

A COMPARATIVE STUDY OF SELF-SUPERVISED SPEECH REPRESENTATIONS IN READ AND SPONTANEOUS TTS

Siyang Wang, Gustav Eje Henter, Joakim Gustafson, Éva Székely

Division of Speech, Music and Hearing, KTH Royal Institute of Technology, Stockholm, Sweden

ABSTRACT

Recent work has explored using self-supervised learning (SSL) speech representations such as wav2vec2.0 as the representation medium in standard two-stage TTS, in place of conventionally used mel-spectrograms. It is however unclear which speech SSL is the better fit for TTS, and whether or not the performance differs between read and spontaneous TTS, the later of which is arguably more challenging. This study aims at addressing these questions by testing several speech SSLs, including different layers of the same SSL, in two-stage TTS on both read and spontaneous corpora, while maintaining constant TTS model architecture and training settings. Results from listening tests show that the 9th layer of 12-layer wav2vec2.0 (ASR finetuned) outperforms other tested SSLs and mel-spectrogram, in both read and spontaneous TTS. Our work sheds light on both how speech SSL can readily improve current TTS systems, and how SSLs compare in the challenging generative task of TTS. Audio examples can be found at <https://www.speech.kth.se/tts-demos/ssr-tts>

Index Terms— speech synthesis, self-supervised speech representation, spontaneous speech

1. INTRODUCTION

Recently, text-to-speech (TTS) researchers have turned to using self-supervised learning (SSL) speech representations as intermediate representations in the popular two-stage TTS scheme [1, 2], with an acoustic model followed by a vocoder. These SSLs take the place of the conventionally used mel-spectrograms as acoustic-modelling targets. They can be learned from audio alone (without transcriptions) thus can scale to vast amounts of data and compute to build better representations of speech audio. They have demonstrated advantages across increasing spectrum of speech technology applications, with recent addition of building high-quality TTS from mixed-quality audio [1].

This research was supported by the Digital Futures project Advanced Adaptive Intelligent Systems (AAIS), the Swedish Research Council projects Connected (VR-2019-05003), Perception of speaker stance (VR-2020-02396), the Riksbankens Jubileumsfond project CAPTivating (P20-0298), and by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

However, SSL-based approaches to TTS are still in their infancy [1, 2, 3]. Although speech SSLs have been compared in various applications such as ASR, speaker identification, and speech conversion [4], it remains unclear which SSL is superior for TTS and why. The demonstrated effectiveness of intermediate layers of the same SSL model for downstream tasks [4] further clouds the picture regarding what representation to use for TTS. Moreover, two-stage TTS with SSL-based intermediate representations has thus far only been applied to speech audio of text read aloud [1], whereas most speech in the wild is spontaneous and unscripted. Spontaneous speech has many unique verbal and nonverbal phenomena (e.g., breathing, disfluencies, discourse markers) often not represented in text, and is severely underrepresented in TTS research today due to these challenges. SSLs have the potential to revolutionise spontaneous TTS as they do not require high-quality audio or transcripts to learn effectively, but to what extent they can live up to this potential is unknown.

We address these questions by building several two-stage TTS systems with different speech SSLs as intermediate representations, and study their objective resynthesis accuracies as well as subjective TTS qualities.¹ The key findings are:

- (1) Not all SSLs are created equal for TTS.
- (2) The final layer of a SSL fine-tuned for ASR is impoverished for TTS, compared to a middle layer.
- (3) The resynthesis error is not a good predictor of the overall TTS quality afforded by using a given SSL.

We also present the first results on applying SSLs as intermediate representations in spontaneous TTS, where we find:

- (4) Similar ordering of resynthesis and TTS performance within SSLs as in read speech.
- (5) SSL features outperform mel-spectrograms more in spontaneous TTS than in read-speech TTS.

2. BACKGROUND AND PRIOR WORK

2.1. Self-supervised speech representations

Learning general-purpose representations from raw audio is a key question in speech processing. Borrowing from the success of self-supervised learning (SSL) in other domains (vision, language), speech SSL learned with only self-

¹Audio examples: <https://www.speech.kth.se/tts-demos/ssr-tts>.

supervised objectives and then finetuned on down-stream tasks have shown promise in applications such as ASR and speaker identification [4]. Examples of these representations are wav2vec [5], wav2vec2.0 [6], HuBERT [7]. Their initial success in ASR has led to research of their usage in generative tasks such as speech encoding/resynthesis [8], speech language modeling [9], and TTS [1].

With the proliferation of SSL approaches, recent work has also focused on understanding the difference between SSL models [4, 10] and the difference between layers of the same SSL model [11, 12]. Our work extends these findings by testing different SSLs and their layers in two-stage TTS.

2.2. Two-stage Text-to-Speech

Neural speech synthesis today is dominated by a so-called “two-stage” approach, where a stage-1/acoustic model first generates intermediate representations which are then vocoded by stage-2/vocoder model. Since the first successful neural two-stage TTS Tacotron2 [13], the focus of research has been on improving either stage-1 or stage-2 architecture and/or training loss, e.g. transformer-based stage-1 model [14, 15], normalizing flow-based stage-1 model [16], GAN-based stage-2 model [17]. Only recently did the intermediate representation start to garner interest with approaches that use SSL instead of the default choice of mel-spectrogram [1, 2].

2.3. Synthesizing spontaneous speech

Modelling spontaneous speech is a critical part of creating more human-like speech synthesis [18]. However, synthesizing spontaneous speech is fundamentally different to read-speech synthesis which is more explored in TTS research today. The differences, which are mostly hindering factors for developing TTS, include more dynamic prosody, disfluencies, various breathing sounds, vocal outbursts, and large variations in voice and recording quality. Many of these characteristics are poorly approximated in the transcript or are completely left out [19], making training TTS difficult. These challenges and spontaneous TTS as a whole are receiving increasing attention from TTS community [20, 21].

3. METHOD

3.1. Representations

Four speech SSLs were selected for testing, as summarized in Table 1. They are different in the training-loss design as well as in the data used for training. We also selected two different layers in the same model wav2vec2.0, as prior work has shown that a middle layer of a ASR-finetuned SSL model contains more prosodic information [11] which could benefit synthesis. We recognize that this set is limited in some aspects and future work could extend by adding more comparisons that probe specific aspects. These SSLs are compared

against standard 80-channel log mel-spectrogram used in current two-stage TTS.

3.2. Data

For read speech TTS, we used the widely benchmarked LJSpeech corpus [22]. For spontaneous TTS, we used the audio recordings of the Trinity Speech-Gesture Dataset (TSGD) [23], comprising 25 impromptu monologues (avg. 10.6m long), spoken by a male actor in colloquial style. To prepare the dataset for TTS, we segmented the corpus into stretches of speech delineated by breath events, and combined these in an overlapping fashion to form an utterance structure. [24].

3.3. System

We largely follow WavThruVec [1] in neural architecture choice and training setup, both for simplicity and because architecture improvement is not the main focus of our work. Our architecture is essentially the single-speaker version of WavThruVec. We also note that WavThruVec used wav2vec2.0-base-asr to demonstrate the effectiveness of the model, while we compare a range of SSL representations with the focus on probing their differences in WavThruVec framework. An overview of the method is shown in Fig. 1.

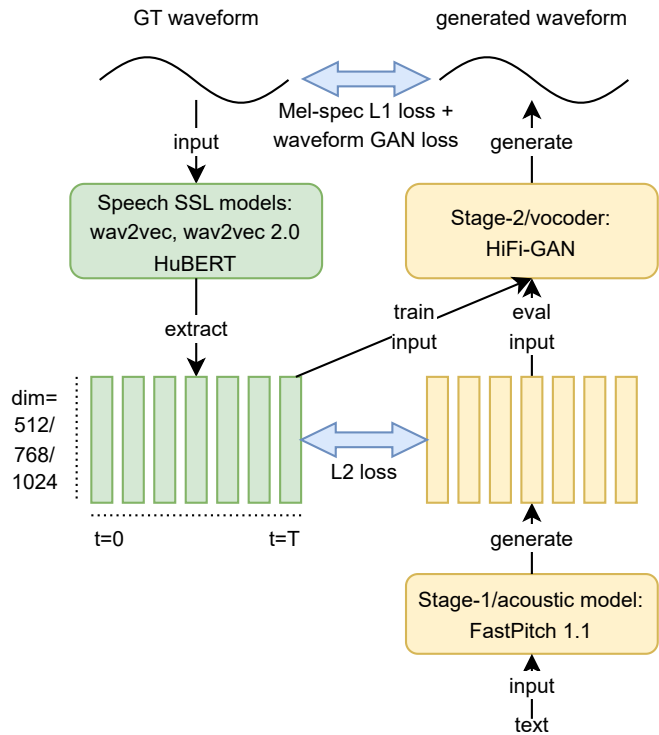


Fig. 1. Diagram of the system.

The stage-1 acoustic model is adapted from FastPitch

Representation Notation (SSL Model)	Pre-training Data (Loss)	ASR Finetuning	Architecture	Layer	Hz	Dim	Resynthesis Error↓ Read Spontaneous
w2v2-L12 (wav2vec2.0-base-asr [6])	LibriSpeech960H (contrastive+diversity)	LibriSpeech960H	7CL→ 12TL	TL-12	50	768	0.7115 0.6724
w2v2-L9 * (wav2vec2.0-base-asr [6])	LibriSpeech960H (contrastive+diversity)	LibriSpeech960H	7CL→ 12TL	TL-9	50	768	0.6089 0.5582
w2v1-c (wav2vec [5])	LibriSpeech960H (contrastive)	-	5CL(z)→ 9CL(c)	c	100	512	0.5253 0.4437
HuBERT-L12 (HuBERT-large-asr [7])	LibriLight60,000H (masked prediction)	LibriSpeech960H	7CL→ 24TL	TL-12	50	1024	0.4406 0.4056

Table 1. Speech SSL tested. In Architecture and Layer columns, “CL” denotes convolutional layer and “TL” denotes transformer layer. wav2vec2.0 and HuBERT share convolutional and transformer layer architecture. Vocoding Error column corresponds to Figure 2. *: w2v2-L9 performs best in TTS perceptual test.

[14], a parallel transformer-based TTS model, with added monotonic alignment [25].² All stage-1 models are trained with identical hyperparameters as in [1], except with a batch size of 96. The models are trained for variable iterations until the validation loss stops improving depending on the corpus, which takes about 200 epochs for most models on LJSpeech. We use an official pre-trained FastPitch [14] as mel-spec baseline.

During initial experiments with the spontaneous corpus TSGD, we found that random weight initialization results in low-quality speech. This is likely due to the smaller size of this corpus, at approximately 4 hours in total. To obtain better results, we instead start training from our LJSpeech-trained first-stage model and finetuned it on the spontaneous corpus for 200 epochs. This methodology is highly effective in allowing neural TTS to be trained on smaller databases [20].

The stage-2 model (vocoder) is adapted from HiFi-GAN [17], trained in a similar setup as in [1]. Training batches consist of 0.5s random audio excerpts, with the batch size varying from 160 to 192 depending on the dimension and sampling rate of the representation, to ensure that per-iteration data throughput is approximately consistent across models. All LJSpeech SSL vocoders are trained for 180k iterations at 22 kHz and we use an official pre-trained LJSpeech mel-spectrogram HiFi-GAN from authors of FastPitch [14] as baseline. All vocoders for spontaneous corpus [23] are trained for 100k iterations at 16 kHz from scratch.

4. RESULTS

4.1. Resynthesis

We first look at resynthesis performance of the SSL representations. We extract SSLs from speech audios in the validation set and then reconstruct the waveforms using the SSLs’ respectively trained HiFi-GAN vocoders. Mel-spectrogram

error (L1) of the reconstructed waveform is used as a measure of resynthesis performance. The results are visualised in Fig. 2 with corresponding values listed in the last column of Table 1.

The ordering of representations is consistent in both read and spontaneous corpora resynthesis performance. Between the two wav2vec2.0 layers, L9 has better resynthesis performance than L12. This is consistent with prior findings that earlier layers of wav2vec2.0-ASR contain more prosodic information than later layers [11] [12].

4.2. Text-to-speech

We conducted comparison-MOS (CMOS) listening tests to compare TTS performance of the SSLs. Each test consists of two speech audio clips synthesized by two models with the same input text, and the listener is asked to choose between 3 scales of preferences towards either audio or a no-preference,




Fig. 2. Resynthesis performance comparison between representations on the two corpora.

²The model is called “Fastpitch 1.1” in the official implementation.

Rep. 1 vs. Rep. 2		CMOS (pref. for Rep. 1)	
		Read (LJSpeech)	Spontaneous (TSGD)
w2v2.0-L9	w2v2.0-L12	1.007*	0.626*
w2v2.0-L9	HuBERT-L12	0.793*	1.012*
w2v2.0-L9	mel-spec.	0.354*	1.877*
HuBERT-L12	mel-spec.	-0.614*	-

Table 2. Comparison-MOS (CMOS) Results from TTS subjective listening tests (*: p-value obtained from Wilcoxon signed-rank test < 0.001).

totaling 7 options. The scores thus range from [-3, 3].³ We conducted all listening tests on the crowdsourcing platform Prolific. Each set of CMOS test was completed by a separate group of at least 30 participants of native English speakers.

4.2.1. Read-speech TTS

We found through subjective listening that w2v2-L9 is superior in read-speech TTS than other representations, thus we conducted four sets of CMOS tests to confirm this finding as shown in the LJSpeech column of Table 2. We did not test w2v1-c in listening tests due to its low output intelligibility and overall bad synthesis quality that is clearly much inferior than other representations. The test input are 40 meaningful sentences chosen from the test set of the corpus prior to testing. The CMOS results (Table 2) show that w2v2-L9 outperforms all other SSLs and mel-spec. HuBERT-L12 does not perform better than mel-spec. These results show that, in a standard read TTS corpus, using SSL can perform better than using mel-spectrogram as representation medium in a two-stage TTS pipeline. However, the performance of different SSLs also differ greatly.

4.2.2. Spontaneous TTS

We conducted same set of CMOS tests as in read-speech TTS to confirm that w2v2-L9 is also superior in spontaneous TTS. We did not test w2v1-c in this setting due to its clearly low output quality similar to that in read-speech TTS.

The original test set from the spontaneous corpus contains largely context-dependent utterances which makes sentence-wise evaluation difficult. We instead fine-tuned a pre-trained generative language model on the training set and used it to generate 20 independently coherent sentences.

The results are shown in TSGD column in Table 2. w2v2-L9 again outperforms all other SSLs and mel-spec. This is consistent with results from read speech TTS (Sec. 4.2.1). Notably the preference for w2v2-L9 over mel-spec is much higher in spontaneous TTS than in read TTS, suggesting that

³This setup of CMOS is as described in ITU-T P808 (“Subjective evaluation of speech quality with a crowdsourcing approach”).

SSL could be especially suitable as intermediate representation in two-stage spontaneous TTS.

5. DISCUSSION

It is clear from Table 1 and 2 that resynthesis performance does not predict overall TTS performance. This suggests that the better SSL for two-stage TTS is not necessarily the one that contains more acoustic information. If we look at other aspects of the tested SSLs as summarized in Table 1, it is also hard to find one single axis of variation that fully explains either resynthesis or TTS performance ranking. A wider comparison with other SSL models and more layers of the same SSL could uncover which SSL and which layer is most suitable for TTS. Another potential addition to further comparison is the use of weighted [4] or stacked SSL layers.

From the perspective of speech SSL research, our study shows that TTS (in both read and spontaneous speech settings) is another feasible downstream task, the performance of which could be taken into consideration when evaluating speech SSLs. However, we note that, (1) the specific TTS models could have bias towards which SSLs they work better with, (2) unlike other more easily benchmarkable downstream tasks, TTS is known to be complex to evaluate quantitatively or subjectively [26]. But regarding the second point, we also note that SSL itself has already emerged as a promising tool for automatic evaluation of TTS systems [27].

6. CONCLUSIONS

We compared several speech SSLs as the representation medium in a standard two-stage TTS pipeline, on both read and spontaneous speech corpora. We used the same models and training setup for both stages of the TTS pipeline for fair comparison. Listening tests were conducted in comparative-MOS (CMOS) form. The results show that one tested SSL, layer 9 of 12-layer wav2vec2.0-ASR, outperforms other tested SSLs as well as the mel-spectrogram on both read and spontaneous TTS. Our results also show how utilizing SSL in TTS is a promising new research direction, and that it effectively addresses challenges in spontaneous speech synthesis. Interesting directions for future work include wider comparisons of SSLs in TTS and further optimizing SSL towards the spontaneous TTS task.

7. REFERENCES

- [1] H. Siuzdak, P. Dura, P. van Rijn, and N. Jacoby, “WavThruVec: Latent speech representation as intermediate features for neural speech synthesis,” in *Proc. Interspeech*, 2022, pp. 833–837.
- [2] C. Du, Y. Guo, X. Chen, and K. Yu, “VQTTs: High-fidelity text-to-speech synthesis with self-supervised

- VQ acoustic feature,” in *Proc. Interspeech*, 2022, pp. 1596–1600.
- [3] S.-H. Lee, S.-B. Kim, J.-H. Lee, E. Song, M.-J. Hwang, and S.-W. Lee, “HierSpeech: Bridging the gap between text and speech by hierarchical variational inference using self-supervised representations for speech synthesis,” in *Proc. NeurIPS*, 2023.
- [4] Shu-wen et al. Yang, “SUPERB: Speech Processing Universal PERFORMANCE Benchmark,” in *Proc. Interspeech*, 2021, pp. 1194–1198.
- [5] S. Schneider, A. Baevski, R. Collobert, and M. Auli, “wav2vec: Unsupervised pre-training for speech recognition,” in *Proc. Interspeech*, 2019, pp. 3465–3469.
- [6] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” in *Proc. NeurIPS*, 2020, vol. 33, pp. 12449–12460.
- [7] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, “HuBERT: Self-supervised speech representation learning by masked prediction of hidden units,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3451–3460, 2021.
- [8] A. Polyak, Y. Adi, J. Copet, E. Kharitonov, K. Lakhotia, W.-N. Hsu, A. Mohamed, and E. Dupoux, “Speech resynthesis from discrete disentangled self-supervised representations,” in *Proc. Interspeech*, 2021, pp. 3615–3619.
- [9] K. Lakhotia, E. Kharitonov, W.-N. Hsu, Y. Adi, A. Polyak, B. Bolte, T.-A. Nguyen, J. Copet, A. Baevski, A. Mohamed, et al., “On generative spoken language modeling from raw audio,” *Transactions of the Association for Computational Linguistics*, vol. 9, pp. 1336–1354, 2021.
- [10] D. Ma, N. Ryant, and M. Liberman, “Probing acoustic representations for phonetic properties,” in *Proc. ICASSP*, 2021, pp. 311–315.
- [11] A. Pasad, J.-C. Chou, and K. Livescu, “Layer-wise analysis of a self-supervised speech representation model,” in *Proc. ASRU*, 2021, pp. 914–921.
- [12] Y. Li, Y. Mohamied, P. Bell, and C. Lai, “Exploration of a self-supervised speech model: A study on emotional corpora,” in *Proc. SLT*, 2022.
- [13] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerrv-Ryan, et al., “Natural TTS synthesis by conditioning wavenet on mel spectrogram predictions,” in *Proc. ICASSP*, 2018, pp. 4779–4783.
- [14] A. Lańcucki, “FastPitch: Parallel text-to-speech with pitch prediction,” in *Proc. ICASSP*, 2021, pp. 6588–6592.
- [15] Y. Ren, C. Hu, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T.-Y. Liu, “FastSpeech 2: Fast and high-quality end-to-end text to speech,” in *Proc. ICLR*, 2021.
- [16] J. Kim, S. Kim, J. Kong, and S. Yoon, “Glow-TTS: A generative flow for text-to-speech via monotonic alignment search,” *Proc. NeurIPS*, vol. 33, pp. 8067–8077, 2020.
- [17] J. Kong, J. Kim, and J. Bae, “HiFi-GAN: Generative adversarial networks for efficient and high fidelity speech synthesis,” *Proc. NeurIPS*, pp. 17022–17033, 2020.
- [18] X. Tan, T. Qin, F. Soong, and T.-Y. Liu, “A survey on neural speech synthesis,” *arXiv preprint arXiv:2106.15561*, 2021.
- [19] É. Székely, G. E. Henter, J. Beskow, and J. Gustafson, “How to train your fillers: uh and um in spontaneous speech synthesis,” in *Proc. SSW*, 2019, vol. 10, pp. 245–250.
- [20] É. Székely, G. E. Henter, J. Beskow, and J. Gustafson, “Spontaneous conversational speech synthesis from found data,” in *Proc. Interspeech*, 2019, pp. 4435–4439.
- [21] L.-W. Chen, S. Watanabe, and A. Rudnicky, “A vector quantized approach for text to speech synthesis on real-world spontaneous speech,” *arXiv preprint arXiv:2302.04215*, 2023.
- [22] K. Ito and L. Johnson, “The LJ Speech dataset,” <https://keithito.com/LJ-Speech-Dataset/>, 2017.
- [23] Y. Ferstl and R. McDonnell, “Investigating the use of recurrent motion modelling for speech gesture generation,” in *Proc. IVA*, 2018, pp. 93–98.
- [24] É. Székely, G. E. Henter, J. Beskow, and J. Gustafson, “Breathing and speech planning in spontaneous speech synthesis,” in *Proc. ICASSP*, 2020, pp. 7649–7653.
- [25] R. Badlani, A. Lańcucki, K. J. Shih, R. Valle, W. Ping, and B. Catanzaro, “One tts alignment to rule them all,” in *Proc. ICASSP*, 2022, pp. 6092–6096.
- [26] P. Wagner, J. Beskow, S. Betz, J. Edlund, J. Gustafson, G. Eje Henter, S. Le Maguer, Z. Malisz, É. Székely, C. Tännander, et al., “Speech synthesis evaluation—state-of-the-art assessment and suggestion for a novel research program,” in *Proc. SSW*, 2019.
- [27] E. Cooper, W.-C. Huang, T. Toda, and J. Yamagishi, “Generalization ability of mos prediction networks,” in *Proc. ICASSP*, 2022, pp. 8442–8446.