



Enhanced ASR Robustness to Packet Loss with a Front-End Adaptation Network

Yehoshua Dissen, Shiry Yonash, Israel Cohen, Joseph Keshet

Faculty of Electrical and Computer Engineering,
Technion – Israel Institute of Technology, Israel

shua.dissen@gmail.com

Abstract

In the realm of automatic speech recognition (ASR), robustness in noisy environments remains a significant challenge. Recent ASR models, such as Whisper, have shown promise, but their efficacy in noisy conditions can be further enhanced. This study is focused on recovering from packet loss to improve the word error rate (WER) of ASR models. We propose using a front-end adaptation network connected to a frozen ASR model. The adaptation network is trained to modify the corrupted input spectrum by minimizing the criteria of the ASR model in addition to an enhancement loss function. Our experiments demonstrate that the adaptation network, trained on Whisper’s criteria, notably reduces word error rates across domains and languages in packet-loss scenarios. This improvement is achieved with minimal affect to Whisper model’s foundational performance, underscoring our method’s practicality and potential in enhancing ASR models in challenging acoustic environments.

Index Terms: Robust speech recognition, packet loss concealment, transformers, Whisper.

1. Introduction

Automatic speech recognition (ASR) has made huge progress in the past few years with the introduction of large pre-trained [1, 2] or weakly supervised transformer models trained on massive amounts of data, such as Whisper [3]. While these models perform well in many domains, there still is room for significant improvement in challenging environments such as noisy or reverberant data. This work focuses on one of the more complex cases, namely, the packet loss scenario, where parts of the audio data are lost or corrupted during transmission.

Improving the models’ robustness under these conditions can be challenging. Finetuning these models can easily overfit to the domain of data you finetune with. For example, if the model is finetuned on English data, the model might “forget” other languages [4]. Even within the same language, the model can improve on read speech while degrading performance on phone call speech. Training a model from scratch is often impractical due to the large size of the model and the vast amount of data required, which can demand excessive computational resources and time.

Another approach, would be using a packet loss concealment (PLC) model. PLC algorithms aim to solve the task of reconstructing missing frames from a signal. This work seeks to create a packet loss concealer directed at the downstream task of improving automatic speech recognition (ASR). Most PLC algorithms aim to improve the speech quality perception, while very few aim to solve a downstream task such as [5], where they work on PLC for speech emotion recognition. Traditionally, concealment was done using some form of Linear Prediction

or interpolation algorithms such as in [6, 7, 8]. Following the proliferation of neural nets, they have become the predominant method [9, 10, 11, 12]. These models generally improve human intelligibility, and usually improve WER as well; however, they can introduce artifacts or distortions in the signal that are not well-received by an ASR model, compromising their efficacy. We further discuss related work in Section 4.

We aim to create a simple method for improving packet loss robustness for foundational ASR models without needing in-domain data. For this, we turned to the ASR model inputs or features. We added a small front-end adaptation model that fills the gaps in the input spectrum before passing it on to the back-end ASR model. We used a U-net [13] architecture with skip connections from popular PLC and inpainting models. However, since our goal is to improve ASR metrics, specifically Word Error Rate (WER) and not the audio quality, instead of using perceptual losses, we utilize the gradients from the ASR model to update the adaptation models weights. We are essentially training a packet loss concealer with ASR objectives while keeping the ASR model frozen. This strategy allows for improving the model’s robustness to packet loss while keeping the ASR model’s initial capabilities without retraining the entire model, a process often constrained by resource limitations or fine-tuning where the model is prone to domain overfitting. A comprehensive series of experiments with the proposed method demonstrated greatly improved robustness to packet-loss corruption compared to the baseline models, fine-tuning, and other PLC methods. This includes performance across domains and languages that differed from those in the training set. Moreover, by maintaining the ASR models’ weights unchanged, the original performance was not compromised, ensuring that the improvements in robustness did not detract from the models existing capabilities.

The main contribution of this work is a method that can improve foundational ASR models’ robustness to packet loss without changing the underlying ASR model or degrading its results using a very lightweight adapter model. Our implementation and trained models are available here ¹.

2. Methods

This study proposes a technique that improves ASR robustness to packet loss scenarios while maintaining the pre-trained ASR architecture and weights. As stated earlier, one option would be to use a PLC module to reconstruct the speech and subsequently apply ASR on the resulting speech. However, this solution is sub-optimal as the PLC model can introduce artifacts detrimental to the ASR model. Here, we would like to consider a different approach when replacing the PLC module with a

¹<https://github.com/MLSpeech/WhisperDenoiser>

module that will adapt the signal explicitly to improve the ASR robustness rather than make the speech sound better. We start by presenting the notation and our general setting.

We denote the speech signal by $X = (x_1, \dots, x_T)$ as a sequence of T frames (here, each x_t denotes a frame of the mel-spectrum). We denote the corrupted speech signal by \tilde{X} , where $\tilde{x}_k \dots \tilde{x}_{k+j}$ are j lost frames starting at the k -th frame. There might be several spans of packet loss within a single utterance. We assume a transcript is associated with the speech signal, which is a sequence of U words or sub-words (tokens). It will be denoted by $Y = (y_1, \dots, y_U)$. Note that T and U differ for each input (and target) sequence. In our setting, we would like to propose a model that receives the corrupted speech \tilde{X} and outputs the target transcription Y as if it had received the original (unobserved) signal X .

Our model has two main components: a front-end *adaptation network*, and a *frozen ASR*. We denote the ASR model g_ϕ with a parameter set ϕ . This function $\hat{Y} = g_\phi(\tilde{X})$ gets as input a speech signal and predicts the word (token) sequence spoken. It is trained with some loss function $L(g_\phi(\tilde{X}), Y)$.

We aim to design the adaptation network f_θ with a parameter set θ . This network gets as input the noisy speech and outputs an adapted version of it $\hat{X} = f_\theta(\tilde{X})$, which is used as input to the ASR, $\hat{Y} = g_\phi(\hat{X})$. Our goal is to have $\text{WER}(\hat{Y}, Y) \leq \text{WER}(\tilde{Y}, Y)$. We note that \hat{X} is generated to improve the ASR performance and might not improve human intelligibility.

This study uses Whisper models as the ASR models and a U-net architecture as the adaptation network. This network is trained with two loss functions. The first loss function is the ASR model’s principal loss function, which for Whisper is cross entropy L_{CE} . This loss function guides the adapter network toward generating a spectrum better suited for greater token classification accuracy. We found that training on the ASR loss alone can sometimes converge in an unstable manner, so we added a second loss function, the L_1 loss component between the original signal X and adapted signal $f_\theta(\tilde{X})$. This loss serves as a form of regularization, i.e.,

$$\min_{\theta} \lambda L_{\text{CE}}(g_\phi(f_\theta(\tilde{X}), Y)) + (1 - \lambda)L_1(X, f_\theta(\tilde{X})) . \quad (1)$$

We emphasize that the minimization is over the adapter network parameters θ , while the Whisper ϕ parameters are fixed. In the evaluation, we show the advantage of our model over fine-tuning Whisper ϕ .

The adaptation network is a fully convolutional network with a U-net architecture and skip connections. The bottleneck consists of residual-blocks. Downscaling is done by maxpooling. Upscaling is done by nearest neighbor resizing followed by a convolutional layer. The input to the Whisper model is a mel-spectrum. Hence, the adapter network is designed to receive the mel-spectrum of the noisy signal \tilde{X} and output an adapted mel-spectrum. This is depicted in Figure 1.

3. Empirical Evaluation

In this section, we describe a set of experiments to demonstrate the proposed method’s effectiveness empirically. We start by defining the datasets used for evaluation.

3.1. Datasets

For training, we use the 960 hours of English LibriSpeech [14]. For evaluation, we use multiple datasets from very different

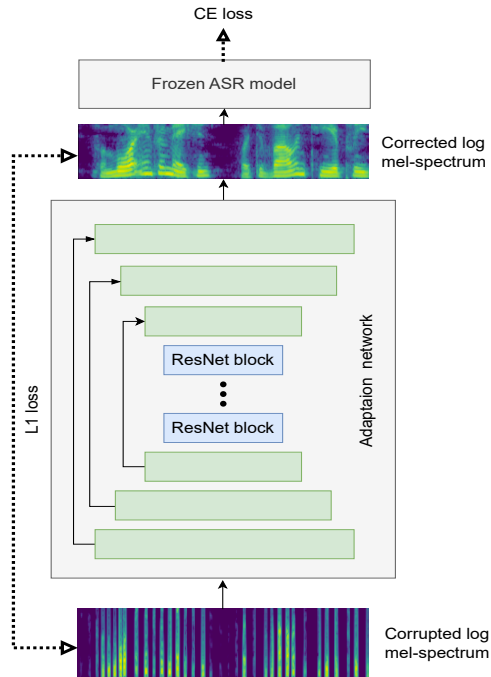


Figure 1: Our model comprises an adaptation network connected to an ASR model. The network is a convolutional U-net that receives the corrupted speech and outputs a mel-spectrum. The ASR model is a trained Whisper model.

domains to showcase the method’s robustness. A subset of ALLSSTAR [15], which is a collection of L1 Mandarin speakers speaking English [16], and Fleurs [17] for testing on multiple languages. We don’t report improvements on LibriSpeech test as they are not interesting since the improvements can come from overfitting to the training domain.

For the packet-loss simulation, we randomly zero out frames based on two probabilities: a drop frequency (the percentage of zeroed frames per utterance) and a probabilistic distribution governing the span of consecutive frame losses. A single utterance can have multiple spans of packet loss. During training, there is a drop frequency distribution, and each sample loaded accordingly gets assigned a drop rate. During inference, for reporting reasons, we duplicate the test set to multiple fixed drop frequencies. Due to the nature of the span length distribution, there might be tiny (up to a tenth of a percent) variations from the fixed rate. When we report packet loss percentage, we mean the total percentage of lost frames in the utterance.

Additionally, we evaluated the models on the blind set from the Interspeech 2022 Audio Deep PLC Challenge [18] for further validation of the results. Naturally, the data for this challenge has a set packet loss, so it’s used as is.

3.2. Experimental setting

Recall that our model consists of two components: a frozen ASR model and an adaptation network. Whisper, which serves as the ASR model, comes in several parameter sizes. In this study, we demonstrated the effectiveness of our method with the multilingual versions of base (74M) and large-v2 (1550M) model sizes. We utilized the same Whisper parameters for all decoding, namely beam size 5, without timestamps, and manually set the language.




Figure 2: WER% of different models on ALLSSTAR with various PLRs. All the decoding is done with Whisper base. tPLCnet refers to Westhausen and Meyer [10].




Figure 3: WER% of different models on ALLSSTAR with various PLRs. All the decoding is done with Whisper large-v2. tPLCnet refers to [10]. BTA refers to adaptation networks trained on Whisper base but connected to Whisper large-v2.

The input to the adaptation network is a mel-spectrum, and the output is an estimated mel-spectrum of the same dimensions. The network is composed of three downsampling and upsampling layers, with skip connections between each equivalently sized layer, with 6 ResNet [19] blocks serving as the bottleneck layers. Additionally, there are single input and single output convolutional layers that retain the same dimensions. This model’s total number of trainable parameters is 7.5M, making this a negligible addition to Whisper. It is trained with the cross entropy and the L_1 loss functions, where the loss functions are scaled by 0.9 for the L_1 and 0.1 for the L_{CE} (chosen on a validation set). We used a learning rate of 0.0005 with a 10% decay rate per epoch.

3.3. Results

In this section, we present the evaluation of the proposed method and analyze the effect of different loss functions on WER. We demonstrate the relative improvement of the proposed method over the unchanged baseline Whisper model and a recently published, open-source PLC model [10]. We then evaluate the model’s robustness to different domains and compare it to fine-tuning Whisper. In all the experiments, the Whisper baselines use *zero-fill* for the dropped frames.

Table 1: Comparison of WER% for Different Models on the Packet Loss Challenge blind set.

Model	WER% (base)	WER% (large-v2)
Whisper	24.0	15.4
tPLCnet [10]	20.4	16.2
Ours	18.1	14.2

Figures 2 and 3 present the performances of Whisper base and Whisper large-v2, respectively, on the original mel-spectrums in comparison with the spectrums generated by our adaptation networks and by a PLC model. The graphs present WER% for various packet loss rate (PLR) values on the ALLSSTAR dataset. We note that the vanilla Whisper large model is more robust to frame loss, only starting to seriously degrade at PLRs larger than 20%, whereas the base model starts degrading immediately.

We present the effect of training the adaptation network with each loss function. Specifically, we compare the performance while training (i) solely using the CE loss function, L_{CE} , where the gradients flow from Whisper (noted as CE only); (ii) solely L_1 loss between the clean and lossy signals without referencing Whisper (noted as L1 only), which can be seen as similar to the TF-Unet in [20]; and (iii) a combined loss of L_{CE} and L_1 (CE + L1).

It can be seen that all methods improve WER over the original whisper model. However, the CE loss improves results more than the L_1 enhancement loss, whereas combining the two losses generates the most significant improvement. In Figure 3, we also presented the performance of the adaptation network trained on the gradients of the base model (the one that is depicted in Figure 2) but connected and evaluated with Whisper large-v2. In Figure 3, we denote these models as *Based Trained Adaptation (BTA)* and labeled them with *Ours: CE (BTA)* and *Ours: CE+L1 (BTA)*). Interestingly, training the model on Whisper base and connecting it to Whisper large-v2 gets better results than the models trained directly using Whisper large-v2. We assume this is because the gradients of the base model are easier to handle and, therefore, more effectively influence the adaptation networks. This suggests that a better training parameters exist for the large model. We defer this issue for further research. This example underscores the broader principle, that applying ASR metrics in PLC model training, can significantly enhance ASR performance across a range of models.

Furthermore, the graph shows the WER% of tPLCnet [10], a time-domain many-to-one RNN model for PLC trained with a combined magnitude and complex mean absolute error loss in the time-frequency domain. We ran the large version of this model on the corrupted files and then decoded them with Whisper (base and large). The graph shows that this method improves the WER. However, the models trained using ASR metrics improve the WER more drastically.

Next, in Table 1, we compare the WER of the baseline whisper models, tPLCnet [10] large and Ours on the blind set from the interspeech 2022 PLC challenge [18]. Here, the PLRs are set by the challenge. As seen, tPLCnet improves the WER over the baseline model but not Whisper large, and Ours performs best in both scenarios.

To further demonstrate the model’s robustness to domains and to showcase that this training method doesn’t harm the original Whisper models as opposed to fine-tuning where training in one domain or language degrades the model’s performance in

Table 2: Comparison of WER% for Different Languages using Whisper base and large-v2.

Whisper Size		Base				Large-V2			
Packet Loss Rate	Model	French	German	Russian	Spanish	French	German	Russian	Spanish
0%	Whisper	24.7	17.2	20.3	10.3	7.2	4.6	6.4	3.7
	Ours	26.3	18.4	21.8	10.7	7.6	4.7	6.4	3.7
5%	Whisper	29.1	20.8	24.6	12.7	7.5	4.7	6.5	3.7
	Ours	27.7	19.6	23.0	11.2	7.8	4.8	6.7	3.9
10%	Whisper	33.8	25.3	28.7	15.8	8.5	5.0	6.8	3.9
	Ours	29.3	20.9	24.7	11.5	8.1	5.0	6.8	3.9
20%	Whisper	48.6	39.0	39.1	24.7	10.3	6.0	7.9	4.2
	Ours	32.6	23.7	27.9	13.2	8.8	5.7	7.5	4.1
30%	Whisper	69.5	60.7	53.7	38.3	16.0	8.3	11.5	5.3
	Ours	38.3	27.2	32.6	15.5	9.9	6.3	8.8	4.4
40%	Whisper	104.9	102.2	75.6	57.2	27.5	13.4	18.2	7.5
	Ours	41.5	32.8	39.0	19.0	12.3	7.6	9.9	5.2
50%	Whisper	126.7	141.2	100.8	89.2	48.3	26.4	34.4	12.7
	Ours	49.6	40.4	46.1	24.4	15.5	10.3	12.7	6.2
60%	Whisper	124.6	140.4	121.1	127.9	71.7	51.8	61.5	26.6
	Ours	61.2	53.2	59.2	33.7	20.6	15.5	19.1	9.4

Table 3: Comparison of WER% for finetuning vs Ours using Whisper base.

Dataset	ALLSSTAR			Spanish		
	0	0.2	0.4	0	0.2	0.4
Whisper	18.4	37.8	70.0	10.3	24.7	57.2
Fine-tune	24.9	27.1	31.8	89.5	91.1	94.5
Ours	18.7	22.1	30.0	10.7	13.2	19.0

other domains or languages, we compare, in Table 2, the WER of the model to the original Whisper models in multiple languages selected at random from the Fleurs dataset [17]. Here, the pattern is similar to the results on the ALLSSTAR dataset, as shown in Figures 2 and 3: the base model starts degrading immediately, and the large only after 20% PLR, whereas Ours, other than a slight degradation in the zero PL scenario, improves results for all PLRs in all languages.

Finally, to establish the advantage of this method over fine-tuning, we also fine-tuned Whisper using the same training data, and the outcomes aligned precisely with our earlier theories. Notably, the performance on the LibriSpeech test set improved substantially. However, this enhancement came at a cost: first, as seen in Table 3, the model lost its multilingual capabilities; second, its performance on ALLSSTAR, an English dataset but from a different domain, experienced a significant decline from the baseline in the no packet loss (clean) case.

4. Related work

PLC. While packet loss concealment is well studied, it primarily focuses on perceptual audio quality, focusing on improving metrics like Perceptual Evaluation of Speech Quality (PESQ) and Short-Time Objective Intelligibility (STOI), whereas this research explicitly targets ASR performance. The current more prevalent framework for neural network-based PLC are encoder decoder speech inpainting methods such as Wang *et al.*, [9] or Pascual *et al.*, [11] who use adversarial training to generate the most natural sounding audio in the gaps. Another standard method is a causal, recurrent, sequence-to-one model where the objective is to predict the next frame given the previous frames.

Westhausen and Meyer [10] whom we compared to, predict the next lost frame with a short context buffer in the time domain, using an RNN. Lin *et al.* [12] frame it as a generative regression problem and utilize a convolutional encoder-decoder with LSTM layers for next frame prediction in the time domain.

Spectral inpainting. Context-based retrieval of missing parts of time-frequency representation of speech using a convolutional U-net was originally demonstrated by Kegler *et al.*[21]. Simon *et al.* [22], used it for correcting mispronunciations in speech by cutting out the mispronunciations and inpainting it correctly with the context.

ASR guided enhancement. There is also precedent for using ASR objectives to improve distorted signals. Subramanian *et al.* [23] used end-to-end speech recognition objectives to train a speech enhancement model. They showed that in addition to improving the WER, they also improved the speech enhancement metrics. Another interesting method, presented by Yang *et al.* [24], uses an auxiliary loss between latent ASR representations of the clean signal and the representations generated by their packet loss concealment model. This generates more natural speech reconstruction.

5. Discussion and future work

This study introduced a novel approach to enhance the reliability of large ASR models in facing packet loss scenarios. The proposed method involves the integration of a smaller model, which is specifically designed to adapt the input features of an ASR model. The study has demonstrated that this integration process leads to significant improvements in the robustness of the ASR models when trained using the gradients of a larger ASR model. The promising results shown by the proposed method opens up new avenues for future research in the field of ASR development. Future research can investigate the applicability of this method in enhancing the robustness of ASR models against various noise types, such as white noise, babble, clipping, and echo suppression. It would also be interesting to examine the generalizability of this approach across different ASR model architectures, such as HuBERT, and evaluate the extent to which this approach can improve traditional intelligibility metrics.

6. References

- [1] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” *Advances in neural information processing systems*, vol. 33, pp. 12 449–12 460, 2020.
- [2] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhota, R. Salakhutdinov, and A. Mohamed, “Hubert: Self-supervised speech representation learning by masked prediction of hidden units,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3451–3460, 2021.
- [3] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, “Robust speech recognition via large-scale weak supervision,” in *International Conference on Machine Learning*. PMLR, 2023, pp. 28 492–28 518.
- [4] G. Sun, X. Zheng, C. Zhang, and P. C. Woodland, “Can Contextual Biasing Remain Effective with Whisper and GPT-2?” in *Proc. INTERSPEECH 2023*, 2023, pp. 1289–1293.
- [5] M. M. Mohamed and B. W. Schuller, “Concealnet: An end-to-end neural network for packet loss concealment in deep speech emotion recognition,” *arXiv preprint arXiv:2005.07777*, 2020.
- [6] E. Gunduzhan and K. Momtahan, “Linear prediction based packet loss concealment algorithm for pcm coded speech,” *IEEE Transactions on Speech and Audio Processing*, vol. 9, no. 8, pp. 778–785, 2001.
- [7] J.-H. Chen, “Packet loss concealment based on extrapolation of speech waveform,” in *2009 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2009, pp. 4129–4132.
- [8] H. Ofir, D. Malah, and I. Cohen, “Audio packet loss concealment in a combined mdct-mdst domain,” *IEEE Signal Processing Letters*, vol. 14, no. 12, pp. 1032–1035, 2007.
- [9] J. Wang, Y. Guan, C. Zheng, R. Peng, and X. Li, “A temporal-spectral generative adversarial network based end-to-end packet loss concealment for wideband speech transmission,” *The Journal of the Acoustical Society of America*, vol. 150, no. 4, pp. 2577–2588, 2021.
- [10] N. L. Westhausen and B. T. Meyer, “tPLCnet: Real-time Deep Packet Loss Concealment in the Time Domain Using a Short Temporal Context,” in *Proc. Interspeech 2022*, 2022, pp. 2903–2907.
- [11] S. Pascual, J. Serrà, and J. Pons, “Adversarial auto-encoding for packet loss concealment,” in *2021 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*. IEEE, 2021, pp. 71–75.
- [12] J. Lin, Y. Wang, K. Kalgaonkar, G. Keren, D. Zhang, and C. Fuegen, “A time-domain convolutional recurrent network for packet loss concealment,” in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 7148–7152.
- [13] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*. Springer, 2015, pp. 234–241.
- [14] 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)*. IEEE, 2015, pp. 5206–5210.
- [15] A. Bradlow, L. Ackerman, L. Burchfield, L. Hesterberg, J. Luque, and K. Mok, “Allstar: Archive of l1 and l2 scripted and spontaneous transcripts and recordings,” in *Proceedings of the International Congress on Phonetic Sciences*, 2010, pp. 356–359.
- [16] S.-E. Kim, B. R. Chernyak, O. Seleznova, J. Keshet, M. Goldrick, and A. R. Bradlow, “Automatic recognition of second language speech-in-noise,” *JASA Express Letters*, vol. 4, no. 2, 2024.
- [17] A. Conneau, M. Ma, S. Khanuja, Y. Zhang, V. Axelrod, S. Dalmia, J. Riesa, C. Rivera, and A. Bapna, “Fleurs: Few-shot learning evaluation of universal representations of speech,” in *2022 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2023, pp. 798–805.
- [18] L. Diener, S. Sootla, S. Branets, A. Saabas, R. Aichner, and R. Cutler, “INTERSPEECH 2022 Audio Deep Packet Loss Concealment Challenge,” in *Proc. Interspeech 2022*, 2022, pp. 580–584.
- [19] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [20] A. A. Nair and K. Koishida, “Cascaded time+ time-frequency unet for speech enhancement: Jointly addressing clipping, codec distortions, and gaps,” in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 7153–7157.
- [21] M. Kegler, P. Beckmann, and M. Cernak, “Deep Speech Inpainting of Time-Frequency Masks,” in *Proc. Interspeech 2020*, 2020, pp. 3276–3280.
- [22] T.-B. Simon, F. Kreuk, F. Awwad, J. T. Cohen, and J. Keshet, “Correcting mispronunciations in speech using spectrogram inpainting,” *Proc. Interspeech 2022*, pp. 1208–1212, 2022.
- [23] A. S. Subramanian, X. Wang, M. K. Baskar, S. Watanabe, T. Taniguchi, D. Tran, and Y. Fujita, “Speech enhancement using end-to-end speech recognition objectives,” in *2019 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*. IEEE, 2019, pp. 234–238.
- [24] D.-H. Yang and J.-H. Chang, “Towards robust packet loss concealment system with asr-guided representations,” in *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2023, pp. 1–8.