

AlignSTS: Speech-to-Singing Conversion via Cross-Modal Alignment

Ruiqi Li¹, Rongjie Huang¹, Lichao Zhang¹, Jinglin Liu², Zhou Zhao^{1*}

¹Zhejiang University

{ruiqili, rongjiehuang, zju_zlc, zhaozhou}@zju.edu.cn

²ByteDance

liu.jinglin@bytedance.com

Abstract

The speech-to-singing (STS) voice conversion task aims to generate singing samples corresponding to speech recordings while facing a major challenge: the alignment between the target (singing) pitch contour and the source (speech) content is difficult to learn in a text-free situation. This paper proposes AlignSTS, an STS model based on explicit cross-modal alignment, we 1) adopt a novel rhythm adaptor to predict the target rhythm representation to bridge the modality gap between content and pitch, where the rhythm representation is disentangled in a simple yet effective way and is quantized into a discrete space; and 2) leverage the cross-modal aligner to re-align the content features explicitly according to the predicted rhythm and conduct a cross-modal fusion for re-synthesis. Experimental results show that AlignSTS achieves superior performance in terms of both objective and subjective metrics. Audio samples are available at <https://alignsts.github.io>.

1 Introduction

Speech-to-singing (STS) voice conversion (Saitou et al., 2004; Cen et al., 2012; Parekh et al., 2020) aims to transfer speech samples into the corresponding singing samples, with the timbre identity and phoneme information unaltered. An STS system takes a speech sample and a target melody as conditions and then generates a high-quality singing sample following the target musical melody. Speech-to-singing research is important for human voice study and useful for practical applications such as computer-aided music production or musical entertainment.

Researchers have developed three major STS approaches: 1) Model-based approaches (Saitou et al., 2004, 2007) use phone-score synchronization information to manually align the phonemes and the target musical notes with artificial control

models. 2) Template-based STS (Cen et al., 2012; Vijayan et al., 2017, 2018) requires an available high-quality reference vocal input, which will be aligned with the input speech for subsequent musical feature extractions. The alignment is the key part and is mostly based on dynamic time warping (DTW). 3) Style transfer approach (Parekh et al., 2020) views STS as a style-transfer problem. This class of methods considers the specific properties that are transformed during conversion as "style". Parekh et al. (2020) stretch the input speech to the same length as the target F0 contour and concatenate the latent features to fuse the asynchronous representations. Wu and Yang (2020) serve as a continuation and extension of their prior work by leveraging boundary-equilibrium GAN (BEGAN) (Berthelot et al., 2017).

Despite their recent success, however, the style information is complex and is composed of multiple entangled features in the time domain (the duration information and the temporal length) and the frequency domain (the pitch information). Simply stretching the representations temporally or applying implicit self-attention can cause alignment problems. In a larger sense, human voice information (such as speech (Huang et al., 2023b) and singing (Huang et al., 2021, 2022a)) is composed of several variance information, each controlling a specific sensory modality. Plenty of arts focus on the decomposition and resynthesis of speech signal (Qian et al., 2020; Chan et al., 2022; Choi et al., 2021, 2022; Huang et al., 2022b; Huang et al.). They managed to roughly decompose speech signals into components like linguistic content, pitch, rhythm, timbre identity, etc. By manipulating any of these components, one can resynthesize a customized speech waveform. In the context of STS, the components that will be manipulated are 1) the frequency modality, namely the pitch information; and 2) the duration modality, or rhythm. The challenge remains, however, since the rhythm infor-

* Corresponding author

mation can be indeterminate given only the content and the target pitch information.

In this paper, we present a novel approach AlignSTS based on modality disentanglement and cross-modal alignment. To tackle the alignment problem, we introduce a novel rhythm representation to bridge the modality gap between content and pitch. A rhythm adaptor is designed to predict the target rhythm representation, which is used to guide the content features to perform the alignment and the cross-modal fusion. Further, we explore speech-to-sing conversion in zero-shot scenarios, where we train the model with singing data in a self-supervised manner and test it on unseen speech data. We categorize AlignSTS as one of the style transfer approaches since we only need a source speech and a target pitch contour for conversion. Experimental results demonstrate significant performance improvements over baseline models. The main contributions of this work are summarized as follows:

- We leverage the temporal modality, rhythm, to bridge the modality gap between the speech content and the target pitch. The rhythm representation is carefully designed and quantized into discrete space to model temporal dynamic states.
- We propose AlignSTS, an STS model based on cross-modal alignment, which predicts the target rhythm representation and uses it to conduct an explicit cross-modal alignment to re-align the content information.
- Experimental results demonstrate that AlignSTS achieves state-of-the-art results in both objective and subjective evaluations. Over baseline models, AlignSTS reaches an absolute improvement of 0.39 in MOS for overall quality and 0.36 for prosody naturalness.

2 Related Works

2.1 Voice Conversion

Voice conversion (VC) focuses on changing the timbre identity of an utterance to a target speaker while keeping the content information (phoneme sequence intact). Inspired by image style transfer, CycleGAN-VC (Kaneko and Kameoka, 2018) combines CycleGAN (Zhu et al., 2017) with gated CNN and identity-mapping loss to capture sequential and hierarchical structures. Similarly, StarGAN-VC (Kameoka et al., 2018) adapts StarGAN (Choi et al., 2018) and manages to train with-

out parallel utterances and pay more attention to real-time processing. Apart from GANs, conditional variational autoencoders (CVAE) are also an important class of approaches. VAE-VC (Hsu et al., 2016) uses the encoder of a VAE to learn the speaker-independent phoneme representations, thus disentangling the timbre identity. ACVAE-VC (Kameoka et al., 2019) notices that VAEs easily ignore the attribute class label input, i.e. the speaker identity, which therefore utilizes an auxiliary speaker classifier. AutoVC (Qian et al., 2019) carefully designs a bottleneck mechanism within a simple autoencoder to achieve zero-shot many-to-many voice conversion with non-parallel data.

2.2 Speech Representation Disentanglement

Human speech (Huang et al., 2023a, 2022d) is a severely complicated and comprehensive information stream, where a number of latent units entangled with each other such as content, pitch, timbre, etc. The disentanglement of the speech signal is an attempt to learn factorized and even interpretable generative factors (Bengio et al., 2013) for further application, like style transfer or domain adaptation. NANSY (Choi et al., 2021) manipulates disentangled factors such as content, pitch, and speed of speech signal, where content and pitch are decomposed by wav2vec 2.0 (Baevski et al., 2020) and Yingram, respectively. Information bottleneck (Qian et al., 2019) is also a popular way to disentanglement. Following AutoVC, SpeechSplit (Qian et al., 2020) introduces three carefully designed information bottlenecks to improve decomposition. VoiceMixer (Lee et al., 2021) leverages a similarity-based information bottleneck and adversarial feedback to disentangle content and voice style. SpeechSplit 2.0 (Chan et al., 2022) alleviates the bottleneck tuning in SpeechSplit by applying efficient signal processing techniques on encoder inputs.

3 AlignSTS

In this section, we first define and formulate the problem of speech-to-singing (STS) voice conversion. We then present the information perturbation methods, the rhythm modality modeling, and the cross-modal fusion mechanism. Finally, we introduce the overall architecture of AlignSTS and the training/inference procedure.

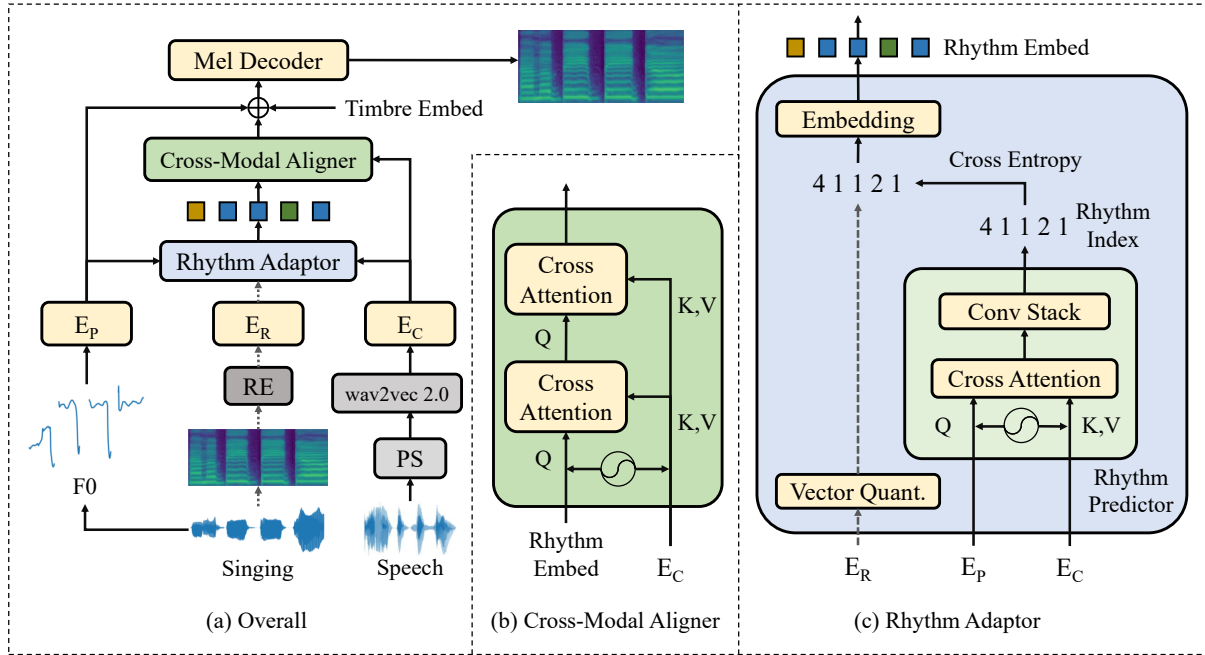


Figure 1: The architecture of the proposed model. Subfigure (a) is the overall structure. "PS" stands for pitch smooth operation. "RE" denotes extracting rhythm information. "E_P", "E_R", and "E_C" stand for three encoders used to encode pitch, rhythm, and content information, respectively. The parts connected by gray arrows on the dotted lines are only involved in the training stage and are deactivated during inference. In subfigures (b) and (c), "E_P", "E_R", and "E_C" denote the output from each encoder. In subfigure (c), a CE loss is applied to the rhythm indices predicted by the rhythm predictor during the training stage.

3.1 Problem Formulation

Let S_{sp} and S_{sg} denote the spectrograms extracted from speech and singing signals. We assume that vocal signals are a comprehensive fusion of several variance information, i.e., content, pitch, rhythm, and timbre. Therefore, this process can be formally defined in Equation 1, where c_{sp} , i_{sp} , f_{sp} , and r_{sp} denote the representations of content, timbre identity, pitch, and rhythm information of speech data. $g(\cdot)$ denotes the multi-modality fusion.

$$S_{sp} = g(c_{sp}, i_{sp}, f_{sp}, r_{sp}) \quad (1)$$

$$\widehat{S}_{sg} = f_{\theta}(c_{sp}, i_{sp}, f_{sg}, r_{sg}(c_{sp}, f_{sg})) \quad (2)$$

The problem of STS is to learn a neural network f_{θ} , such that if we switch the pitch component f_{sp} to target pitch feature f_{sg} , the corresponding singing spectrogram \widehat{S}_{sg} will be generated while preserving the content and timbre intact, as shown in Equation 2. Note that the rhythm modality implies temporal information, which will also be influenced and should be adapted according to the pitch and content. r_{sg} is the adapted rhythm representation and is generated conditioned on f_{sg} and c_{sp} .

3.2 Method Overview

AlignSTS treats speech and singing signals as a comprehensive fusion of several variance information, which can be further regarded as different sensory modalities. Pitch and rhythm features are the main modalities to convert in STS, which can be disentangled in parallel. However, the synthesis logic of variance information needs to be carefully designed during conversion. A phoneme sequence and a pitch contour seem to be uncorrelated with each other at first glance, yet it is highly possible for a human to create a suitable alignment and produce a singing melody. The mechanism behind the human behavior is: (1) find an appropriate sequence of onset and offset timings of phonemes and notes, or as known as rhythm, according to the lyrics and the melody; and (2) place the time-stretched phonemes in order according to the rhythm and combine them with the melody to produce the singing result. Inspired by this mechanism, AlignSTS:

1. Decompose the input speech signal into several disentangled variance information.
2. According to the altered speech component, i.e., the pitch contour, adapt the speech component that controls temporal duration information,

namely the rhythm representation.

3. Perform a cross-modal alignment to re-align the content according to the adapted rhythm representations and carry out a modality fusion to combine the variance information.

3.3 Information Perturbation

Speech and singing samples are fully complex and can be decomposed into variance information such as content, pitch, rhythm, and timbre. Each feature needs to be extracted and disentangled from the other, which can be achieved by certain information perturbations.

- **Linguistic Content.** We use a wav2vec 2.0 (Baevski et al., 2020) model pre-trained and fine-tuned on 960 hours of Librispeech on 16kHz sampled speech audio * to extract the linguistic content. It is shown that extracted features from speech SSP models such as wav2vec 2.0 can be applied to downstream tasks like ASR, indicating that the extracted features should provide rich isolated linguistic information. Prior works like SpeechSplit (Qian et al., 2020) and AutoPST (Qian et al., 2021) utilize random sampling to perturb the rhythm information within the content, given their input and output audios are the same speech. Using the paired speech and singing data, we do not need extra random sampling but leverage the natural discrepancy between speech and singing to perturb the rhythm within the content information.
- **Pitch.** We extract the fundamental frequency contour F0 of singing data as pitch information. The fundamental frequency contour is then quantized to 256 possible values uniformly. The F0 contour contains minimum rhythm information, given a common singing situation where one phoneme corresponds to several musical notes (or vice versa), making the rhythm theoretically indeterminate.
- **Rhythm.** Rhythm is a crucial speech component that controls the overall speed, the duration of each phoneme, and the pattern of the onset and offset of syllables. Therefore, the rhythm modality provides the duration information for both content and pitch. A good rhythm representation creates a "fill in the blank" mechanism (Qian et al., 2020) for content information to re-align. Meanwhile, the patterns of the rhythms

of speech and singing voices differ significantly, the intensity of singing voices is generally more fluctuating and distinct than speech. Therefore, we directly utilize the time-domain energy contour e_t of singing samples, which is computed by taking the L2-norm of all the frequencies for each time step. To eliminate the relative fluctuation and leave only the rhythm information, we further normalize the energy contour e_t using the Sigmoid function $\sigma(\cdot)$:

$$r_t = \sigma \left(\beta \times \frac{e_t - \text{mean}(e)}{\text{std}(e) + \epsilon} \right) \quad (3)$$

where ϵ is used to avoid division-by-zero error and β is a hyperparameter used to control the normalizing effect. r_t is the resulting rhythm representation. We further stabilize the representation by applying Gaussian filtering with a standard deviation σ of 1.0. This scalar representation disentangles the rhythm information the most thoroughly while preserving the duration information. A visualization is shown in Figure 2.

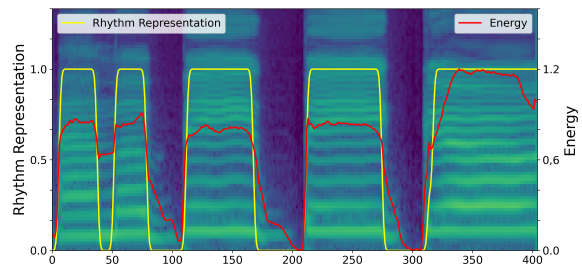


Figure 2: A visualization of rhythm representation with $\beta = 400$. The mel-spectrogram is extracted from a singing utterance "I'm a big big girl". A clear duration pattern is revealed and the relative intensity difference across phonemes is eliminated.

- **Timbre.** We leverage an open-source speaker identity encoding API Resemblyzer † to extract the timbre representations. Resemblyzer generates a 256-dimensional vector that summarizes the characteristics of the given voice spoken.

3.4 Rhythm Modality Modeling

The rhythm modality of the singing signal can be considered as a series of discrete temporal dynamic states, such as attack, sustain, transition, silence, etc. In automatic music transcription (AMT), the recognition of states refers to onset/offset detection. (Chang and Lee, 2014; Fu and Su, 2019) Similarly, traditional ASR methods like Kaldi (Povey

*<https://huggingface.co/facebook/wav2vec2-base-960h>

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

et al., 2011) model the intra- and inter-states of phonemes to improve distinguishability. From this perspective, the rhythm modality is a "softened" version of duration with more intermediate states. The onset and offset states (intermediate states) of singing may last longer than that of speech (consider the fade-out effect). Inspired by this, we adopt a **Vector Quantization** (VQ) module to quantize the continuous rhythm features to model these temporal states and form an information bottleneck (Van Den Oord et al., 2017). The designed discrete latent embedding space can be denoted as $e \in \mathbb{R}^{K \times D}$ where K is the number of clustered categories and D is the dimension of each latent embedding vector $e_k \in \mathbb{R}^D, k \in 1, 2, \dots, K$. A commitment loss (Van Den Oord et al., 2017) is used to constrain the input representation to commit to a discrete embedding:

$$\mathcal{L}_C = \|z_e(x) - \text{sg}[e]\|_2^2 \quad (4)$$

where $z_e(\cdot)$ denotes the VQ module and sg denotes the stop-gradient operator.

To model the target rhythm representations conditioned on both input content information and target pitch contour, we design a **Cross-Attention** module by adopting the Scaled Dot-Product Attention (Vaswani et al., 2017). The encoded pitch representation X_P is used as the query, and the encoded content representation X_C is used as both the key and the value. A positional encoding embedding is added to both representations before the attention module. The attention mechanism can be formulated as:

$$\begin{aligned} \text{Attention}(Q, K, V) &= \text{Attention}(X_P, X_C, X_C) \\ &= \text{Softmax}\left(\frac{X_P X_C^T}{\sqrt{d}}\right) X_C \end{aligned} \quad (5)$$

where X_P and X_C are first projected to the query, the key, and the value representations in practice.

Since the target rhythm features are not available during inference, we only activate the VQ module and use the discrete embeddings generated from it for downstream modules during the training stage. The cross-attention module and a stack of convolution layers are combined to serve as the **Rhythm Predictor**, which takes the generated discrete rhythm embeddings from the VQ module as the training target. A cross-entropy (CE) loss is applied to train the rhythm predictor so it can generate

the desired rhythm embeddings during inference:

$$\mathcal{L}_R = -\frac{1}{T} \sum_{t=1}^T \sum_{c=1}^K y_{t,c} \log(\hat{y}_{t,c}) \quad (6)$$

where T denotes the number of time frames. $y_{t,c}$ denotes the rhythm embedding indices generated from the VQ module, where $y_{t,c} = 1$ if $c = k_t$ and k_t is the rhythm index at time step t . $\hat{y}_{t,c}$ denotes the predicted rhythm indices.

3.5 Cross-Modal Alignment

With the target rhythm sequence generated, we design a cross-modal aligner to place the linguistic features along the time axis according to the target rhythm. This aligner uses rhythm information to bridge the gap between content and pitch modalities. We simply use a stack of two cross-attention layers mentioned before to complete this task. The rhythm embedding X_R generated from the rhythm adaptor is used as the query, and the encoded content representation X_C is again used as both the key and value.

The alignment using the cross-attention mechanism can be considered a soft-selection operation over the linguistic content representations X_C along the time axis at each target time step. Ideally, the resulting attention weight matrix should show a monotonic pattern and the alignment path should be concentrated and nearly diagonal. This requires extra constraints and regulations to make sure the model does not bypass the attention mechanism and simply interpolate the input content to the same length as the required rhythm representation. We mainly apply two techniques: windowing (Chorowski et al., 2015) and guided attention loss (Tachibana et al., 2018), which is described in detail in Appendix B.

With each variance information re-aligned, the content, the rhythm, and the pitch representations should have the same temporal length. We apply the inter-modality fusion across these features by element-wise vector addition. Furthermore, we involve the timbre information by adding the timbre embeddings extracted earlier.

3.6 Architecture

AlignSTS mainly consists of four modules: the encoders, the rhythm adaptor, the cross-modal aligner, and the mel decoder. The overall architecture of AlignSTS is presented in Figure 1. The detailed

description and the overall hyperparameter setting are listed in [Appendix A](#).

To accelerate the re-synthesis process and improve the audio quality, we adopt the teacher model of ProDiff ([Huang et al., 2022c](#)), a 4-step generator-based diffusion model, to be the mel decoder. A generator-based diffusion model parameterizes the denoising model by directly predicting the clean data, instead of estimating the gradient of data density. Therefore, the generator-based method bypasses the difficulty of predicting sample x_t using a single network at different diffusion steps t , which allows us to train the decoder with a reconstruction loss. We use two objectives to be the reconstruction loss:

- **Mean Absolute Error (MAE).** We apply MAE at each random term of the diffusion step t :

$$\mathcal{L}_{\text{MAE}} = \left\| \mathbf{x}_\theta \left(\alpha_t \mathbf{x}_0 + \sqrt{1 - \alpha_t^2} \epsilon \right) - \mathbf{x}_0 \right\| \quad (7)$$

where α_t denotes a derived constant that $\alpha_t = \prod_{i=1}^t \sqrt{1 - \beta_i}$, in which β_t is the predefined fixed noise schedule at diffusion step t . ϵ is randomly sampled and $\epsilon \in \mathcal{N}(0, \mathbf{I})$. \mathbf{x}_0 denotes the clean data and \mathbf{x}_θ denotes the denoised data sample predicted by the denoising neural networks θ .

- **Structural Similarity Index (SSIM) Loss.** We adopt SSIM ([Wang et al., 2004](#)), one of the state-of-the-art perceptual metrics to measure image quality, to tackle the problem of over-smoothness ([Ren et al., 2022](#)):

$$\mathcal{L}_{\text{SSIM}} = 1 - \text{SSIM} \left(\mathbf{x}_\theta \left(\alpha_t \mathbf{x}_0 + \sqrt{1 - \alpha_t^2} \epsilon \right), \mathbf{x}_0 \right) \quad (8)$$

where $\text{SSIM}(\cdot)$ is the SSIM function and is between 0 and 1.

3.7 Training and Inference

The final loss of AlignSTS consists of the following loss terms: 1) the commitment loss of VQ module \mathcal{L}_C ; 2) the CE loss of rhythm predictor \mathcal{L}_R ; 3) the guided loss for cross-attention $\mathcal{L}_{\text{attn}}$; 4) the MAE reconstruction loss \mathcal{L}_{MAE} ; and 5) the SSIM loss $\mathcal{L}_{\text{SSIM}}$. During the training stage, the cross-modal aligner takes the rhythm embeddings generated from the VQ module as input directly. At the same time, the rhythm predictor is trained to

predict the correct indices of quantized vectors for rhythm embeddings using CE loss. During the inference stage, the VQ module is deactivated and the predicted rhythm indices from the rhythm predictor are used to look up the embeddings for subsequent modules.

4 Experiments

4.1 Experimental Setup

4.1.1 Dataset

We utilize a subset of the PopBuTFy database ([Liu et al., 2022](#)) as our dataset. PopBuTFy is originally used for the singing voice beautifying (SVB) task, which consists of paired amateur and professional singing recordings. Additionally, we collected and annotated the speech version of a subset of PopBuTFy to create a paired speech and singing dataset. In all, the dataset consists of 152 English pop songs (~ 5.5 hours in total) and the respective speech recordings (~ 3.7 hours in total) from 16 singers. More details are listed in [Appendix C](#).

4.1.2 Implementation Details

We use mel-spectrograms extracted from singing samples to be the training target. We transform the raw waveform with the sampling rate of 24000 Hz into mel-spectrograms with window size 512 and hop size 128. We extract the fundamental frequency contour F_0 using Parselmouth ([Jadoul et al., 2018](#); [Boersma and Weenink, 2021](#)) as pitch information. In addition, we remove all the silent frames to accelerate the training process. The output mel-spectrograms are transformed into audio waveforms using a HiFi-GAN vocoder ([Kong et al., 2020](#)) trained with singing data in advance. More details are listed in [Appendix A](#).

4.1.3 Training and Evaluation

We train the proposed model on a single NVIDIA GeForce RTX 3090 GPU with a batch size of 20 sentences for 200k steps. The performance evaluation consists of two parts, objective and subjective evaluations, respectively.

- **Objective evaluation.** Following ([Parekh et al., 2020](#)), we use log-spectral distance (LSD) and F0 raw chroma accuracy (RCA) using `mir_eval` ([Raffel et al., 2014](#)) as the objective metrics. LSD is computed by taking the average of the Euclidean distance between the predicted log-spectrogram and the ground truth recordings. For RCA, we set the maximum tolerance deviation as 50 cents.

In addition, we design a new evaluation metric, rhythm representation distance (RRD), to measure rhythm reconstruction performance by computing the Euclidean distance between the rhythm representations described in Equation 3.

- **Subjective evaluation.** We conducted crowd-sourced mean opinion score (MOS) listening tests for subjective evaluation. Specifically, MOS-Q indicates the overall quality of the audio and MOS-P indicates the naturalness and coherence of prosody. The metrics are rated from 1 to 5 and reported with 95% confidence intervals.

4.1.4 Baseline Models

We compare the quality of the generated audio samples of AlignSTS with other approaches, including 1) *GT Mel*, in which we first convert the reference audio into mel-spectrograms and then convert them back to audio using HiFi-GAN; 2) (Parekh et al., 2020), an STS model based on encoder-decoder framework; 3) *SpeechSplit 2.0 (w/o SE)* (Chan et al., 2022), where we train the model only conditioned on the target pitch contour, while the rhythm input is implemented by interpolating the source spectral envelope (from speech) to match the length of pitch contour (in both *SpeechSplit 2.0* baselines, we interpolate the content input to match the target length); 4) *AlignSTS (zero-shot)*, where we conduct zero-shot STS and test the model on unseen speech samples; and 5) *AlignSTS (GAN)*, where we change the diffusion mel decoder to the decoder of FastSpeech 2 (Ren et al., 2020), one of the SOTA approaches of non-autoregressive text-to-speech (NAR-TTS), combining a multi-window discriminator (Wu and Luan, 2020). In addition, since *SpeechSplit 2.0* is not originally designed for STS tasks, we conduct extra experiments on baseline *SpeechSplit 2.0 (w/ SE)* in Appendix D, where we involve the target rhythm representations.

4.2 Main Results

The results are shown in Table 1. The quality of *GT Mel* is the upper limit of STS systems. In both objective and subjective evaluations, AlignSTS outperforms the baseline systems by a large margin. In both objective and subjective evaluations, the results of RCA and MOS-P are better than LSD and MOS-Q, which indicates that the condition of the target pitch contour possesses rich melody and prosody information, making the melody transfer much more effortless than phoneme modeling.

The main challenge remains in the coherence and recognizability of phonemes. The results of RRD show the reconstruction performance of the rhythm representation, demonstrating that AlignSTS with explicit rhythm modeling does the best job. In addition, the results indicate that RRD can be a valid metric for rhythm modeling.

A visualization of output mel-spectrograms is shown in Figure 3. The effectiveness of rhythm modality modeling is clearly demonstrated: 1) (Parekh et al., 2020) barely expresses the dynamic of the melody and the phonemes are scarcely distinguished; 2) the voiced parts are clustered temporally in *SpeechSplit 2.0 (w/o SE)*, but the phonemes are still unrecognizable and the formants are disordered; and 3) AlignSTS successfully re-align the phonemes in order. However, all of them reconstruct the pitch information in various degrees.

We set an additional baseline *AlignSTS (GAN)* by changing the diffusion decoder to a GAN-based decoder. Results indicate superior performance of diffusion models in singing voice synthesis.

4.3 Ablation Study

As shown in Table 2, we conduct ablation studies to demonstrate the effectiveness of several designs in AlignSTS, including the rhythm adaptor, the cross-modal alignment, and the skip-connection of pitch representation. We conduct CMOS-Q (comparative mean opinion score of quality) and CMOS-P (comparative mean opinion score of prosody) evaluations on each setting: 1) *w/o RA*: we stretch the speech rhythm representations defined in Equation 3 using linear interpolation and use it for the subsequent cross-modal alignment. 2) *w/o CM*: we drop the cross-modal alignment operation and stretch the content representation to the same length as the pitch contour and the adapted rhythm and simply fuse them together using element-wise addition. 3) *w/o FO*: we cut off the skip-connection of pitch representation in the fusion, i.e., we only combine the adapted rhythm and the aligned content to the mel decoder. The results demonstrate a significant loss of performance when dropping any module. Specifically, it indicates that the rhythm adaptor plays an important role in modeling phonemes, replacing the rhythm adaptor will end up with a singing melody with unintelligible syllables. Removing the cross-modal alignment operation witnesses the degradation of both the audio quality and the prosody naturalness, but the pres-




Figure 3: Visualization of the output mel-spectrogram of each system. (a) (Parekh et al., 2020); (b) SpeechSplit 2.0 (w/o SE); (c) AlignSTS (ours); (d) GT Mel

Method	LSD ↓	RCA ↑	RRD ↓	MOS-Q ↑	MOS-P ↑
GT Mel	2.8974	0.9959	0.3693	4.04±0.10	4.18±0.17
(Parekh et al., 2020)	7.3613	0.9218	0.7865	2.86±0.12	2.91±0.11
SpeechSplit 2.0 (w/o SE)	5.7681	0.9870	0.8262	3.19±0.06	3.45±0.13
AlignSTS (GAN)	5.4926	0.9875	0.5709	3.41±0.11	3.76±0.17
AlignSTS (ours)	5.0129	0.9934	0.5366	3.58±0.19	3.81±0.09
AlignSTS (zero-shot)	5.6607	0.9871	0.5693	3.29±0.08	3.46±0.11

Table 1: The Objective and Subjective evaluation results of STS systems.

Setting	CMOS-Q	CMOS-P
AlignSTS	0.00	0.00
w/o RA	-0.34	-0.25
w/o CM	-0.27	-0.22
w/o F0	-0.37	-0.76

Table 2: Ablation study results. RA denotes the rhythm adaptor, CM denotes the cross-modal alignment, F0 denotes the final skip-connection of pitch representation.

ence of the adapted rhythm representation allows the mel decoder to implicitly re-align the content as thoroughly as possible, resulting in a slightly smaller quantity of loss compared to w/o RA. As expected, The removal of F0 skip-connection drastically drops the prosody naturalness.

4.4 Zero-Shot Speech-to-Singing Conversion

We conduct extensional experiments on zero-shot STS given only the singing samples and test the model on unseen speech samples. Specifically, we use the identical model architecture and the training pipeline, but carry out a singing-to-singing task and train the model to reconstruct singing samples from the proposed dataset. During the inference stage, we input the unseen speech signal and the target pitch contour to generate the corresponding singing samples. The results are also shown in Table 1 and the task is denoted as AlignSTS (zero-shot). AlignSTS demonstrates great potential in zero-shot STS.

5 Conclusion

In this work, we presented AlignSTS, a speech-to-singing model based on modality disentanglement and cross-modal alignment. To achieve better voice quality and improve interpretability, we decomposed the input speech waveform into four variance information and proposed a novel cross-modal fusion mechanism. Specifically, we designed a rhythm adaptor to adjust the rhythm representation to deal with the altered pitch modality, and a cross-modal aligner to re-align the content representation. Finally, we conduct a cross-modal fusion to combine the different components together.

Experimental results demonstrated that AlignSTS achieved superior quality and naturalness compared with several baselines. Ablation studies demonstrated that each design in AlignSTS was effective. Extensive experiments showed the great potential of AlignSTS in zero-shot STS. We envisage that our work could serve as a basis for future STS studies.

6 Limitations

Research on the speech-to-singing conversion is important for human voice study and useful for practical applications such as computer-based music productions or entertainment. However, current STS approaches require an input condition of a fine-grained target F0 contour, which is always unavailable. In addition, the F0 contour of a singing utterance often possesses rich speaker-related infor-

mation, which still needs further disentanglement. Finetuning F0 contours in real applications brings significant extra work. One of our future directions is to simplify the input conditions, such as musical scores. Furthermore, the preliminary attempt at the zero-shot STS task may lead to a better perspective.

Except for positive applications, STS systems may face ethical concerns. With the development of speech/singing voice synthesis technology, the cost of faking an utterance of a specific individual is gradually declining. Researchers need further consideration of the regulation and recognition of speech/singing voice synthesis.

7 Acknowledgement

This work was supported in part by the National Key R&D Program of China under Grant No.2022ZD0162000, National Natural Science Foundation of China under Grant No.62222211, Grant No.61836002 and Grant No.62072397.

References

- Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. 2020. wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in Neural Information Processing Systems*, 33:12449–12460.
- Yoshua Bengio, Aaron Courville, and Pascal Vincent. 2013. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828.
- David Berthelot, Thomas Schumm, and Luke Metz. 2017. Began: Boundary equilibrium generative adversarial networks. *arXiv preprint arXiv:1703.10717*.
- Paul Boersma and David Weenink. 2021. Praat: doing phonetics by computer [Computer program]. Version 6.1.38, retrieved 2 January 2021 <http://www.praat.org/>.
- Ling Cen, Minghui Dong, and Paul Chan. 2012. Template-based personalized singing voice synthesis. In *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4509–4512. IEEE.
- Chak Ho Chan, Kaizhi Qian, Yang Zhang, and Mark Hasegawa-Johnson. 2022. Speechsplit2. 0: Unsupervised speech disentanglement for voice conversion without tuning autoencoder bottlenecks. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6332–6336. IEEE.
- Sungkyun Chang and Kyogu Lee. 2014. A pairwise approach to simultaneous onset/offset detection for singing voice using correntropy. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 629–633. IEEE.
- Hyeong-Seok Choi, Juheon Lee, Wansoo Kim, Jie Lee, Hoon Heo, and Kyogu Lee. 2021. Neural analysis and synthesis: Reconstructing speech from self-supervised representations. *Advances in Neural Information Processing Systems*, 34:16251–16265.
- Hyeong-Seok Choi, Jinhyeok Yang, Juheon Lee, and Hyeongju Kim. 2022. Nansy++: Unified voice synthesis with neural analysis and synthesis. *arXiv preprint arXiv:2211.09407*.
- Yunje Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. 2018. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8789–8797.
- Jan K Chorowski, Dzmitry Bahdanau, Dmitriy Serdyuk, Kyunghyun Cho, and Yoshua Bengio. 2015. Attention-based models for speech recognition. *Advances in neural information processing systems*, 28.
- Zih-Sing Fu and Li Su. 2019. Hierarchical classification networks for singing voice segmentation and transcription. In *Proceedings of the 20th International Society for Music Information Retrieval Conference (ISMIR 2019)*, pages 900–907.
- Chin-Cheng Hsu, Hsin-Te Hwang, Yi-Chiao Wu, Yu Tsao, and Hsin-Min Wang. 2016. Voice conversion from non-parallel corpora using variational auto-encoder. In *2016 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA)*, pages 1–6. IEEE.
- Rongjie Huang, Feiyang Chen, Yi Ren, Jinglin Liu, Chenye Cui, and Zhou Zhao. 2021. Multi-singer: Fast multi-singer singing voice vocoder with a large-scale corpus. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 3945–3954.
- Rongjie Huang, Chenye Cui, Feiyang Chen, Yi Ren, Jinglin Liu, Zhou Zhao, Baoxing Huai, and Zhefeng Wang. 2022a. Singgan: Generative adversarial network for high-fidelity singing voice generation. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 2525–2535.
- Rongjie Huang, Jiawei Huang, Dongchao Yang, Yi Ren, Luping Liu, Mingze Li, Zhenhui Ye, Jinglin Liu, Xiang Yin, and Zhou Zhao. 2023a. Make-an-audio: Text-to-audio generation with prompt-enhanced diffusion models. *arXiv preprint arXiv:2301.12661*.
- Rongjie Huang, Max WY Lam, Jun Wang, Dan Su, Dong Yu, Yi Ren, and Zhou Zhao. 2022b. Fastdiff: A fast conditional diffusion model for high-quality speech synthesis. *arXiv preprint arXiv:2204.09934*.

- Rongjie Huang, Mingze Li, Dongchao Yang, Jia-tong Shi, Xuankai Chang, Zhenhui Ye, Yuning Wu, Zhiqing Hong, Jiawei Huang, Jinglin Liu, et al. 2023b. Audiogpt: Understanding and generating speech, music, sound, and talking head. *arXiv preprint arXiv:2304.12995*.
- Rongjie Huang, Yi Ren, Jinglin Liu, Chenye Cui, and Zhou Zhao. Generspeech: Towards style transfer for generalizable out-of-domain text-to-speech. In *Advances in Neural Information Processing Systems*.
- Rongjie Huang, Zhou Zhao, Huadai Liu, Jinglin Liu, Chenye Cui, and Yi Ren. 2022c. Prodiff: Progressive fast diffusion model for high-quality text-to-speech. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 2595–2605.
- Rongjie Huang, Zhou Zhao, Jinglin Liu, Huadai Liu, Yi Ren, Lichao Zhang, and Jinzheng He. 2022d. Transpeech: Speech-to-speech translation with bilateral perturbation. *arXiv preprint arXiv:2205.12523*.
- Yannick Jadoul, Bill Thompson, and Bart de Boer. 2018. [Introducing Parselmouth: A Python interface to Praat](#). *Journal of Phonetics*, 71:1–15.
- Hirokazu Kameoka, Takuhiro Kaneko, Kou Tanaka, and Nobukatsu Hojo. 2018. Stargan-vc: Non-parallel many-to-many voice conversion using star generative adversarial networks. In *2018 IEEE Spoken Language Technology Workshop (SLT)*, pages 266–273. IEEE.
- Hirokazu Kameoka, Takuhiro Kaneko, Kou Tanaka, and Nobukatsu Hojo. 2019. Acvae-vc: Non-parallel voice conversion with auxiliary classifier variational autoencoder. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 27(9):1432–1443.
- Takuhiro Kaneko and Hirokazu Kameoka. 2018. Cyclegan-vc: Non-parallel voice conversion using cycle-consistent adversarial networks. In *2018 26th European Signal Processing Conference (EUSIPCO)*, pages 2100–2104. IEEE.
- Jungil Kong, Jaehyeon Kim, and Jaekyoung Bae. 2020. Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis. *Advances in Neural Information Processing Systems*, 33:17022–17033.
- Sang-Hoon Lee, Ji-Hoon Kim, Hyunseung Chung, and Seong-Whan Lee. 2021. Voicemixer: Adversarial voice style mixup. *Advances in Neural Information Processing Systems*, 34:294–308.
- Jinglin Liu, Chengxi Li, Yi Ren, Zhiying Zhu, and Zhou Zhao. 2022. [Learning the beauty in songs: Neural singing voice beautifier](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7970–7983, Dublin, Ireland. Association for Computational Linguistics.
- Jayneel Parekh, Preeti Rao, and Yi-Hsuan Yang. 2020. Speech-to-singing conversion in an encoder-decoder framework. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 261–265. IEEE.
- Daniel Povey, Arnab Ghoshal, Gilles Boulianne, Lukas Burget, Ondrej Glembek, Nagendra Goel, Mirko Hannemann, Petr Motlicek, Yanmin Qian, Petr Schwarz, Jan Silovsky, Georg Stemmer, and Karel Vesely. 2011. The kaldi speech recognition toolkit. In *IEEE 2011 Workshop on Automatic Speech Recognition and Understanding*. IEEE Signal Processing Society. IEEE Catalog No.: CFP11SRW-USB.
- Kaizhi Qian, Yang Zhang, Shiyu Chang, Mark Hasegawa-Johnson, and David Cox. 2020. Unsupervised speech decomposition via triple information bottleneck. In *International Conference on Machine Learning*, pages 7836–7846. PMLR.
- Kaizhi Qian, Yang Zhang, Shiyu Chang, Jinjun Xiong, Chuang Gan, David Cox, and Mark Hasegawa-Johnson. 2021. [Global prosody style transfer without text transcriptions](#). In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8650–8660. PMLR.
- Kaizhi Qian, Yang Zhang, Shiyu Chang, Xuesong Yang, and Mark Hasegawa-Johnson. 2019. Autovc: Zero-shot voice style transfer with only autoencoder loss. In *International Conference on Machine Learning*, pages 5210–5219. PMLR.
- Colin Raffel, Brian McFee, Eric J Humphrey, Justin Salamon, Oriol Nieto, Dawen Liang, Daniel PW Ellis, and C Colin Raffel. 2014. mir_eval: A transparent implementation of common mir metrics. In *Proceedings of the 15th International Society for Music Information Retrieval Conference, ISMIR*. Cite-seer.
- Yi Ren, Chenxu Hu, Xu Tan, Tao Qin, Sheng Zhao, Zhou Zhao, and Tie-Yan Liu. 2020. Fastspeech 2: Fast and high-quality end-to-end text to speech. *arXiv preprint arXiv:2006.04558*.
- Yi Ren, Xu Tan, Tao Qin, Zhou Zhao, and Tie-Yan Liu. 2022. Revisiting over-smoothness in text to speech. *arXiv preprint arXiv:2202.13066*.
- Takeshi Saitou, Masataka Goto, Masashi Unoki, and Masato Akagi. 2007. Speech-to-singing synthesis: Converting speaking voices to singing voices by controlling acoustic features unique to singing voices. In *2007 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, pages 215–218. IEEE.
- Takeshi Saitou, Naoya Tsuji, Masashi Unoki, and Masato Akagi. 2004. Analysis of acoustic features affecting "singing-ness" and its application to singing-voice synthesis from speaking-voice. In *Eighth International Conference on Spoken Language Processing*.

- Hideyuki Tachibana, Katsuya Uenoyama, and Shunsuke Aihara. 2018. Efficiently trainable text-to-speech system based on deep convolutional networks with guided attention. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4784–4788. IEEE.
- Aaron Van Den Oord, Oriol Vinyals, et al. 2017. Neural discrete representation learning. *Advances in neural information processing systems*, 30.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Karthika Vijayan, Minghui Dong, and Haizhou Li. 2017. A dual alignment scheme for improved speech-to-singing voice conversion. In *2017 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, pages 1547–1555. IEEE.
- Karthika Vijayan, Xiaoxue Gao, and Haizhou Li. 2018. Analysis of speech and singing signals for temporal alignment. In *2018 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, pages 1893–1898. IEEE.
- Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612.
- Da-Yi Wu and Yi-Hsuan Yang. 2020. Speech-to-singing conversion based on boundary equilibrium gan. *arXiv preprint arXiv:2005.13835*.
- Jie Wu and Jian Luan. 2020. Adversarially trained multi-singer sequence-to-sequence singing synthesizer. *arXiv preprint arXiv:2006.10317*.
- Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232.

A Architecture

The architecture and hyperparameters are listed in Table 4. We use three encoders E_P , E_R , and E_C to encode the target F0, the target rhythm, and the source content features:

$$\mathbf{x}_P = E_P(\mathbf{f}_0), \mathbf{x}_R = E_R(\mathbf{r}), \mathbf{x}_C = E_C(\mathbf{x}) \quad (9)$$

where all encoders are stacks of several convolution layers.

These encoded features are fed into the rhythm adaptor, to both generate the discrete rhythm embeddings and train the rhythm predictor. The

rhythm predictor consists of the cross-attention module and several convolutional layers, predicting the target discrete rhythm embeddings conditioned on the content and pitch features. The target rhythm embeddings are then used to re-align the source content representations. Finally, we carry out an inter-modality fusion across four different representations for re-synthesis.

Each encoder (E_R , E_P , and E_C) consists of two 1-D convolutional layers, where the kernel sizes are 7, 5, and 3, respectively. The linguistic hidden features extracted from wav2vec 2.0 are 32-dimensional, which are widely used for downstream tasks. The hidden size of all the model components is 256. The size of the codebook in the VQ module is set to 6. The cross-attention layer is implemented with one multi-head layer and a feed-forward layer, where the latter consists of a 1-D convolution layer and a fully-connected layer.

B Attention and Alignment Regulation

B.1 Attention Windowing

Attention windowing is a widely used technique that controls the "field of view" at each time step. Only a subsequence of key $\hat{\mathbf{x}} = [\mathbf{x}_{p_t-w}, \dots, \mathbf{x}_{p_t+w}]$ are considered at each query time step t , where w is the window width and p_t is the middle position of the window along the key time axis. Specifically, we replace all the values outside the window with -10^8 before $\text{Softmax}(\cdot)$ so that the contribution outside the window is reduced significantly.

B.2 Guided Attention Loss

To make sure that the attention weight matrix is nearly diagonal and monotonic, we adopt guided attention loss. Let $\alpha_{t,n}$ denotes the attention weight at query time step t that attends to key time step n , a guided attention loss can be defined as:

$$\mathcal{L}_{\text{attn}} = \frac{1}{TN} \sum_{t=1}^T \sum_{n=1}^N \alpha_{t,n} w_{t,n}, \text{ where} \quad (10)$$

$$w_{t,n} = 1 - \exp\left(-\frac{\left(\frac{n}{N} - \frac{t}{T}\right)^2}{2g^2}\right) \quad (11)$$

where T and N are the lengths of the query and the key, respectively. $w_{t,n}$ is the weight distribution of the constraint. g is a hyperparameter used to control the concentration degree, which is set $g = 0.1$ in practice. If $\alpha_{t,n}$ is far from the diagonal, meaning that the key representations (like the content

Method	LSD ↓	RCA ↑	MOS-Q ↑	MOS-P ↑
<i>GT Mel</i>	2.8974	0.9959	4.04±0.10	4.18±0.17
<i>SpeechSplit 2.0 (w/o SE)</i>	5.7681	0.9870	3.19±0.06	3.45±0.13
<i>SpeechSplit 2.0 (w/ SE)</i>	4.4871	0.9848	3.65±0.14	3.88±0.05
<i>AlignSTS (ours)</i>	5.0129	0.9934	3.58±0.19	3.81±0.09

Table 3: Experimental results involving *SpeechSplit 2.0 (w/ SE)*.

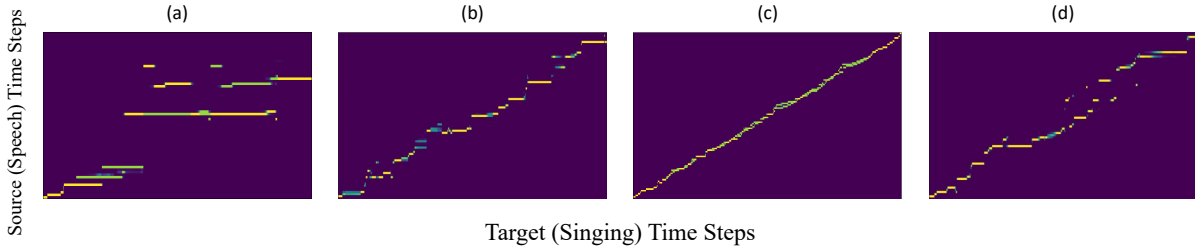


Figure 4: Visualization of Attention weights. (a) Cross-attention without alignment regulation; (b) *AlignSTS (ours)*; (c) *AlignSTS (zero-shot)* during training stage; (d) *AlignSTS (zero-shot)* during inference stage.

features) are placed in a random order, it provides a strong penalty.

B.3 Visualization

A visualization of the attention weights is shown in Figure 4. All the attention weights are extracted from the last layer of the cross-modal aligner and averaged across attention heads.

Comparing subfigure (a) the attention weights without alignment regulation and (b) that with regulation, the importance of attention path regulation is clearly demonstrated. Without alignment regulation, skips and non-monotonic situations occur, causing disordered or even indistinguishable phonemes.

As for the zero-shot scenario, we compare the attention weights during the training stage and during the inference stage, which are shown in subfigures (c) and (d), respectively. Since the training is conducted in a self-supervised manner, the attention pattern demonstrates a perfectly linear pattern, as expected. The cross-modal aligner learns how to uniformly interpolate and stretch the input wav2vec 2.0 features to the same length as the target representations. However, in inference, the aligner is still capable of predicting the specific duration information of each linguistic unit of unseen speech data to a certain degree. As shown in subfigure (d), *AlignSTS* demonstrates its generalizability to unseen speech data and the ability to explore modality interaction in self-supervised pre-training.

C Dataset

We collected and annotated the speech version of a subset of PopBuTFy to create a paired speech and singing dataset. During the collection, the private information of the speakers was protected. The qualified speakers are requested to read the lyrics of songs that have been sung by themselves. The personal vocal timbre is kept still during the recording process. We carefully select a subset of the collected recordings to create a high-quality dataset. In all, the dataset consists of 152 English pop songs (~5.5 hours in total) and the respective speech recordings (~3.7 hours in total) from 16 singers. All the audio files are recorded in a professional recording studio by professional singers, male and female. The recordings are sampled at 22050 Hz with 16-bit quantization. We randomly pick 111 pieces for validation and testing.

D Extensional Experiments

SpeechSplit 2.0 is originally designed for aspect-specific voice conversion, not STS tasks. Only manipulating the pitch component of *SpeechSplit 2.0* input for STS may cause severe alignment problems, since the latent rhythm information is influenced. To lower the training difficulty and explore the importance of rhythm information, we add a new baseline *SpeechSplit 2.0 (w/ SE)* that involves the target ground truth rhythm information, i.e., the perturbed spectral envelope, in the training and inference procedure. The performances are listed in Table 3. The results show a great improvement

brought by this "information leak", in that the target rhythm information should not be available in a real situation. Also, the perturbed spectral envelope may still carry residual linguistic information for phoneme reconstruction.

Hyperparameter		AlignSTS
Pitch Encoder	Encoder Kernel	5
	Encoder Layers	3
	Encoder Hidden	256
Content Encoder	Encoder Kernel	3
	Encoder Layers	2
	Encoder Hidden	256
Rhythm Encoder	Encoder Kernel	7
	Encoder Layers	2
	Encoder Hidden	256
Rhythm Adaptor	Attention Hidden	256
	Attention Heads	1
	Attention Layers	1
	Attention Window Size	0.5
	Attention Guided g	0.1
	Attention FFN Kernel	9
	Conv1D Kernel	3
	Conv1D Layers	2
	Conv1D Hidden	256
	Conv1D Dropout	0.8
	VQ Embeddings	6
VQ Hidden	256	
Cross-Modal Aligner	Attention Hidden	256
	Attention Heads	2
	Attention Layers	2
	Attention Dropout	0.1
	Attention Window Size	0.4
	Attention Guided g	0.1
	Attention FFN Kernel	9
Diffusion Decoder	Denoiser Layers	20
	Denoiser Hidden	256
	Time Steps	4
	Noise Schedule Type	VPSDE
Total Number of Parameters		26M

Table 4: Hyperparameters of AlignSTSmodules.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
section 6, after the conclusion
- A2. Did you discuss any potential risks of your work?
section 6
- A3. Do the abstract and introduction summarize the paper’s main claims?
The introduction is in section 1.
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

section 3 and 4

- B1. Did you cite the creators of artifacts you used?
section 3 and 4
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
The license of a GPU can be easily found online and is known to all.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
Section 4.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
section 4.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
The documentation of a GPU can be easily found online and is known to all.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Left blank.

C Did you run computational experiments?

section 4

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
section 4

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?
section 4
- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?
section 4
- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?
section 3 and 4
- D** **Did you use human annotators (e.g., crowdworkers) or research with human participants?**
section 4
- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?
It's not relevant
- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?
It's not relevant and this information is redundant.
- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?
There is no space for that, but yes, our instructions explain how the data would be used and the participants agreed with the terms.
- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?
We hire a company to collect the data.
- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?
section 4