
CONTENTVEC: An Improved Self-Supervised Speech Representation by Disentangling Speakers

Kaizhi Qian^{*1} Yang Zhang^{*1} Heting Gao² Junrui Ni² Cheng-I Jeff Lai³ David Cox¹
Mark Hasegawa-Johnson² Shiyu Chang⁴

Abstract

Self-supervised learning (SSL) in speech involves training a speech representation network on a large-scale unannotated speech corpus, and then applying the learned representations to downstream tasks. Since the majority of the downstream tasks of SSL learning in speech largely focus on the content information in speech, the most desirable speech representations should be able to disentangle unwanted variations, such as speaker variations, from the content. However, disentangling speakers is very challenging, because removing the speaker information could easily result in a loss of content as well, and the damage of the latter usually far outweighs the benefit of the former. In this paper, we propose a new SSL method that can achieve speaker disentanglement without severe loss of content. Our approach is adapted from the HuBERT framework, and incorporates disentangling mechanisms to regularize both the teachers (masked prediction labels) and the students (learned representations). We evaluate the benefit of speaker disentanglement on a set of content-related downstream tasks, and observe a consistent and notable performance advantage of our speaker-disentangled representations.¹

1. Introduction

Over the recent years, self-supervised learning (SSL) has emerged as a state-of-the-art solution to many speech processing problems with relatively few annotated data. The

^{*}Equal contribution ¹MIT-IBM Watson AI Lab ²University of Illinois at Urbana-Champaign ³Massachusetts Institute of Technology ⁴University of California, Santa Barbara. Correspondence to: Kaizhi Qian <kqian@ibm.com>, Yang Zhang <yang.zhang2@ibm.com>.

basic idea of SSL in speech is to train a speech representation network on large-scale *unannotated* corpora, with an objective to capture and elicit meaningful speech structures and information. The resulting speech representation is then applied to the training of downstream tasks with a small amount of annotated data. Since the speech representation is already well-structured, it reduces the dependency of downstream task training on large-scale datasets.

While speech SSL has demonstrated advantages in a surprisingly wide range of tasks, one of the primary foci of speech SSL is on tasks that process the *content* of speech, such as speech recognition/phone classification, speech content generation, *etc.* For these tasks, the most desirable speech representations should be the ones that can disentangle content information in speech from other interfering variations, such as speaker variations. However, among the most widely-used existing speech representations, few can achieve a reasonable disentanglement of speaker variations. For example, the HUBERT representation (Hsu et al., 2021) can achieve a speaker identification accuracy of up to 81.4% on the SUPERB benchmark (Yang et al., 2021). This observation suggests that there may still be room for performance gain for SSL on content-related speech processing tasks, if the disentanglement of speaker is adequately addressed.

However, it has been widely acknowledged that disentangling speakers is very challenging. Since no text annotations are accessible during the training of the speech representation network, any attempt to remove speaker variations from speech representation could easily lead to a loss of content information (Choi et al., 2021). In most content-related downstream tasks, the cost of losing content information far outweighs the advantage in disentangling speakers.

In this paper, we seek to investigate the following two research questions. First, is there a way to disentangle speaker variations during SSL training *without* significant content loss? To this end, we propose CONTENTVEC, an SSL framework that is adapted from the HUBERT training paradigm. The key idea of HUBERT is that by having some relatively poor speech representations, such as MFCC, serve as the teacher labels for the masked prediction task, one can derive speech representations (which are sometimes referred to

as *students*) that are far better than the teachers in many aspects, including content preservation. This inspires us that by combining HUBERT’s teacher-student framework with speaker disentanglement techniques, we could potentially restore the content loss caused by the latter.

This has led us to the design of CONTENTVEC, which incorporates into HUBERT three disentangling mechanisms - *disentanglement in teachers* and *disentanglement in students*, and *speaker conditioning*. Specifically, disentanglement in teachers refers to removing the speaker information from the teacher labels. Disentanglement in students refers to introducing a regularization loss that directly enforces speaker invariance on the speech representations. Speaker conditioning refers to inputting speaker information to the masked prediction task, so that the need for the speech representation to encode speaker information is relieved. As we will show, all three modules are essential in shaping the speaker information flow across the speech representation network layers, and thereby achieve a superior disentanglement quality while keeping the content information intact.

The second research question we would like to explore is: How much performance gain, if any, can speaker disentanglement in SSL features contribute to downstream tasks? Our extensive evaluation shows that speaker disentanglement can achieve a consistent performance advantage over the baseline speech representations on content-related applications. The findings of this paper can shed some light on next-generation speech representations that can supply more targeted information to the downstream tasks and enable more powerful content processing directly on speech.

2. Related Work

Voice Conversion Voice conversion is among the first research areas where speaker disentanglement is explored. The general trend follows the analysis-synthesis framework, where the analysis stage learns a speaker-independent speech representation that only preserves the content, and the synthesis stage uses the speaker-independent speech representation and the speaker-related variations to synthesize the conversion results. Much research focuses on learning better linguistic representations during the analysis stage and/or injecting speaker variations better during the synthesis stage. VAE-VC (Hsu et al., 2016) is an early attempt of directly using VAE for voice conversion. Afterward, Chou et al. (2018) disentangles more speaker variations from the latent representation by discouraging the latent representation to be classified as the source speaker using an auxiliary speaker classifier on the latent representation. In contrast, ACVAE-VC (Kameoka et al., 2019) indirectly encourages more speaker disentanglement by encouraging the conversion output to be correctly classified as the source speaker. Inspired by image style transfer, StarGAN-

VC (Kameoka et al., 2018), StarGAN-VC2 (Kaneko et al., 2019b), CycleGAN-VC (Kaneko & Kameoka, 2018), and CycleGAN-VC2 (Kaneko et al., 2019a) adapted StarGAN (Choi et al., 2018) and CycleGAN (Zhu et al., 2017) respectively for voice conversion. AutoVC (Qian et al., 2019) disentangles speakers and content by directly tuning the bottleneck dimensions of a vanilla autoencoder. The following AutoVC-F0 (Qian et al., 2020a) improves pitch disentanglement by conditioning the synthesis stage on pitch representations. VoiceMixer (Lee & Kim, 2021) improves the content loss of AutoVC using similarity-based downsampling as the bottleneck. AdaIN-VC (Chou et al., 2019) uses instance normalization to normalize out the speaker variations in the analysis stage, and AGAIN-VC (Chen et al., 2021) additionally uses an activation function to constrain the speaker variations from flowing into the synthesis stage. Instead of pursuing extreme speaker disentanglement, another slightly different track of research encourages the synthesis stage to use the supplied speaker variations by using partially disentangled content representations combined with speaker variations that are easier for the synthesis stage to utilize. SpeechSplit (Qian et al., 2020b), AutoPST (Qian et al., 2021), and NANSY (Choi et al., 2021) perturb the speaker variations during the analysis stage to encourage the synthesis stage to use the supplied more stable speaker representations. In particular, Polyak et al. (2021) and NANSY start with the self-supervised speech representations as the partially disentangled content representation.

Self-supervised Learning in Speech Learning self-supervised speech representation usually encodes the speech feature into context representations followed by pretext tasks to extract content information, which mainly has two tracks. The first track is generative learning. Chung et al. (2019); Chung & Glass (2020) uses Autoregressive Predictive Coding (APC) for self-supervised representation learning. Mockingjay (Liu et al., 2020) learns speech representation by predicting the current frame given both the past and future contexts. TERA (Liu et al., 2021) learns speech representation by reconstructing acoustic frames from their altered counterparts. DeCoAR 2.0 (Ling & Liu, 2020) reconstructs the frames from their vector-quantized counterparts. Wang et al. (2020) reconstructs masked frames. The second track is discriminative. van den Oord et al. (2018) uses contrastive predictive coding to learn multi-modal representations including speech. Wav2vec (Schneider et al., 2019) learns to predict the future samples from distractors. Wav2vec 2.0 (Baevski et al., 2020b), an end-to-end version of vq-wav2vec (Baevski et al., 2020a), learns to identify the true vq-quantized frame among the distractors using contrastive loss. Kharitonov et al. (2021) significantly improves CPC-based SSL with speech data augmentation. Zhang et al. (2020) pushes the limits of SSL using noisy student training by giant Conformer models pre-trained using wav2vec 2.0. Hubert (Hsu et al., 2021) predicts masked frames pre-

quantized using k-means. ILS-SSL (Wang et al., 2021) further improves Hubert by adding masked prediction loss on intermediate layers. Besides, there are also research using multiple tasks (Pascual et al., 2019; Ravanelli et al., 2020; Chung et al., 2021) or using both labeled and unlabeled data (Wang et al., 2020).

3. Approach

In this section, we will introduce of our approach. We use upper-cased letters, X and \mathbf{X} , to represent random scalars and vectors, respectively, and lower-cased letters, x and \mathbf{x} , to represent deterministic scalars and vectors, respectively.

3.1. Problem Formulation

Denote $\mathbf{X} = [X_1, \dots, X_T]$ as the sequence of a speech features, where X_t is the speech feature vector at frame t , and T is the total number of frames. Our goal is to learn a speech representation network $\mathbf{R} = f(\mathbf{X})$, where $\mathbf{R} = [R_1, \dots, R_T]$ and R_t is the representation for frame t . \mathbf{R} should desirably satisfy the following two properties.

- \mathbf{R} should preserve as much content information as possible, and the content information roughly corresponds to the phonetic/text transcriptions of the utterance.
- \mathbf{R} should be invariant across speaker variations.

As mentioned, the pursuit of one goal can easily compromise another. In the following, we will describe our method to strike a better balance and discuss the rationale behind.

3.2. The General Framework

The CONTENTVEC framework builds upon the mask-prediction framework of HUBERT. Specifically, there are three components in the HUBERT framework: 1) the speech representation network $f(\cdot)$, 2) the predictor $p(\cdot)$, and 3) the teacher label generator $g(\cdot)$.

During training, the speech representation network takes the partially masked speech utterance, $\tilde{\mathbf{X}}$, as the input, and produces a representation for the masked speech sequence, $\tilde{\mathbf{R}} = f(\tilde{\mathbf{X}})$. On the other hand, the teacher label generator generates a label sequence $\mathbf{L} = g(\mathbf{X})$ from the *unmasked* speech. The goal of the predictor is to predict the teacher labels \mathbf{L} from the masked speech representation $\tilde{\mathbf{R}}$. The teacher label generator $g(\cdot)$ is usually predefined and fixed during training. The other two modules, $f(\cdot)$ and $p(\cdot)$, are trained jointly to minimize the following prediction loss:

$$\mathcal{L}_{pred} = \mathbb{E}[\ell_m(p \circ f(\tilde{\mathbf{X}}), g(\mathbf{X}))], \quad (1)$$

where ℓ_m denotes the cross-entropy loss computed over the *masked* frames only. To make our description more intuitive, we will refer to $f(\tilde{\mathbf{X}})$ as *students*, and $g(\mathbf{X})$ as *teachers*.

It has been reported (Hsu et al., 2021) that even if the HUBERT teacher is poor (e.g., losing content), the student can still preserve the content far better than the teacher, thanks to the masked prediction mechanism. This observation inspires us to test the hypothesis that one can combine speaker disentanglement techniques (potentially causing loss of content) with the masked prediction framework, and in this way, preserve content more faithfully than using a speaker disentanglement algorithm on its own. Since teachers, students, and the predictor are three major components of the masked prediction, CONTENTVEC introduces three disentanglement mechanisms, *disentanglement in teachers*, *disentanglement in students*, and *speaker conditioning*, to tackle the three components respectively, as shown in Figure 1.

3.3. Disentanglement in Teachers

Disentanglement in teachers aims to remove the speaker information in the teacher labels. Recently, there has been marked progress in unsupervised voice conversion systems, which can now significantly obscure the source speaker information without losing too much content (Polyak et al., 2021). Inspired by this, we adopt a voice conversion model to convert all utterances to the same speaker before generating the teacher labels.

Specifically, as shown in Figure 1(c), the teacher labels, $\mathbf{L} = g(\mathbf{X})$, are generated via the following three steps. First, all the utterances \mathbf{X} in the training set are converted to a single speaker using a competent unsupervised voice conversion system. Second, the converted utterances are passed through a pre-trained unsupervised speech representation network, in our case HUBERT, to generate a set of speech representations, which should contain very little speaker information. Finally, the speech representations are quantized to discrete teacher labels using k -means clustering.

It is worth noting that although the teacher speech representation described above already achieves speaker disentanglement, its content preservation is not satisfactory because any voice conversion systems sometimes (for some speakers) cause a non-negligible content loss (Choi et al., 2021). In order to ameliorate this shortcoming of modern voice conversion, we use voice conversion as a teacher to train better students, instead of directly applying its output to downstream tasks.

3.4. Disentanglement in Students

Disentanglement in students enforces speaker-invariant student representations, which can be achieved with SIMCLR (Chen et al., 2020), a contrastive-learning-based algorithm.

Specifically, as shown in Figure 1(a), each speech utterance, \mathbf{X} , is passed into two random transformations that alter only the speaker information, before it is masked. Denote the two

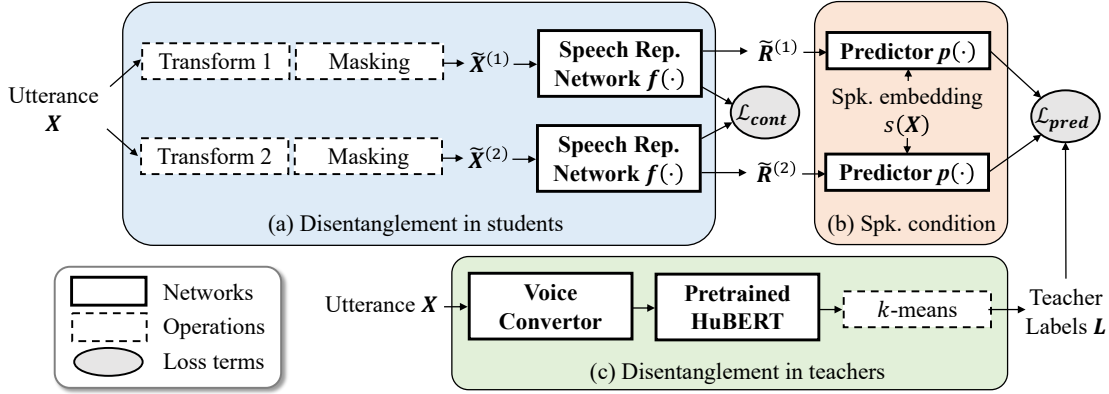


Figure 1. The overall structure of CONTENTVEC. Each colored block represents one disentanglement module in CONTENTVEC.

masked, transformed copies of \mathbf{X} as $\tilde{\mathbf{X}}^{(1)}$ and $\tilde{\mathbf{X}}^{(2)}$. Then, this pair of utterances are passed through the speech representation network, $f(\cdot)$, to generate the representations $\mathbf{R}^{(1)}$ and $\mathbf{R}^{(2)}$, and the following contrastive loss is introduced to penalize dissimilarity between $\mathbf{R}^{(1)}$ and $\mathbf{R}^{(2)}$:

$$\begin{aligned} \mathcal{L}_{contr} = & \sum_{t=1}^T \frac{\exp(\text{cossim}(\mathbf{R}_t^{(1)}, \mathbf{R}_t^{(2)})/k)}{\sum_{\tau \in \{t\} \cup \mathcal{I}_t} \exp(\text{cossim}(\mathbf{R}_t^{(1)}, \mathbf{R}_\tau^{(1)})/k)} \\ & + \sum_{t=1}^T \frac{\exp(\text{cossim}(\mathbf{R}_t^{(2)}, \mathbf{R}_t^{(1)})/k)}{\sum_{\tau \in \{t\} \cup \mathcal{I}_t} \exp(\text{cossim}(\mathbf{R}_t^{(2)}, \mathbf{R}_\tau^{(2)})/k)}, \end{aligned} \quad (2)$$

where $\text{cossim}(\cdot, \cdot)$ denotes the cosine similarity, and \mathcal{I}_t denotes a set of random time indices at which the representations are chosen as the negative examples for time t . The contrastive loss consists of two terms so that it is symmetric with respect to $\mathbf{R}^{(1)}$ and $\mathbf{R}^{(2)}$. According to Equation (2), the negative examples for the utterance pair, $(\mathbf{R}_t^{(1)}, \mathbf{R}_t^{(1)})$, are uniformly randomly drawn from the remaining frames within the *same* utterances. As an extension to Equation (2), the contrastive loss can be applied to an intermediate layer, instead of the final layer, of $f(\cdot)$. Section 3.6 will discuss how the choice of layer in which the contrastive loss is imposed would affect the disentanglement behavior.

The biggest challenge of applying the contrastive loss is how to design a random transformation that only alters the speaker identity of the utterance with minimal changes in the other aspects. To this end, we adopt the random transformation algorithm proposed by Choi et al. (2021). Specifically, the algorithm consists of three steps of transformations. First, all the formant frequencies within an utterance are scaled by a factor of ρ_1 ; second, F0 in every frame is scaled by a factor of ρ_2 ; finally, a random equalizer is applied to accommodate any channel effects. ρ_1 and ρ_2 are both randomly drawn from the uniform distribution $\mathcal{U}([1, 1.4])$, and then flipped to their reciprocals with probability 0.5. Since the majority of voice information resides in the formant frequency and F0 frequency ranges (e.g., (Eide & Gish, 1996)),

while content information resides in the relative formant frequency ratios (Stevens, 1987), uniform scaling of all the formant and F0 tends to change the speaker information while retaining the content.

To further strengthen the invariance, the same random transformations are also applied to the student representations in the masked prediction task, *i.e.*, Equation (1) is modified as

$$\mathcal{L}_{pred} = \mathbb{E}[\ell_m(p \circ f(\tilde{\mathbf{X}}^{(1)}), g(\mathbf{X})) + \ell_m(p \circ f(\tilde{\mathbf{X}}^{(2)}), g(\mathbf{X}))]. \quad (3)$$

Again, the masked prediction loss is applied to both $f(\tilde{\mathbf{X}}^{(1)})$ and $f(\tilde{\mathbf{X}}^{(2)})$ for symmetry.

3.5. Speaker Conditioning

Although disentanglement in teacher can remove the majority of the speaker information from the teacher labels, certain speaker information would remain. As a result, the student representations are undesirably forced to carry the same amount of speaker information as the teachers do in order to reasonably predict the teacher labels. To break this entailment between the speaker information in students and in teachers, we feed the speaker embeddings to the predictor. Speaker embeddings are produced by a speaker embedding network, in our case a pre-trained GE2E (Wan et al., 2018), which takes a speech utterance as input and outputs a vector summarizing the speaker information in the utterance. Therefore, by conditioning the predictor on the speaker embedding, we can supply whatever speaker information is needed for the mask prediction task, so that the students do not have to carry the speaker information themselves.

Formally, the masked prediction loss now becomes

$$\begin{aligned} \mathcal{L}_{pred} = & \mathbb{E}[\ell_m(p(f(\tilde{\mathbf{X}}_1), s(\mathbf{X})), g(\mathbf{X})) \\ & + \ell_m(p(f(\tilde{\mathbf{X}}_2), s(\mathbf{X})), g(\mathbf{X}))], \end{aligned} \quad (4)$$

where $s(\mathbf{X})$ denotes the speaker embeddings. The final loss is the superposition of the prediction and contrastive losses:

$$\mathcal{L} = \mathcal{L}_{pred} + \lambda \mathcal{L}_{contr}. \quad (5)$$

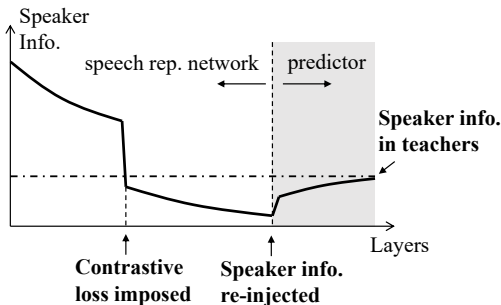


Figure 2. A conceptual curve of how speaker information changes through the network layers.

As can be observed, although CONTENTVEC requires speaker labels to identify speaker information, speaker labels are only used in pre-training the speaker embedding network. The training of CONTENTVEC itself only requires speaker embeddings, not speaker labels. Since the speaker embedding network is pre-trained on a separate dataset, and can well generalize to unseen speakers, the training set for CONTENTVEC does not need to contain any speaker labels.

3.6. An Information Flow Perspective

To provide an intuitive illustration of how the aforementioned modules work collaboratively towards disentangling speakers, Figure 2 shows a conceptual curve of how the amount of speaker information changes along different layers of the speech representation network $f(\cdot)$ and the predictor $p(\cdot)$. The vertical axis denotes the amount of speaker information, and the horizontal axis denotes the number of layers. The white area denotes the speech representation network layers, and the grey area denotes the prediction layers, which are on top of the speech representation network. To the left, the speaker information is equal to the full speaker information in the input utterance. To the right, the speaker information should be roughly equal to the speaker information in the teacher labels, which is much lower than that in the input but is still not zero. Due to the information processing inequality, the speaker information is monotonically decreasing as the layer progresses, except for the predictor layers where speaker information is re-injected.

As can be observed, there are two places where the speaker information undergoes abrupt changes. The first is where the contrastive loss (Equation (2)) is imposed, and the speaker information is largely reduced. The second is where the speaker information is re-injected, and the speaker information slightly increases. As a result, the speaker information should reach its minimum at the intersection between the speech representation network and the predictor. Figure 2 shows that all the modules in CONTENTVEC are essential to a successful speaker disentanglement.

4. Experiments

In this section, we will evaluate CONTENTVEC on an extensive set of content-related tasks. In particular, we would like to investigate whether disentangling speakers has benefit in real-world tasks, and how large the benefit would be. Further experimental details can be found in Appendix B.

4.1. Configurations

Implementation Details The speech representation network of CONTENTVEC has the same architecture as HUBERT. According to Section 3.6, the speaker-disentanglement is optimal at the output of the speech representation network, so we select the output-layer representation as the CONTENTVEC features. The contrastive loss is imposed at the last but five layer, the temperature k is set to 0.1, and the contrastive loss weight $\lambda = 1e-5 * num_train_steps$, which linearly increases to 10 when training for 100k steps. A parameter sensitivity evaluation will be provided in Section 4.7.

The predictor of CONTENTVEC contains three transformer layers *without* layer drop. The frame masking scheme and prediction logit generation for the masked prediction task are the same as HUBERT.

To generate the teacher labels, we use the voice converter proposed by Polyak et al. (2021). We re-train the voice converter on a subset of 200 speakers chosen from the Librispeech dataset (Panayotov et al., 2015), using the publicly-released HUBERT base model² with 100 clusters. The model checkpoint that gives the highest average target speaker similarity is selected. The teacher utterances are then generated by converting the entire Librispeech-960h to the voice of that selected target speaker with the selected model. After passing the converted utterance to the pretrained HUBERT, the seventh layer feature representation is chosen, because compared with the commonly chosen sixth, the seventh layer has a lower speaker classification accuracy and yet comparable teacher quality (Hsu et al., 2021). The number of clusters for quantizing final teacher labels is 100.

Baselines and Dataset The following baselines are included in the evaluation.

- WAV2VEC 2.0 (Baevski et al., 2020b): Following Lakhota et al. (2021), the 14th layer representation is chosen.
- HUBERT (Hsu et al., 2021): We adopt the publicly-released pretrained model in Ott et al. (2019). Following Hsu et al. (2021), the sixth layer representation is chosen.
- HUBERT-ITER: Since CONTENTVEC is guided by a

²<https://github.com/pytorch/fairseq/tree/main/examples/hubert>

Table 1. Results on zero-resource content probe and language modeling.

Model	ABX(w) ↓	ABX(a) ↓	Lexical ↓	Syntactic ↓	PPX ↓	VERT ↓	AUC ↓
CONTENTVEC	5.13	6.32	33.27	43.95	650.04	46.05	45.01
HUBERT-ITER	6.01	7.20	34.00	44.36	739.12	47.55	53.28
HUBERT	6.06	7.37	36.19	46.48	790.17	54.35	75.23
WAV2VEC 2.0	8.70	10.34	35.93	46.40	840.34	58.59	88.83

pretrained HUBERT as teachers, for a fair comparison, we introduce another HUBERT baseline that is trained by the same pretrained HUBERT as teachers (except that no voice conversion is performed). This baseline controls the benefit of iterative training. We performed the same teacher quality evaluation as in Hsu et al. (2021), and identified layer eight as the best performing layer.

CONTENTVEC and all the baselines are trained on the Librispeech dataset (Panayotov et al., 2015). If the evaluation task requires discrete representations, all the representations will be quantized to 100 clusters by k -means. Otherwise, the continuous representations will be used.

4.2. Zero-shot Content Probe

The first set of experiments we would like to evaluate is the set of zero-shot probing tasks proposed in the Zero-Resource Speech Challenges (Rivière & Dupoux, 2021; Dunbar et al., 2021), because they require a high alignment between the discrete representation and phonetic content. For some of these tasks, a language model trained on the discrete speech representations is needed. We use the same language model and hyperparameter setting as in Lakhotia et al. (2021), which is the *transformer LM big architecture* implemented in fairseq (Ott et al., 2019). We evaluate on four tasks.

- **ABX(w):** Given a pair of words with one difference in phoneme and a test word containing the same phoneme as one of the two words, ABX measures the probability that the test phoneme representation is closer to the representation of the correct phoneme in the word pair than to that of the incorrect phoneme. ‘(w)’ indicates that the comparison is ‘within’ the same speaker.
- **ABX(a):** Same as ABX(w), except that the test utterance is uttered by a different speaker. ‘(a)’ indicates that the comparison is ‘across’ different speakers.
- **Spot the Word (Lexical):** Spot the word measures the accuracy of identifying the correct word from a pair of real/fake words, based on the perplexity of the language model.
- **Acceptability Judgment (Syntactic):** Acceptability judgement measures the accuracy of identifying the syntactically correct sentence from a pair of correct/incorrect sentences, based on the perplexity of the language model.

Table 1 (left four columns) shows the results of the zero-shot probing tasks. There are three key observations. First, CONTENTVEC achieves consistent advantage on all four metrics, demonstrating that speaker disentanglement does help in these tasks. However, the size of the performance gain varies across different tasks and across different number of clusters. The benefit is largest for the phonetic-level tasks, ABX(w) and ABX(a). On the other hand, for lexical- and semantic-level tasks, the benefit is smaller. We believe that this is because the performance in these tasks not only depends on the quality of the speech representation, but also on the language model. Our second observation is that HUBERT-ITER consistently outperforms HUBERT, which confirms that there is a benefit in iterative training, and this is why it is very important to include HUBERT-ITER as a baseline for a fair comparison. Finally, note that there is a slight difference between our results for HUBERT and the results reported in Lakhotia et al. (2021). This is likely because the publicly-released model that we use is different from the model used in Lakhotia et al. (2021), in terms of number of clusters in teachers, batch size, the number of GPUs *etc.* However, since CONTENTVEC and HUBERT-ITER are both derived from HUBERT, we expect a similar performance gap among these three methods if a different HUBERT model is used.

4.3. Language Modeling

Language models built directly on the discrete speech representations can be applied to many content-related speech generation and analysis tasks. We would like to explore whether a disentangled speech representation can contribute to a language model with higher quality. To this end, we use the same language models as in Section 4.2 to generate random speech representation sequences under different temperatures, and resynthesize speech from these sequences using TACOTRON as in Lakhotia et al. (2021). For each temperature level, we compute the perplexity (PPX) and variety (VERT) score of the transcript as proposed in Lakhotia et al. (2021). All the PPX-VERT pairs at different temperature levels form a curve depicting the quality-variety trade-off of each language model. We report the following three metrics.

- **PPX at oracle VERT:** The perplexity score when the VERT score equals the VERT score of true text.
- **VERT at oracle PPX:** The VERT score when the PPX

Table 2. Results on SUPERB tasks within “Content” and “Semantics” categories.

Tasks Metrics	PR PER ↓	ASR WER ↓	KS ACC ↑	QbE MTWV ↑	IC ACC ↑	SF F1 ↑ CER ↓	
CONTENTVEC	0.049	5.7	0.964	0.0590	0.991	0.896	0.236
HUBERT-ITER	0.052	6.5	0.963	0.0891	0.983	0.886	0.259
HUBERT	0.054	6.4	0.963	0.0736	0.983	0.885	0.256

score equals the PPX score of the true text.

- **AUC:** Area under the perplexity-VERT curve.

Table 1 (right columns) shows the results. As can be seen, CONTENTVEC achieves a significantly lower PPX score than all the baselines and a slightly lower VERT score. This indicates that the improvement in speaker disentanglement contributes to the correctness of the language model. This observation shows that speaker disentanglement significantly helps in improving the speech generation quality. Some qualitative results are shown in Appendix A.

4.4. SUPERB Experiments

To extend our evaluation to supervised tasks, we use SUPERB (Yang et al., 2021), a benchmark dataset containing an extensive list of supervised speech processing tasks. We select the subset of tasks belonging to the categories of “content” and “semantic.” These tasks include phone recognition (PR), automatic speech recognition (ASR), keyword spotting (KS), Query by Example Spoken Term Detection (QbE), intent classification (IC), and slot filling (SF). Detailed descriptions of these tasks can be found in Yang et al. (2021). During the training of these tasks, the speech representation networks are frozen.

Unlike the tasks discussed in the previous sections, the SUPERB tasks use the *continuous* representations rather than the discrete ones. Therefore, the HUBERT baseline, which is trained with a 500-class teacher and prolonged training iterations, is expected to provide more information than do CONTENTVEC and HUBERT-ITER, which are both trained with a 100-class teacher. Therefore, we retrain HUBERT and HUBERT-ITER with matched number of teacher clusters and training iterations for a fair comparison. Table 2 lists the results on the SUPERB tasks. We can observe that CONTENTVEC generally outperforms both HUBERT-ITER and HUBERT. This observation verifies that the benefit of speech disentanglement can generalize to content-related supervised tasks.

4.5. Speaker & Accent Classification

Speaker identification serves as a proxy for speaker disentanglement. In addition, we are also interested in whether our speaker disentanglement algorithm removes regional ac-

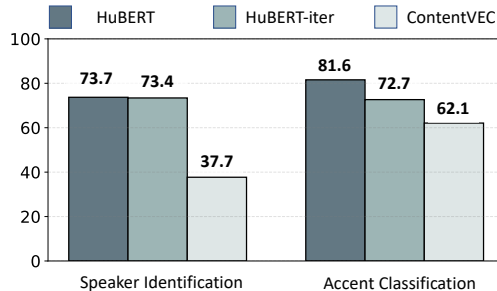


Figure 3. Bar plots of SID accuracy (left) and accent classification accuracy (right)

cents or not: is regional accent primarily communicated by phoneme content, or by speaker-dependent style? We therefore evaluate the quality of speaker disentanglement on the speaker identification task (SID) in the SUPERB benchmark, and an accent classification task using the L2-ARCTIC dataset, which contains 7 accent groups.

Figure 3 plots the accuracy of the two classification tasks. As can be seen, CONTENTVEC sharply reduces the accuracy in both tasks. In the SID task, the reduction is as high as 36% compared to the HUBERT-ITER, indicating that the speaker disentanglement mechanisms in CONTENTVEC are very effective. HUBERT-ITER has a slightly lower SID accuracy than HUBERT, which shows that iterative training reduces the amount of speaker information. In the accent classification task, the reduction is also significant, which verifies that disentangling speaker information also reduces accent information to some degree.

4.6. Voice Conversion

The existing mainstream voice conversion systems follow an encoder-decoder paradigm, where the encoder derives a speech representation with speaker information disentangled and the decoder synthesizes the speech signal conditional on speaker embedding/labels. Polyak et al. (2021) further shows that a decoder built directly on top of self-supervised speech representations suffices to produce state-of-the-art voice conversion results. To evaluate whether the advantage of speaker disentanglement in CONTENTVEC translate to better voice conversion, we run the voice conversion model in Polyak et al. (2021) on CONTENTVEC and other baseline speech representations and evaluate the speaker similarity of the converted speech to the target speaker. Instead of

Table 3. Average cosine similarity (\uparrow) of the d-vectors between the converted speech based on different speech representations and the target speakers.

Setting	C2C	O2C	C2O	O2O
CONTENTVEC	0.9316	0.9277	0.9150	0.9257
HUBERT-ITER	0.9286	0.9243	0.9050	0.9215
HUBERT	0.9029	0.8982	0.8848	0.9036

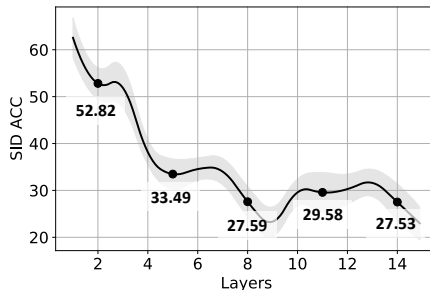
using discretized speech representations as in Polyak et al. (2021), we consider a much more challenging setting where the continuous speech representations are used. The models are trained and evaluated on Librispeech. Our test set contains the same speakers as the training set. To reduce the evaluation complexity, while the source speakers are from the entire test set, the target speakers are from a subset of 20 speakers, 10 are from the ‘clean’ subset and 10 are from the ‘other’ subset. Our evaluation is thus divided into four scenarios, clean to clean (C2C), clean to other (C2O), other to clean (O2C), and other to other (O2O).

Table 3 shows the average cosine similarity of the d-vectors (Heigold et al., 2016) between the converted speech and the target speakers. As can be observed, thanks to the inherent speaker disentanglement capability of the neural decoder, the HUBERT-based voice conversion model already has a decent speaker similarity, which is consistent with the findings by Polyak et al. (2021). However, CONTENTVEC is able to further advance the performance by a significant margin, further verifying the advantage of CONTENTVEC’s speaker disentanglement quality. We also observe that HUBERT-ITER, which has a slightly better speaker disentanglement property according to Section 4.5, also improves over HUBERT in terms of speaker similarity in this task.

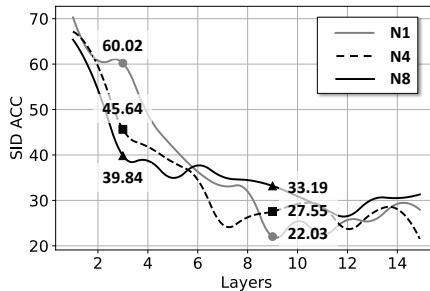
4.7. Ablation Studies

This section evaluates different variants of CONTENTVEC to understanding the contribution of each model design choice. Since there are a large number of models, we only select the following three efficient yet representative metrics: the phone normalized mutual information between the discrete representations and ground truth phonetic unit (PNMI) proposed in Hsu et al. (2021), ABX(w) accuracy, and ABX(a) accuracy. For each model variant, we report the best-layer results under each respective metric.

Contribution of Each Disentanglement Module To measure how much each of our three major disentanglement mechanisms (disentanglement in teachers, disentanglement in students, and speaker conditioning of the predictor) contribute to the overall performance, we build variants of CONTENTVEC with each individual mechanism removed, named NO-DTEACHERS, NO-DSTUDENTS, and NO-COND, re-



(a) The CONTENTVEC model



(b) Comparison of CONTENTVEC-N1, CONTENTVEC-N4, and CONTENTVEC-N8

Figure 4. SID accuracies as functions of network layers

spectively. Specifically, NO-DTEACHERS means that no voice conversion module is introduced; NO-DSTUDENTS means that no transformation or contrastive loss is imposed in the student module; NO-COND means that no speaker embeddings are fed to the predictor. Their respective performance is reported in Table 4. As can be seen, all three models perform significantly worse than CONTENTVEC. We can conclude that all three modules are essential to CONTENTVEC.

Speaker Information Flow In order to verify if our conceptual curve in Figure 2 is correct, we evaluate the SID accuracy (as a proxy of speaker information) on every layer of CONTENTVEC. For each experiment, we select the best SID model based on validation accuracy for every 10k iterations and compute the accuracy achieved by all these SID models to gauge statistical significance. Figure 4(a) plots the SID accuracy as a function of network layers. By comparing Figures 2 and 4(a), we find that our key hypotheses are verified. First, the SID monotonically decreases before layer 9, and there is a significance decrease before layer seven, which is roughly where the contrastive loss is imposed.³ This observation is consistent with the sharp decrease in Figure 2. Second, there is a slight increase towards the last layers, which is consistent with the increase in Figure 2, although it occurs at the lower level than we expect.

³The contrastive loss is always imposed at the last but fifth layer. Due to layer drop, the actual layer at which the contrastive loss is imposed is randomly distributed before layer seven, with the major probability mass in layers five to seven.

Table 4. Results of ablation studies.

	CONTENTVEC	Removing each module			Position of \mathcal{L}_{cont}			Weight of \mathcal{L}_{cont} (λ)			
		NO-DT	NO-DS	NO-COND	N1	N4	N8	1e-6	5e-6	2e-5	5e-5
ABX(w) ↓	5.13	5.67	5.62	5.76	5.76	5.19	5.14	5.18	5.10	5.18	5.22
ABX(a) ↓	6.32	6.92	6.93	7.08	7.03	6.43	6.40	6.36	6.39	6.51	6.58
PNMI ↑	0.590	0.576	0.577	0.593	0.545	0.587	0.584	0.586	0.586	0.584	0.582

A likely reason is that our SID is performed on the discrete representations, whereas our conceptual curve is based on continuous representation. Nevertheless, these observations confirm the rationales behind our model design.

Position of Contrastive Loss To illustrate how the position of contrastive loss influences the speaker information flow, we run three variants of CONTENTVEC, denoted as CONTENTVEC-N1, CONTENTVEC-N4, CONTENTVEC-N8. “N1”, “N4” and “N8” mean that the contrastive loss is imposed in last, last but three, and last but seven layers, respectively (instead of last but five as in our standard CONTENTVEC). Figure 4 compares the SID accuracy against the layers for these variants. We can observe a similar decreasing and then increasing trend in all these curves. More importantly, there is an interesting correlation between the SID accuracy drop positions and the layer at which contrastive loss is imposed. CONTENTVEC-N8 drops earliest, followed by CONTENTVEC-N4 and then CONTENTVEC-N1. We can also observe that the model that starts dropping late would drop to a lower point.

We also evaluate the performance of these model variants with respect to the aforementioned metrics, and the results are listed in Table 4. As can be observed, all the variants achieve very competitive results, outperforming all the baselines. These results show that the competitiveness of CONTENTVEC does not hinge on specific choices of this hyperparameter.

Contrastive Loss Weight We also evaluate different contrastive loss weights (the λ in Equation (5)). Recall that our standard CONTENTVEC uses a weight coefficient of 1e-5. The following four weight coefficient are tested, 1e-6, 5e-6, 2e-5, and 5e-5, and the results are listed in Table 4. We can observe that CONTENTVEC performs competitively across all the models, showing a low hyperparameter sensitivity.

5. Conclusions

In this paper, we propose CONTENTVEC, which is a speech representation learning network that aims to remove speaker information while preventing loss of content information. CONTENTVEC builds upon the HUBERT framework and introduces three key disentanglement components: disentanglement in teachers, disentanglement in students, and teacher conditioning of the predictor. Our empirical analy-

ses confirm that all three modules are essential to the success of CONTENTVEC. We also verified that a successful speaker disentanglement does help with a wide range of content-related speech processing tasks. Meanwhile, CONTENTVEC still have some limitations, e.g. slight content loss and the lack of hyperparameter selection method. Finding solutions to these problems will be our future directions.

References

Baevski, A., Schneider, S., and Auli, M. vq-wav2vec: Self-supervised learning of discrete speech representations. *ArXiv*, abs/1910.05453, 2020a.

Baevski, A., Zhou, H., rahman Mohamed, A., and Auli, M. wav2vec 2.0: A framework for self-supervised learning of speech representations. *ArXiv*, abs/2006.11477, 2020b.

Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pp. 1597–1607. PMLR, 2020.

Chen, Y.-H., Wu, D.-Y., Wu, T.-H., and yi Lee, H. Again-vc: A one-shot voice conversion using activation guidance and adaptive instance normalization. *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5954–5958, 2021.

Choi, H.-S., Lee, J., Kim, W., Lee, J. H., Heo, H., and Lee, K. Neural analysis and synthesis: Reconstructing speech from self-supervised representations. In *Advances in Neural Information Processing Systems*, 2021.

Choi, Y., Choi, M.-J., Kim, M. S., Ha, J.-W., Kim, S., and Choo, J. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8789–8797, 2018.

Chou, J.-C., chieh Yeh, C., yi Lee, H., and Lee, L.-S. Multi-target voice conversion without parallel data by adversarially learning disentangled audio representations. In *INTERSPEECH*, 2018.

Chou, J.-C., chieh Yeh, C., and yi Lee, H. One-shot voice conversion by separating speaker and content representations with instance normalization. In *INTERSPEECH*, 2019.

- Chung, Y.-A. and Glass, J. R. Generative pre-training for speech with autoregressive predictive coding. *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 3497–3501, 2020.
- Chung, Y.-A., Hsu, W.-N., Tang, H., and Glass, J. R. An unsupervised autoregressive model for speech representation learning. In *INTERSPEECH*, 2019.
- Chung, Y.-A., Zhang, Y., Han, W., Chiu, C.-C., Qin, J., Pang, R., and Wu, Y. W2v-bert: Combining contrastive learning and masked language modeling for self-supervised speech pre-training. *ArXiv*, abs/2108.06209, 2021.
- Dunbar, E., Bernard, M., Hamilakis, N., Nguyen, T. A., de Seyssel, M., Roz'e, P., Rivière, M., Kharitonov, E., and Dupoux, E. The zero resource speech challenge 2021: Spoken language modelling. *Interspeech 2021*, 2021.
- Eide, E. and Gish, H. Parametric approach to vocal tract length normalization. In *ICASSP*, pp. 346–9, 1996.
- Heigold, G., Moreno, I., Bengio, S., and Shazeer, N. M. End-to-end text-dependent speaker verification. *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5115–5119, 2016.
- Hsu, C.-C., Hwang, H.-T., Wu, Y.-C., Tsao, Y., and Wang, H.-M. Voice conversion from non-parallel corpora using variational auto-encoder. In *Signal and Information Processing Association Annual Summit and Conference (APSIPA), 2016 Asia-Pacific*, pp. 1–6. IEEE, 2016.
- Hsu, W.-N., Bolte, B., Tsai, Y.-H. H., Lakhotia, K., Salakhutdinov, R., and Mohamed, A. Hubert: Self-supervised speech representation learning by masked prediction of hidden units. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:3451–3460, 2021.
- Ito, K. and Johnson, L. The lj speech dataset. <https://keithito.com/LJ-Speech-Dataset/>, 2017.
- Kameoka, H., Kaneko, T., Tanaka, K., and Hojo, N. Stargan-vc: non-parallel many-to-many voice conversion using star generative adversarial networks. *2018 IEEE Spoken Language Technology Workshop (SLT)*, pp. 266–273, 2018.
- Kameoka, H., Kaneko, T., Tanaka, K., and Hojo, N. Acvae-vc: Non-parallel voice conversion with auxiliary classifier variational autoencoder. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 27(9):1432–1443, 2019. doi: 10.1109/TASLP.2019.2917232.
- Kaneko, T. and Kameoka, H. Cyclegan-vc: Non-parallel voice conversion using cycle-consistent adversarial networks. *2018 26th European Signal Processing Conference (EUSIPCO)*, pp. 2100–2104, 2018.
- Kaneko, T., Kameoka, H., Tanaka, K., and Hojo, N. Cyclegan-vc2: Improved cyclegan-based non-parallel voice conversion. *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6820–6824, 2019a.
- Kaneko, T., Kameoka, H., Tanaka, K., and Hojo, N. StarGAN-VC2: Rethinking conditional methods for StarGAN-based voice conversion. *Proc. Interspeech 2019*, pp. 679–683, 2019b.
- Kharitonov, E., Rivière, M., Synnaeve, G., Wolf, L., Mazar'e, P.-E., Douze, M., and Dupoux, E. Data augmenting contrastive learning of speech representations in the time domain. *2021 IEEE Spoken Language Technology Workshop (SLT)*, pp. 215–222, 2021.
- Kong, J., Kim, J., and Bae, J. Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. F., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 17022–17033. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/c5d736809766d46260d816d8dbc9eb44-Paper.pdf>.
- Lakhotia, K., Kharitonov, E., Hsu, W.-N., Adi, Y., Polyak, A., Bolte, B., Nguyen, T., Copet, J., Baeovski, A., Mohamed, A. B., and Dupoux, E. On generative spoken language modeling from raw audio. *Transactions of the Association for Computational Linguistics*, 9:1336–1354, 2021.
- Lee, S.-H. and Kim, J.-H. Voicemixer: Adversarial voice style mixup. 2021.
- Ling, S. and Liu, Y. Decoar 2.0: Deep contextualized acoustic representations with vector quantization. *ArXiv*, abs/2012.06659, 2020.
- Liu, A. T., wen Yang, S., Chi, P.-H., Hsu, P.-C., and yi Lee, H. Mockingjay: Unsupervised speech representation learning with deep bidirectional transformer encoders. *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6419–6423, 2020.
- Liu, A. T., Li, S.-W., and yi Lee, H. Tera: Self-supervised learning of transformer encoder representation for speech. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:2351–2366, 2021.

- Ott, M., Edunov, S., Baevski, A., Fan, A., Gross, S., Ng, N., Grangier, D., and Auli, M. fairseq: A fast, extensible toolkit for sequence modeling. In *NAACL*, 2019.
- Panayotov, V., Chen, G., Povey, D., and Khudanpur, S. Librispeech: An asr corpus based on public domain audio books. *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5206–5210, 2015.
- Pascual, S., Ravanelli, M., Serrà, J., Bonafonte, A., and Bengio, Y. Learning problem-agnostic speech representations from multiple self-supervised tasks. In *INTERSPEECH*, 2019.
- Polyak, A., Adi, Y., Copet, J., Kharitonov, E., Lakhota, K., Hsu, W.-N., Mohamed, A., and Dupoux, E. Speech Resynthesis from Discrete Disentangled Self-Supervised Representations. In *Proc. Interspeech 2021*, 2021.
- Qian, K., Zhang, Y., Chang, S., Yang, X., and Hasegawa-Johnson, M. Autovc: Zero-shot voice style transfer with only autoencoder loss. In *International Conference on Machine Learning*, pp. 5210–5219. PMLR, 2019.
- Qian, K., Jin, Z., Hasegawa-Johnson, M., and Mysore, G. J. F0-consistent many-to-many non-parallel voice conversion via conditional autoencoder. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6284–6288. IEEE, 2020a.
- Qian, K., Zhang, Y., Chang, S., Hasegawa-Johnson, M., and Cox, D. Unsupervised speech decomposition via triple information bottleneck. In *International Conference on Machine Learning*, pp. 7836–7846. PMLR, 2020b.
- Qian, K., Zhang, Y., Chang, S., Xiong, J., Gan, C., Cox, D., and Hasegawa-Johnson, M. A. Global prosody style transfer without text transcriptions. In *ICML*, 2021.
- Ravanelli, M., Zhong, J., Pascual, S., Swietojanski, P., Monteiro, J., Trmal, J., and Bengio, Y. Multi-task self-supervised learning for robust speech recognition. *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6989–6993, 2020.
- Rivière, M. and Dupoux, E. Towards unsupervised learning of speech features in the wild. *2021 IEEE Spoken Language Technology Workshop (SLT)*, pp. 156–163, 2021.
- Schneider, S., Baevski, A., Collobert, R., and Auli, M. wav2vec: Unsupervised pre-training for speech recognition. In *INTERSPEECH*, 2019.
- Stevens, K. N. Relational properties as perceptual correlates of phonetic features. In *Proc. Eleventh Int. Conf. Phonetic Sciences*, volume 4, pp. 352–356, Tallinn, Estonia, 1987.
- van den Oord, A., Vinyals, O., and Kavukcuoglu, K. Neural discrete representation learning. In *NIPS*, 2017.
- van den Oord, A., Li, Y., and Vinyals, O. Representation learning with contrastive predictive coding. *ArXiv*, abs/1807.03748, 2018.
- Wan, L., Wang, Q., Papir, A., and Moreno, I. L. Generalized end-to-end loss for speaker verification. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 4879–4883. IEEE, 2018.
- Wang, C., Wu, Y., Chen, S., Liu, S., Li, J., Qian, Y., and Yang, Z. Self-supervised learning for speech recognition with intermediate layer supervision. *ArXiv*, abs/2112.08778, 2021.
- Wang, W., Tang, Q., and Livescu, K. Unsupervised pre-training of bidirectional speech encoders via masked reconstruction. *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6889–6893, 2020.
- Yang, S.-w., Chi, P.-H., Chuang, Y.-S., Lai, C.-I. J., Lakhota, K., Lin, Y. Y., Liu, A. T., Shi, J., Chang, X., Lin, G.-T., et al. Superb: Speech processing universal performance benchmark. *arXiv preprint arXiv:2105.01051*, 2021.
- Zhang, Y., Qin, J., Park, D. S., Han, W., Chiu, C.-C., Pang, R., Le, Q. V., and Wu, Y. Pushing the limits of semi-supervised learning for automatic speech recognition. *ArXiv*, abs/2010.10504, 2020.
- Zhu, J.-Y., Park, T., Isola, P., and Efros, A. A. Unpaired image-to-image translation using cycle-consistent adversarial networks. *2017 IEEE International Conference on Computer Vision (ICCV)*, pp. 2242–2251, 2017.
- Zhu, Y., Lu, S., Zheng, L., Guo, J., Zhang, W., Wang, J., and Yu, Y. Texus: A benchmarking platform for text generation models. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, SIGIR '18*, pp. 1097–1100, New York, NY, USA, 2018. Association for Computing Machinery. ISBN 9781450356572. doi: 10.1145/3209978.3210080. URL <https://doi.org/10.1145/3209978.3210080>.

A. Additional Experiment Results

A.1. Visualizing Speech Generation Cluster Sequences

We compare HUBERT-ITER and CONTENTVEC by visualizing their speech labels on a small subset of Librispeech and calculating the average distance between male and female speakers speaking the same content.

For each k -means label, we compute its ratio of occurrence between male and female speakers in the dev-clean split of Librispeech. We then rank the labels by the occurrence ratio of female; the more frequently the label occurs in female utterances, the larger rank it gets. We randomly selected 100 female utterances from train-clean-100 split of Librispeech and convert them to male speeches using the voice converter. Then for each pair of male and female utterances with the same content, we map their label sequences to ranked label sequences. Since each pair of utterance shares the same content, their corresponding ranked label should ideally be the same. Figure 5 shows the ranked label sequences on an example male-female pair, where their differences are highlighted with blue blocks. We observe that there are fewer blue blocks in CONTENTVEC plot and therefore the male and female curves align better using CONTENTVEC labels. This observation implies that CONTENTVEC disentangles more speaker information from the content compared to HUBERT-ITER.

To confirm this observation, we apply additional dynamic time warping to align the utterance pairs and then compute the average L0 distance in their ranked labels sequences over frames. Dynamic time warping is necessary as we find that the converted audio still have slight mismatch in speech rate. The result is shown in Table 5; CONTENTVEC does have lower average L0 distance, indicating the gender information is hardly preserved in the CONTENTVEC features.

Table 5. Results of average L0 distance between male and female speech pairs

	HUBERT-ITER	CONTENTVEC
L0 Distance	0.165	0.118

A.2. ASR Transcriptions of Example Generated Speech

In Table 6 we show several example transcripts of speech generated by HUBERT-ITER and CONTENTVEC, using the first three seconds of utterances from the test-clean split of the Librispeech dataset.

It seems the speech representation from CONTENTVEC is better at reconstructing the prompts than that from the HUBERT-ITER. Both models are able to generate meaningful speech within a few words after the prompt. However,

neither of them is capable of generating meaningful long sentences.

Table 6. Example transcripts of resynthesized speech using prompt. H denotes the transcripts from HUBERT-ITER and C denotes those from CONTENTVEC. The highlighted texts are resynthesized from the given prompt and the non-highlighted texts are from labels generated by language models

Prompt		after early nightfall the yellow lamps
H :		after early nightfall the yellow lamps would light i double loud in my hands and a couple of dances told me we
H :		ah after early nightfall meat you lamps would light in that six thous heaven when happens
C :		after early nightfall the yellow lamps would light and they were upon their liberty until morning when they were
C :		after early nightfall the yellow lamps would light and chill in some dim spectacle lancers in the sun
Prompt		if a layman in giving baptism pour the water
H :		if alane and giving baptism for the waters which dry even stream formed the foreheads of
H :		fish iffen men and giving baptism for the water of flame and he came on a great
C :		if a layman in giving baptism pour the water forth to the knot and come out without any
C :		if the laymen in giving baptism pour the water in a shower of sparkle which made these sure

A.3. Voice Converter Quality

Since the voice converter in teacher is a key mechanism of CONTENTVEC, we would like to investigate how the quality of the voice converter impact the quality of CONTENTVEC. To this end, we performed an ablation study where we replaced the voice converter with a compromised version trained with fewer number of steps (10,000 steps instead of 40,000 steps). Accordingly, the best-layer ABX(w), ABX(a) and PNMI degrade to 6.05, 7.78, and 0.5616 respectively (The best layer for these three metrics are 8, 12, 9, respectively). This result shows that the quality of the voice converter is very essential to the performance of CONTENTVEC. A poor voice converter would produce an even worse performance than the variant with no voice converter at all, as indicated by the ‘NO-DT’ results in Table 4.

A.4. Contrastive Loss on Multiple Layers

So far, we have only evaluated on CONTENTVEC models where the contrastive loss is imposed on only one layer. To investigate how the performance will change if the contrastive loss is imposed on multiple layers, we perform an ablation study where the contrastive loss is simultaneously imposed on the last, last but third, and last but seventh layers. To keep the scale consistent, the weight on each contrastive

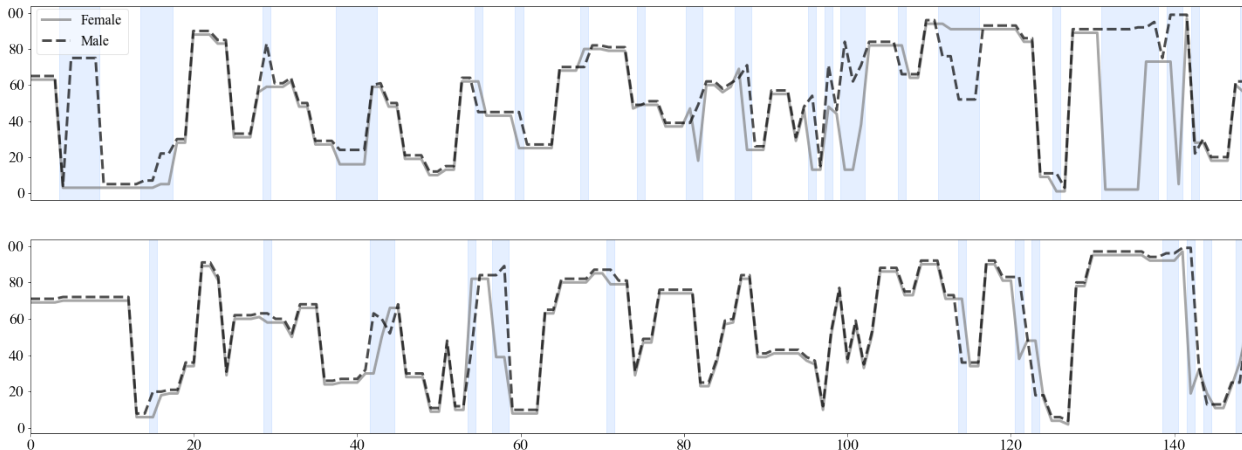


Figure 5. Example label rank sequences generated using HUBERT-ITER (upper) and CONTENTVEC (lower). The mismatched frames are highlighted using blue.

loss term is reduced to 1/3 of its original value. The best-layer ABX(w), ABX(a) and PNMI become 5.19, 6.53, and 0.5684 respectively (The best layer for these three metrics are all 12), which is slightly worse than the original CONTENTVEC model. In other words, there is no evidence that imposing contrastive loss on multiple layers can improve performance of CONTENTVEC.

B. Experiment Details

B.1. CONTENTVEC Implementation and Training

The speech representation network of CONTENTVEC has the same architecture as the HUBERT, which has 7 temporal convolutional feature extraction blocks followed by 12 layers of transformer layers of model dimension 768. During training, each layer is independently dropped with a probability of 0.05. According to Section 3.6, the speaker-disentanglement is optimal at the output of the speech representation network, so we select the output-layer representation as the CONTENTVEC features. The contrastive loss is imposed at the last but five layer, the temperature k is set to 0.1, and the contrastive loss weight $\lambda = 1e-5 * num_train_steps$, which linearly increases to 10 when training for 100k steps.

The masking strategy is the same as in Wav2Vec 2.0 (Baevski et al., 2020b), with the masking probability set to 0.08. The masks of the two transformation paths are the same. It is also worth mentioning that the masking is only imposed for the masked prediction loss, \mathcal{L}_{pred} . No masking is imposed for the contrastive loss, \mathcal{L}_{cont} .

The predictor of CONTENTVEC contains three transformer layers *without* layer drop. Each transformer layer uses the

conditional layer normalization to inject speaker information, of which the scale and bias are learnable linear mappings of the conditioned speaker embedding. The frame masking scheme and prediction logit generation for the masked prediction task are the same as HUBERT.

Our model is trained for 100k steps using 36 GPUs, with a batch size of at most 76 seconds of audio per GPU, which takes about 19 hours to complete. The best model is selected based on the lowest validation masked prediction loss.

To generate the teacher labels, we use the voice converter proposed by Polyak et al. (2021), which is trained by reconstructing speech from quantized speech representations. We directly use the publicly available implementation⁴, and all hyperparameters are kept the same unless explicitly mentioned. We re-train the resynthesis model on a subset of 200 speakers chosen from the Librispeech dataset (Panayotov et al., 2015). Of the 200 speakers, 100 of them are from the train-clean-100 or the train-clean-360 subset, while the other 100 speakers are from the train-other-500 subset. The ratio of male and female speakers is 1:1. These speakers are chosen based on the amount of audio available as listed in the metadata. The input consists of three components: the output from a speech-to-unit model, the output from a pitch-to-unit model, and speaker embedding. We choose the publicly-released HUBERT-BASE model⁵ as the speech-to-unit model. To extract the discrete speech units, we train a k-means model with 100 clusters on the features extracted

⁴<https://github.com/facebookresearch/speech-resynthesis>

⁵<https://github.com/pytorch/fairseq/tree/main/examples/hubert>

from the seventh layer of the HUBERT-BASE model, using the `train-clean-100` subset. The pitch-to-unit model is a VQ-VAE (van den Oord et al., 2017) with a convolutional encoder plus a bottleneck, trained to reconstruct per-speaker mean-variance normalized pitch contours. The speaker embedding is based on d-vectors (Heigold et al., 2016), and we use an existing implementation⁶. To resynthesize speech, the inputs are fed into a decoder based on HiFi-GAN (Kong et al., 2020), except that the input format is modified. The decoder is trained for a total of 460k steps, and each step consists of alternating updates between the generator and the discriminators. After training, we randomly chose eight unseen speakers from the `dev-clean` and the `dev-other` subsets, with an equal number of male and female speakers. For each saved checkpoint, we convert the utterances from those eight unseen speaker to the voices of ten seen speakers from the 100 clean speakers used for training. The conversion is done simply by changing the speaker embedding to those of the target speakers. For each target speaker, the speaker similarity is calculated as the cosine similarity between the speaker embedding obtained with the converted utterances, and the ground-truth target speaker embedding extracted using the target speaker’s own utterances. The saved iteration that gives the highest average target speaker similarity on the ten seen speakers is selected. After that, we further select a target speaker for the teacher utterances from the 100 clean speakers seen during training. This is obtained by converting the utterances from the unseen speakers into each of the 100 clean speakers with the selected model iteration, and then calculating the same cosine similarities as stated above. The teacher utterances are then generated by converting the entire `Librispeech-960` to the voice of that selected target speaker (which is a male speaker), with the selected model iteration.

B.2. Baselines and Dataset

- WAV2VEC 2.0 (Baeovski et al., 2020b): Following Lakhota et al. (2021), the 14th layer representation is chosen.
- HUBERT (Hsu et al., 2021): We adopt the publicly-released pretrained model in Ott et al. (2019). Following Hsu et al. (2021), the sixth layer representation is chosen.
- HUBERT-ITER: Since CONTENTVEC is guided by a pretrained HUBERT as teachers, for a fair comparison, we introduce another HUBERT baseline that is trained by the same pretrained HUBERT as teachers (except that no voice conversion is performed). This baseline controls the benefit of iterative training. We performed the same teach quality evaluation as in Hsu et al. (2021), and identified layer eight as the best performing layer.

⁶<https://github.com/resemble-ai/Resemblyzer>

CONTENTVEC and all the baselines are trained on the full 960 hours of `Librispeech` dataset (Panayotov et al., 2015). The teacher quality evaluations are conducted on the `dev-clean` and `dev-other` partition of `Librispeech`, and we use publicly available phone alignment using the Montreal Aligner.⁷ If the evaluation task requires discrete representations, all the representations will be quantized to 100 clusters by k -means except for SUPERB, which uses 500 clusters. Otherwise, the continuous representations will be used.

B.3. Zero-shot Content Probe

We evaluate our models on set of zero-shot probing tasks proposed in the Zero-Resource Speech Challenges (Rivière & Dupoux, 2021; Dunbar et al., 2021), because they require a high alignment between the discrete representation and phonetic content. For some of these tasks, a language model trained on the discrete speech representations is needed. We use the same language model and hyperparameter setting as in Lakhota et al. (2021), which is the *transformer LM big architecture* implemented in `fairseq` (Ott et al., 2019). This transformer model has 12 layers with 16 attention heads, embedding size of 1024, FFN size of 4096 and dropout probability of 0.1. The model is trained on eight 32-GB GPUs for 20000 updates the using distributed training with a dynamic batch size that contains at most 4096 tokens. The learning rate is set to 5×10^{-4} . We select the model that has the lowest loss on the validation set. In this experiment, we also test the models with 200 clusters.

B.4. Language Modeling

Language models built directly on the discrete speech representations can be applied to many content-related speech generation and analysis tasks. We would like to explore whether a disentangled speech representation can contribute to a language model with higher quality. To this end, we use the same language models as in Section 4.2 to generate random speech representation sequences under different temperatures, resynthesize speech from these sequences and transcribe the audio using a pretrained speech recognizer.

We adopt the same TACOTRON-based speech synthesizer as in Lakhota et al. (2021), and finetune it on `LJ Speech` dataset (Ito & Johnson, 2017) that contains only one female voice. The model is trained on eight 32-GB GPUs with a batch size of 32 for 20000 iterations. The learning rate is set to 0.001. We use the model from the last iteration for speech resynthesis.

We use WAV2VEC2.0-LARGE model pretrained on the `Libri-Light` dataset without finetuning to generate tran-

⁷<https://github.com/CorentinJ/librispeech-alignments>

Table 7. Best layers of the ablation models

	Removing each module			Position of \mathcal{L}_{cont}			Weight of \mathcal{L}_{cont} (λ)			
	NO-DT	NO-DS	NO-COND	N1	N4	N8	1e-6	5e-6	2e-5	5e-5
ABX(w)	12	12	10	12	12	12	12	12	12	12
ABX(a)	12	12	10	12	12	12	12	12	12	12
PNMI	12	11	11	12	12	12	12	12	12	12

criptions. Note that the TACOTRON model was pretrained on 22050 Hz speech and works best with 22050 Hz. We need to resample the generated speech to 16000 Hz so that WAV2VEC can work correctly.

For each temperature level, we sample speech sequences from language model, resynthesize audio using these sequences. We compute the perplexity (PPX) and VERT score of the transcript. VERT score is the geometric mean of the self- and auto-BLEU as proposed in Lakhotia et al. (2021).

The perplexity is evaluated using a pretrained English language model *transformer_lm.wmt19.en* from fairseq. BLEU scores measures the similarity between two sentences and self-BLEU score (Zhu et al., 2018) measures the similarity between sentences generated using the same prompt. Higher self-BLEU score means less diversity of the generated speech.

Auto-BLEU score (Lakhotia et al., 2021) measures within-sentence diversity and is computed as the ratio of k -grams that are repeated at least once:

$$\text{AUTO-BLEU}(u, k) = \frac{\sum_s \mathbb{I}[s \in (\text{NG}_k(u) \setminus s)]}{|\text{NG}_k(u)|}, \quad (6)$$

where $\text{NG}_k(n)$ is the set of k -grams of utterance u .

All the perplexity-VERT pairs at different temperature levels form a curve depicting the quality-variety trade-off of each language model. We use the perplexity and VERT scores of the ground-truth text as an anchor point to compute AUC: the area above the anchor point and under the perplexity-VERT curve. We use AUC as the tradeoff measure; lower AUC means the model is closer to the anchor point, therefore closer to ground-truth text.

B.5. Voice Conversion

For the voice conversion experiments, we use the same voice conversion model proposed by Polyak et al. (2021) and follow the same re-synthesis training described in Appendix B.1, except for the input, the input layer and train/dev/test split. In order to compare how well our CONTENTVEC system disentangles speaker information with the two baseline models (HUBERT and HUBERT-ITER), we extract the continuous speech unit of the best performing layer of all three models (sixth layer for HUBERT, eighth layer for HUBERT-ITER and CONTENTVEC).

The 200-speaker sub-corpus created from LibriSpeech in Appendix B.1 is further splitted so that around 70% of the utterances for each speaker serve as the training set, around 10% serve as the development set, and around 20% serve as the test set. The decoder is trained for a total of 440k steps, and we choose the best performing checkpoint (over the last 200k steps) by re-synthesizing the development set and then averaging the per-speaker cosine similarities. The per-speaker cosine similarity is calculated as dot-product between the length-normalized d-vector (Heigold et al., 2016) extracted from re-synthesized development-set utterances of that speaker and the speaker’s original (length-normalized) d-vector used in training. Voice conversion is carried out by choosing 10 seen speakers that belong to the `train-clean-100` or `train-clean-360` subsets of LibriSpeech and another 10 seen speakers that belong to the `train-clean-other` subset. All test utterances are converted to the 20 target speakers (the utterances from those 20 speakers are only converted to the other 19 target speakers). We separate the test set into a ‘clean’ subset and an ‘other’ subset, and we separately calculate the average-speaker cosine similarity for the 10 ‘clean’ target speakers and the 10 ‘other’ target speakers. For each of the four scenarios, the average cosine similarity is obtained by averaging the per-speaker cosine similarities over the 10 target speakers in that scenario. The per-speaker cosine similarity is calculated as the dot product between a length-normalized d-vector obtained from all the converted speech of a specific target speaker and that speaker’s original d-vector used in training.

B.6. Ablation Studies

For each metric, we report the best-layer result. The best layers of each ablation model and each metric shown in Table 7.