different two sequences of Mel cepstral are. The idea is, the smaller the MCD is between synthesized and natural Mel cepstral sequences, the closer the synthesized speech is to the natural speech. This is by no means a perfect measurement, but it can be a useful indicator, especially combined with other metrics.

log-F0 RMSE Log-F0 Root-Mean-Square Error is a measure which computes the root-mean-square error of the log-F0 of the signal. Where F0 is the fundamental frequency which is closely related to the pitch.

For both, the MCD and log-F0 RMSE dynamic time-warping (DTW) is applied to match the different length of the ground-truth and the generated signal.

CER and WER The Character Error Rate (CER) and the Word Error Rate (WER) are calculating the modifications needed to transform the ground-truth text to the generated text. A modification is defined as either a substitution (S), a deletion (D) or an insertion (I). When there is no modification needed, we call it a Word/Character Hit (H) [18]. For CER the calculation is done on a character basis:

$$CER = \frac{S + D + I}{H + S + D} \qquad (1)$$

$H + S + D$ is the number of characters in the ground-truth text.

For WER the calculation is done on a word basis, the same equation can be used, but the modifications are counted per word and not per character.

$$WER = \frac{S + D + I}{H + S + D} \qquad (2)$$

Before calculating the error rate, both texts are preprocessed:

- Convert the text to lower case
- Remove multiple spaces
- Trim white spaces at the start and end
- Remove the three characters: «, » and ,
- Remove punctuation

BLEU The Bilingual Evaluation Understudy score counts the matching n-grams in the ground-truth text to the generated text. A perfect match results in a score of 1.0, whereas a perfect mismatch results in a score of 0.0.

MCD and log-F0 RMSE reflect speaker, prosody, and phonetic content similarities, whereas CER, WER and BLEU the intelligibility.

Vocoder Discriminator Score A further method – developed in this work – to measure the quality of the speech synthesis is to use a trained vocoder which uses Generative Adversarial Networks (GAN). The discriminator tries to discriminate fake from real waveforms. The intuition is, the better our synthesized speech can fool the discriminator, the better the model performed. To evaluate this, we can take the discriminator or discriminators of the trained model and let them predict if the speech input is fake or real. The discriminators will output a value for each waveform step between 0 and 1 where a bigger number means the discriminator is more confident that the waveform step is real. For getting the score, the discriminator outputs are concatenated to $\hat{y}$, then the mean squared error between $\hat{y}$ and a zero vector with the same dimension is calculated. The resulting $loss$ will be transformed to a $score$ by subtracting the $loss$ from 1.
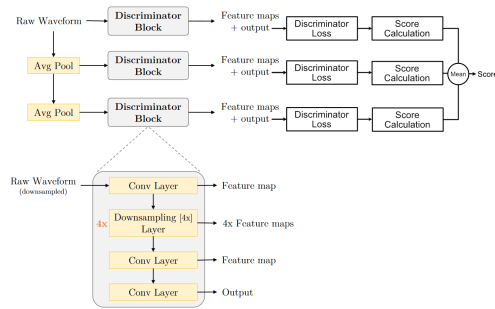
$$loss = \frac{1}{n} \sum_{i=1}^{n} (\hat{y}_i - 1)^2$$
$$score = 1 - loss \qquad (3)$$

Two models were trained for the evaluation, a Parallel WaveGAN [19] model and a MelGAN [20] model; where the Parallel WaveGAN has one discriminator and the MelGAN has three. For the MelGAN we calculate the score for each discriminator and then we take the mean of the three scores. Booth models were trained on the SDS-200 and on the SwissDial corpus.
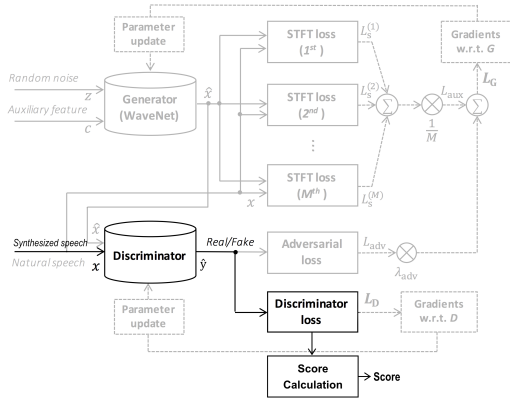
(a) Parallel WaveGAN score calculation. Base on the illustration form [19, S.2].

(b) MelGAN score calculation. Base on the illustration form [20, S.4].

SDS-200	
Metric Name	Error / Score
↓ MCD	7.50 ± 1.63
↓ log-F0 RMSE	0.22 ± 0.09
↑ PGAN Score	0.73
↑ MelGAN Score	0.35
↓ CER	0.27
↓ WER	0.43
↑ BLEU	0.39

SwissDial	
Metric Name	Error / Score
↓ MCD	9.29 ± 1.11
↓ log-F0 RMSE	0.31 ± 0.11
↑ PGAN Score	0.75
↑ MelGAN Score	0.69
↓ CER	0.15
↓ WER	0.28
↑ BLEU	0.54

Slowsoft	
Metric Name	Error / Score
↓ MCD	8.27 ± 0.66
↓ log-F0 RMSE	0.25 ± 0.08
↑ PGAN Score	0.74
↑ MelGAN Score	0.57

6. Results

In this section we present the results of the pipeline. For this we show the scores of the T5 model and the different scores of the VITS models trained on the three data sets.

6.1. T5 - Results

The T5 model achieves a BLEU score of 50.47 on the test split of the SwissDial corpus over all the eight dialects, which is very promising. However, the BLEU score is not fully representative. Thus, we add a qualitative analysis. Table 3 shows three examples of the Swiss German texts generated from the standard German input. We note that the past tense is handled correctly in these cases. Since Swiss German does not use the simple past, all the past tenses need to be rewritten to past progressive tense. In the example, "In Gegenrichtung wurde jedoch eine durchgehende Leitplanke installiert."

is translated correctly to "In Gegarichtig isch jedoch a durchgegehndi Leitplanka installiart worda." Similarly the last example changes "führte" to "hät ... igfüert". Most discrepancies between the ground truth and the generated Swiss German sentences stem from the fact that different spellings were used. For instance, in the first example, the ZH groundtruth would have expected to start with "Ide Gägerichtig..." instead of "In Gägerichtig...". However, the generated utterance is not wrong. The evaluation is challenging since Swiss German has no writing rules and the specific spelling is, thus, very individual. There are also real mistakes, such as in the last example the LU translation uses the word "Wiisscreme", which is clearly wrong.

6.2. VITS - Results

For evaluating CER, WER and BLEU an internally provided ASR fairseq model was used. The model reached on SDS-200 with GT speech a BLEU of 0.698, a WER of 0.18 and a CER of 0.12. Because the Slowsoft model uses phonemes, it is not possible to evaluate the CER, WER and BLEU metrics.

Dialect	Utterance
DE	In Gegenrichtung wurde jedoch eine durchgehende Leitplanke installiert.
ZH	In Gägerichtig isch aber e durchgehendi Litplank installiert worde.
SG	In Gegenrichtig isch aber e durchgehendi Leitplank installiert worde.
BE	In Gägärichtig isch aber ä durchgehendi Litplanki inschtalliert wordä.
GR	In Gegarichtig isch jedoch a durchgegehndi Leitplanka installiart worda.
VS	In Gägurichtig isch aber en durchgehendi Leitplanka installiert wordu.
BS	In Gegerichtig isch aber ä durchgenedi Leitplank installiert worde.
AG	E Gägerechtig esch aber en dorchgehendi Leifantworte installiert worde.
LU	En Gägesiitig esch aber en döregehendi Liitplank installiert worde.
DE	Zunächst kommen aber ohnehin die Rezepte des Bundesrates ins Parlament.
ZH	Zerst chömed aber sowieso d Rezäpt vom Bundesrat is Parlamänt.
SG	Zersch chömmed aber trotzdem dRefekt vom Bundesrot is Parlament.
BE	Zersch chömä aber trotzdäm d Projekt vom Bundesrat is Parlamänt.
GR	Zerst kömmen aber sowieso d Rezept vum Bundesrot ins Parlament.
VS	Zersch chämunt aber sowieso d'Reschtoff vam Bundesrat ins Parlamänt.
BS	Zerst kömme aber sowieso d Rezäkt vom Bundesroot ins Parlament.
AG	Zerscht chömme aber sowieso d'Reschpondänt fom Bondesrot es Parlamänt.
LU	Zerscht chömed aber sowieso d Rezäpt vom Bondesrot is Parlamänt.
DE	Er führte auch das Servieren von Eiscreme bei nachmittäglichen Empfängen ein.
ZH	Er hät au s Serviere vo Glacé bi nahmittägliche Empfäng igfüert.
SG	Er het au sServiere vo Iiscreme bi nochemtägliche Empfäng iigführt.
BE	Är het o s Servierä vo Glacä bi nachmittäglechä Empfäng iigfüehrt.
GR	Er het au ds Serviara vu Iiscreme bi nochmittäglicha Empfäng igfüahrt.
VS	Är het öi z'Servieru va Iischcreme bi nahmittaglichu Empfängi ihgfiehrt.
BS	Är het au s Serviere vo Glace bi nochmittägliche Empfäng ihgfüehrt.
AG	Er het au s'Serviere fo Iscreme be nachhertägliche Emfäng ihgfüehrt.
LU	Er hed au s Serviere vo Wiisscreme be nochmettägliche Emfäng igfüehrt.

Table 3
Three randomly picked examples of standard German text translated to each of the dialects in the Siwss Dial corpus.

The SDS-200 model performed best on the MCD and log-F0 RMSE metrics. The corpus SDS-200 contains some noise which is an easy task for the model to reproduce correctly. We think this is also the reason the model performed the best on the MCD and log-F0 RMSE metrics which compares the ground-truth signal with the synthesized. But on all other metrics, the SwissDial model performed best. The Slowsoft model has significantly less training data than the other two corpora, but does not show a significant drop in performance. It even achieves better results on MCD and log-F0 RMSE compared to the SwissDial model.

6.3. Human Evaluation - MOS

Since automated evaluation is flawed, we also run a human evaluation. For this, we randomly select 15 samples, which we synthesize into waveforms using the three different VITS models. We then let human rate the audio on a scale from 1 to 5 with the following classification:
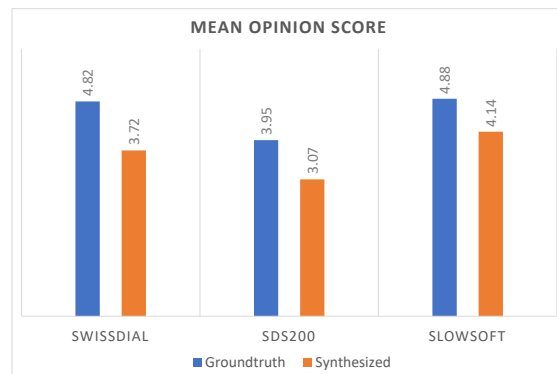
Figure 2: MOS Results of the TTS human evaluation.

5 : Excellent: Neural speech synthesis for Swiss German.

4 : Good: Through attentive listening, speech can

be perceived without effort.

3 : Fair: Speech can be perceived with slight effort.

2 : Poor: It takes great concentration and effort to understand the speech.

1 : Bad: Despite great effort, you cannot understand the speech.

We let each sample be rated by 7 native speakers. Alongside the synthesized audio, we also let the groundtruth audio be rated as a reference. In Figure 2 the results of the human evaluation are shown. We see that the groundtruth audio are still significantly higher rated than the synthesized audios. Consistent with our observations, the groundtruth SDS-200 audio samples were rated worse compared to the high-quality SwissDial and Slowsoft samples. The synthesized samples show the same trends. The VITS model trained with the high-quality Slowsoft samples performs best with an average MOS score of 4.1, which is higher than the groundtruth samples of the SDS-200 corpus. Analogous to above, the VITS model trained on SDS-200 samples performs worst, with a score of 3.1.

In summary, the results show that TTS already works very well for Swiss German, provided that the training material is of high quality. Thus, it is best to use a small amount of data, which is of highest quality, than using a large dataset of low-quality samples. Furthermore, the human evaluation shows that the automated metrics are not indicative of the final quality of the generated audio.

7. Conclusion

In this work, we compared three VITS models trained on different data sets to synthesize Swiss German speech from standard German texts. For this, we build a two-step pipeline where we first use a fine-tuned T5 model to translate the standard German text to Swiss German text, and then the VITS model that synthesizes Swiss German audio from the text. We evaluated the quality of the audio with automated metrics and a human evaluation. The results showed that the VITS model trained on the highest quality data performs best. Furthermore, the quality is good enough to be considered high quality audio samples, even surpassing the groundtruth samples of another corpus. This shows that TTS methods are applicable to the case of Swiss German provided that the training data is of high quality and a text translation from standard German to Swiss German text is available.

Acknowledgments

References

[1] P. Dogan-Schönberger, J. Mäder, T. Hofmann, Swissdial: Parallel multidialectal corpus of spoken swiss german, 2021. URL: https://arxiv.org/abs/2103.11401. doi:doi:10.48550/ARXIV.2103.11401.

[2] M. Plüss, M. Hürlimann, M. Cuny, A. Stöckli, N. Kapotis, J. Hartmann, M. A. Ulasik, C. Scheller, Y. Schraner, A. Jain, J. Deriu, M. Cieliebak, M. Vogel, Sds-200: A swiss german speech to standard german text corpus, 2022. URL: https://arxiv.org/abs/2205.09501. doi:doi:10.48550/ARXIV.2205.09501.

[3] SlowSoft GmbH, SlowSoft GmbH | Über SlowSoft, 2019. URL: https://slowsoft.ch/deu/company.html, [Online; accessed 27. Dec. 2022].

[4] J. Kim, J. Kong, J. Son, Conditional variational autoencoder with adversarial learning for end-to-end text-to-speech, CoRR abs/2106.06103 (2021). URL: https://arxiv.org/abs/2106.06103. arXiv:2106.06103.

[5] Swiss German Text To Speech, 2022. URL: https://www.narakeet.com/languages/swiss-german-text-to-speech, [Online; accessed 27. Dec. 2022].

[6] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. J. Skerry-Ryan, R. A. Saurous, Y. Agiomyrgiannakis, Y. Wu, Natural TTS synthesis by conditioning wavenet on mel spectrogram predictions, CoRR abs/1712.05884 (2017). URL: http://arxiv.org/abs/1712.05884. arXiv:1712.05884.

[7] Y. Ren, C. Hu, X. Tan, T. Qin, S. Zhao, Z. Zhao, T.-Y. Liu, Fastspeech 2: Fast and high-quality end-to-end text to speech, 2020. URL: https://arxiv.org/abs/2006.04558. doi:doi:10.48550/ARXIV.2006.04558.

[8] E. Casanova, J. Weber, C. Shulby, A. C. Júnior, E. Gölge, M. A. Ponti, Yourtts: Towards zero-shot multi-speaker TTS and zero-shot voice conversion for everyone, CoRR abs/2112.02418 (2021). URL: https://arxiv.org/abs/2112.02418. arXiv:2112.02418.

[9] C. Wang, S. Chen, Y. Wu, Z. Zhang, L. Zhou, S. Liu, Z. Chen, Y. Liu, H. Wang, J. Li, L. He, S. Zhao, F. Wei, Neural codec language models are zero-shot text to speech synthesizers, 2023. URL: https://arxiv.org/abs/2301.02111. doi:doi:10.48550/ARXIV.2301.02111.

[10] SlowSoft GmbH, SlowSoft GmbH | Home, 2019. URL: https://slowsoft.ch/eng/index.html, [Online; accessed 21. Dec. 2022].

[11] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, P. J. Liu, Exploring the limits of transfer learning with a unified text-to-text transformer, Journal of Machine Learning Research 21 (2020) 1–67. URL: http://jmlr.org/papers/

v21/20-074.html.

[12] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. De-langue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Fun-towicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. L. Scao, S. Gugger, M. Drame, Q. Lhoest, A. M. Rush, Transformers: State-of-the-art natural language processing, in: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, Association for Computational Linguistics, Online, 2020, pp. 38–45. URL: https://www.aclweb.org/anthology/2020.emnlp-demos.6.

[13] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishi-toba, Y. Unno, N. Enrique Yalta Soplin, J. Heymann, M. Wiesner, N. Chen, A. Renduchintala, T. Ochiai, ESPnet: End-to-end speech processing toolkit, in: Proceedings of Interspeech, 2018, pp. 2207–2211. URL: http://dx.doi.org/10.21437/Interspeech.2018-1456. doi:doi:10.21437/Interspeech.2018-1456.

[14] eSpeak: Speech Synthesizer, 2013. URL: https://espeak.sourceforge.net, [Online; accessed 27. Dec. 2022].

[15] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, S. Khudanpur, X-vectors: Robust dnn embed-dings for speaker recognition, in: 2018 IEEE In-ternational Conference on Acoustics, Speech and Signal Processing (ICASSP), 2018, pp. 5329–5333. doi:doi:10.1109/ICASSP.2018.8461375.

[16] K. Papineni, S. Roukos, T. Ward, W.-J. Zhu, Bleu: a Method for Automatic Evaluation of Machine Translation, in: Proceedings of the 40th Annual Meeting of the Association for Computational Lin-guistics, Association for Computational Linguistics, Philadelphia, Pennsylvania, USA, 2002, pp. 311–318. URL: https://www.aclweb.org/anthology/P02-1040. doi:doi:10.3115/1073083.1073135.

[17] R. Kubichek, Mel-cepstral distance measure for objective speech quality assessment, in: Proceedings of IEEE Pacific Rim Conference on Communications Computers and Signal Pro-cessing, volume 1, 1993, pp. 125–128 vol.1. doi:doi:10.1109/PACRIM.1993.407206.

[18] A. C. Morris, V. Maier, P. Green, From WER and RIL to MER and WIL: improved evaluation measures for connected speech recognition, ResearchGate (2004). doi:doi:10.21437/Interspeech.2004-668.

[19] R. Yamamoto, E. Song, J.-M. Kim, Parallel wavegan: A fast waveform generation model based on gen-erative adversarial networks with multi-resolution spectrogram, 2019. URL: https://arxiv.org/abs/1910.11480. doi:doi:10.48550/ARXIV.1910.11480.

[20] K. Kumar, R. Kumar, T. de Boissiere, L. Gestin, W. Z. Teoh, J. Sotelo, A. de Brebisson, Y. Ben-gio, A. Courville, Melgan: Generative adversar-ial networks for conditional waveform synthe-sis, 2019. URL: https://arxiv.org/abs/1910.06711. doi:doi:10.48550/ARXIV.1910.06711.