

SCRAPS: Speech Contrastive Representations of Acoustic and Phonetic Spaces

Ivan Valles-Perez^{1,*}, Grzegorz Beringer¹, Piotr Bilinski^{1,2}, Gary Cook² and Roberto Barra-Chicote²

¹ Alexa AI, Amazon

² University of Warsaw

Abstract. Numerous examples in the literature proved that deep learning models have the ability to work well with multimodal data. Recently, *CLIP* has enabled deep learning systems to learn shared latent spaces between images and text descriptions, with outstanding zero- or few-shot results in downstream tasks. In this paper we explore the same idea proposed by *CLIP* but applied to the speech domain, where the phonetic and acoustic spaces usually coexist. We train a *CLIP*-based model with the aim to learn shared representations of phonetic and acoustic spaces. The results show that the proposed model is sensible to phonetic changes, with a 91% of score drops when replacing 20% of the phonemes at random, while providing substantial robustness against different kinds of noise, with a 10% performance drop when mixing the audio with 75% of *Gaussian* noise. We also provide empirical evidence showing that the resulting embeddings are useful for a variety of downstream applications, such as intelligibility evaluation and the ability to leverage rich pre-trained phonetic embeddings in speech generation task. Finally, we discuss potential applications with interesting implications for the speech generation and recognition fields.

1 Introduction

Recent advances in deep learning have catapulted computer vision, improving the domain across multiple axes [5]. One of the most important breakthroughs is *CLIP* (Contrastive Language-Image Pre-training) [24], a deep learning model trained on vast amounts of multimodal data. In this case, by means of contrastive learning [19], *CLIP* showed the ability to learn to match images with written descriptions, and vice-versa. The representation power of this model at zero-shot regimes is such that it is able to surpass supervised models specifically trained for a determined task. The embeddings generated with *CLIP* have also been proved useful for other downstream tasks involving image and/or descriptions, in form of pretrained embeddings with rich latent representations [6, 27, 16, 31]. In these settings, reusing *CLIP* embeddings becomes a cornerstone for problems with small datasets. The extraordinary success of *CLIP* and the simplicity of its approach motivates exploring other applications to other fields where multimodal data is essential.

The field of speech technologies has come a long way in recent years, with deep learning models achieving impressive performance on a wide range of tasks [32, 28, 10]. However, there are still many challenges and opportunities in this domain [11], particularly when

it comes to exploiting large amounts of data. On the speech generation side, one of the main difficulties is to build a model that correctly aligns the phonetic and acoustic sequences, leading to a natural prosody with fluent speech and high intelligibility, while still capturing the prosody variations [25]. On the opposite side, automatic speech recognition systems struggle with long-tail words recognition [35], and speech vs background disentanglement [18].

This work explores the use of *CLIP* models for learning shared phonetic-acoustic latent spaces that can be later exploited for further downstream tasks. The methodology described in this study, which we name *SCRAPS* (Speech Contrastive Representation of Acoustic and Phonetic Spaces) offers a promising approach for learning such common representations. By learning phonetic-acoustic representations, we aim to enable the use of rich latent spaces in a variety of downstream tasks, such as cheap intelligibility evaluation, rich pre-trained embeddings, a differential intelligibility loss and on-the-fly training data filtering. This last example is specially relevant at the moment given the vast amounts of found data available online.

We believe that this approach has the potential to help advancing the state of the art in speech generation and recognition towards learning from large-scale data, looking at the impact that *CLIP* model has had in the computer vision domain. In the following sections, we will provide a detailed description of our proposed methodology, as well as experiments and results that backup our hypotheses, demonstrating its effectiveness. We also discuss potential downstream applications and use cases where the *SCRAPS* model may be useful, together with some initial experiments, and suggest future research lines.

2 Related work

There are various examples in the literature with notable similarity to the work presented in this paper, one of the most relevant being *CLAP* [8]. *CLAP* stands for Contrastive Language-Audio Pretraining, and it is a *CLIP* based model that is trained on text description-audio pairs with the aim of producing a shared multimodal representation of acoustic and semantic domains. Although arguably similar to this work, *CLAP* uses the text from a semantic perspective, as the dataset used to train this model contents pairs of audio samples with corresponding textual descriptions. *SCRAPS* is trained over pairs of audio clips containing speech utterances and their corresponding transcribed phonetic sequences. The objective of *SCRAPS* is to learn a shared latent space between the phonetics and its acoustic realizations, hence the nature of the problem is orthogonal to *CLAP*.

* Corresponding Author Email: ivallesp@amazon.com.

Work conducted when all authors were at Alexa AI, Amazon.

Another example of *CLIP*-based approach that operates in the speech domain is *SpeechCLIP* [29]. In that work, the authors build a system that learns shared representations of images with spoken captions. By means of *CLIP*, the authors establish a third alignment with text descriptions. In *M-SpeechCLIP* [2], the authors study the usefulness of English pre-trained *SpeechCLIP* to other languages and also propose multilingual versions of it. *Wav2CLIP* [36] is another work with the same idea as *SpeechCLIP*, with the difference that the authors use the *CLIP* image encoder to distill an audio encoder that produces latent representations compatible with the images. *Audio-CLIP* [13] is another example where the original *CLIP* model gets extended to the audio domain. The authors first pre-train the audio encoder using the frozen pre-trained *CLIP* encoders, and later unfreeze all the modules to fine-tune the whole system together.

These approaches differ from ours in the following points:

1. The approaches described above operate with text descriptions and at word level to infer semantics (as they rely on *CLIP* text encoder), while *SCRAPS* operates in the phonemes space, which brings the problem closer to the speech generation/recognition field, and makes it more practical for those use-cases.
2. The models found in the literature (aside from CLAP) rely on the pre-trained *CLIP* models (*SpeechCLIP* at inference time and *Wav2CLIP* and *Audio2CLIP* at training time), while our solution is an end-to-end model, and all the models are directly optimized to work with phonetic and acoustic spaces.
3. *SpeechCLIP* and *Wav2CLIP* do not directly optimize the speech-text shared space, but use images as a common intermediate space. In practice, this means that the error inherent in the original *CLIP* model (which operates on the image-text multimodality) becomes effectively irreducible, and it is added to the image-audio model error.

To the best of our knowledge, our work is the first attempt to use a *CLIP*-based strategy to learn joint phonetic and acoustic spaces. For this reason, we believe this work is of high interest for the speech generation/recognition community.

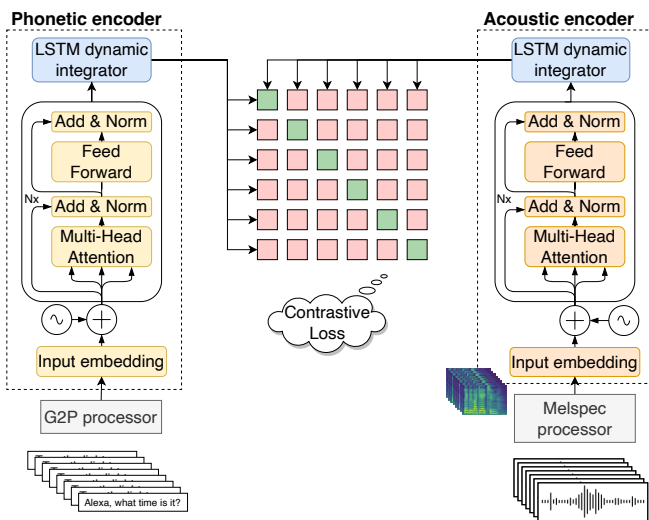


Figure 1. *SCRAPS* model architecture, with a phonetic encoder in the left hand side, and an acoustic encoder in the right hand side.

3 Methods

3.1 CLIP

CLIP is composed of two encoders, one for the images and one for the text. Each one receives the corresponding input (*i.e.* the image is fed to the visual encoder and its description to the text encoder) and generates a real valued vector, one derived from the image $\vec{u} \in \mathbb{R}^D$, another derived from the text $\vec{v} \in \mathbb{R}^D$, both with the same dimensionality D . The objective of the network is to make both vector spaces compatible, such that corresponding image-text pairs produce, ideally, the same vectors – *e.g.* (\vec{u}_1, \vec{v}_1) – and non-corresponding pairs produce dissimilar vectors (\vec{u}_1, \vec{v}_2) . This is achieved by maximizing the dot product between corresponding pairs of vectors, which matrix is represented as \mathbf{L} , with shape $N \times M$, *i.e.* N images times M descriptions, and minimizing it for non-corresponding pairs. In the original form, $M = N$ as the positive and negative pairs are constructed using all the pairwise combinations of the elements in the *minibatch*.

This is done by using a symmetric cross-entropy loss (cross-entropy across text and across images), as described in Equation 2, where \mathbf{L} represents the logits matrix resulting from the dot product of the outputs of the two encoders, and \mathbf{P}_1 and \mathbf{P}_2 are the two *softmax*-normalized matrices across the two possible axes (T represents the *softmax* temperature). The pairing of images and descriptions is done using the elements of the *minibatch*.

$$\mathbf{P}_1(i, j) = \frac{e^{\mathbf{L}(i, j)/T}}{\sum_{k=1}^N e^{\mathbf{L}(k, j)/T}}, \quad \mathbf{P}_2(i, j) = \frac{e^{\mathbf{L}(i, j)/T}}{\sum_{k=1}^M e^{\mathbf{L}(i, k)/T}}. \quad (1)$$

$$\mathcal{L} = \sum_{j=1}^M \sum_{k=1}^N \mathbf{I}(k, j) \log \mathbf{P}_1(k, j) + \sum_{i=1}^N \sum_{k=1}^M \mathbf{I}(i, k) \log \mathbf{P}_2(i, k). \quad (2)$$

3.2 SCRAPS

SCRAPS is an adaption of the original *CLIP* idea to the speech domain. Consequently, the description encoder included in the original *CLIP* architecture is replaced by a phonetic encoder, which receives a sequence of phonemes as input, and produces a vector as output. The image encoder is replaced by an acoustic encoder, which receives a mel-spectrogram as input and produces a vector as output.

The architecture used for this study is represented in Figure 1. As it can be seen in the figure, the model is built of two blocks: a phonetic encoder (left hand side) and an acoustic decoder (right hand side). Both modules take their corresponding input and produce a D dimensional vector (time independent): \mathbb{R}^D . As shown in the diagram, the blocks contain two modules: a transformer backbone (inspired on transformer TTS [20, 33]), and an *LSTM* [14] on top. The role of the *LSTM* is to integrate the representations coming from the transformers whose inputs have dynamic lengths. It produces a forward-rolling integration of the outputs vectors from the transformers. The last state of the *LSTM* is used as the final *SCRAPS* vector. The *LSTMs* of the two encoders share weights to propagate the latent spaces compatibility back to the output of the transformer outputs (see section 4.7 as an empirical proof of this assumption). That helps the time-dependent vectors live in the same space for both encoders, while still maintaining a time-dependent latent representation which may be used for future downstream tasks where the either of the *SCRAPS* encoders is used as a pre-trained module.

As it is shown in the middle of Figure 1, the *SCRAPS* model is trained with a contrastive loss where the objective is to maximize

the scores of the matching pairs, while minimizing the scores of the non-matching pairs. This is achieved by computing the dot product of $\mathbf{T} \in \mathbb{R}^{B \times D}$, i.e. *SCRAPS* vectors of the B phonetic sequences of the *minibatch*, and $\mathbf{A} \in \mathbb{R}^{B \times D}$, i.e. *SCRAPS* vectors of the B mel-spectrograms in the *minibatch*: $\mathbf{L} = \mathbf{T}\mathbf{A}^T$; $\mathbf{L} \in \mathbb{R}^{N \times N}$. \mathbf{L} is referred as the *SCRAPS* scores matrix, hereafter. *SCRAPS* is trained with the same loss as in the original *CLIP* model (Equation 2).

4 Experiments

4.1 Setup

We have trained a *SCRAPS* model on a large proprietary dataset containing 60,000 hours of de-identified speech data in US English. The data is composed of recordings of untrained speakers with variety of background noises, unnormalized pauses, and in some cases, even some samples with concurrent speech where one of the speakers dominates over the others. Each recording is accompanied with its corresponding transcription. The phonetic sequences are derived from the text transcriptions by using a simplified grapheme to phoneme processor (G2P) based on a dictionary of pronunciations. The *X-Sampa* phoneme set [34] has been used along all the experiments. 80-mel spectrograms have been extracted from the recordings (using 50ms window length, 12.5ms skips and 1024-point FFTs), and then they have been standardized. We held out a set composed of 5000 utterances, for development purposes. Along the following experiments, we describe various test sets. These test sets are always disjoint from the training set described above. We have trained that model for 2.5M training steps using *Adam* optimizer [17] with default hyper-parameters, a learning rate of $5 \cdot 10^{-4}$, a *minibatch* size of 128 and using a machine with 8 *Nvidia V100 16GB Tensor Core GPUs*.

Both the phonetic and the acoustic encoder transformer backbones are built with 3 layers of 8 attention heads each and a dropout rate of 0.1. The encoders operate with hidden and output spaces of size \mathbb{R}^{256} . The *LSTM* integrator head, which is shared between both encoders, is a single layer *LSTM* with 1024 units, which corresponds to the size of the output embeddings (similar to *CLIP*). The last output state of the *LSTM* is used as the final *SCRAPS* vector. We used a temperature of $T = 1.0$ during all the training and concluded that our training process is stable with that setting.

Contrastive models are difficult to evaluate in an objective manner [26]. In our case, just taking an unseen set of utterances, computing the *SCRAPS* vectors from a random sample of the spectrograms and phonetic sequences and measuring how well the model discriminates between matching and non-matching pairs is not informative, due to the fact that the examples in the *minibatches* are very dissimilar among them (i.e. when the samples are chosen at random to build the *minibatch* it may be too easy to discriminate between positive and negative pairs). We have conducted this exercise with a held-out test set of around 6000 samples, and achieved an Equal-Error-Rate (EER) in the order of 10^{-4} and Area Under the Receiver Operating Curve (*AUC-ROC*) in the order of $1 - 10^{-7}$. Handcrafting a test set by enrolling challenging negative samples is a possible alternative [4], but it relies on a criterion of difficulty that is not easy to define and that is prone to subjective biases. This difficulty motivates the exploration of more creative methods for understanding the performance of our model.

We propose a sensitivity and robustness analysis to study how does the model react against perturbation of the input data (referred hereafter as corruption). In this analysis, we corrupt either the spectrogram or the phonetic sequence and measure how the results vary due

to that perturbation. We use the following methods to corrupt spectrograms or phonetic sequences:

- **Random phonemes substitution:** substitute a portion of phonemes in the phonetic sequence at random.
- **Spectrogram *Gaussian* noise:** add different amounts of *Gaussian* noise to spectrogram. The noise is added according to Equation 3, where \mathbf{M} is the spectrogram, \mathbf{N} is the noise and $\alpha \in [0, 1]$ is the varying parameter that controls the amount of noise.
- **Spectrogram mix:** analogous to the *Gaussian* noise method, but replacing the noise with another random sampled spectrogram from the *minibatch*. Equation 3 is used to mix the original and the noise sources.

$$\hat{\mathbf{M}} = (1 - \alpha) \cdot \mathbf{M} + \alpha \cdot \mathbf{N}. \quad (3)$$

4.2 Sensitivity

We take a random sample of 1152 matching phonetic sequences and spectrograms from the test set, and calculate their *SCRAPS* score. We apply the corruption methods described above, one by one, and calculate their corresponding *SCRAPS* scores. Finally, for each of the tested corruption proportions we calculate the percentage of cases where the *SCRAPS* score of the corrupted examples *drops* or *lifts* by any amount, with respect to the original score (without corruption). High sensitivity to phonetic corruption is desirable to spot if utterances and transcriptions match. Hence, a good model would have an early raise of % of drops, and none/very few lifts. However, for acoustic corruption, high sensitivity is bad. Corrupting the acoustics to change the content is not an easy task. We instead noised the utterances with *Gaussian* noise, or by mixing in another utterance. With those methods a good model would have low sensitivity to acoustic corruption (i.e., ability to match phonetic and acoustic content in adverse conditions).

The results are represented visually in Figure 2-top and 2-middle. Table 1 shows the percentage of *drops* and *lifts* in *SCRAPS* scores after corruption for a set of specific corruption amounts.

In Table 1 we observe that in the 91.06% of the examples where at least 20% of the phonemes are corrupted, the *SCRAPS* model reacts by decreasing the score of the match, while in 2.62% of the cases the model score increases.

Finally, we also study the effect that the number of phonemes in the sequences has in the *SCRAPS* score (see figure 3), observing that the distribution of *SCRAPS* score remains constant when the number of phonemes is greater than 20.

4.3 Robustness

Another important feature to test on the *SCRAPS* model is how robust it is against perturbations. In the worst case scenario, this model would only work on conditions that are very similar to those of the training data, and completely fail if the data characteristics change slightly, which is undesirable. To measure the robustness, we have calculated the *AUC-ROC* score over the *SCRAPS* probabilities of each *minibatch*, considering the matching pairs as positive examples and the non-matching pairs as negative examples. We have calculated that score after corrupting the inputs in different quantities. These measurements have been performed over 1152 random sampled examples from the test set.

Table 1. Robustness (*AUC-ROC*) and sensitivity (*drop* and *lift*) results of *SCRAPS* at different levels of corruption. The second row of the table indicates the method used to corrupt the inputs in each case. All the results are represented as the mean \pm the 95% confidence interval.

Space	Acoustic						Phonetic		
Method	Mix spectrograms			Gaussian noise			Random substitution		
Metric	<i>AUC-ROC</i>	<i>Drop</i> (%)	<i>Lift</i> (%)	<i>AUC-ROC</i>	<i>Drop</i> (%)	<i>Lift</i> (%)	<i>AUC-ROC</i>	<i>Drop</i> (%)	<i>Lift</i> (%)
Corruption (%)									
5.0	1.00 \pm 0.00	46.01 \pm 2.88	53.99 \pm 2.88	1.00 \pm 0.00	43.49 \pm 2.86	56.51 \pm 2.86	1.00 \pm 0.00	52.78 \pm 2.88	8.94 \pm 1.65
10.0	1.00 \pm 0.00	49.31 \pm 2.89	50.69 \pm 2.89	1.00 \pm 0.00	45.14 \pm 2.87	54.86 \pm 2.87	1.00 \pm 0.00	76.91 \pm 2.43	7.20 \pm 1.49
20.0	1.00 \pm 0.00	53.21 \pm 2.88	46.79 \pm 2.88	1.00 \pm 0.00	47.83 \pm 2.88	52.17 \pm 2.88	1.00 \pm 0.00	91.06 \pm 1.65	2.60 \pm 0.92
40.0	1.00 \pm 0.00	79.43 \pm 2.33	20.57 \pm 2.33	1.00 \pm 0.00	63.54 \pm 2.78	36.46 \pm 2.78	0.97 \pm 0.01	96.44 \pm 1.07	1.22 \pm 0.63
60.0	0.76 \pm 0.02	95.23 \pm 1.23	4.77 \pm 1.23	1.00 \pm 0.00	92.10 \pm 1.56	7.90 \pm 1.56	0.85 \pm 0.02	98.44 \pm 0.72	0.87 \pm 0.54
80.0	0.54 \pm 0.03	99.57 \pm 0.38	0.43 \pm 0.38	0.80 \pm 0.02	99.22 \pm 0.51	0.78 \pm 0.51	0.71 \pm 0.03	99.48 \pm 0.42	0.35 \pm 0.34
90.0	0.52 \pm 0.03	99.65 \pm 0.34	0.35 \pm 0.34	0.65 \pm 0.03	99.13 \pm 0.54	0.87 \pm 0.54	0.62 \pm 0.03	99.57 \pm 0.38	0.35 \pm 0.34
95.0	0.51 \pm 0.03	99.65 \pm 0.34	0.35 \pm 0.34	0.57 \pm 0.03	99.39 \pm 0.45	0.61 \pm 0.45	0.59 \pm 0.03	99.83 \pm 0.24	0.17 \pm 0.24




Figure 2. Sensitivity and robustness results. Top and middle charts represent the percentage of *drops* (top) and *lifts* (middle) in *SCRAPS* score as a function of different levels of corruption in the input phonemes score. Bottom chart shows the robustness evaluation results using *Gaussian* noise and spectrogram mix methods. Dashed lines represent the 95% confidence intervals.

The results of this experiment for the two types of noise are represented in Figure 2-bottom where the amount of noise has been represented in the x axis and the *AUC-ROC* in the y axis. Table 1 provides the *AUC-ROC* scores achieved for a set of manually selected noise amounts (α in Equation 3). As we can observe, the *SCRAPS* model is robust against different levels of noise depending on its nature. For *Gaussian* noise, the performance degrades a 10% or more when adding noise with $\alpha \geq 0.75$. However, for spectrogram mixing, the performance drops at least a 10% when $\alpha \geq 0.52$.




Figure 3. Distribution of *SCRAPS* scores with respect to the lengths of the phonetic sequences.

An interesting observation arises when looking at Figure 2-bottom: the *Gaussian* noise curve (green) never reaches the random performance level (even for $\alpha = 1.0$, point in which the spectrograms have no information at all, the *AUC-ROC* is significantly greater than 0.5). We hypothesize that this is due to the fact that despite not providing explicit information to the acoustic transformer, it extracts information from the positional encoding to later determine if the length of the acoustic input is feasible for the phonetic sequence provided.

4.4 Covariate shift

Looking at the results of subsection 4.2, one could think that the observed effects may be due to covariate shift induced by the corruption methods described. Four additional experiments have been conducted. These experiments involved more sophisticated corruption techniques such as word swaps, cut-offs, accent switching, and denoising. In all cases, the *SCRAPS* scores were affected by the corruptions, hence the likelihood of covariate shift being the cause of the initial drop in scores was determined to be minimal. Appendix A, which we included in the preprint version of this manuscript¹ due to length constraints, describes the findings of these experiments more in detail.

4.5 Open source evaluation baselines

With the aim of facilitating future benchmarking and reproducibility, we have included two extra evaluations using the *Common Voice* test set. In the first experiment we measure the correlation between *SCRAPS* scores and *Whisper* base perplexity using the invalidated

¹ <https://arxiv.org/abs/2307.12445>

subset of the *Common Voice* dataset, observing a Pearson correlation of -0.60. This indicates a significant level of agreement between *SCRAPS* and *Whisper*, which establishes a first reproducible benchmark.

The second experiment replicates the sensitivity analysis against natural noise (denoising) described previously (Appendix A). This time, the open source *Common Voice* test set is used, which contains approximately 16K utterances. After denoising, the *SCRAPS* scores increase by 2.94 points, with 72.79% of the utterances showing a score increase, confirming that *SCRAPS* scores significantly increase after denoising the utterances.

Due to length restrictions, these results have been included in Appendix B of the preprint¹.

4.6 Examples

Aside from quantitative results, qualitative ones are often informative. This section provides a set of cherry-picked examples where we use *SCRAPS* to retrieve the spectrograms with highest score in our test set given an input transcription. For that, we score a random sample of the test set containing 7124 pairs (spectrogram, phonetic sequence) and calculate the pairwise dot product between acoustic and phonetic vectors. The resulting matrix has shape 7124×7124 . Then we normalize that matrix and do a lookup to find the closest spectrograms given a sentence. In the paragraphs below, the first sentence highlighted in bold represents the query sentence, and the ones below are the transcriptions corresponding with the top 3 spectrograms with highest probability. The numerical probability value, rounded to the 2nd decimal place, is indicated between parentheses before the sentences. In the cases where one of the top 3 closest spectrograms corresponds to the actual matching spectrogram for the reference sentence, we have highlighted its transcription with a star (*).

Astro dance.

*(0.97) Astro dance.
(0.03) Astro, dance.
(0.00) Astro dance for me.

Alexa, end this call on my phone.

*(1.00) Alexa, end this call on my phone.
(0.00) Alexa, unpair my fire phone.
(0.00) Alexa, end this meeting for all.

Echo, keep reading.

*(0.74) Echo, keep reading.
(0.25) Hey Alexa keep reading.
(0.01) Echo, keep going.

Alexa, disconnect my fire phone.

*(0.88) Alexa, disconnect my fire phone.
(0.04) Alexa disconnect the phone.
(0.03) Alexa disconnect my cell phone.

Alexa, do i have any appointments on June nineteenth.

*(0.98) Alexa, do i have any appointments on June nineteenth.
(0.01) Alexa do i have any appointments on February twenty fifth.
(0.00) Alexa, do i have any appointments at seven fifty five AM.

One of the most important observations by looking at the examples above is that in some cases the commas in the transcriptions (corresponding to pauses in the associated spectrograms, given that the commas were annotated through forced alignment [22]) have a notable effect in the *SCRAPS* probability. We have analysed that effect in more detail and observed that when removing the commas from a pair that originally had a comma, the *SCRAPS* score drops a $5.49\% \pm 0.38\%$ with respect to the original value (as per the 95% confidence interval over the mean, calculated with 768 samples). Through an informal listening test, we have confirmed that the

commas usually correspond to long pauses in the spectrogram.




Figure 4. Word Error Rate (WER) of the TTS models, relative to the WER over the test set recordings. The two lines represent how the intelligibility of the model improves as the model trains for the architectures with pretrained *SCRAPS* and with *E2E* phonetic encoder. The WER has been measured with a pretrained ASR model over a set of 512 sentences unseen during training, synthesized every 25,000 training steps. The WER represented in the figure is relative to the recordings WER.

4.7 Downstream application 1: pretrained phonetic embeddings for speech generation

For this application, we have trained an autoregressive multispeaker text-to-speech model with attention, similar to Transformer TTS. We used a proprietary dataset with 180,000 hours of de-identified en-US recordings. Around 2/3 of the dataset is automatically annotated by an ASR model, the rest of utterances have been manually annotated. The evaluation set, used to generate the figure 4, is composed of 10,000 sentences, randomly held out from the training set.

The baseline architecture consists of three modules: a phonetic encoder, an acoustic decoder and a speaker embedding. The acoustic decoder is autoregressive and attends to the output of the phonetic encoder and the speaker embedding. Then we have substituted the phonetic encoder by a *SCRAPS* pretrained phonetic encoder (only the transformer backbone, not the *LSTM* integrator). The test metrics evolution during training are represented in Figure 4. In the figure, we observe that both architectures get a very similar final performance, but when using *SCRAPS*, the model converges much faster (the initial error drop happens 250,000 steps earlier in the case of the model that uses the pretrained *SCRAPS*). Apart from the clear faster convergence, this observation allows us concluding that the *SCRAPS* transformer latent space is rich enough to be used as a pretrained module to speed up Text-To-Speech (TTS) model training.

4.8 Downstream application 2: text-less intelligibility evaluation for voice conversion systems

Intelligibility performance is a key metric for Voice Conversion (VC) systems [21]. Usually, WER is used to assess such performance. However, it is relatively costly: apart from ground-truth texts, one needs to transcribe VC audio using humans or ASR. In case of using an ASR system, it needs to be trained on the target language with sufficient amount of data in order to be robust enough to bypass the need for human annotations. For these reasons, there is a need for cheaper alternatives that, if possible, do not require any text.

In this application we study the use of *SCRAPS* as a text-less intelligibility metric. For that, we used a VC model trained over a proprietary dataset with 60k hours of de-identified en-US recordings, manually annotated by humans. Each of these recordings are paired with

a synthetically generated audio of a target speaker, to train the many-to-one voice conversion model. These syntheses have been generated using a similar technique of the one described in this paper [9]. As mentioned in the section 4.6 of the manuscript, the evaluation set is composed of 3500 utterances, which were held out from the training set. Although *SCRAPS* is trained to match a sequence of phonemes to the corresponding audio, at inference time it can also be used to compute correspondence between two audio files without requiring any text. In this scenario, the *SCRAPS* score is computed between vectors of synthetic audio (VC) and source audio (pre-conversion). We refer to this metric as *SCRAPS(VC, source)*.

We evaluate usability of *SCRAPS(VC, source)* as an intelligibility metric by testing how it relates to WER computed over human annotations (h-WER). We run an en-US any-to-one VC system (similar to [7]) over a test set of 3532 recordings. Then, we compute h-WER by gathering human transcriptions on VC audio and comparing it to ground-truths. Finally, *SCRAPS(VC, source)* is computed using only VC and source (pre-conversion) audios. We compute the *Spearman* correlation between *SCRAPS(VC, audio)* and h-WER (Table 2). For comparison, two alternative metrics are computed: (1) ASR-WER, where we use a proprietary ASR system to replace human transcriptions, (2) *SCRAPS(VC, text)* that uses ground-truth text instead of input audio. The proposed metric obtains a negative correlation of -0.1938 , which is significantly lower (in absolute values) than *SCRAPS(VC, text)* (-0.3232) and ASR (0.6197). Since h-WER has a bounded distribution at 0, with a large proportion of samples (69%) having exactly h-WER=0, while *SCRAPS* scores are unbounded, we also compute correlation using only samples with h-WER>0 (third column). As can be observed, the differences in correlation between metrics are then smaller, with no perceivable difference between *SCRAPS(VC, source)* and *SCRAPS(VC, text)* scores. This suggests that *SCRAPS(VC, source)* more often incorrectly assigns low score to high-intelligibility (h-WER=0) samples than *SCRAPS(VC, text)*, which could originate from *SCRAPS* not being trained with pairs of audio and being less robust to slight differences between them, even if content in both is the same. This experiment shows there is a statistically significant correlation between *SCRAPS* scores and h-WER, although using a robust ASR is still the strongest automatic solution, despite the higher computational cost.

We also test if *SCRAPS(VC, source)* can be used to find high-WER VC samples, which could serve as a post-filter to remove low-intelligibility VC outputs. The results can be seen in Figure 5, where four h-WER thresholds were used to binarize samples into "low" and "high" h-WER samples. While *SCRAPS(VC, source)* is not robust enough to classify h-WER>0 samples without supervision (AUC=0.580), it might be sufficient enough to act as a filter for removing the worst cases, *i.e.* h-WER>0.75 (AUC=0.830), thus reducing the computational cost of the filtering process.

Table 2. Spearman correlation between automatic intelligibility metrics and human WER (h-WER). Results on all utterances (second column) and utterances with h-WER>0 (third column) are shown.

Metric	Correlation (all)	Correlation (h-WER>0)
ASR-WER	0.6197	0.7064
<i>SCRAPS(VC, text)</i>	-0.3232	-0.4812
<i>SCRAPS(VC, source)</i>	-0.1938	-0.4740




Figure 5. ROC of h-WER>T classification using *SCRAPS(VC, source)* scores for different thresholds T.

4.9 Ablation study

In this subsection, we question some of the design decisions of the *SCRAPS* architecture and test the impact of ablating them in terms of sensitivity and robustness. The aspects included in this ablation study are the following ones:

- **Head parameters sharing:** this feature was included to try to bias the acoustic and phonetic transformer towards producing compatible representations. However, there is a risk that by doing that, the representation power of the encoders decreases.
- **LSTM integrator modules:** they may be superfluous. In such case, the output of the acoustic and phonetic encoders is the vector corresponding to the sequence length at the output of the transformer.

To do these experiments we trained two new models with the described ablated components, for the same number of training steps as the full model described before. The results, included in Table 3, show that in the majority of metrics there is no significant difference between the baseline and the ablated versions. More specifically, from the table we can conclude that sharing parameters of the *LSTM* integrators is safe, as it does not have a significant effect in the performance, compared with the baseline. However, the *SCRAPS* version without integrators seems to be significantly less robust to *Gaussian* noise than the rest, showing large differences in *AUC-ROC* score. This may imply that the transformer is not large enough to handle noisy inputs and that the *LSTM* module at the output of the transformer helps making the output more robust.

5 Discussion

The *SCRAPS* model has shown the ability to model a rich common latent space between phonetic and acoustic spaces, which is interesting for speech generation and recognition tasks. This task is not trivial, as it involves implicitly learning the underlying allophonic and phonetic joint probability distribution, which is the root difficulty of speech technologies, while learning to ignore non-related information such as prosody, background, non-verbal sounds and other stationary aspects such as speaker id or different accents.

We have proved that the *SCRAPS* model has high sensitivity to small changes in phonetic sequences, and at the same time it is robust against high levels of noise, independently of its nature (stationary

Table 3. Results of the ablation study, showing the robustness (*AUC-ROC*) and sensitivity (*drop* and *lift*) numbers of the ablated versions of *SCRAPS* at different levels of corruption. All the results are represented as the mean \pm the 95% confidence interval. The cases where the results differ significantly from the originally defined version of *SCRAPS* have been highlighted in bold. They have been colored in blue in case of representing a negative deviation over the baseline (ablated < baseline) or in red if the deviation is positive (ablated > baseline).

	Space	Acoustic						Phonetic		
	Method	Mix spectrograms			Gaussian noise			Random substitution		
Ablation	Metric Corruption (%)	<i>AUC-ROC</i>	<i>Drop (%)</i>	<i>Lift (%)</i>	<i>AUC-ROC</i>	<i>Drop (%)</i>	<i>Lift (%)</i>	<i>AUC-ROC</i>	<i>Drop (%)</i>	<i>Lift (%)</i>
Parameters sharing	5.0	1.00 \pm 0.00	49.57 \pm 2.89	50.43 \pm 2.89	1.00 \pm 0.00	45.57 \pm 2.88	54.43 \pm 2.88	1.00 \pm 0.00	52.00 \pm 2.89	9.72 \pm 1.71
	10.0	1.00 \pm 0.00	49.74 \pm 2.89	50.26 \pm 2.89	1.00 \pm 0.00	45.31 \pm 2.87	54.69 \pm 2.87	1.00 \pm 0.00	75.95 \pm 2.47	8.16 \pm 1.58
	20.0	1.00 \pm 0.00	56.77 \pm 2.86	43.23 \pm 2.86	1.00 \pm 0.00	47.57 \pm 2.88	52.43 \pm 2.88	1.00 \pm 0.00	90.36 \pm 1.70	3.30 \pm 1.03
	40.0	1.00 \pm 0.00	81.51 \pm 2.24	18.49 \pm 2.24	1.00 \pm 0.00	61.20 \pm 2.81	38.80 \pm 2.81	0.96 \pm 0.01	96.18 \pm 1.11	1.48 \pm 0.70
	60.0	0.77 \pm 0.02	96.53 \pm 1.06	3.47 \pm 1.06	1.00 \pm 0.00	90.89 \pm 1.66	9.11 \pm 1.66	0.85 \pm 0.02	98.52 \pm 0.70	0.78 \pm 0.51
	80.0	0.54 \pm 0.03	99.57 \pm 0.38	0.43 \pm 0.38	0.79 \pm 0.02	98.96 \pm 0.59	1.04 \pm 0.59	0.71 \pm 0.03	99.39 \pm 0.45	0.43 \pm 0.38
	90.0	0.52 \pm 0.03	99.74 \pm 0.29	0.26 \pm 0.29	0.62 \pm 0.03	99.39 \pm 0.45	0.61 \pm 0.45	0.63 \pm 0.03	99.65 \pm 0.34	0.26 \pm 0.29
	95.0	0.52 \pm 0.03	99.65 \pm 0.34	0.35 \pm 0.34	0.57 \pm 0.03	99.31 \pm 0.48	0.69 \pm 0.48	0.60 \pm 0.03	99.83 \pm 0.24	0.17 \pm 0.24
<i>LSTM</i>	5.0	1.00 \pm 0.00	48.00 \pm 2.89	52.00 \pm 2.89	1.00 \pm 0.00	51.30 \pm 2.89	48.70\pm2.89	1.00 \pm 0.00	51.56 \pm 2.89	10.16 \pm 1.74
	10.0	1.00 \pm 0.00	49.74 \pm 2.89	50.26 \pm 2.89	1.00 \pm 0.00	49.74 \pm 2.89	50.26 \pm 2.89	1.00 \pm 0.00	75.26 \pm 2.49	8.85 \pm 1.64
	20.0	1.00 \pm 0.00	54.77 \pm 2.87	45.23 \pm 2.87	1.00 \pm 0.00	59.29\pm2.84	40.71\pm2.84	1.00 \pm 0.00	90.62 \pm 1.68	3.04 \pm 0.99
	40.0	1.00 \pm 0.00	80.03 \pm 2.31	19.97 \pm 2.31	1.00 \pm 0.00	92.97\pm1.48	7.03\pm1.48	0.98 \pm 0.01	96.44 \pm 1.07	1.22 \pm 0.63
	60.0	0.76 \pm 0.02	98.09\pm0.79	1.91\pm0.79	0.77\pm0.02	99.22\pm0.51	0.78\pm0.51	0.90\pm0.02	98.52 \pm 0.70	0.78 \pm 0.51
	80.0	0.54 \pm 0.03	99.74 \pm 0.29	0.26 \pm 0.29	0.54\pm0.03	99.57 \pm 0.38	0.43 \pm 0.38	0.76 \pm 0.02	99.57 \pm 0.38	0.26 \pm 0.29
	90.0	0.53 \pm 0.03	99.74 \pm 0.29	0.26 \pm 0.29	0.53\pm0.03	99.48 \pm 0.42	0.52 \pm 0.42	0.68 \pm 0.03	99.57 \pm 0.38	0.35 \pm 0.34
	95.0	0.53 \pm 0.03	99.91 \pm 0.17	0.09 \pm 0.17	0.52\pm0.03	99.57 \pm 0.38	0.43 \pm 0.38	0.64 \pm 0.03	99.65 \pm 0.34	0.35 \pm 0.34

and non-stationary). One of the most surprising findings has been *SCRAPS* robustness in the face of concurrent speech, showing an *AUC-ROC* of 1.00 up to 60%-40% overlap of primary and secondary spectrograms, respectively (see Figure 2-bottom).

The *SCRAPS* model can be seen as an implicit speech generation-recognition system. Given a mel-spectrogram, one can search for the sentence that better matches that spectrogram, transforming the system in an automatic speech recognizer. On the other side, given a sentence, one could look for the spectrogram that better matches the input sentence, thus converting the system into a speech generator. Obviously the search processes makes these approaches very computationally expensive, but this interpretation of *SCRAPS* helps devise further applications.

Finally, on the practical side, *SCRAPS* brings its value when used as a tool for a downstream application. There are many potential use cases for this model, two of which we explored in sections 4.7 and 4.8. Further research lines could explore other applications. Below, we state some of the most interesting examples, but there may be many more.

- Transcription quality assurance: machine learning models are sensitive to the quality of the training data. In speech generation and recognition applications, the datasets often contain transcription problems at different levels: from mistaken words to poor annotation (e.g. not containing hesitations or background events). *SCRAPS* can be trained with high quality data and used to score other datasets to spot transcription issues.
- Grapheme to phoneme mapping enhancement: speech datasets are usually composed of pairs of audio clips and sentences. The sentences are mapped to the phonetic sequences using grapheme to phoneme (G2P) systems, which often rely on a word-level dictionary of pronunciations [30]. Although this approach usually works in practice, it is prone to errors when a speaker in the data set uses non-normative pronunciations, or in the presence of homographs [23]. *SCRAPS* could potentially be used to find the most probable phoneme, that better matches the allophones present in the

associated recording, allowing to have a higher quality automatic annotation and hence a potentially better model.

- Intelligibility optimization: intelligibility is an important metric for speech synthesis systems, although it is not the only one (other examples are signal quality, speaker id match or naturalness). Usually, these models do not directly optimize these objectives, but others such as the maximum likelihood [3], that indirectly optimize the previously mentioned ones. It has been repeatedly proven that the those indirect objectives do not always correlate with the perceptual ones [12]. *SCRAPS* model, as a consequence of being a fully differentiable system, offers the possibility of being used as a loss function to directly optimize the match between the target transcription and the content of the spectrogram, or in other words, to directly maximize intelligibility. This could be achieved by plugging the *SCRAPS* encoders (frozen) in a speech generation architecture and maximizing the dot product between the target phonetic sequence and the synthesized spectrogram. This could be done in combination with the traditional objectives. Alternatively, one could also use *SCRAPS* to monitor the evolution of intelligibility during training in order to, for instance, implement early stopping mechanisms.

6 Conclusions

We showed how by using a *CLIP*-like architecture we are able to model a latent space that connects phonetic and acoustic domains. The experiments provide evidence that the embeddings are robust to different kinds of noise, with no degradation up to 60% of corruption when using *Gaussian* noise and 40% when dealing with concurrent speech. At the same time, we observed that *SCRAPS* reacts decreasing the score, in 91.06% of the examples where at least 20% of the phonemes are corrupted. Finally, we provided two examples of downstream applications that show that *SCRAPS* is a useful component in several downstream applications, and discussed future research lines.

References

- [1] Rosana Ardila, Megan Branson, Kelly Davis, Michael Kohler, Josh Meyer, Michael Henretty, Reuben Morais, Lindsay Saunders, Francis Tyers, and Gregor Weber, 'Common voice: A massively-multilingual speech corpus', in *Proceedings of the 12th Language Resources and Evaluation Conference*, pp. 4218–4222, Marseille, France, (May 2020). European Language Resources Association.
- [2] Layne Berry, Yi-Jen Shih, Hsuan-Fu Wang, Heng-Jui Chang, Hung-yi Lee, and David Harwath, 'M-speechclip: Leveraging large-scale, pre-trained models for multilingual speech to image retrieval', in *arXiv Preprint, submitted to ICASSP 2023*, (2022).
- [3] Christopher M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)*, Springer-Verlag New York Inc., April 2011.
- [4] Rui Cao, Yihao Wang, Yuxin Liang, Ling Gao, Jie Zheng, Jie Ren, and Zheng Wang, 'Exploring the impact of negative samples of contrastive learning: A case study of sentence embedding', in *Findings of the Association for Computational Linguistics: ACL 2022*, pp. 3138–3152, Dublin, Ireland, (May 2022). Association for Computational Linguistics.
- [5] Junyi Chai, Hao Zeng, Anming Li, and Eric W.T. Ngai, 'Deep learning in computer vision: A critical review of emerging techniques and application scenarios', *Machine Learning with Applications*, **6**, 100134, (2021).
- [6] Marcos V. Conde and Kerem Turgutlu, 'Clip-art: Contrastive pre-training for fine-grained art classification', in *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pp. 3951–3955, (2021).
- [7] Rohan Doshi, Youzheng Chen, Liyang Jiang, Xia Zhang, Fadi Biadsy, Bhuvana Ramabhadran, Fang Chu, Andrew Rosenberg, and Pedro J. Moreno, 'Extending parrottron: An end-to-end, speech conversion and speech recognition model for atypical speech', in *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6988–6992, (2021).
- [8] Benjamin Elizalde, Soham Deshmukh, Mahmoud Al Ismail, and Huaming Wang, 'CLAP: Learning Audio Concepts From Natural Language Supervision', *arXiv e-prints*, arXiv:2206.04769, (June 2022).
- [9] Adam Gabryś, Goeric Huybrechts, Manuel Sam Ribeiro, Chung-Ming Chien, Julian Roth, Giulia Comini, Roberto Barra-Chicote, Bartek Perz, and Jaime Lorenzo-Trueba, 'Voice filter: Few-shot text-to-speech speaker adaptation using voice conversion as a post-processing module', in *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 7902–7906, (2022).
- [10] Basi Garcia, Brendan Shillingford, Hank Liao, Olivier Siohan, Otavio de Pinho Forin Braga, Takaki Makino, and Yannis Assael, 'Recurrent neural network transducer for audio-visual speech recognition', in *Proceedings of IEEE Automatic Speech Recognition and Understanding Workshop*, (2019).
- [11] Kallirroi Georgila, *Speech Synthesis: State of the Art and Challenges for the Future*, 257–272, Cambridge University Press, 2017.
- [12] Ian Goodfellow, Joshua Bengio, and Aaron Courville, *Deep Learning*, MIT Press, November 2016.
- [13] Andrey Guzhov, Federico Raue, Jörn Hees, and Andreas Dengel, 'Audioclip: Extending clip to image, text and audio', in *ICASSP 2022*, pp. 976–980, (2022).
- [14] Sepp Hochreiter and Jürgen Schmidhuber, 'Long Short-Term Memory', *Neural Computation*, **9**(8), 1735–1780, (November 1997).
- [15] Umut Isik, Ritwik Giri, Neerad Phansalkar, Jean-Marc Valin, Karim Helwani, and Arvindh Krishnaswamy, 'Poconet: Better speech enhancement with frequency-positional embeddings, semi-supervised conversational data, and biased loss', in *Interspeech 2020*, (2020).
- [16] Apoorv Khandelwal, Luca Weihs, Roozbeh Mottaghi, and Aniruddha Kembhavi, 'Simple but effective: Clip embeddings for embodied ai', in *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 14809–14818, (2022).
- [17] Diederik P. Kingma and Jimmy Ba, 'Adam: A Method for Stochastic Optimization.', in *Proceedings of the 3rd International Conference of Learning Representations (ICLR)*, (December 2014).
- [18] Keisuke Kinoshita, Tsubasa Ochiai, Marc Delcroix, and Tomohiro Nakatani, 'Improving noise robust automatic speech recognition with single-channel time-domain enhancement network', in *Proceedings of the IEEE ICASSP*. arXiv, (2020).
- [19] Phuc H. Le-Khac, Graham Healy, and Alan F. Smeaton, 'Contrastive representation learning: A framework and review', *IEEE Access*, **8**, 193907–193934, (2020).
- [20] Naihan Li, Shujie Liu, Yanqing Liu, Sheng Zhao, and Ming Liu, 'Neural speech synthesis with transformer network', in *Proceedings of the 31st AAAI Conference on Artificial Intelligence, AAAI'19/IAAI'19/EAAI'19*. AAAI Press, (2019).
- [21] Jaime Lorenzo-Trueba, Junichi Yamagishi, Tomoki Toda, Daisuke Saito, Fernando Villavicencio, Tomi Kinnunen, and Zhenhua Ling, 'The voice conversion challenge 2018: Promoting development of parallel and nonparallel methods, 2018.
- [22] Michael McAuliffe, Michaela Socolof, Sarah Mihuc, Michael Wagner, and Morgan Sonderegger, 'Montreal Forced Aligner: Trainable Text-Speech Alignment Using Kaldi', in *Proceedings of the Interspeech conference*, pp. 498–502, (2017).
- [23] Artem Ploujnikov and Mirco Ravanelli, 'Soundchoice: Grapheme-to-phoneme models with semantic disambiguation', in *Interspeech 2022, 23rd Annual Conference of the International Speech Communication Association, Incheon, Korea, 18-22 September 2022*, eds., Hanseok Ko and John H. L. Hansen, pp. 486–490. ISCA, (2022).
- [24] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever, 'Learning transferable visual models from natural language supervision', in *Proceedings of the 38th International Conference on Machine Learning*, eds., Marina Meila and Tong Zhang, volume 139 of *Proceedings of Machine Learning Research*, pp. 8748–8763. PMLR, (18–24 Jul 2021).
- [25] Yi Ren, Xu Tan, Tao Qin, Zhou Zhao, and Tie-Yan Liu, 'Revisiting over-smoothness in text to speech', in *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1)*, pp. 8197–8213, Dublin, Ireland, (May 2022). Association for Computational Linguistics.
- [26] Joshua Robinson, Ching-Yao Chuang, Suvrit Sra, and Stefanie Jegelka, 'Contrastive learning with hard negative samples', in *Proceedings of the International Conference of Learning Representations (ICLR)*, (2021).
- [27] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer, 'High-resolution image synthesis with latent diffusion models', in *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10674–10685, (2022).
- [28] Jonathan Shen, Ruoming Pang, Ron J. Weiss, Mike Schuster, Navdeep Jaitly, Zongheng Yang, Zhifeng Chen, Yu Zhang, Yuxuan Wang, Rj Skerrv-Ryan, Rif A. Saurous, Yannis Agiomvrgiannakis, and Yonghui Wu, 'Natural tts synthesis by conditioning wavenet on mel spectrogram predictions', in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 4779–4783, (2018).
- [29] Yi-Jen Shih, Hsuan-Fu Wang, Heng-Jui Chang, Layne Berry, Hung-yi Lee, and David Harwath, 'Speechclip: Integrating speech with pre-trained vision and language model', in *Proceedings of the IEEE Workshop on Spoken Language Technology 2022*, (2022).
- [30] Paul Taylor, *Text-to-Speech Synthesis*, Cambridge University Press, 2009.
- [31] Federico Urli, Emiliano Versini, and Lauro Snidaro, 'Fusion of sentence embeddings for news retrieval', in *2022 25th International Conference on Information Fusion (FUSION)*, pp. 1–7, (2022).
- [32] Aäron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew W. Senior, and Koray Kavukcuoglu, 'WaveNet: A Generative Model for Raw Audio', in *Proceedings of the 9th International Speech Communication Association (ISCA), Speech Synthesis Workshop, Sunnyvale, CA, USA*, p. 125. ISCA, (September 2016).
- [33] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin, 'Attention is All you need', in *Proceedings of the 31st Neural Information Processing Systems conference (NIPS), NIPS'17*, p. 6000–6010, Red Hook, NY, USA, (2017). Curran Associates Inc.
- [34] J. C. Wells. Computer-coding the IPA: a proposed extension of SAMPA, 1995. [Online resource; accessed 23-July-2023].
- [35] Genta Indra Winata, Guangsen Wang, Caiming Xiong, and Steven C. H. Hoi, 'Adapt-and-adjust: Overcoming the long-tail problem of multilingual speech recognition', in *Proceedings of the Interspeech conference*, volume abs/2012.01687, (2021).
- [36] Ho Hsiang Wu, Prem Seetharaman, Kundan Kumar, and Juan Pablo Bello, 'Wav2clip: Learning robust audio representations from clip', in *2022 IEEE International Conference on Acoustics, Speech, and Signal Processing, ICASSP 2022 - Proceedings*, pp. 4563–4567, (2022).