

A PROOFS

Theorem A.1 (Theorem 3.1). Let $\boldsymbol{\mu}_v^{(-ui)} = \frac{1}{N_v} \sum_{k=1}^{N_v} \mathbf{s}_{vk}$ if $u \neq v$; and $\boldsymbol{\mu}_u^{(-ui)} = \frac{1}{N_u-1} \sum_{j \neq i} \mathbf{s}_{uj}$. Then,

$$\mathcal{I}(\mathbf{u}; \mathbf{s}) \geq \mathbb{E} \left[\frac{1}{N} \sum_{u=1}^M \sum_{i=1}^{N_u} \left[-\|\mathbf{s}_{ui} - \boldsymbol{\mu}_u^{(-ui)}\|^2 - \frac{e^{-1}}{N} \sum_{v=1}^M [N_v \exp(-\|\mathbf{s}_{ui} - \boldsymbol{\mu}_v^{(-ui)}\|^2)] \right] \right]. \quad (12)$$

Proof of Theorem 3.1. By the condition in the Theorem, we have the given sample pairs $\{(u, \mathbf{s}_{ui})\}_{1 \leq u \leq M, 1 \leq i \leq N_u}$. Note that each pair of speaker identity and style embedding, (u, \mathbf{s}_{ui}) , can be regarded as a sample from the joint distribution $p(\mathbf{u}, \mathbf{s})$. To clarify the proof, we change the notation of random variables \mathbf{u} and \mathbf{s} to \mathbf{U} and \mathbf{S} , which are distinct to samples $\{(u, \mathbf{s}_{ui})\} \sim p(\mathbf{U}, \mathbf{S})$.

For a sample pair (u, \mathbf{s}_{ui}) , by the NWJ lower bound, we have

$$\begin{aligned} \mathcal{I}(\mathbf{U}; \mathbf{S}) &\geq \mathbb{E}_{p(\mathbf{U}, \mathbf{S})} [f(\mathbf{U}, \mathbf{S})] - e^{-1} \mathbb{E}_{p(\mathbf{U})p(\mathbf{S})} [e^{f(\mathbf{U}, \mathbf{S})}] \\ &= \mathbb{E}_{p(\mathbf{S})} \left[\mathbb{E}_{p(\mathbf{U}|\mathbf{S})} [f(\mathbf{U}, \mathbf{S})] - e^{-1} \mathbb{E}_{p(\mathbf{U})} [e^{f(\mathbf{U}, \mathbf{S})}] \right] \\ &= \mathbb{E}_{\mathbf{s}_{ui} \sim p(\mathbf{S})} \left[\mathbb{E}_{p(\mathbf{U}|\mathbf{S}=\mathbf{s}_{ui})} [f(\mathbf{U}, \mathbf{S}=\mathbf{s}_{ui})] - e^{-1} \mathbb{E}_{p(\mathbf{U})} [e^{f(\mathbf{U}, \mathbf{S}=\mathbf{s}_{ui})}] \right], \end{aligned} \quad (13)$$

with a score function $f(\mathbf{U}, \mathbf{S})$. Given $\mathbf{S} = \mathbf{s}_{ui}$, $\mathbb{E}_{p(\mathbf{U}|\mathbf{S}=\mathbf{s}_{ui})} [f(\mathbf{U}, \mathbf{S}=\mathbf{s}_{ui})]$ has an unbiased estimation $f(\mathbf{U} = u, \mathbf{S} = \mathbf{s}_{ui})$; $\mathbb{E}_{p(\mathbf{U})} [e^{f(\mathbf{U}, \mathbf{S}=\mathbf{s}_{ui})}]$ has an unbiased estimation by taking average of all possible values $v \sim p(\mathbf{U})$ in samples $\{(u, \mathbf{s}_{ui})\}$,

$$\mathbb{E}_{p(\mathbf{U})} [e^{f(\mathbf{U}, \mathbf{S}=\mathbf{s}_{ui})}] = \mathbb{E} \left[\sum_{v=1}^M \frac{N_v}{N} e^{f(\mathbf{U}=v, \mathbf{S}=\mathbf{s}_{ui})} \right]. \quad (14)$$

With the two estimations, (13) becomes

$$\mathcal{I}(\mathbf{U}; \mathbf{S}) \geq \mathbb{E} \left[f(u, \mathbf{s}_{ui}) - e^{-1} \sum_{v=1}^M \frac{N_v}{N} e^{f(v, \mathbf{s}_{ui})} \right]. \quad (15)$$

Specifically, we select score function $f(\mathbf{U} = v, \mathbf{S} = \mathbf{s}) = -\|\mathbf{s} - \boldsymbol{\mu}_v^{(-ui)}\|^2$, then (15) becomes

$$\mathcal{I}(\mathbf{U}; \mathbf{S}) \geq \mathbb{E} [\hat{\mathcal{I}}_{ui}] = \mathbb{E} \left[-\|\mathbf{s}_{ui} - \boldsymbol{\mu}_u^{(-ui)}\|^2 - e^{-1} \sum_{v=1}^M \frac{N_v}{N} e^{-\|\mathbf{s}_{ui} - \boldsymbol{\mu}_v^{(-ui)}\|^2} \right]. \quad (16)$$

Since the selection of index ui is arbitrary, we take an average on all $\hat{\mathcal{I}}_{ui}$,

$$\begin{aligned} \mathcal{I}(\mathbf{U}; \mathbf{S}) &\geq \frac{1}{N} \sum_{u=1}^M \sum_{i=1}^{N_u} \mathbb{E}_{(u, \mathbf{s}_{ui}) \sim p(\mathbf{U}, \mathbf{S})} [\hat{\mathcal{I}}_{ui}] = \mathbb{E} \left[\frac{1}{N} \sum_{u=1}^M \sum_{i=1}^{N_u} \hat{\mathcal{I}}_{ui} \right] \\ &= \mathbb{E} \left[\frac{1}{N} \sum_{u=1}^M \sum_{i=1}^{N_u} \left[-\|\mathbf{s}_{ui} - \boldsymbol{\mu}_u^{(-ui)}\|^2 - \frac{e^{-1}}{N} \sum_{v=1}^M [N_v \exp(-\|\mathbf{s}_{ui} - \boldsymbol{\mu}_v^{(-ui)}\|^2)] \right] \right], \end{aligned} \quad (17)$$

where the right-hand side of equation (7) is derived. \square

Theorem A.2 (Theorem 3.2). Assume that given $\mathbf{s} = \mathbf{s}_u$, samples $\{(\mathbf{x}_{ui}, \mathbf{c}_{ui})\}_{i=1}^{N_u}$ are observed. With a variational distribution $q_\phi(\mathbf{x}|\mathbf{s}, \mathbf{c})$, we have $\mathcal{I}(\mathbf{x}; \mathbf{c}|\mathbf{s}) \geq \mathbb{E}[\hat{\mathcal{I}}]$, where

$$\hat{\mathcal{I}} = \frac{1}{N} \sum_{u=1}^M \sum_{i=1}^{N_u} \left[\log q_\phi(\mathbf{x}_{ui}|\mathbf{c}_{ui}, \mathbf{s}_u) - \log \left(\frac{1}{N_u} \sum_{j=1}^{N_u} q_\phi(\mathbf{x}_{uj}|\mathbf{c}_{ui}, \mathbf{s}_u) \right) \right]. \quad (18)$$

Proof of Theorem 3.2. Given $\mathbf{s} = \mathbf{s}_u$, we observe sample pair $\{(\mathbf{x}_{ui}, \mathbf{c}_{ui})\}_{i=1}^{N_u}$. By the InfoNCE lower bound (Oord et al., 2018), with a score function f , we have

$$\mathcal{I}(\mathbf{x}; \mathbf{c}|\mathbf{s} = \mathbf{s}_u) \geq \mathbb{E} \left[\frac{1}{N_u} \sum_{i=1}^{N_u} \left[f(\mathbf{x}_{ui}, \mathbf{c}_{ui}) - \log \left(\frac{1}{N_u} \sum_{j=1}^{N_u} e^{f(\mathbf{x}_{uj}, \mathbf{c}_{ui})} \right) \right] \right]. \quad (19)$$

We select $f(\mathbf{x}, \mathbf{c}) = \log q_\phi(\mathbf{x}|\mathbf{c}, \mathbf{s} = \mathbf{s}_u)$, then

$$\mathcal{I}(\mathbf{x}; \mathbf{c}|\mathbf{s} = \mathbf{s}_u) \geq \mathbb{E} \left[\frac{1}{N_u} \sum_{i=1}^{N_u} \left[\log q_\phi(\mathbf{x}_{ui}|\mathbf{c}_{ui}, \mathbf{s}_u) - \log \left(\frac{1}{N_u} \sum_{j=1}^{N_u} q_\phi(\mathbf{x}_{uj}|\mathbf{c}_{ui}, \mathbf{s}_u) \right) \right] \right]. \quad (20)$$

Taking expectation of \mathbf{s} on both sides, we derive

$$\mathcal{I}(\mathbf{x}; \mathbf{c}|\mathbf{s}) \geq \mathbb{E} \left[\frac{1}{N} \sum_{u=1}^M \sum_{i=1}^{N_u} \left[\log q_\phi(\mathbf{x}_{ui}|\mathbf{c}_{ui}, \mathbf{s}_u) - \log \left(\frac{1}{N_u} \sum_{j=1}^{N_u} q_\phi(\mathbf{x}_{uj}|\mathbf{c}_{ui}, \mathbf{s}_u) \right) \right] \right]. \quad (21)$$

□

Theorem A.3 (Theorem 3.3). *If $p(\mathbf{s}|\mathbf{c})$ provides the conditional distribution between variables \mathbf{s} and \mathbf{c} , then*

$$\mathcal{I}(\mathbf{s}; \mathbf{c}) \leq \mathbb{E} \left[\frac{1}{N} \sum_{u=1}^M \sum_{i=1}^{N_u} \left[\log p(\mathbf{s}_{ui}|\mathbf{c}_{ui}) - \frac{1}{N} \sum_{v=1}^M \sum_{j=1}^{N_v} \log p(\mathbf{s}_{uj}|\mathbf{c}_{vj}) \right] \right]. \quad (22)$$

Proof of Theorem 3.3. By the upper bound in Cheng *et al.* (Cheng *et al.*, 2020b), we have

$$\mathcal{I}(\mathbf{s}; \mathbf{c}) \leq \mathbb{E}_{p(\mathbf{s}, \mathbf{c})} [\log p(\mathbf{s}|\mathbf{c})] - \mathbb{E}_{p(\mathbf{s})p(\mathbf{c})} [\log p(\mathbf{s}|\mathbf{c})]. \quad (23)$$

With embedding samples $\{\mathbf{s}_{ui}, \mathbf{c}_{ui}\}_{1 \leq u \leq M, 1 \leq i \leq N_u}$, the right-hand side of (23) can be estimated by

$$\mathcal{I}(\mathbf{s}; \mathbf{c}) \leq \mathbb{E} \left[\frac{1}{N} \sum_{u=1}^M \sum_{i=1}^{N_u} \left[\log p(\mathbf{s}_{ui}|\mathbf{c}_{ui}) - \frac{1}{N} \sum_{v=1}^M \sum_{j=1}^{N_v} \log p(\mathbf{s}_{uj}|\mathbf{c}_{vj}) \right] \right]. \quad (24)$$

□

Discussion on variational approximation As mentioned in Section 3.3, we approximate $p(\mathbf{s}|\mathbf{c})$ with a variational distribution $q_\theta(\mathbf{s}|\mathbf{c})$ in equation (10), since the closed form of $p(\mathbf{s}|\mathbf{c})$ is unknown. We claim that with $q_\theta(\mathbf{s}|\mathbf{c})$ as a good approximation of $p(\mathbf{s}|\mathbf{c})$, equation (10) remains a MI upper bound. We calculate the difference between $\mathcal{I}(\mathbf{s}; \mathbf{c})$ and the approximated version of (23):

$$\begin{aligned} \Delta &:= \mathcal{I}(\mathbf{s}; \mathbf{c}) - [\mathbb{E}_{p(\mathbf{s}, \mathbf{c})} [\log q_\theta(\mathbf{s}|\mathbf{c})] - \mathbb{E}_{p(\mathbf{s})p(\mathbf{c})} [\log q_\theta(\mathbf{s}|\mathbf{c})]] \\ &= \mathbb{E}_{p(\mathbf{s}, \mathbf{c})} [\log p(\mathbf{s}|\mathbf{c}) - \log p(\mathbf{s})] - \mathbb{E}_{p(\mathbf{s}, \mathbf{c})} [\log q_\theta(\mathbf{s}|\mathbf{c})] + \mathbb{E}_{p(\mathbf{s})p(\mathbf{c})} [\log q_\theta(\mathbf{s}|\mathbf{c})] \\ &= \left[\mathbb{E}_{p(\mathbf{s}, \mathbf{c})} [\log p(\mathbf{s}|\mathbf{c})] - \mathbb{E}_{p(\mathbf{s}, \mathbf{c})} [\log q_\theta(\mathbf{s}|\mathbf{c})] \right] - \left[\mathbb{E}_{p(\mathbf{s})} [\log p(\mathbf{s})] - \mathbb{E}_{p(\mathbf{s})p(\mathbf{c})} [\log q_\theta(\mathbf{s}|\mathbf{c})] \right] \\ &= \mathbb{E}_{p(\mathbf{s}, \mathbf{c})} \left[\log \frac{p(\mathbf{s}|\mathbf{c})}{q_\theta(\mathbf{s}|\mathbf{c})} \right] - \mathbb{E}_{p(\mathbf{s})p(\mathbf{c})} \left[\log \frac{p(\mathbf{s})}{q_\theta(\mathbf{s}|\mathbf{c})} \right] \\ &= \text{KL}(p(\mathbf{s}|\mathbf{c}) \| q_\theta(\mathbf{s}|\mathbf{c})) - \text{KL}(p(\mathbf{s}) \| q_\theta(\mathbf{s}|\mathbf{c})). \end{aligned}$$

When $q_\theta(\mathbf{s}|\mathbf{c})$ is a good approximation to $p(\mathbf{s}|\mathbf{c})$, the divergence $\text{KL}(p(\mathbf{s}|\mathbf{c}) \| q_\theta(\mathbf{s}|\mathbf{c}))$ can be smaller than $\text{KL}(p(\mathbf{s}) \| q_\theta(\mathbf{s}|\mathbf{c}))$. Then Δ remains negative, which indicates $[\mathbb{E}_{p(\mathbf{s}, \mathbf{c})} [\log q_\theta(\mathbf{s}|\mathbf{c})] - \mathbb{E}_{p(\mathbf{s})p(\mathbf{c})} [\log q_\theta(\mathbf{s}|\mathbf{c})]]$ still be an MI upper bound.

B EXPERIEMENTS

More ablation study on bottleneck design We kept the same bottleneck design as AUTOVC to have a fair comparison for the effectiveness of the proposed disentangled learning scheme. To further provide evidence of effectiveness of IDE-VC, we also conducted an ablation study in which the bottleneck is widened in Table 5. Specifically, we use set sampling rate as 4 and conduct experiments under the zero-shot setup. The results demonstrated that the bottleneck design has little impact on the disentanglement ability of the proposed model.

Table 5: Ablation study with bottleneck design for zero-shot VST. We set sampling rate as 4. Performance is measured by objective metrics.

	Distance	Verification[%]
AUTOVC	6.59	41
IDE-VC	6.24	80

More ablation study on visualization We further provide t-SNE visualization for content embedding from IDE-VC and AUTOVC in Figure 3 and Figure 4 under same hyperparameter setups. Comparing between the two t-SNE plottings, the content embeddings generated with IDE-VC are more indistinguishable for different speakers than the ones from AUTOVC, which proves that the proposed model has stronger ability to eliminate speaker-related information in content embedding.

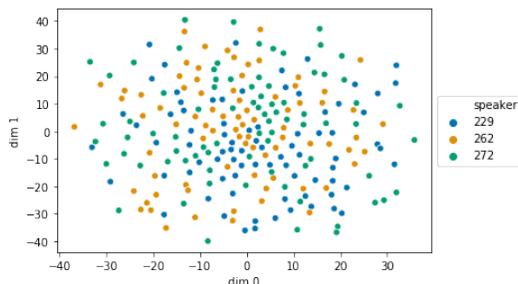


Figure 3: t-SNE visualization for content embedding from IDE-VC . The embeddings are extracted from the voice samples of 3 different speakers.

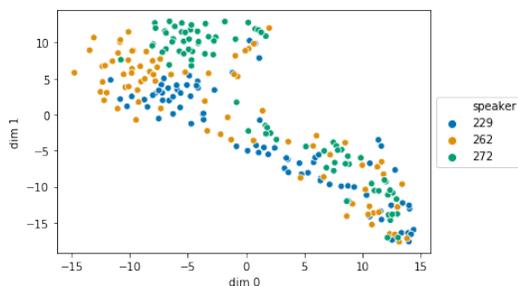


Figure 4: t-SNE visualization for content embedding from AUTOVC. The embeddings are extracted from the voice samples of 3 different speakers.

Speaker encoder pretraining Our speaker encoder is pretrained with GE2E loss on a combination of VoxCeleb1 (Nagrani et al., 2017) and Librispeech (Panayotov et al., 2015) datasets, in total of 3549 speakers.

Implementation details In our experiments, we use official implementation of AdaIN-VC⁴, AUTOVC⁵ and Blow⁶. Specifically, same pretrained speaker encoder is used in AUTOVC (Qian et al., 2019) and our model for fair comparison. Blow model is trained with 100 epochs and suggested hyperparameters, the training takes over 10 GPU days on Nvidia V100 in comparison with 1 GPU day on Nvidia Xp for our model. For StarGAN-VC, we use an open source implementation⁷, which achieves better performance according to multiple previous works (Qian et al., 2019; Serrà et al., 2019). All above models are trained on all 109 speakers in VCTK dataset, and same splits are used for testing and validation.

For our model, we use loss on validation set to conduct grid search on hyperparameter β , and we use $\beta = 5$ in final experiment. The other hyperparameters are set as the same as in AUTOVC (Qian et al., 2019).

Sample speeches We also provide several sample conversed speeches on <https://idevc.github.io/>.

⁴https://github.com/jjery2243542/adaptive_voice_conversion

⁵<https://github.com/auspicious3000/autovc>

⁶<https://github.com/joansj/blow>

⁷<https://github.com/liusongxiang/StarGAN-Voice-Conversion>

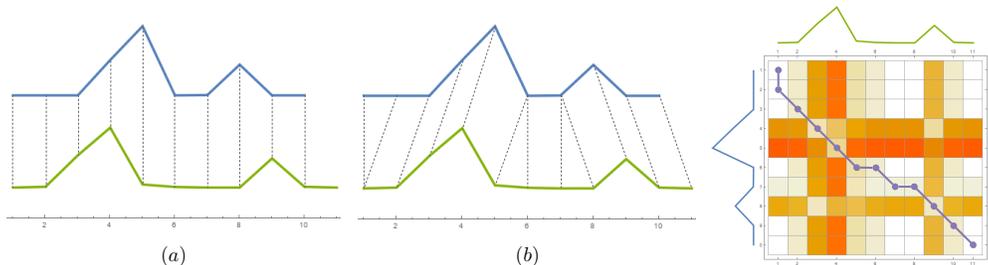


Figure 5: Left: (a) The naive Euclidean distance matching between two time series might miss the important voice patterns because of the time shift and different speakers’ speaking speeds; (b) DTW automatically find the optimal matching that captures important voice patterns. Right: DTW converts the time sequence matching problem into the minimal cost path searching on the distance matrix.

C EVALUATION DETAILS

Verification score In details, in Resemblyzer, a pre-trained speaker encoder is provided $\mathbf{s} = E_{sr}(\mathbf{x})$. The voice profile for each speaker u is first computed as $\mathbf{s}_u = \frac{1}{N_t} \sum_n E_{sr}(\mathbf{x}_{un})$, in which N_t represents the number of speeches each speaker has in testing set, \mathbf{x}_{un} represents the n -th speech of speaker u in testing set. For each speech converted from i -th speech of speaker u to j -th speech of speaker v , represented as $\hat{\mathbf{x}}_{u_i \rightarrow v_j}$, the speaker embedding is computed with Resemblyzer: $\hat{\mathbf{s}}_{u_i \rightarrow v_j} = E_{sr}(\hat{\mathbf{x}}_{u_i \rightarrow v_j})$. Dot product is used to compute similarity between the speaker embedding and the voice profile. If among all speakers in testing set, the speaker embedding of the converted speech has highest similarity score with the target speaker’s profile \mathbf{s}_v , we view it as a success conversion. The portion of success conversion among all conversion trials is reported as verification score.

Details in evaluation for zero-shot VST Based on setting in AdaIN-VC (Chou & Lee, 2019), subjective evaluation is performed on converted voice between male to male, male to female, female to male and female to female speakers. To reduce variance, 3 speakers are selected for each gender, thus, in total 36 pairs of speakers. The speakers of these pairs were unseen during training. Following the setting in AdaIN-VC (Chou & Lee, 2019), the converted result of each pair was transferred from our proposed model with only one source utterance and one target utterance.

Dynamic Time Wrapping When evaluating the voices generated by neural networks from latent embeddings, the mismatching problem in time alignment occurs due to the time shift and different speaking speeds in the generation. Important voice patterns may be neglected when directly calculating the Euclidean distance between the generated voice and the ground-truth. One effective solution to the sequential time-alignment problem is the Dynamic Time Wrapping (DTW) algorithm (Berndt & Clifford, 1994), which has been widely applied in speech recognition and matching (Muda et al., 2010; Chapaneri, 2012; Dhingra et al., 2013). The DTW algorithm seeks the optimal matching path $\mathbf{P}^* \in \mathcal{P}(T, S)$ that minimizes the sequential matching cost between two time series $\mathbf{x} = (x^1, x^2, \dots, x^T)$ and $\mathbf{y} = (y^1, y^2, \dots, y^S)$ (e.g., the purple path in the right of Figure 5). A consecutive matching path $\mathbf{P} \in \mathcal{P}(T, S)$ denotes a sequence of index pairs $\mathbf{P} = (\mathbf{p}^1, \mathbf{p}^2, \dots, \mathbf{p}^L)$, in which each pair $\mathbf{p}^l = (t_l, s_l)$ matches x^{t_l} and y^{s_l} following the time order (i.e., $\mathbf{p}^1 = (1, 1)$, $\mathbf{p}^L = (T, S)$, $0 \leq t_{l+1} - t_l \leq 1$, and $0 \leq s_{l+1} - s_l \leq 1$). The DTW score is $\mathcal{S}_{\text{DTW}}(\mathbf{x}, \mathbf{y}) = \min_{\mathbf{P} \in \mathcal{P}(T, S)} \sum_{l=1}^L d(x^{t_l}, y^{s_l})$, where $d(\cdot, \cdot)$ is a ground distance measuring the dissimilarity of any two points in time series. The optimization needed for calculating the DTW score can be solved efficiently by dynamic programming.

Human Evaluation Following Wester *et al.* (Wester et al., 2016), we use the naturalness of the speech and the similarity of the transferred speech to target identity as subjective metrics. Figure 6 and Figure 7 shows the contents of the two human evaluation webpage layouts respectively.

Qualifications Required: HIT Approval Rate (%) for all Requesters' HITs greater than 90 , Location is US , Number of HITs Approved greater than 500

Instructions ×

[View full instructions](#)

[View tool guide](#)

Some of the samples may sound somewhat degraded / distorted. Please try to listen **beyond the distortion** and concentrate on identifying the speech.

Do you think these two samples could have been produced by the same speaker?

Sample A:

▶ 0:00 / 0:00 🔊

Sample B:

▶ 0:00 / 0:00 🔊

Select an option

Same - Absolutely sure	1
Same - Not sure	2
Different - Not sure	3
Different - Absolutely sure	4

Figure 6: Human evaluation: similarity

Qualifications Required: HIT Approval Rate (%) for all Requesters' HITs greater than 90 , Location is US , Number of HITs Approved greater than 500

Instructions ×

[View full instructions](#)

[View tool guide](#)

Listen to the sample of computer generated speech and assess the quality of the audio based on how close it is to natural speech.

How natural (i.e. human-sounding) is this speech?

▶ 0:00 🔊

Select an option

Excellent - Completely natural speech	1
Good - Mostly natural speech	2
Fair - Equally natural and unnatural speech	3
Poor - Mostly unnatural speech	4
Bad - Completely unnatural speech	5

Figure 7: Human evaluation: naturalness

D DATA PROCESSING INEQUALITY

Theorem D.1. *If three variables $x \rightarrow y \rightarrow z$ follow a markov chain, then $\mathcal{I}(x; y) \geq \mathcal{I}(x; z)$.*