



# Iteratively Improving Speech Recognition and Voice Conversion

Mayank Kumar Singh<sup>1</sup>, Naoya Takahashi<sup>2</sup>, Naoyuki Onoe<sup>1</sup>

<sup>1</sup>Sony Research India, India, <sup>2</sup>Sony Group Corporation, Japan  
{mayank.a.singh, naoya.takahashi, naoyuki.onoe}@sony.com

## Abstract

Many existing works on voice conversion (VC) tasks use automatic speech recognition (ASR) models for ensuring linguistic consistency between source and converted samples. However, for the low-data resource domains, training a high-quality ASR remains to be a challenging task. In this work, we propose a novel iterative way of improving both the ASR and VC models. We first train an ASR model which is used to ensure content preservation while training a VC model. In the next iteration, the VC model is used as a data augmentation method to further fine-tune the ASR model and generalize it to diverse speakers. By iteratively leveraging the improved ASR model to train VC model and vice-versa, we experimentally show improvement in both the models. Our proposed framework outperforms the ASR and one-shot VC baseline models on English singing and Hindi speech domains in subjective and objective evaluations in low-data resource settings.

**Index Terms:** iterative training, automatic speech recognition, voice conversion, data augmentation

## 1. Introduction

Speech processing technologies such as voice conversion (VC) and automatic speech recognition (ASR) have dramatically improved in the past decade owing to the advancements in deep learning technologies. However, the task of training these models remains challenging on low resource domains as they suffer from over-fitting and do not generalize well for practical applications. As VC models often rely on ASR model for extracting content features or imposing content consistency loss, [1, 2, 3], degradation of ASR directly affects the quality of VC models. On the other hand, to improve the generalization capability of ASR model, a variety of data augmentation techniques have been proposed with voice conversion being one of them [4].

This creates a causality dilemma, wherein poor quality of ASR model affects VC model training which in turn leads to low quality data augmentation for training ASR models. Conversely, improving the ASR model should lead to better VC models, which should produce better data augmentation samples for improving the ASR models. Motivated from this, in this work we propose to iteratively improve the ASR model by using the VC model as a data augmentation method for training the ASR and simultaneously improve the VC model by using the ASR model for linguistic content preservation.

We validate the proposed framework on two low-resource domains, namely Hindi speech voice and English singing domains, and show that our proposed iterative framework improves the WER of ASR model as well as improves the subjective and objective metrics of the VC model. We chose

English singing voice as it has low amount of public resources for the task of VC as well as low amount of manual annotations for training an ASR. Acapella singing voice without annotations is also scarce, making the training of unsupervised ASR models unfeasible. For showing the applicability of our approach across languages different from English, we chose the Hindi speech domain in low resource settings.

The contributions of our work are summarized as follows:

1. We propose an iterative framework to improve ASR models by using VC models for data augmentation and VC models by using ASR models as speech consistency loss.
2. We evaluate the proposed framework on Hindi speech domain under low data resource settings and successfully show its superiority over the baseline models for both ASR and VC.
3. We evaluate the proposed framework for the one-shot singing voice conversion and singing voice recognition task using NUS-48E and NHSS datasets and show that the proposed framework significantly improves the word error rate of the ASR model and quality of samples generated using the VC model compared to models trained without the proposed framework.

The audio samples are available at our website<sup>1</sup>

## 2. Related Work

The problem of converting the speaker identity of a voice while preserving the linguistic content, such as speech voice conversion (converting the speaker identity of speech signal) [2, 5, 6, 7, 8], singing voice conversion (converting the speaker identity of a singing signal) [9, 1, 3, 10] and emotional voice conversion [11] (converting the emotion of expressive speech signals) has been actively studied for decades. Recent advances in deep learning have been inspired from techniques like auto-encoder [5, 6, 9, 10] and generative adversarial networks (GAN) [2, 3]. Recently, ASRs have been used to improve the linguistic content preservation of VC models, [1, 2, 3]. These techniques have dramatically improved VC models, yet the performance drops significantly when we consider the challenging and realistic scenario of working on low data resource domains.

ASR performance in the supervised case has seen rapid improvements in the last few years due to advances in the model architecture [12, 13, 14], training methods [15] and increasing data size [16]. Recently, unsupervised ASR has gained popularity [17] and shown remarkable efficacy in utilizing large sets of unlabelled data. Yet, the accuracy of the ASR models drop

<sup>1</sup>[https://demosamplesites.github.io/IterativeASR\\_VC/](https://demosamplesites.github.io/IterativeASR_VC/).




Figure 1: VC Model training with Speech Consistency Loss.

significantly when we try to generalize them across domains, for e.g. from speech to singing, across languages, or train them on low data resource domains. It becomes more important to address the limitations of current ASR models because of their usage in domains having wide variety of applications like voice conversion [1, 3, 4], source separation [18] and text to speech generation [19].

### 3. VC Data Augmentation in ASR training

Data augmentation has been effectively used for training ASR models by modifying the pitch [20], applying spectral augmentation and speed perturbation [21]. Although these augmentation strategies have been robustly shown to improve ASR models, recently, voice conversion has been shown to outperform these methods [22, 4]. The effectiveness of using VC to augment the dataset in very low resource domains and superiority of using VC over spectral augmentation is established in [22]. However, they do not make use of an ASR model during the training of the VC model, hence their VC model cannot be shown to leverage the improved ASR model which is trained using VC data augmentation. In [4], it is shown that even in the case of high data resource domain, VC is effective as a data augmentation method for generalizing the ASR model to unseen speakers. They replace the encoder of the VC model with a pre-trained ASR encoder and show the effect of using such an encoder on the quality and speaker identity conversion. This motivated us to adopt an iterative approach of using the VC as a data augmentation tool for improving the ASR and using the ASR to improve the VC model. Different from [4], which uses the ASR encoder to disentangle the speaker identity and the linguistic content, we use the ASR encoder as an auxiliary loss to preserve the linguistic content. We further detail this in Section 4.

### 4. Speech Consistency Loss

The effectiveness of using a pre-trained ASR model for preserving the linguistic content during voice conversion has been shown in [1, 2, 4]. One way of using ASR features to preserve the linguistic features is by applying the speech consistency loss [2, 1]. In [2], the speech consistency loss is defined as

$$\mathcal{L}_{asr}(\theta) = \mathbb{E}_{X,s} \left[ \| |h_{asr}(X) - h_{asr}(V_{\theta}(X, s))| \|_1 \right] \quad (1)$$

where  $X$  denotes mel-spectrogram of the source,  $h_{asr}(\cdot)$  denotes the encoder features of a differentiable pre-trained ASR model,  $s$  is the speaker embedding of the target speaker, and  $V$  is the voice conversion model which is conditioned upon  $X$  and  $s$ .  $\theta$  denotes the VC model parameters. It is assumed that the output of the encoder of an ASR model is independent of

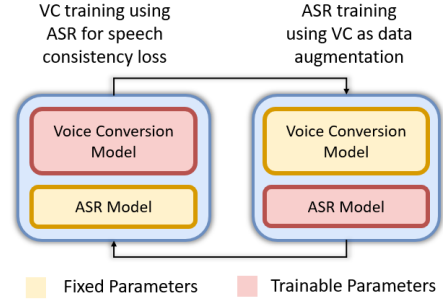


Figure 2: Iterative training for VC and ASR

speaker identity and pitch and contains only the linguistic content. Thus minimizing the  $L_{asr}$  loss emphasises the preservation of the linguistic content after voice conversion and does not hinder speaker identity conversion nor pitch conversion.

The voice conversion training process using an ASR model for preserving the linguistic content is shown in Figure 1.

## 5. Iterative Training Framework

### Algorithm 1 Iterating training of ASR and VC

**Require:** Training data set  $\tau$ , ASR fine-tuning routine  $F(A_i, Dataset)$ , VC training routine  $T(Dataset, A_i)$

- 1: Train  $A_0$  using standard ASR training
- 2: Train  $V_0$  by utilizing  $A_0$  for minimizing  $L_{asr}$
- 3:  $i = 0$
- 4: **while** not done **do**
- 5:  $\hat{\tau}_{i+1} = V_i(\tau)$  ▷ Apply VC data augmentation
- 6:  $A_{i+1} = F(A_i, \{\tau, \hat{\tau}_{i+1}\})$  ▷ Fine-tune ASR Model
- 7:  $V_{i+1} = T(\tau, A_{i+1})$  ▷ Train VC Model
- 8:  $i \leftarrow i + 1$
- 9: **end while**

The proposed iterative training framework is illustrated in Figure 2 and its pseudo code is shown in Algorithm 1. In the first step of the iterative training, an ASR model is trained on the low resource domain dataset,  $\tau$ , which we call as  $A_0$ . Next, we train a VC model which uses  $A_0$  for minimizing  $L_{asr}$ . We call this VC model as  $V_0$ . In the next step, we use  $V_0$  to apply data augmentation to  $\tau$  and get  $\hat{\tau}_1$ . For applying the VC data augmentation, we sample both the source and reference from the train dataset and provide them to  $V_0$  to artificially increase the diversity of the training samples in terms of speaker style. We further fine-tune  $A_0$  on  $\tau$  and  $\hat{\tau}_1$  and refer to the fine-tuned model as  $A_1$ . Using  $A_1$ , for minimizing  $L_{asr}$ , we again train a VC model which we denote as  $V_1$ . This process of iterative refinement of ASR and VC models is repeated until the WER of the ASR model converges on the validation set.

The motivation of our approach is as follows. For training a VC model, using an ASR model which is able to extract better linguistic features and generates speaker independent encoder features will lead to better content preservation and clarity. Similarly, for training an ASR model, using a VC model which preserves more linguistic content and produces less artefacts will generate more realistic samples for data augmentation. By iteratively training the ASR and VC model we leverage the improved models for training the next iteration and thus achieve better performance on the objective and subjective metrics.

Table 1: Hindi ASR WER across iterations

Iter	0	1	2	3
WER ↓	27.6 ± 0.13	26.2 ± 0.17	25.8 ± 0.16	<b>25.7 ± 0.15</b>

Although in this work we use speech consistency loss to leverage the ASR model for improving the linguistic preservation capacity of the VC models, another approach of using the ASR models is by providing the ASR features extracted from  $X_{src}$  as an input to the VC model. Such an approach is used in [1], which is an auto-encoder based VC model that takes the pitch, loudness and ASR features as input. Although this approach has the advantage of utilizing non-differentiable ASR models for training the VC model, it strongly relies on the ASR features for extracting the linguistic content that can lead to poor generalization of the VC model to out-of-domain samples.

## 6. Experiments

To evaluate the proposed framework in a realistic scenario of low resource domains we chose Hindi speech and English singing domains for the task of voice conversion and automatic speech recognition.

### 6.1. Datasets

For the Hindi speech dataset, we use an internal data consisting of nine speakers with a total duration of three hours. Out of the nine speakers, five are male and four are female. The speakers were given non-overlapping transcripts in the Hindi language and were asked to speak them in a neutral tone. The recordings were conducted in a studio setup. As the transcripts are in Hindi, we use a transliteration package [23] to convert the Devanagiri characters to Latin characters. For training, we use a subset of seven speakers and use two speakers for evaluation. We train an ASR model for WER objective evaluation of the Hindi speech VC models. For this, we use the publicly available MUCS2021 dataset [24] having 95.05 hours of training audio and 5.55 hours of testing audio.

For the English singing dataset, we use the publicly available NUS-48E [25] and NHSS [26] datasets. The NUS-48E dataset consists of twelve singers with a total of two hours of singing data and one hour of speech data. The NHSS dataset consists of ten singers with a total of 4.75 hours of singing data and 2.25 hours of speech data. For the experiments we use only the singing data from NUS-48E and NHSS datasets. For training, we use a subset of ten singers from NUS-48E dataset and eight singers from NHSS dataset which has a duration of four hours and use the remaining speakers for evaluation.

For both the Hindi speech and English singing domain task, the ASR model is trained on a combination of the target domain and LibriSpeech dataset, which is a large English dataset consisting of 960 hours of data [27] and is publicly available. The LibriSpeech corpus is incorporated during training to improve generalization of the ASR model.

For the VC and ASR training, all the data is first re-sampled to 24kHz and an 80 bin mel-spectrogram representation is calculated as in [2].

### 6.2. Model Architecture

For the ASR model, we choose the ESPNet framework [28] having a hybrid CTC-Attention model [29]. For  $h_{asr}$ , we use the output of the encoder with 512 dimensions. As the ASR

Table 2: Hindi VC Results.

Iterations	WER ↓	pMOS ↑	Identity ↑
0	41.93 ± 0.07	3.79 ± 0.01	0.7410 ± 0.0007
1	38.87 ± 0.07	3.89 ± 0.01	0.7388 ± 0.0007
2	38.85 ± 0.07	3.91 ± 0.01	0.7282 ± 0.0007
3	<b>38.35 ± 0.07</b>	<b>3.94 ± 0.01</b>	<b>0.7498 ± 0.0007</b>

model needs to extract linguistic features from the VC model’s converted samples, we modify the mel-spectrogram extraction hyper-parameters to match those of the VC model.

For the VC model we adopted [3], which is a recently proposed modification of [2], a GAN-based network. A brief overview of the VC model architecture is provided in Figure 1. The implementation details are the same as in [3].

### 6.3. Training Details

**ASR** For training  $A_0$  we first train an ASR on the LibriSpeech dataset, which we call as  $A_{base}$  and then fine-tune it on the low resource domain. The  $A_{base}$  model was trained for 35 epochs with a batch size having cumulative duration of 30 minutes. We use Adam [30] optimizer with a learning rate of 0.0025 with a warm-up scheduler of 40k iterations. For each iteration in the iterative training, we fine-tune the ASR model for a total of 100k iterations with a learning rate of 0.00025. During fine-tuning, we keep the distribution of the LibriSpeech dataset to that of the target dataset as 1:3. The training details remain the same across iterations.

**VC** We train the VC model with the same hyper-parameters as in [3], except for the weight for the  $L_{asr}$  which we set as 100 for all the experiments. We use the hierarchical prior grad vocoder, proposed in [31], to synthesize waveform from mel-spectrogram. All the experiments are conducted using 2 A-100 Nvidia GPUs having runtime of 24 hour.

### 6.4. Evaluations

For showcasing the effectiveness of our proposed framework, we consider  $A_0$  as a baseline model for the ASR evaluations and  $V_0$  as a baseline model for the VC evaluations. We evaluate our ASR models using the WER metric. For evaluating the VC models, we use both objective as well as subjective metrics.

**Objective metric for VC.** We use NISQA-TTS [32] for predicted mean opinion score (pMOS). For evaluating the speaker conversion of the VC models, we use the Identity metric, which is the cosine similarity of d-vectors extracted using Resemblyzer [33], a speaker verification model. For calculating how well the VC model preserves linguistic content, we calculate the WER using an ASR model which is different from the ASR model used for training the VC model. For the Hindi speech domain, we train an ASR model on the MUCS 2021 dataset which is publicly available. For English singing domain, we use an ASR model trained using iterative framework and having different distribution of speech and song data during training. For the objective and subjective test, we use 346 and 420 samples for the Hindi speech and English singing domain, respectively. The samples have a duration between 2 and 12 seconds.

**Subjective metric for VC.** We ask 17 audio engineers for the Hindi speech and 23 audio engineers for the English singing domain domain evaluation. The evaluators have no known hearing impairment. For the Hindi speech AB and SMOS test, the evaluators evaluate a total of 170 samples per model. For the

Table 3: Singing ASR WER across iterations. Note that we stop training the ASR after the WER converges.

Iter	0	1	2	3
WER ↓	6.7 ± 0.07	6.0 ± 0.1	5.0 ± 0.1	<b>4.9 ± 0.1</b>

Table 4: Singing VC Results.

Iterations	WER ↓	pMOS ↑	Identity ↑
0	20.83 ± 0.22	3.11 ± 0.01	0.7701 ± 0.001
1	19.12 ± 0.21	3.11 ± 0.01	0.7794 ± 0.001
2	<b>17.55 ± 0.21</b>	<b>3.38 ± 0.01</b>	0.7807 ± 0.001
3	17.56 ± 0.20	3.36 ± 0.01	<b>0.7808 ± 0.001</b>

English singing AB and SMOS test, the evaluators evaluate a total of 230 utterances per model.

## 7. Results

We present the results of the objective and subjective evaluation along with their 95% confidence interval.

### 7.1. Objective Evaluation

#### 7.1.1. Hindi speech voice dataset

**ASR** The WER of the Hindi ASR model is reported in Table 5. The Hindi ASR model achieves the best WER of 25.8 for iteration 3 with a relative improvement over the ASR model trained without iterative training of 6.85%.

**VC** Objective test results for Hindi speech VC model is reported in Table ?? which shows a decrease in WER and increase in pMOS across iterations.

#### 7.1.2. English singing voice dataset

**ASR** We evaluate our proposed framework on the English singing voice domain and report the WER in Table 3. Similar to the Hindi speech domain we observe that after 3 iterations the WER of the ASR model converges. We validate our proposed framework and show that the WER converges to 4.9 after 3 iterations with a relative improvement over the ASR model trained without our proposed framework of 26.8%.

**VC** For the VC evaluation of our proposed framework, we report the results in Table 2. We validate our proposed iterative training framework and show a relative decrease in WER of 12.63% and relative increase in the pMOS of 8.83%. Across iterations, the Identity metric remains nearly unchanged.

### 7.2. Subjective Evaluation

For the Hindi speech domain, we compare the  $V_3$  model with the baseline  $V_0$  model based on the results of the objective evaluation. Similarly, for the English singing domain, we compare the  $V_2$  model with the baseline  $V_0$  model. For subjective evaluation we consider AB quality test and Similarity MOS (SMOS). The AB quality test is constructed to compare the quality of the samples generated using the proposed model with the samples generated using the baseline model. During evaluation, we present one audio sample each from the baseline and proposed model to the evaluator and ask them to rate their preference in terms of quality. The two audio samples are generated using the same source and reference samples. The SMOS test is constructed to evaluate the speaker similarity of the con-

Table 5: Subjective Test Results for SMOS.

Model	ASR Domain	SMOS ↑
Oracle Hindi	-	4.68 ± 0.14
Hindi	Hindi Iteration 0	2.93 ± 0.16
Hindi	Hindi Iteration 3	<b>3.01 ± 0.18</b>
Oracle Singing	-	4.70 ± 0.09
Singing	Singing Iteration 0	3.06 ± 0.13
Singing	Singing Iteration 2	<b>3.17 ± 0.14</b>




Figure 3: Subjective AB Test results for Hindi Speech Iteration 3 vs Iteration 0 and English Singing Iteration 2 vs Iteration 0

verted sample to that of the reference sample. The evaluators are presented the reference sample and the voice converted sample and asked to rate the similarity of the voices between the samples on a scale of 1-5. We also evaluate the SMOS for oracle samples in which the reference and converted are real samples from the same speaker. The results of the subjective test for the VC models are summarized in Table ???. As we can observe, the proposed iterative training outperforms the baseline models on AB quality test with a preference rate of 0.66 to 0.34 on Hindi speech VC and 0.61 to 0.39 on English singing VC while achieving better SMOS.

### 7.3. Observations

By using the improved ASR for training the VC models we expected a decrease in the WER of the samples after conversion which is validated in Table 2 and ??. However, we also observe that the quality of the converted samples improve via the proposed framework due to decrease in artefacts and clearer pronunciation of words. This further demonstrates that with decrease in WER of the ASR model, the encoder features better capture the linguistic content which is essential for improving the quality of VC models. Because of the improvement in the quality of the converted samples because of less artefacts and clearer pronunciation, we expected the VC models to act as better data augmentation tools for training the ASR model which is validated from the results in 3 and 5.

## 8. Conclusions

We present a novel iterative framework for improving voice conversion models and automatic speech recognition models on low resource domains and verify its applications on the Hindi speech domain and English singing domain. We show improved speech preservation and MOS quality of the converted samples on voice conversion tasks as well as improved the word error rate on ASR tasks using this framework. Future work includes further improving the content preservation of the one-shot VC models so as to bring WER of the VC converted samples closer to the WER on the ground truth samples which would also lead to better MOS quality. We would also like to investigate combining the ASR and VC training in an end-to-end system.

## 9. References

- [1] A. Polyak, L. Wolf, Y. Adi, and Y. Taigman, “Unsupervised Cross-Domain Singing Voice Conversion,” in *Proc. Interspeech 2020*, 2020, pp. 801–805.
- [2] Y. A. Li, A. Zare, and N. Mesgarani, “StarGANv2-VC: A Diverse, Unsupervised, Non-Parallel Framework for Natural-Sounding Voice Conversion,” in *Proc. Interspeech 2021*, 2021, pp. 1349–1353.
- [3] N. Takahashi, M. Singh, and Y. Mitsufuji, “Robust one-shot singing voice conversion,” in *arXiv*, 2022. [Online]. Available: <https://arxiv.org/abs/2210.11096>
- [4] G. Wang, A. Rosenberg, B. Ramabhadran, F. Biadsy, J. Emond, Y. Huang, and P. J. Moreno, “Non-Parallel Voice Conversion for ASR Augmentation,” in *Proc. Interspeech 2022*, 2022, pp. 3408–3412.
- [5] K. Qian, Y. Zhang, S. Chang, X. Yang, and M. Hasegawa-Johnson, “AutoVC: Zero-shot voice style transfer with only autoencoder loss,” in *Proceedings of the 36th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, K. Chaudhuri and R. Salakhutdinov, Eds., vol. 97. PMLR, 09–15 Jun 2019, pp. 5210–5219.
- [6] J. chieh Chou and H.-Y. Lee, “One-Shot Voice Conversion by Separating Speaker and Content Representations with Instance Normalization,” in *Proc. Interspeech 2019*, 2019, pp. 664–668.
- [7] Y.-H. Chen, D.-Y. Wu, T.-H. Wu, and H.-y. Lee, “Again-vc: A one-shot voice conversion using activation guidance and adaptive instance normalization,” in *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 5954–5958.
- [8] D.-Y. Wu, Y.-H. Chen, and H. yi Lee, “VQVC+: One-Shot Voice Conversion by Vector Quantization and U-Net Architecture,” in *Proc. Interspeech 2020*, 2020, pp. 4691–4695.
- [9] N. Takahashi, M. K. Singh, and Y. Mitsufuji, “Hierarchical disentangled representation learning for singing voice conversion,” in *2021 International Joint Conference on Neural Networks (IJCNN)*, 2021, pp. 1–7.
- [10] C. Wang, Z. Li, B. Tang, X. Yin, Y. Wan, Y. Yu, and Z. Ma, “Towards high-fidelity singing voice conversion with acoustic reference and contrastive predictive coding,” in *Proc. Interspeech 2022*, 2022, pp. 4287–4291.
- [11] N. Shah, M. Kumar, Singh, N. Takahashi, and N. Onoe, “Nonparallel emotional voice conversion for unseen speaker-emotion pairs using dual domain adversarial network & virtual domain pairing,” in *ICASSP 2023*, 2023.
- [12] L. Dong, S. Xu, and B. Xu, “Speech-transformer: A no-recurrence sequence-to-sequence model for speech recognition,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 5884–5888.
- [13] G. Synnaeve, Q. Xu, J. Kahn, T. Likhomanenko, E. Grave, V. Pratap, A. Sriram, V. Liptchinsky, and R. Collobert, “End-to-end ASR: from supervised to semi-supervised learning with modern architectures,” in *ICML 2020 Workshop on Self-supervision in Audio and Speech*, 2020. [Online]. Available: <https://openreview.net/forum?id=OSVxDDc360z>
- [14] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu, and R. Pang, “Conformer: Convolution-augmented Transformer for Speech Recognition,” in *Proc. Interspeech 2020*, 2020, pp. 5036–5040.
- [15] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 12 449–12 460.
- [16] D. Galvez, G. Diamos, J. Torres, K. Achorn, J. Cerón, A. Gopi, D. Kanter, M. Lam, M. Mazumder, and V. Janapa Reddi, “The people’s speech: A large-scale diverse english speech recognition dataset for commercial usage,” in *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, J. Vanschoren and S. Yeung, Eds., vol. 1, 2021.
- [17] A. Baevski, W.-N. Hsu, A. Conneau, and M. Auli, “Unsupervised speech recognition,” in *Advances in Neural Information Processing Systems*, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, Eds., 2021.
- [18] N. Takahashi, M. K. Singh, S. Basak, P. Sudarsanam, S. Ganapathy, and Y. Mitsufuji, “Improving voice separation by incorporating end-to-end speech recognition,” in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 41–45.
- [19] J. Ni, L. Wang, H. Gao, K. Qian, Y. Zhang, S. Chang, and M. Hasegawa-Johnson, “Unsupervised Text-to-Speech Synthesis by Unsupervised Automatic Speech Recognition,” in *Proc. Interspeech 2022*, 2022, pp. 461–465.
- [20] S. Basak, S. Agarwal, S. Ganapathy, and N. Takahashi, “End-to-end lyrics recognition with voice to singing style transfer,” in *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 266–270.
- [21] T. Ko, V. Peddinti, D. Povey, and S. Khudanpur, “Audio augmentation for speech recognition,” in *Proc. Interspeech 2015*, 2015, pp. 3586–3589.
- [22] M. Baas and H. Kamper, “Voice Conversion Can Improve ASR in Very Low-Resource Settings,” in *Proc. Interspeech 2022*, 2022, pp. 3513–3517.
- [23] In-rolls, “Indicate: Hindi transliteration,” <https://github.com/in-rolls/indicate>, 2023.
- [24] A. Diwan, R. Vaideeswaran, S. Shah, A. Singh, S. Raghavan, S. Khare, V. Unni, S. Vyas, A. Rajpuria, C. Yarra, A. Mittal, P. Ghosh, P. Jyothi, K. Bali, V. Seshadri, S. Sitaram, S. Bharadwaj, J. Nanavati, R. Nanavati, and K. Sankaranarayanan, “Mucs 2021: Multilingual and code-switching asr challenges for low resource indian languages,” in *InterSpeech2021*, 08 2021, pp. 2446–2450.
- [25] Z. Duan, H. Fang, B. Li, K. C. Sim, and Y. Wang, “The nus sung and spoken lyrics corpus: A quantitative comparison of singing and speech,” in *2013 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference*, 2013, pp. 1–9.
- [26] B. Sharma, X. Gao, K. Vijayan, X. Tian, and H. Li, “NHSS: A speech and singing parallel database,” *CoRR*, vol. abs/2012.00337, 2020. [Online]. Available: <https://arxiv.org/abs/2012.00337>
- [27] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: An asr corpus based on public domain audio books,” in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015, pp. 5206–5210.
- [28] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishitoba, Y. Unno, N. E. Y. Soplin, J. Heymann, M. Wiesner, N. Chen, A. Renduchintala, and T. Ochiai, “Espnet: End-to-end speech processing toolkit,” *CoRR*, vol. abs/1804.00015, 2018. [Online]. Available: <http://arxiv.org/abs/1804.00015>
- [29] S. Watanabe, T. Hori, S. Kim, J. R. Hershey, and T. Hayashi, “Hybrid ctc/attention architecture for end-to-end speech recognition,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, no. 8, pp. 1240–1253, 2017.
- [30] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, Y. Bengio and Y. LeCun, Eds., 2015. [Online]. Available: <http://arxiv.org/abs/1412.6980>
- [31] N. Takahashi, M. Kumar, Singh, and Y. Mitsufuji, “Hierarchical diffusion models for singing voice neural vocoder,” in *ICASSP 2023*, 2023.
- [32] G. Mittag and S. Möller, “Deep Learning Based Assessment of Synthetic Speech Naturalness,” in *Proc. Interspeech 2020*, 2020, pp. 1748–1752.
- [33] R. AI, “Resemblyzer,” <https://github.com/resemble-ai/Resemblyzer>, 2023.