



EXPRESSO: A Benchmark and Analysis of Discrete Expressive Speech Resynthesis

Tu Anh Nguyen*, Wei-Ning Hsu*, Antony D’Avirro*, Bowen Shi*, Itai Gat, Maryam Fazal-Zarani, Tal Remez, Jade Copet, Gabriel Synnaeve, Michael Hassid, Felix Kreuk, Yossi Adi⁺, Emmanuel Dupoux⁺

Meta AI

{ntuanh, adiyoss, dpx}@meta.com

Abstract

Recent work has shown that it is possible to resynthesize high-quality speech based, not on text, but on low bitrate discrete units that have been learned in a self-supervised fashion and can therefore capture expressive aspects of speech that are hard to transcribe (prosody, voice styles, non-verbal vocalization). The adoption of these methods is still limited by the fact that most speech synthesis datasets are read, severely limiting spontaneity and expressivity. Here, we introduce EXPRESSO, a high-quality expressive speech dataset for textless speech synthesis that includes both read speech and improvised dialogues rendered in 26 spontaneous expressive styles. We illustrate the challenges and potentials of this dataset with an *expressive resynthesis benchmark* where the task is to encode the input in low-bitrate units and resynthesize it in a target voice while preserving content and style. We evaluate resynthesis quality with automatic metrics for different self-supervised discrete encoders, and explore tradeoffs between quality, bitrate and invariance to speaker and style. The dataset, evaluation metrics and baseline models will be open sourced¹.

Index Terms: Speech synthesis evaluation, self-supervised speech representations, expressive synthesis

1. Introduction and related work

Speech synthesis has been traditionally approached as a mapping between text and speech. This has a series of limiting consequences: text is an impoverished representation of language, that does not specify many expressive dimensions: rhythm, intonation, emotion, emphasis, and often fails to encode non-verbals like laughter, cries, lip smacks, etc. As a result, speech synthesizers typically resort to standard read speech as the main target, which severely limits the expressivity of AI systems.

Everything changed with the advent of Self Supervised Learning (SSL) speech models [1, 2], which enabled to build discrete representations for speech without needing any textual annotation [3, 4]. Because such representations can be learned from much more diverse audio than read speech (conversational, casual speech), this opens up the possibility to build more expressive systems based on SSL units instead of text. Recently, [5, 6] utilize neural audio codecs [7, 8] to encode speech features into codes and generate natural speech from textual input. However, these models still rely on large-scale read speech that lacks expressivity. One of the roadblocks in building such expressive systems is the lack of datasets that are sufficiently expressive and of high audio quality for learning a synthesizer. Most existing expressive datasets (e.g., EmoV [9]) have used expressive reading, where voice actors read a (neutral) sentence

in different expressions (happy, sad, etc.). This method puts the voice actors in an artificial situation, resulting in not very plausible rendering of these expressions. In this work, we also reproduced this protocol to create an expressive speech dataset but added a section based on conversational improvisation. The two actors are prompted with a situation and a character (e.g., two drivers involved in a car accident, a parent and a child, etc.) and improvise a dialogue impersonating their characters. This yields much more realistic and casual speech, with spontaneous hesitations, laughter, etc., that would be extremely hard to transcribe accurately, but, which can be in principle captured by SSL units.

Next, we illustrate the potential of this dataset by setting up the task of discrete expressive resynthesis. As in discrete resynthesis [10], it consists in taking audio as input, encoding it in low bitrate discrete units, and synthesizing the same content with a different target voice. In this work, we additionally include the task of preserving expressive style. We introduce automatic metrics for content and style preservation and evaluate a HuBERT [1] encoder trained on a masked prediction objective followed by k-means clustering, which we compare to Encodec [7], a compression model which acts as a high bitrate baseline. We compare different pretraining sources for the HuBERT units based on public datasets of read speech and/or spontaneous speech. Synthesis is done with a units-based HiFi-GAN vocoder [10] conditioned on speaker or speaker and style.

2. The EXPRESSO dataset

The EXPRESSO dataset consists of 47 hours of expressive speech from 4 speakers of North American English. The dataset is divided into two main sections: an *expressive reading* section (37% of the corpus) where actors read short prompts in a parallel fashion in 7 different styles, with additional long-form and emphasis material (see 2.1), and an *improvised dialog* section (72% of the corpus) where pairs of actors are prompted to improvise a conversation in a fictive setting that illustrates one of 25 specified styles (see 2.2). In a small additional singing section, actors sing a few of their favorite songs.

The 26 different styles were chosen for their universality/recognizability (i.e. common emotions like happiness, sadness), utility for current/anticipated speech applications (i.e. whispered, enunciated speech), and to elicit the large range of possible vocalizations of the human voice (including addressing or imitating a child or an animal, and non-verbals like grunting, coughing, whistling, etc., see Table 1 for a full list).

Data was recorded in a professional recording studio with minimal background noise at 48kHz/24bit. The files for read speech and singing are in a mono wav format; and for the dialog section in stereo (one channel per actor), where the original flow of turn-taking is preserved.

⁺Core contribution as first and last authors.

¹<https://speechbot.github.io/expresso/>

Table 1: EXPRESSO’s expressive styles. * singing is the only improvised style that is not in dialogue format.

Style	Read (min)	Improvised (min)	total hours
angry	-	82	1.4
animal	-	27	0.4
animal_directed	-	32	0.5
awe	-	92	1.5
bored	-	92	1.5
calm	-	93	1.6
child	-	28	0.4
child_directed	-	38	0.6
confused	94	66	2.7
default	133	158	4.9
desire	-	92	1.5
disgusted	-	118	2.0
enunciated	116	62	3.0
fast	-	98	1.6
fearful	-	98	1.6
happy	74	92	2.8
laughing	94	103	3.3
narration	21	76	1.6
non_verbal	-	32	0.5
projected	-	94	1.6
sad	81	101	3.0
sarcastic	-	106	1.8
singing*	-	4	.07
sleepy	-	93	1.5
sympathetic	-	100	1.7
whisper	79	86	2.8

2.1. Expressive reading

Seven of the styles (confused, default, enunciated, happy, laughing, sad, whisper) were applied in a parallel fashion to the same set of prompts, so content did not necessarily reflect the emotion being conveyed, and we relied on the actor’s expertise to convey the desired style. Written instructions were delivered for each style describing the upper and lower performative bounds of the style, some including video examples.

Aside from a small corpus of shared essential lines (greetings, common phrases, numbers, letters), each speaker had a unique script in order to maximize linguistic diversity across the entire dataset. In total, the written corpus contains roughly 21,000 words over 2,400 unique lines. Material was scraped from open-source datasets like Wikipedia and commissioned datasets containing voice-assistant style utterances and then proofread and scrubbed for PII. Although not specifically balanced for phonetic coverage, the corpus was tuned both overall and per-speaker for a desired ratio of statements, questions, exclamations, jokes, etc.

Contrastive emphasis. We enclose certain words/spans in asterisks to denote emphasis, designed to convey contrastive focus in the reading of an utterance. Actors were trained to read this syntax with the desired prosodic effect. These occur in isolation throughout the read-speech corpus, but also in a subset of each speaker’s lines labeled “emphasis” where the same line is repeated 2 to 4 times with contrastive emphasis placed on different words/spans.

Long-form material. To capture longer-range prosodic dependencies, each speaker’s script contains one news article (read in the “default” style) and one long-form narrative piece (read in the “narration” style), roughly totalling 100 lines per speaker.

2.2. Improvised dialogs section

Dialogs were elicited via a set of situational prompts designed to evoke the desired styles or emotions. Some prompts resemble

Table 2: Encoder and Tokenizer used for our discrete units. Bitrate is $\log_2(\text{codebook size}) \times n$ units per sec in BPS. Mean within- and across-speaker ABX discrimination scores, resp., on 1-hot vectors and units’ centroids. PNMI is mutual information between units and phonemes. For Encodec RVQ8, we concatenate multiple codebooks for ABX. PNMI is not available for this tokenizer.

Model	Tokenizer	BPS	ABX 1-hot (\downarrow)		ABX-centr. (\downarrow)		PNMI (\uparrow)	
			LS	Fisher	LS	Fisher	LS	Fisher
HuBERT (LS960)	KM500 (LS960)	450	8.98	15.45	5.28	10.84	67.57	48.42
	KM2000 (Expr)	550	11.00	17.81	4.32	9.07	73.23	56.33
HuBERT (Mix1)	KM2000 (Mix1)	550	10.27	14.96	4.92	8.93	72.36	55.77
	KM2000 (Expr)	550	10.50	15.27	4.39	8.04	74.37	59.14
Encodec	RVQ1	500	45.47	44.29	26.10	29.57	21.42	13.17
	RVQ8	4000	41.38	40.50	20.20	25.60	-	-

voice-application domains such as reporting the weather, navigation, information retrieval, while others are more open-ended scenarios; some of them were proposed by the actors.

Tradeoffs were made to capture usable data while preserving the feel of natural conversation, i.e. actors were recorded in separate booths but watching each other over video conference. A small number of dialogs (<10) were interrupted by the studio director to provide notes. These dialogs were edited to remove the pause and maintain a more natural flow.

2.3. Singing section

Each speaker recorded several versions of popular nursery rhymes and public domain songs. The original recording had more data (93 minutes total) but could not be shared because the songs turned out not to be in the public domain.

2.4. Data preparation

The raw dataset contains mostly pre-cut segments of 3-4 seconds for the read section, except for long-form ones, and long waveforms ranging from 2 to 10 minutes for dialogs section. For the purpose of speech synthesis, we cut long files into segments of 15 seconds. We split the dataset into train/dev/test subsets such that each speaker-style contains roughly 60s in dev/test splits, resulting in 1.5 hours of speech in each subset.

3. Method

Additional datasets. We refer to the read section of EXPRESSO as Exp-R and the improvised section as Exp-I. To train the units and vocoder, and to evaluate the results, we use additional open source datasets: LJspeech [11] (LJ), VCTK [12], Librispeech dev-other [13] (LS), Fisher [14] and EmoV-DB (EmoV) [9].

Evaluation metrics For evaluating the discrete units, we compute their bitrate, ABX discrimination (the probability that the DTW distance between minimally different triphones like /bit/ versus /bet/ are more distant to one another than two instances of the same triphone [15]) both on the 1-hot representations (as in [16]), and using the dense embedding corresponding to the centroid of the units, and PNMI (phone-normalized mutual information between units z and ground truth phonemes y : $I(y; z)/H(y)$ as in [1]).

For evaluating the quality of resynthesized speech, we build automatic metrics for the preservation of content, pitch, and expressivity. As in [10], content preservation is evaluated by running a publicly available Automatic Speech Recognition (ASR) model [17] on the resynthesized sentence and computing the

Table 3: Content preservation evaluation: WERs (%) of speech resynthesized by our models. We bold the best HuBERT or best Encodec model within each column. We denote speaker conditioning as S, and speaker with expression conditioning as S_E. Results are reported for Espresso (E), VCTK (V) and Fisher (Fish).

				In Domain Source						Out of Domain Source								
				Same Speaker			Swapped Speaker			LS		Fish		LS		Fish		
Model	Tokenizer	Data	Src. Tgt.	LJ	V	E	V	E	LS	Fish	LS	Fish	LS	Fish	LS	Fish		
				LJ	V	E	V	E	LJ	LJ	V	V	E	E	Orig Voices			
				WER														
Original audio				2.04	1.74	14.76	-	-	-	-	-	-	-	-	-	-	3.55	30.26
HuBERT (LS960)	KM500 (LS960)	LJ+V	S	3.12	6.85	-	7.20	-	11.56	50.13	10.83	48.26	-	-	-	-	-	
		E+LJ+V	S	3.22	7.19	24.21	6.80	24.34	11.61	49.24	10.42	46.63	10.93	47.18	-	-	-	
		E+LJ+V	S_E	3.65	6.76	23.65	7.79	24.82	12.00	50.49	10.75	47.72	10.57	46.57	-	-	-	
	KM2000 (E)	LJ+V	S	2.98	7.15	-	7.09	-	10.95	47.44	9.98	47.06	-	-	-	-	-	
		E+LJ+V	S	3.41	7.09	21.64	7.01	22.80	10.80	46.90	10.20	45.79	10.61	46.77	-	-	-	
		E+LJ+V	S_E	2.83	6.48	22.35	6.56	21.92	10.34	45.93	9.72	45.08	9.52	43.13	-	-	-	
HuBERT (Mix1)	KM2000 (Mix1)	LJ+V	S	2.60	6.98	-	7.60	-	9.60	41.78	8.34	40.53	-	-	-	-		
		E+LJ+V	S	2.80	6.84	21.25	7.20	22.52	9.17	40.61	8.38	38.91	9.76	42.87	-	-	-	
		E+LJ+V	S_E	2.85	7.17	20.36	7.33	20.81	9.50	41.09	8.92	40.82	8.39	38.47	-	-	-	
	KM2000 (E)	LJ+VCTK	S	2.77	5.60	-	5.89	-	9.48	41.42	8.39	40.81	-	-	-	-		
		E+LJ+V	S	2.95	4.85	20.64	5.07	21.01	9.04	39.62	7.91	38.45	8.46	39.84	-	-	-	
		E+LJ+V	S_E	3.05	5.48	19.52	5.59	20.27	9.20	38.79	7.75	37.48	8.00	36.67	-	-	-	
Encodec (RVQ-1)	-	None	5.52	17.46	34.36	-	-	-	-	-	-	-	-	-	18.88	60.68		
Encodec (RVQ-8)	-	None	2.20	2.52	16.85	-	-	-	-	-	-	-	-	-	4.62	35.64		

Word Error Rate (WER) relative to the transcription of the input sentence². We run this on in-domain inputs (LJ, VCTK and Exp-R) and out-of-domain inputs (LS and Fisher). Pitch preservation is evaluated by computing F0 Frame Error (FFE), which measures the percentage of frames with a deviation of more than 20% in pitch value between the input and resynthesized output. Expressivity preservation is computed by training an expressive style classifier³ on the train set of EXPRESSO and applying it to resynthesized versions of its dev set. These classifiers are also run on the original data for comparison.

4. Models

4.1. Unit encoding

For unit encoding, we compared three models: two HuBERT-based, one Encodec-based. The HuBERT models use the same architecture (HuBERT base with 12 Transformer layers), but are trained on different corpora. HuBERT-LS960 was trained on LibriSpeech 960 as in [3]. We used the available model in textless-lib [21]. We use HuBERT-Mix1 from [22], which was trained with a more varied mixture of datasets: an 8 language subset of VoxPopuli [23] (167K h), Common Voice [24] (4K h) and Multilingual LibriSpeech (MLS) [25] (50K h), totalling 221K hours. For quantization, we trained k-means models on HuBERT features either on a subset of HuBERT pre-training dataset or on EXPRESSO, with k=500 on LS960 or k=2000 on other datasets. The Encodec model is from [7], we used two models with 1 and 8 codebooks of cardinality 1024, which were trained on VoxPopuli400k English, People’s Speech [26], LibriSpeech 960, LibriLight [27] and Spotify [28]. Training hyperparameters were identical to [7] except for having no audio normalization and using zero padding instead of reflect.

4.2. Vocoder

For HuBERT units, we produce the waveform using HiFiGAN, which we train on the units presented above. For each set of units, we train either on LJ and VCTK, or on LJ, VCTK and EXPRESSO. The vocoder is either conditioned on speaker ID (using a look-up table), or on speaker ID and expression ID (also using a look-up table). We distinguish the read and improvised versions of the expressions, yielding a total of 34 ex-

pressions. For Encodec units, we used the Encodec decoder to directly produce the waveform, and compared systems with 1 and 8 codebooks.

5. Results

Table 2 shows the phonetic quality metrics across different SSL units. The HuBERT encoders trained on the larger and noisier corpus (Mix1) tend to have overall better results than when trained on LS960 only, especially when tested on Fisher. The ABX-centroid and PNMI metrics gave better results when k-means clustering was run on EXPRESSO (a small high quality, high diversity dataset) than on the large dataset used to train HuBERT itself. This was not the case, however with the ABX 1-hot metric, so further study is necessary to confirm this result. The Encodec units gave poor results, which is not surprising given that Encodec units are generic representations trained for audio compression that also encodes non-phonetic variations whereas HuBERT units are trained with a masked language modeling objective.

Table 3 shows the result of content preservation for the resynthesis task (WER), distinguishing the case where the input sentences are drawn from the same distribution as the vocoder training sets (In-domain Source: LJ, VCTK and Exp-R) and when the input sentences are from different datasets (Out-of-domain Source: LS and Fisher). For in-domain, we further distinguish the cases where the input voice is the same as the target voice (Same Speaker) and when the target voice is randomly sampled from the same training set (Swapped Speaker; not applicable to LJ speech). We find that swapping speakers costs on average only a small decrement in WER (3% relative), suggesting that the units are well (although not totally) disentangled from speaker information. On average, the HuBERT units trained on Mix1 give better performances (about 10% relative) than units trained on LS960, a result consistent with the phonetic quality metrics. As for the vocoder, the training voices and conditioning (either speaker alone or speaker+style) did not give systematic results. Overall, the best HuBERT results on the same speaker showed a drop in performance compared to the original audio files (between 30% relative to more than double the error), and compared to the Encodec-8 units. Regarding out-of-domain resynthesis, the HuBERT-Mix1 units again generally outperform the HuBERT-LS960 units (19% relative). In addition, the tokenizer trained on EXPRESSO tend to be better by a small margin (3% relative). The best resynthesis models suffer from a large drop in WER compared to original audio and

²https://dl.fbaipublicfiles.com/fairseq/wav2vec/wav2vec_vox_960h_pl.pt

³We fine-tune the wav2vec2 base model [18] on a 26 style classes audio classification task (as in the SUPERB benchmark [19]) using Huggingface transformers library [20]).

Table 4: Expressive style classification accuracy using a pre-trained emotion classification model. We denote speaker conditioning as *S*, and speaker with expression conditioning as *S.E*. Results are reported for Expresso (*E*) and VCTK (*V*).

Model	Tokenizer	Data	Src. Tgt. Cond.	Same		Swapped		Zero-shot				Out of dom.
				E.R	E.I	E.R	E.I	E.R	E.I	E.R	E.I	EmoV
				E	E	E	E	LJ	LJ	V	V	E
<i>Original Audio</i>				92.47	75.69	-	-	-	-	-	-	27.46
HuBERT (LS960)	KM500 (LS960)	LJ+V	S	-	-	-	-	13.36	9.14	26.42	12.43	11.56
		E+LJ+V	S	33.18	14.99	29.80	13.53	8.45	10.79	28.26	10.97	
		E+LJ+V	S.E	81.57	58.68	81.26	56.12	23.96	36.93	56.07	28.15	
	KM2000 (E)	LJ+V	S	-	-	-	-	5.22	10.42	27.04	11.70	9.25
		E+LJ+V	S	39.17	23.95	33.95	19.38	11.67	12.07	27.34	14.63	
		E+LJ+V	S.E	78.34	62.16	76.96	54.11	22.12	39.31	55.76	32.54	
HuBERT (Mix1)	KM2000 (Mix1)	LJ+V	S	-	-	-	-	7.99	9.32	27.50	11.52	11.27
		E+LJ+V	S	25.81	17.73	28.57	15.72	5.53	8.96	27.65	11.33	
		E+LJ+V	S.E	78.80	61.06	81.41	58.14	31.64	40.77	40.86	31.63	
	KM2000 (E)	LJ+V	S	-	-	-	-	7.68	10.24	27.80	12.07	14.45
		E+LJ+V	S	37.02	16.82	35.33	16.09	5.99	9.69	26.73	11.52	
		E+LJ+V	S.E	72.81	62.16	73.12	55.76	31.18	39.85	55.61	28.52	
Encodec (RVQ-1)	-	-	None	57.76	44.42	-	-	-	-	-	-	22.25
Encodec (RVQ-8)	-	-	None	78.65	64.53	-	-	-	-	-	-	26.88

Table 5: F0 Frame Error (FFE). Bold best absolute scores. We denote speaker conditioning as *S*, and speaker with expression conditioning as *S.E*. Results are reported for Expresso (*E*), Expresso Read (*E.R*), Expresso Improvised (*E.I*), LJ, VCTK (*V*), and EMOV.

Model	Tokenizer	Data	Src. Tgt. Cond.	Same		Swapped		Zero-shot				OOD
				E.R	E.I	E.R	E.I	E.R	E.I	E.R	E.I	EMOV
				E	E	E	E	LJ	LJ	V	V	E
FFE												
HuBERT (LS960)	KM500 (LS960)	E+LJ+V	S	0.31±0.10	0.33±0.12	0.37±0.12	0.38±0.13	0.34±0.10	0.35±0.12	0.38±0.11	0.38±0.14	0.27±0.08
		E+LJ+V	S.E	0.27±0.13	0.30±0.13	0.34±0.16	0.36±0.15	0.33±0.15	0.35±0.14	0.34±0.16	0.36±0.15	0.25±0.09
		E+LJ+V	S	0.31±0.12	0.33±0.13	0.36±0.13	0.36±0.13	0.35±0.09	0.35±0.11	0.38±0.11	0.38±0.14	0.26±0.09
	KM2000 (E)	E+LJ+V	S.E	0.26±0.13	0.29±0.13	0.34±0.16	0.36±0.15	0.31±0.14	0.33±0.13	0.33±0.16	0.35±0.15	0.26±0.09
		E+LJ+V	S	0.32±0.11	0.33±0.13	0.38±0.12	0.37±0.13	0.36±0.09	0.35±0.11	0.38±0.12	0.38±0.14	0.26±0.09
		E+LJ+V	S.E	0.28±0.14	0.29±0.13	0.34±0.16	0.36±0.15	0.32±0.15	0.34±0.14	0.34±0.16	0.36±0.15	0.27±0.09
HuBERT (Mix1)	KM2000 (Mix1)	E+LJ+V	S	0.32±0.11	0.33±0.13	0.38±0.12	0.37±0.13	0.36±0.09	0.35±0.11	0.38±0.12	0.38±0.14	0.26±0.09
		E+LJ+V	S.E	0.31±0.10	0.32±0.12	0.37±0.12	0.37±0.13	0.35±0.09	0.35±0.11	0.38±0.11	0.37±0.14	0.26±0.08
		E+LJ+V	S	0.27±0.13	0.30±0.13	0.34±0.16	0.36±0.14	0.32±0.14	0.34±0.14	0.34±0.16	0.36±0.15	0.25±0.09
	KM2000 (E)	E+LJ+V	S	0.08±0.04	0.11±0.07	-	-	-	-	-	-	0.09±0.06
		E+LJ+V	S.E	0.04±0.02	0.05±0.03	-	-	-	-	-	-	0.04±0.02
		E+LJ+V	S	-	-	-	-	-	-	-	-	-

Encodec-8 for LS (twice the errors) but a much smaller drop for Fisher. Encodec-1 consistently underperform HuBERT-based resynthesis, indicating that one Encodec codebook of size 1024 is not enough to fully capture linguistic content.

Tables 4 and 5 show the results on style and pitch preservation, respectively; these experiments are exclusively ran on EXPRESSO inputs (in-domain) or EmoV inputs (out-of-domain). We group the in-domain results in 3 conditions. The first two are similar to the same-speaker and swapped speaker conditions discussed above. Unsurprisingly, for both style and pitch, the results are uniformly better when the vocoder was conditioned on speaker+expression than in speaker alone. Note, though, that even without expression conditioning, the style classification score is much higher than the chance level (3.8%), suggesting that style is partly transmitted through the units. The results on swapped speakers are slightly worse than on same speaker (cost around 10% relative accuracy on style and 20% rel. on pitch error). Globally, the style scores of Encodec-8 are on par or better than the best HuBERT resynthesis models, but much better for pitch preservation (6-fold). Next, we explored whether expressivity could be transferred to voices that were only trained in the default read speech style (VCTK and LJ). This change cuts the style score in half, still remaining way above chance, and better than the EXPRESSO voices not conditioned by expression, showing that expressive styles can, to a certain extent, generalize to untrained voices in a zero-shot fashion. Finally, we tested expressive resynthesis could be applied out-of-domain, using input data from EmoV. This dataset uses 5 expressions (neutral, amused, angry, sleepy, disgust) which we mapped based on the description to EXPRESSO’s neutral, laughing, angry, sleepy and

disgusted. The performance of the classifier was much lower than in the in-domain case, and unexpectedly higher for resynthesis than the original file. Inspection of the errors pattern showed some systematic style confusions across datasets. For instance, original EmoV voices in angry style were classified as “projected” by our classifier, but as “angry” once resynthesized with the “angry” conditioning. This suggests discrepancies in style rendering across the two datasets for identical labels (e.g., anger rendered as shouting in one versus cold rage in the other). This also suggests that the style-conditioned vocoder can to a certain extent remap input styles to different ones (confirmed by a style swapping experiment not reported in the table). More research and expressive datasets are needed to develop a dataset-independent and speaker-independent expressive style classifiers. Despite these limitations, the results were congruent with in the in-domain case, with better than chance performance and improved performance with expression-conditioned vocoders.

6. Conclusion

We presented a new dataset for expressive discrete resynthesis, and analysed content and style preservation for several baseline discrete SSL models. We showed that Encodec systems which are designed for general audio compression are generally better for resynthesis, although they lack the controllability in output voices and style made possible by the fact that HuBERT units are disentangled from speaker identity and (partially) from expressivity. Further work is needed to improve HuBERT-based expressive resynthesis, and reach the quality of Encodec units, while retaining controllability. In particular, improved performance on pitch preservation could be obtained by conditioning the vocoder on discrete pitch units, as in [10].

7. References

- [1] 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 Trans. Audio, Speech, Language Process. (TASLP)*, vol. 29, pp. 3451–3460, 2021.
- [2] Y. Chung, Y. Zhang, W. Han, C. Chiu, J. Qin, R. Pang, and Y. Wu, "W2v-bert: Combining contrastive learning and masked language modeling for self-supervised speech pre-training," *CoRR*, vol. abs/2108.06209, 2021. [Online]. Available: <https://arxiv.org/abs/2108.06209>
- [3] K. Lakhotia, E. Kharitonov, W. Hsu, Y. Adi, A. Polyak, B. Bolte, T. A. Nguyen, J. Copet, A. Baevski, A. Mohamed, and E. Dupoux, "Generative spoken language modeling from raw audio," *CoRR*, vol. abs/2102.01192, 2021. [Online]. Available: <https://arxiv.org/abs/2102.01192>
- [4] Z. Borsos, R. Marinier, D. Vincent, E. Kharitonov, O. Pietquin, M. Sharifi, O. Teboul, D. Grangier, M. Tagliasacchi, and N. Zeghidour, "Audiolm: a language modeling approach to audio generation," 2022. [Online]. Available: <https://arxiv.org/abs/2209.03143>
- [5] C. Wang, S. Chen, Y. Wu, Z. Zhang, L. Zhou, S. Liu, Z. Chen, Y. Liu, H. Wang, J. Li, L. He, S. Zhao, and F. Wei, "Neural codec language models are zero-shot text to speech synthesizers," 2023.
- [6] E. Kharitonov, D. Vincent, Z. Borsos, R. Marinier, S. Girgin, O. Pietquin, M. Sharifi, M. Tagliasacchi, and N. Zeghidour, "Speak, read and prompt: High-fidelity text-to-speech with minimal supervision," 2023.
- [7] A. Défossez, J. Copet, G. Synnaeve, and Y. Adi, "High fidelity neural audio compression," *arXiv preprint arXiv:2210.13438*, 2022.
- [8] N. Zeghidour, A. Luebs, A. Omran, J. Skoglund, and M. Tagliasacchi, "Soundstream: An end-to-end neural audio codec," *IEEE/ACM Trans. Audio, Speech and Lang. Proc.*, vol. 30, p. 495–507, nov 2021. [Online]. Available: <https://doi.org/10.1109/TASLP.2021.3129994>
- [9] A. Adigwe, N. Tits, K. E. Haddad, S. Ostadabbas, and T. Dutoit, "The emotional voices database: Towards controlling the emotion dimension in voice generation systems," *arXiv preprint arXiv:1806.09514*, 2018.
- [10] A. Polyak, Y. Adi, J. Copet, E. Kharitonov, K. Lakhotia, W. Hsu, A. Mohamed, and E. Dupoux, "Speech resynthesis from discrete disentangled self-supervised representations," *CoRR*, vol. abs/2104.00355, 2021. [Online]. Available: <https://arxiv.org/abs/2104.00355>
- [11] K. Ito and L. Johnson, "The lj speech dataset," <https://keithito.com/LJ-Speech-Dataset/>, 2017.
- [12] C. Veaux, J. Yamagishi, and K. MacDonald, "Cstr vetk corpus: English multi-speaker corpus for cstr voice cloning toolkit," 2017.
- [13] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An asr corpus based on public domain audio books," in *ICASSP*, 2015, pp. 5206–5210.
- [14] C. Cieri, D. Miller, and K. Walker, "The fisher corpus: a resource for the next generations of speech-to-text," in *LREC*, 2004.
- [15] T. Schatz, V. Peddinti, F. Bach, A. Jansen, H. Hermansky, and E. Dupoux, "Evaluating speech features with the Minimal-Pair ABX task: Analysis of the classical MFC/PLP pipeline," in *INTERSPEECH*, 2013. [Online]. Available: <https://hal.archives-ouvertes.fr/hal-00918599>
- [16] T. A. Nguyen, M. de Seyssel, P. Rozé, M. Rivière, E. Kharitonov, A. Baevski, E. Dunbar, and E. Dupoux, "The zero resource speech benchmark 2021: Metrics and baselines for unsupervised spoken language modeling," 2020.
- [17] Q. Xu, A. Baevski, T. Likhomanenko, P. Tomasello, A. Conneau, R. Collobert, G. Synnaeve, and M. Auli, "Self-training and pre-training are complementary for speech recognition," in *ICASSP*. IEEE, 2021, pp. 3030–3034.
- [18] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, vol. 33, 2020, pp. 12 449–12 460.
- [19] S. wen Yang, P.-H. Chi, Y.-S. Chuang, C.-I. Lai, K. Lakhotia, Y. Y. Lin, A. T. Liu, J. Shi, X. Chang, G.-T. Lin, T. hsien Huang, W.-C. Tseng, K. tik Lee, D.-R. Liu, Z. Huang, S. Dong, S.-W. Li, S. Watanabe, A. rahman Mohamed, and H. yi Lee, "Superb: Speech processing universal performance benchmark," in *Inter-speech*, 2021.
- [20] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. L. Scao, S. Gugger, M. Drame, Q. Lhoest, and A. M. Rush, "Transformers: State-of-the-art natural language processing," in *Proceedings of Empirical Methods in Natural Language Processing: System Demonstrations*, 2020, pp. 38–45. [Online]. Available: <https://www.aclweb.org/anthology/2020.emnlp-demos.6>
- [21] E. Kharitonov, J. Copet, K. Lakhotia, T. A. Nguyen, P. Tomasello, A. Lee, A. Elkahky, W.-N. Hsu, A. Mohamed, E. Dupoux, and Y. Adi, "textless-lib: a library for textless spoken language processing," 2022.
- [22] W.-N. Hsu, T. Remez, B. Shi, J. Donley, and Y. Adi, "Re-verse: Self-supervised speech resynthesis with visual input for universal and generalized speech enhancement," *arXiv preprint arXiv:2212.11377*, 2022.
- [23] C. Wang, M. Riviere, A. Lee, A. Wu, C. Talnikar, D. Haziza, M. Williamson, J. Pino, and E. Dupoux, "VoxPopuli: A large-scale multilingual speech corpus for representation learning, semi-supervised learning and interpretation," in *Proceedings of Association for Computational Linguistics*, Aug. 2021, pp. 993–1003. [Online]. Available: <https://aclanthology.org/2021.acl-long.80>
- [24] R. Ardila, M. Branson, K. Davis, M. Kohler, J. Meyer, M. Henretty, R. Morais, L. Saunders, F. Tyers, and G. Weber, "Common voice: A massively-multilingual speech corpus," in *Proceedings of the Twelfth Language Resources and Evaluation Conference*. Marseille, France: European Language Resources Association, May 2020, pp. 4218–4222. [Online]. Available: <https://aclanthology.org/2020.lrec-1.520>
- [25] V. Pratap, Q. Xu, A. Sriram, G. Synnaeve, and R. Collobert, "Mls: A large-scale multilingual dataset for speech research," *ArXiv*, vol. abs/2012.03411, 2020.
- [26] D. Galvez, G. Damos, J. Torres, K. Achorn, J. Cerón, A. Gopi, D. Kanter, M. Lam, M. Mazumder, and V. Janapa Reddi, "The people's speech: A large-scale diverse english speech recognition dataset for commercial usage," in *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, J. Vanschoren and S. Yeung, Eds., vol. 1, 2021.
- [27] J. Kahn, M. Rivière, W. Zheng, E. Kharitonov, Q. Xu, P. E. Mazaré, J. Karadayi, V. Liptchinsky, R. Collobert, C. Fuegen, T. Likhomanenko, G. Synnaeve, A. Joulin, A. Mohamed, and E. Dupoux, "Libri-light: A benchmark for asr with limited or no supervision," in *ICASSP*, 2020, pp. 7669–7673, <https://github.com/facebookresearch/libri-light>.
- [28] A. Clifton, S. Reddy, Y. Yu, A. Pappu, R. Rezapour, H. Bonab, M. Eskevich, G. Jones, J. Karlgren, B. Carterette, and R. Jones, "100,000 podcasts: A spoken English document corpus," in *Proceedings of the 28th International Conference on Computational Linguistics*, Dec. 2020, pp. 5903–5917. [Online]. Available: <https://aclanthology.org/2020.coling-main.519>