

TRAMBA: A Hybrid Transformer and Mamba Architecture for Practical Audio and Bone Conduction Speech Super Resolution and Enhancement on Mobile and Wearable Platforms

YUEYUAN SUI*, Northwestern University, United States

MINGHUI ZHAO*, Columbia University, United States

JUNXI XIA, Northwestern University, United States

XIAOFAN JIANG, Columbia University, United States

STEPHEN XIA, Northwestern University, United States

We propose TRAMBA, a hybrid transformer and Mamba architecture for acoustic and bone conduction speech enhancement, suitable for mobile and wearable platforms. Bone conduction speech enhancement has been impractical to adopt in mobile and wearable platforms for several reasons: (i) data collection is labor-intensive, resulting in scarcity; (ii) there exists a performance gap between state-of-art models with memory footprints of hundreds of MBs and methods better suited for resource-constrained systems. To adapt TRAMBA to vibration-based sensing modalities, we pre-train TRAMBA with audio speech datasets that are widely available. Then, users fine-tune with a small amount of bone conduction data. TRAMBA outperforms state-of-art GANs by up to 7.3% in Perceptual Evaluation of Speech Quality (PESQ) and 1.8% in Short-Time Objective Intelligibility (STOI), with an order of magnitude smaller memory footprint and an inference speed up of up to 465 times. We integrate TRAMBA into real systems and show that TRAMBA (i) improves battery life of wearables by up to 160% by requiring less data sampling and transmission; (ii) generates higher quality voice in noisy environments than over-the-air speech; (iii) requires a memory footprint of less than 20.0 MB.

CCS Concepts: • Computer systems organization → Embedded and cyber-physical systems; • Human-centered computing → Ubiquitous and mobile computing; • Computing methodologies → Artificial intelligence.

Additional Key Words and Phrases: Audio Super Resolution, Speech Enhancement, Bone Conduction Enhancement, Transformer, Mamba, Earable Computing

ACM Reference Format:

Yueyuan Sui, Minghui Zhao, Junxi Xia, Xiaofan Jiang, and Stephen Xia. 2024. TRAMBA: A Hybrid Transformer and Mamba Architecture for Practical Audio and Bone Conduction Speech Super Resolution and Enhancement on Mobile and Wearable Platforms. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 8, 4, Article 205 (December 2024), 29 pages. <https://doi.org/10.1145/3699757>

1 INTRODUCTION

Wearables have revolutionized the way we interact with technology and have enabled continuous and real-time monitoring of our health and wellness. Since the early 2000s, wearables have continued to play an increasingly

*Both authors contributed equally to this work.

Authors' addresses: Yueyuan Sui, Northwestern University, Evanston, Illinois, United States, yueyuansui2024@u.northwestern.edu; Minghui Zhao, Columbia University, New York, New York, United States, mz2866@columbia.edu; Junxi Xia, Northwestern University, Evanston, Illinois, United States, junxia2024@u.northwestern.edu; Xiaofan Jiang, Columbia University, New York, New York, United States, jiang@ee.columbia.edu; Stephen Xia, Northwestern University, Evanston, Illinois, United States, stephen.xia@northwestern.edu.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2024 Copyright held by the owner/author(s).

ACM 2474-9567/2024/12-ART205

<https://doi.org/10.1145/3699757>

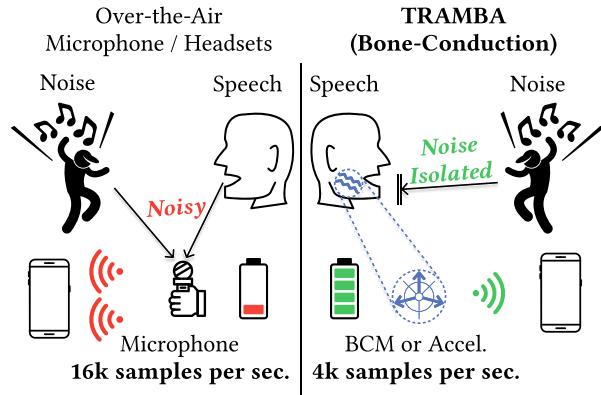


Fig. 1. TRAMBA enhances vibration-based speech, which is naturally insensitive to ambient noises.

important role in our lives. The wearables market is expected to grow from 70 billion USD in 2023 to 230 billion USD by 2032 [39]. Head-worn wearables, including earphones and glasses, continue to be one of the fastest growing segments (71 billion USD in 2023 to 172 billion USD by 2030 [13, 14]). This market has been spurred on, in part, by the increasing importance and prevalence of earables, augmented reality (AR), and virtual reality (VR).

One unique signal that can only be sensed by wearables near or directly on the head is speech. Traditionally, speech is captured by over-the-air (OTA) microphones that convert variations in air pressure into an electrical signal, which can be used for a variety of important applications (e.g., phone conversations and voice commands). OTA microphones are commonly found on earbuds or headphones and are generally positioned near the user's mouth. However, OTA microphones easily pick up background sounds, especially in noisier environments, which can degrade speech quality.

To solve this problem, there have been many works that explore denoising, sound source separation, and speech enhancement methods to separate speech from background noise [23, 36, 37, 44, 49, 55]. However, this approach is challenging because 1) the model is “blind” to the types of background noises that may occur. There could be any number of different sounds, or even another person’s speech, that the model needs to account for without prior knowledge, while preserving the wearer’s speech. 2) It is not uncommon to be in areas where background noise can easily interfere and overpower the user’s speech (e.g., a busy cafeteria or construction), which makes speech extraction and enhancement extremely difficult.

Unlike OTA microphones, bone conduction microphones (BCM) are placed in direct contact with the head. BCMs are sensitive to vibrations produced by the skin and skull as the person speaks and are not sensitive to changes in air pressure, making them naturally robust against ambient noise. Other sensors that can measure vibration and forces of motion, such as accelerometers (ACCEL), have shown sensing capabilities of speech at lower fidelities [20, 31] and of general facial movements for other important applications, such as authentication [52]. Bone conduction sensing methods are naturally not sensitive to ambient background noises compared to OTA microphones. However, vibration-based acoustic modalities see *severe attenuation at higher frequencies, which significantly degrades intelligibility and speech quality*.

There are generally two subproblems in speech enhancement: noise reduction and quality improvement. Noise reduction generally focuses on removing noises from external sources (e.g., source separation [55]), while quality improvement generally focuses on making frequencies and other aspects of speech more pronounced (e.g., super

resolution [38]). Because vibration-based modalities are naturally insensitive to ambient noises, this work focuses on speech super resolution for bone-conduction modalities, namely BCMs and ACCELS.

BCMs are commonly used in commercial and consumer products, such as bone conduction headsets, sports and outdoor equipment, or medical devices like hearing aids, to allow voice communication in noisy environments. ACCELS are much more commonly seen than BCMs due to their prevalence as motion trackers in a wide range of wearables, such as smartwatches and fitness trackers. While their primary use is not as a microphone because they are less sensitive to bone-conduction voice, ACCELS are much more widespread and possess the capability to sense voice that we also consider them in this work.

There are several works that explore vibration and bone-conduction super resolution methods to reconstruct higher frequencies and improve speech quality [20, 30], but do not consider practical aspects for real-time mobile, wearable, and earable systems, as detailed below:

- (1) **Heavy Processing:** State-of-art speech super resolution models, namely generative adversarial networks (GANs), require at least tens of millions of parameters (several hundred MB). This makes loading, running, and finetuning the model less practical.
- (2) **Performance Gap:** Methods with smaller memory footprints that target mobile platforms with less computational resources (e.g., many U-Net architectures [30, 33]) see a large gap in performance compared to top performing models, which TRAMBA addresses.
- (3) **Lack of data:** Training a super resolution model for vibration-based sensors requires paired vibration (BCM or ACCEL) and OTA microphone data that captures the higher frequencies. There is a dearth of this data publicly available. Additionally, collecting this data is very labor intensive because it requires building a device that records both BCM/ACCEL data synced with a microphone that captures the higher frequencies. (i) Most super resolution works focus on OTA microphones, where low-fidelity signals can be generated simply by low pass filtering clean speech, for which there are many datasets for. Applying a model trained on such a dataset performs poorly on speech captured by vibration-based sensors. (ii) Other works that focus on bone-conduction settings make commendable efforts to painstakingly collect data from a limited set of volunteers or artificially generate data from clean speech audio. Training a model with limited real samples does not generalize well. We demonstrate both of these shortcomings in Section 5.
- (4) **System Optimization:** Considerations such as sampling rate and where to perform computation can have significant impact on inference time and battery life of the wearable.

Leveraging recent advances in attention and state space models, we propose TRAMBA, a hybrid transformer and Mamba-based model for acoustic speech enhancement and super resolution. TRAMBA outperforms state-of-art GANs by up to 7.3% in PESQ and 1.8% in STOI, with an order of magnitude smaller memory footprint and an inference speed up of up to 465 times. Compared to state-of-art methods, namely GANs, TRAMBA only requires 5.2 million parameters, compared to a minimum of tens of billions of parameters. This yields a model size in the order of MBs compared to hundreds of MBs or more when implemented on a mobile phone. Figure 1 highlights the advantages of TRAMBA compared to OTA methods.

To adapt TRAMBA to bone conduction sensing modalities, we first pre-train TRAMBA for super resolution on speech collected from OTA microphones, before fine-tuning the model with a small amount of data collected from the user. Building the model in this way allows us to take advantage of the vast amount of standard audio datasets that are publicly available, instead of relying on a large amount of paired audio and vibration data, which is both scarcely available and labor intensive to collect. In real systems, this data collection can occur only one time upon box opening. We show that TRAMBA outperforms state-of-art speech enhancement GANs on BCMs and accelerometers across different users and placements on the face by up to 5.9% for PESQ, 1.8% for STOI, while requiring almost 20 times less time to fine-tune on only 15 minutes of data collected from an individual. The

amount of data required is lower, if not on-par, with other voice generation applications that need to collect voice samples from users, such as Apple’s public offering [4].

We incorporate TRAMBA into an end-to-end wearable and mobile platform and demonstrate, through user studies, that TRAMBA improves word error rate by up to 75% in noisy environments compared to approaches that aim to suppress noise in OTA speech. Additionally, we demonstrate that the lack of high-frequency components in bone conduction-based sensing modalities allows us to significantly reduce both the sampling rate of the sensor and the transmission rate, which can improve the battery life of wearables by up to 160%.

This work introduces a novel speech enhancement method that outperforms existing state-of-art methods, with a memory footprint that is orders of magnitude smaller. To the best of our knowledge, this is also the first work that reconstructs intelligible speech using only a single wearable accelerometer. Our contributions are summarized as follows:

- We propose TRAMBA, a hybrid transformer and Mamba-based architecture for speech super resolution. We demonstrate that TRAMBA outperforms existing state-of-art super resolution methods (U-Net and GAN architectures) by up to 109.1% on standard intelligibility and quality metrics, with a model size 5.2 million parameters, compared to state-of-art GANs that require at least tens of millions of parameters.
- We demonstrate that TRAMBA is generalizable to multiple acoustic modalities, including over-the-air microphones and bone/vibration-based modalities (BCMs and accelerometers). To adapt TRAMBA to contact and vibration-based sensing modalities, we fine-tune TRAMBA with only 15 minutes of labeled data collected from the user and demonstrate up to 92.7% improvements across a variety of different sensor placements, compared to the state-of-art. This work is also the first to sense intelligible speech using only a single head-worn accelerometer.
- We integrate TRAMBA into a wearable and mobile platform to demonstrate real-time speech super resolution and more than a 50% reduction in power consumption from sampling and data transmission. Through user studies, we show that vibration-based sensing modalities enhanced by TRAMBA generate significantly higher quality speech, with less interference from background noise, than systems that leverage over-the-air microphones with noise reduction algorithms to remove background sounds.
- We open-source all our code, designs, and example clips: <https://imec-northwestern.github.io/TRAMBAPage/>.

2 RELATED WORKS

2.1 Audio Super Resolution

Audio super-resolution, also known as bandwidth expansion, refers to the restoration of high resolution signals from low resolution signals. Specifically, it employs the low frequency components of the low resolution signal to predict and generate the high-frequency components of the signal to improve quality and generate high-resolution audio. Deep learning has shown great promise for audio super-resolution and speech enhancement. U-Net architectures [34] are one of the most common types of deep learning architectures for audio and speech super resolution, which generally consist of a series of “contracting” layers to learn latent features that could be represented with lower dimensionality than speech signals. Then, a series of “expanding” layers upsample the embeddings to the final high resolution audio signal, producing high resolution speech. This type of architecture has shown good performance because speech is very structured, with high energy concentrated around specific frequency bands depending on the user’s voice and what is spoken.

Kuleshov et al. proposed Audio-UNet [27], the first architecture employing convolutional networks for audio super-resolution tasks. Birnbaum et al. integrated Feature-wise Linear Modulation (FiLM) [46] with Audio-UNet, resulting in the Temporal FiLM (TFiLM) [5], which significantly improved the performance of Audio-UNet. Rakotonirina incorporated recurrent neural networks (RNNs) with self-attention [48] to further enhance perceptual quality. Other works leverage frequency domain representations as input, such as the three-layer fully connected

network proposed by [28]. To take advantage of the strengths of both time-domain and frequency-domain processing, several works propose multi-pathway architectures to process both representations simultaneously, such as TFNet [32].

Generative adversarial networks (GANs), such as Seanet [29], EBEN [19], and Aero [38], have recently shown exceptional performance in generating high quality and natural sounding speech from low resolution audio. However, GANs typically employ both a generator and discriminator network, which often exhibits high complexity with numerous associated parameters (on order of tens of billions or more) and long training and fine-tuning times (on order of days to weeks). For these reasons GANs are difficult to incorporate practically into mobile and wearable systems where memory is constrained and real-timeliness is valued.

All of the works mentioned above focus on over-the-air audio super resolution. For speech collected by OTA microphones, it is common for ambient noises to pollute and interfere, often reducing its quality and intelligibility. Moreover, low resolution signals used to train these methods are often generated through low pass filtering clean high resolution speech, before downsampling or decimating. Signals generated in this way are easier for the model to learn because aliasing can be avoided. However, this still requires the system to sample at high frequency. A system looking to leverage the benefits of low resolution signals by reducing the sampling rate to lower power consumption, cannot take advantage of this scheme.

2.2 Multi-modal and Vibration-based Speech Enhancement

Speech enhancement aims to restore high quality speech from degraded speech signals that is often contaminated by noise, echoes, or other interference. The primary objective of speech enhancement is to improve speech clarity, intelligibility, and user satisfaction. Speech super resolution, which aims to reconstruct high frequency formants of a person’s voice to improve audio quality, is often viewed as a class of speech enhancement. For OTA audio, ambient background noises can often be heard on the microphone. Another line of speech enhancement focuses on denoising or sound source separation to remove background sounds with minimal impact to the sound quality of the target sound [7, 22, 35–37, 49, 57–60]. Other works, combine information from multiple sensing modalities, such as a camera or ultrasonic sensors, to extract higher quality speech [11, 21, 42, 63].

Head-worn bone conduction and vibration-based sensors, such as accelerometers, IMUs, and BCMs, are naturally insensitive to sounds propagating over-the-air. However, high frequency speech components attenuate much more severely through a user’s bone and skin than over-the-air. As such, speech enhancement for vibration-based acoustic modalities primarily consists of estimating high frequency formants of speech. Recent works have employed noise-free, but low resolution, speech information from bone conduction sensors or accelerometers in conjunction with high resolution signals from OTA microphones, that may experience interference from ambient sounds, to extract speech and remove background noises. [20] leverages bone conduction sensors to enhance high resolution signals from OTA microphones, while [56] leverages an IMU to do the same. This work leverages vibration-based sensing modalities, that naturally attenuate background noise, to enhance speech, and focuses on generating and enhancing speech with only a single speech sensing modality at a time.

There are several works that explore generating high quality speech signals using only vibration-based sensors on mobile and wearable platforms [20, 30, 41]. However, there is a large performance gap between the methods proposed for these platforms (encoder-decoder or U-Net architectures) and state-of-art GANs, which are not as practical to deploy on mobile and wearable platforms. TRAMBA addresses this performance gap.

3 AUDIO SUPER RESOLUTION ARCHITECTURE DESIGN

3.1 Opportunities and Challenges

Figure 2 compares the performance and efficiency between TRAMBA and six different state-of-art audio super resolution methods, which can broadly be categorized into two approaches: 1) U-Net models and 2) GAN models.

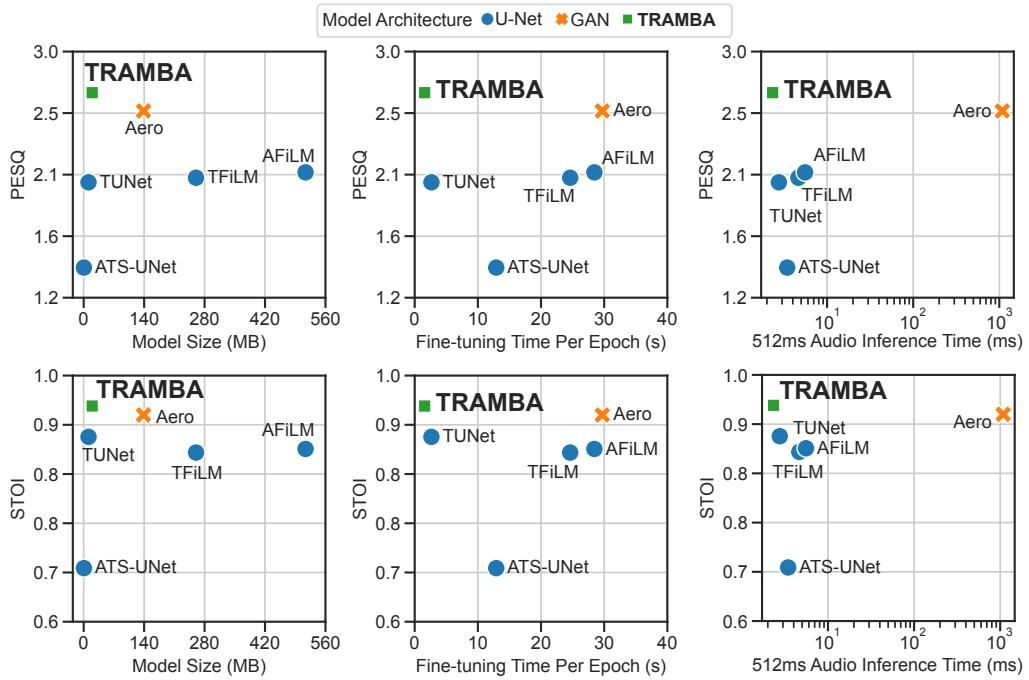


Fig. 2. Comparison of performance (PESQ and STOI) vs efficiency (memory footprint, fine-tuning time, inference time) of TRAMBA and state-of-art audio super resolution methods.

U-Net models consist of contracting layers, typically containing convolutional layers, that compress the input into a lower dimensional latent space before expanding back out. This allows the network to learn important low dimensional features common across all training samples to generate higher frequency content in the expansive layers. GANs consist of a generator network that attempts to learn and generate audio waveforms that are indistinguishable from real data by a discriminator network. From Figure 2, we see several trends and challenges impeding practical adoption of speech super resolution for bone conduction sensors, which TRAMBA addresses.

3.1.1 Heavy Performance and Performance Gap. The general trend is that the best performing models (typically GANs) require orders of magnitude more memory (tens of billions vs. billions of parameters and hundreds of MBs vs MBs), inference time (hundreds of ms vs. ms), training time (days or weeks vs. minutes), and fine-tuning time than U-Net architectures. However, pure U-Net architectures that are more amenable to resource constrained mobile and wearable platforms see significantly worse performance (PESQ and STOI in Figure 2).

3.1.2 Data Scarcity. Training TRAMBA or any bone conduction-based super resolution method requires inputs collected from the BCM or accelerometer, paired with speech collected from OTA microphones that can capture higher frequency formants and act as ground truth. This paired data is not readily available, unlike pure speech or audio collected that can be collected with only a single microphone. Many works that specifically focus on speech enhancement for bone conduction or vibration-based microphones make commendable efforts to collect some amount of data from several volunteers to train/test models [30], which is hardly enough to train a generalizable model. Other works artificially simulate bone conduction microphones and learn transformations

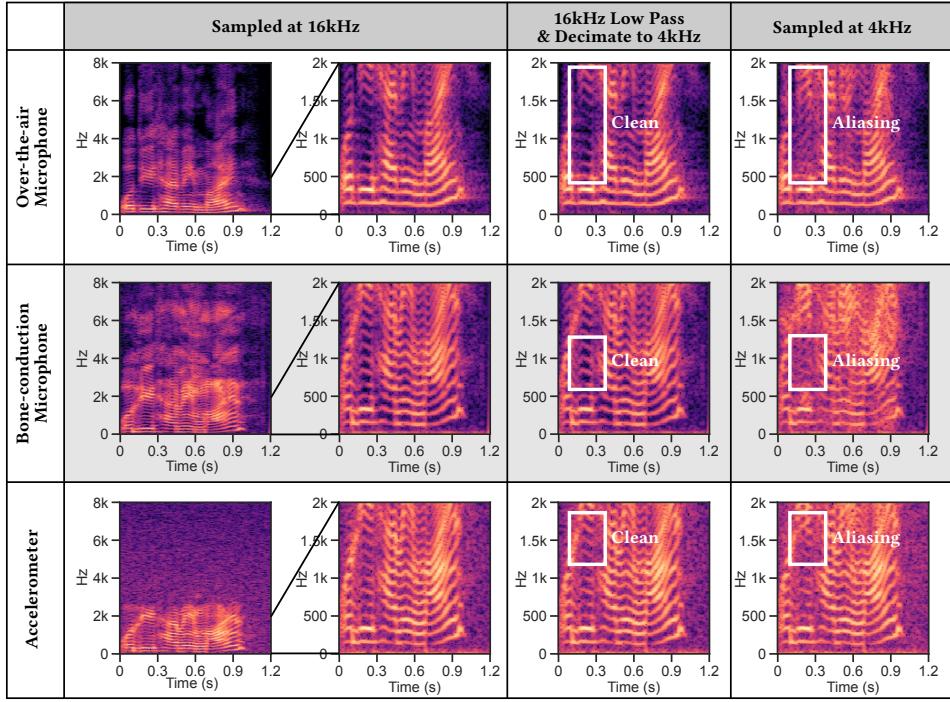


Fig. 3. Comparison of OTA microphone, BCM, and accelerometer recorded audio under different sampling and filtering schemes. In the low pass filter + decimate scenario, a 100 order Butterworth filter with a 2kHz cut-off frequency was applied.

to convert large datasets of audio collected by OTA microphones into signals that you might observe on an IMU or BCM [20], which is used to train the model. However, such methods need to be carefully tuned to generate data that appropriately captures the characteristics of signal observed by an ACCEL or BCM at different sensor locations and different people. Methods trained in this way have difficulty generalizing to new users and sensor placements.

On the other hand, works that focus on over-the-air audio super resolution can train only on standard speech and audio datasets [19]. These works can easily generate realistic inputs by manually attenuating higher frequency components by applying low pass filters, allowing them to take advantage of the vast amount of publicly available audio. We find that directly applying models trained in this way on OTA speech signals produces poor super resolution results (Section 5.3.6).

3.1.3 System Design Opportunities. There are several system design opportunities that we explore, while integrating TRAMBA into a mobile and wearable platform.

Real-Timeliness. The right column of Figure 2 shows the performance of each method vs. the inference time on a 512ms window on an NVIDIA L40 graphics processing unit (GPU). We see that the past best performing models (GANs) require two orders of magnitude longer to perform inference. Some of the models require longer than 512ms to process the input, meaning it cannot be integrated into a real-time system. TRAMBA sees the best performance, while requiring an inference time that is on par with the smallest and fastest running models.

Power Consumption. The general trend in energy-efficient systems research is to perform computation as close to the edge as possible (e.g., on the wearable or earable). The outputs in many detection, classification, and

recognition tasks is much lower in dimension than the inputs (high resolution raw audio), which significantly reduces the amount of data that needs to be transmitted. Conversely, in this work, we aim to produce high resolution audio (16 kHz). Running the model directly on the wearable would require it to transmit full high resolution audio to the mobile platform, meaning power consumption from wireless transmission cannot be reduced. Moreover, current general purpose microcontrollers (MCU) and system-on-chip (SoC) platforms, that are often used to implement earables and wearables, cannot support models larger than tens or hundreds of KB, which limits the performance of models that can be incorporated [1, 2].

Sampling Rate. The left-most column in Figure 3 shows spectrogram plots of speech recorded from a microphone, BCM, and accelerometer at 16kHz. We see that most of the higher frequencies above 2kHz for both the BCM and accelerometer have extremely low energy. This suggests that we can aggressively reduce the sampling rate from the ACCEL and BCM to reduce power consumption of sampling. If model inference and fine-tuning occurs on the mobile platform, the wearable can also reduce the activity of its wireless radio and significantly improve battery life. Moreover, a mobile platform like a smartphone typically has more compute and battery life than a wearable, allowing us to leverage bigger and more powerful models.

However, audio super resolution works operate on audio sampled at the full sampling rate (e.g., 16 kHz), and attempt to fill in higher frequencies [30]. Works that operate on lower sampling frequency inputs tend to low pass filter before decimating the signal. This still requires sampling at a high data rate, which does not improve power consumption on the sensing front. To achieve savings on both sampling and transmission fronts, we aim for robust acoustic super resolution on signals that are directly sampled at a lower frequency rather than filtered signals that have been decimated from higher sampling rates. The right three columns of Figure 3 shows the difference between these two schemes when decimating to or sampling at 4kHz. Although the higher frequencies are not as prevalent in vibration-based sensing modalities, aliasing still occurs. The filter + decimate scheme removes aliasing, but is unrealistic in scenarios aiming to reduce power by directly sampling at a lower rate. We show in Section 5.3 that TRAMBA handles aliasing better than other methods.

3.2 Deep Neural Network Architecture

Figure 4 shows our super resolution architecture. At a high level our architecture adopts a modified U-Net architecture and incorporates self-attention in the downsampling contracting and expanding layers, as well as Mamba in the narrow bottleneck layer, as shown in Figure 4.

Our architectural choices are motivated by the need to balance performance and efficiency in processing time-series data. The U-Net structure allows for capturing both local and global features, while the incorporation of transformer [54] and Mamba [16] enables better handling of long-range dependencies. The transformer’s multi-head attention mechanism is designed to capture complex patterns in high-dimensional data. Residual connections and normalization techniques help maintain stable gradient flow, preventing vanishing or exploding gradients. Mamba introduces a selection mechanism that enables the model to prioritize relevant input, similar to the attention mechanisms in transformers. This makes Mamba particularly well-suited for tasks involving audio processing. Additionally, Mamba has been shown to achieve similar performance as transformers while requiring only half the memory.

3.2.1 Preprocessing. TRAMBA processes 512ms windows of single-channel audio. An accelerometer measures acceleration across three axes. To preprocess acceleration, we subtract the DC offset such that each axis is zero mean. Then, we average the magnitudes of the three axes together to use as input.

3.2.2 Down-Sampling Block. Our model comprises a total of three downsampling blocks, each of which contains a 1D convolutional layer that feeds into a layer of LeakyReLU activations, similar to other U-Net models. Each downsampling block contains 2^{5+b} convolutional filters and a stride of 4. The convolution kernel sizes of each

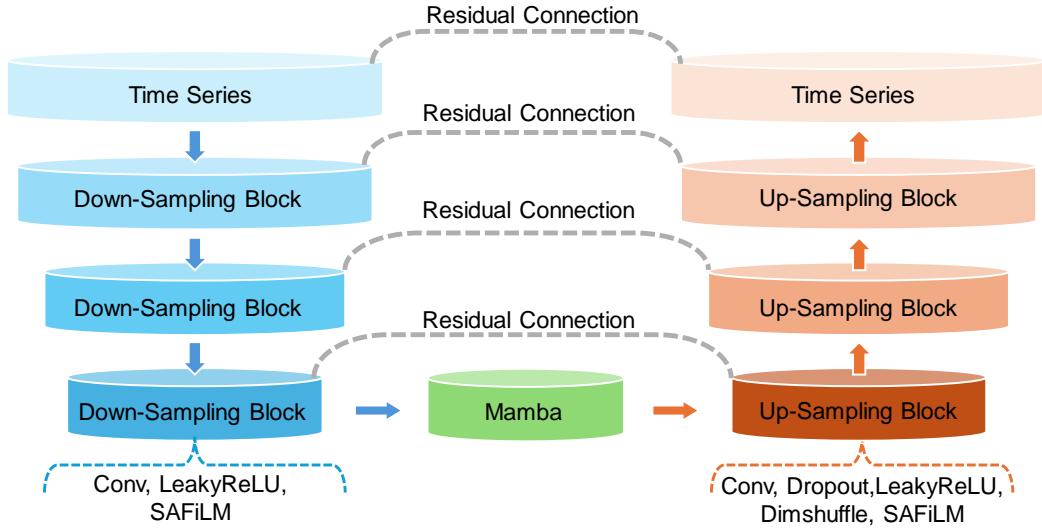


Fig. 4. Super resolution and enhancement architecture.

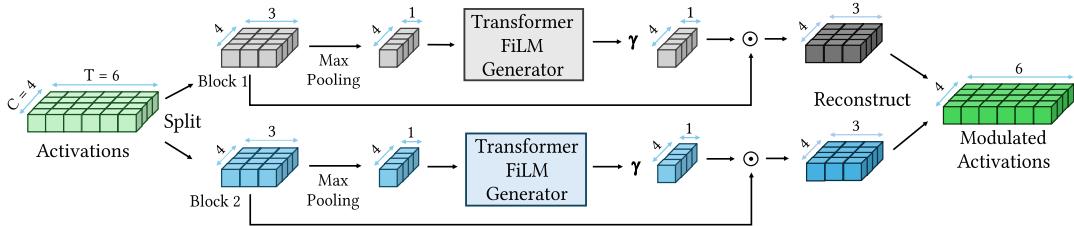


Fig. 5. SAFiLM architecture.

downsampling block are 65, 17, and 7 respectively. Unlike previous works, we apply a novel conditioning layer at the end of each block, called Scale-only Attention-based Feature-wise Linear Modulation (SAFiLM), which we introduce next. The progressive increase in the number of filters allows the network to capture increasingly complex features at each level of abstraction. The decreasing kernel sizes are chosen to gradually reduce the receptive field, focusing on more localized features as we go deeper into the network.

3.2.3 Scale-only Attention-based Feature-wise Linear Modulation (SAFiLM). General purpose conditioning, which aims to learn an affine transformation consisting of a scaling factor (γ) and a shift (β) to apply on intermediate layers in a neural network, has been shown to greatly improve the performance of many tasks. Feature-wise Linear Modulation [46] is one of the most commonly used methods, which uses a recurrent neural network (RNN) to learn these parameters. AFiLM [48] replaced the RNN with a transformer block, containing self-attention that better captures long-range dependencies. The transformer's multi-head attention mechanism enables it to learn intricate characteristics of high dimensional input data, while the residual connections and normalization strategies allow gradients to more easily propagate through the layers without vanishing or exploding, which

facilitates faster training and better performance even with deeper architectures compared to traditional sequence models like RNNs.

However, transformers are memory intensive. To reduce the memory footprint, SAFiLM (Figure 5) only learns and applies the scaling factor, γ , compared to AFiLM, which learns and applies both a scale and shift across blocks. T and C are the total feature length and number of channels resulting from the 1D convolution. The initial activations F , shaped as $\mathbb{R}^{T \times C}$, are segmented into B individual blocks. Each block is represented as F_{block} and undergoes dimensionality reduction along its temporal axis via max pooling to form F_{pool} . Scaling for block b in SAFiLM is expressed as in Equation 1. The activations are then modulated as in Equation 2, where t and c are the feature and channel index resulting from the 1D convolution respectively.

$$\gamma[b, :] = \text{TransformerBlock}(F_{\text{pool}}[b, :]) \quad (1)$$

$$\text{SAFiLM}(F_{\text{block}}[b, t, c]) = \gamma[b, c] \times F_{\text{block}}[b, t, c] \quad (2)$$

3.2.4 Bottleneck. Recently, state space sequence models (SSMs) [18], especially structured state space sequence models (S4) [17], have shown state-of-art performance in a variety of tasks and modalities, including long sequence time series. The Mamba model [16] builds on top of S4 by incorporating a selection mechanism that allows the model to choose relevant information based on the input, which mimics attention mechanisms found in transformers. This enhancement makes Mamba particularly adept at handling pattern-intensive tasks like language processing. Moreover, Mamba matches the performance of transformers with at least two times less memory. Motivated by these strengths, we decide to incorporate Mamba as our bottleneck rather than a transformer.

Given the benefits of Mamba over transformer, it is natural to replace all transformer blocks (SAFiLM) with Mamba. The incorporation of Mamba in the bottleneck allows us to effectively process long sequences and capture complex patterns while maintaining a low memory footprint. This choice contributes significantly to our model’s performance across various metrics. However, we find that doing so often causes the gradients to vanish, so we keep transformers in the downsampling and upsampling blocks. We plan to explore reasons and solutions for this behavior in future work.

3.2.5 Up-Sampling Block. As with the downsampling blocks, our model also comprises three upsampling blocks, each of which performs convolution, dropout, LeakyReLU, pixelshuffle, and SAFiLM. The upsampling block preceding the final output only performs convolution and pixelshuffle.

We implement a one-dimensional version of [51], like in [27] for the pixelshuffle layer.

For pixelshuffle layer, we also implement a one-dimensional version of [51] like [27]. If the input of the pixelshuffle layer is $(C_{\text{in}}, T_{\text{in}})$ and the scaling factor is 4, the output is $(C_{\text{in}} \div 4, T_{\text{in}} \times 4)$. The three upsampling blocks contain convolution filters of size 512, 256, and 4, with a stride of 4. The convolution kernel sizes of each convolution layer are 7, 17, and 65, respectively. The upsampling blocks are designed to mirror the downsampling blocks, gradually increasing the spatial resolution while decreasing the number of channels. This symmetry helps in reconstructing the enhanced audio signal while preserving learned features from earlier stages of the network.

3.2.6 Residual Connection. Many U-Net architectures leverage a skip connection between respective downsample and upsample layers that concatenates the outputs from both. We instead leverage a residual connection, which adds the outputs from both. This allows gradients to more easily flow through layers, facilitating training, reducing the risk of vanishing gradients, and enabling the optimization of deeper networks.

3.3 Fine-tuning

As mentioned in Section 3.1, it is labor intensive to obtain large amounts of paired bone conduction data and clean audio to train a robust model for vibration-based speech enhancement. Rather than take this approach, we instead pre-train our model using the VCTK [61] speech dataset before fine-tuning with a small amount of paired data collected from the user at box opening. This procedure is common in many AI voice generation applications that typically require around 15 minutes of voice samples from the user [4]. To pre-train the model, we downsample clean speech to use as input. During fine-tuning, the clean speech can be captured by the user’s mobile phone, while the wearable records from the BCM or ACCEL simultaneously.

4 EVALUATION: PRE-TRAINING AUDIO SPEECH ENHANCEMENT AND SUPER RESOLUTION

First, we discuss the performance of TRAMBA in context of over-the-air speech super resolution. We compare against several state-of-art audio super resolution methods, including four U-Net architectures and two GANs (Section 4.2).

4.1 Training, Validation, Testing Procedure and Metrics

Dataset. For all methods, we trained, tested, and validated on the VCTK dataset [61], which contains a corpus of spoken language comprising 109 native English speakers of various regional accents. Each speaker reads out approximately 400 different sentences, totaling approximately 44 hours of data. We use audio from 100 individuals for training and the remaining 9 individuals for testing.

Preprocessing and Training. For OTA speech enhancement, we downsample high resolution clean audio to 4kHz to use as input to learn high resolution 16kHz audio. Training data was divided into smaller segments with a window size of 512ms and 50% overlap. For all other methods, we applied the preprocessing steps mentioned the respective papers. We train TRAMBA for 30 epochs. After this time period, we notice that the training, testing, and validation loss changes very little. We trained all other methods using their default training parameters specified in their respective papers and code.

Loss Function. In our experiments, we found that using mean squared error (MSE) as the loss function generated lower quality audio than mean absolute error (MAE). As such, we adopted MAE as our primary loss function. However, leveraging MAE, or any type of sample-by-sample distance metric (e.g., MSE), does not necessarily guarantee perceptual quality [40]. As such, we also incorporate multi-resolution STFT loss [62]. We leverage the same multi-resolution STFT loss parameters as [10, 38]: FFT bins $\in \{512, 1024, 2048\}$, hop lengths $\in \{50, 120, 240\}$, and window sizes $\in \{240, 600, 1200\}$. Our final loss is the sum between MAE and the mutli-resolution STFT loss.

Performance Metrics. The primary metrics we use to compare TRAMBA for this section and throughout the rest of the paper are discussed next.

Signal to Noise Ratio (SNR): SNR is a ratio between the energy of a target signal and background noise, calculated using the formula:

$$\text{SNR}(x, y) = 10 \log_{10} \left(\frac{\sum x^2}{\sum (x - y)^2} \right)$$

x represents the magnitude spectrum of the original (high resolution) signal, and y denotes the magnitude spectrum of the processed signal.

Perceptual Evaluation of Speech Quality (PESQ): In the context of audio super resolution, PESQ [50] serves as an objective measurement to evaluate the quality of reconstructed high-resolution speech signals as perceived by human listeners. This metric compares the generated audio against a high-quality reference, with scores typically ranging from -0.5 to 4.5. A higher PESQ score indicates a closer distance to the original audio’s quality.

Log-Spectral Distance (LSD): This metric measures the distance between the spectrum of the generated speech and the clean speech on the log scale. The difference is measured between log spectra because human hearing follows this scale [15]. Scores are calculated over all frequency bins, with lower values indicating a closer match to the reference spectrum, as shown below.

$$\text{LSD}(x, y) = \frac{1}{T} \sum_{t=1}^T \sqrt{\frac{1}{K} \sum_{k=1}^K (\log X(t, k) - \log \hat{X}(t, k))^2}$$

$X(t, k)$ denotes the magnitude spectrum of the original signal at time frame t and frequency k . $\hat{X}(t, k)$ represents the magnitude spectrum of the super-resolved signal at the same time frame and frequency. T indicates the total number of time frames, and K is the total number of frequency bins.

Short-Time Objective Intelligibility (STOI): STOI [53] compares the clarity and intelligibility of speech enhanced from a lower to a higher resolution against a clean reference, with scores ranging from 0 to 1. A higher STOI score, closer to 1, signifies better intelligibility, indicating that the enhanced speech is easier to understand.

4.2 Models Compared

We compare against six additional models. Four of the models are UNet [34] structures, while the remaining two employ GAN [12] structures.

- TFiLM [5] is a UNet model that uses a recurrent neural network to alter the activations of the convolutional layers in the standard U-Net.
- AFiLM [48] replaces the RNNs in TFiLM with self-attention.
- TUNet [43] uses the Performer [9] architecture as the bottleneck of the UNet model.
- ATS-UNet [30] is a lightweight UNet model that reduces the number of parameters in the convolutional layers.
- EBEN [19] is a GAN for audio speech super resolution with a lightweight generator model.
- Aero [38] is a GAN that operates directly on the complex-valued frequency representation of audio. It employs two independent channels to process amplitude and phase information, circumventing the Spectral discontinuity and phase mismatch associated with high and low frequency splicing that are inherent to traditional methods.

4.3 Performance Summary

	Params (M)	Model Size (MB)	Inference Time (ms)	SNR	PESQ	STOI	LSD
Raw Audio (4kHz)	-	-	-	20.8635	1.8674	0.8195	2.8487
TFiLM (4kHz→16kHz) [5]	68.2	260.3	4.6423	13.7079	2.4349	0.8475	1.8667
AFiLM (4kHz→16kHz) [48]	134.7	513.9	5.5552	11.3508	2.4658	0.8689	1.8665
TUNet (4kHz→16kHz) [43]	2.9	11.2	2.7550	22.3046	2.5783	0.9382	0.9429
ATS-UNet (4kHz→16kHz) [30]	0.1	0.5	3.4419	12.1807	1.5471	0.6695	1.5624
EBEN (4kHz→16kHz) [19]	29.7	113.3	242.1341	18.9196	2.5800	0.8912	1.0079
Aero (4kHz→16kHz) [38]	36.3	138.7	1084.5392	22.5075	3.0136	0.9386	0.8102
TRAMBA (4kHz→16kHz)	5.2	19.7	2.3309	23.2426	3.2344	0.9478	0.8271

Table 1. Summary of performance of TRAMBA compared to existing audio super resolution methods. Inference time is measured on an NVIDIA L40 GPU processing one window of 512 ms.

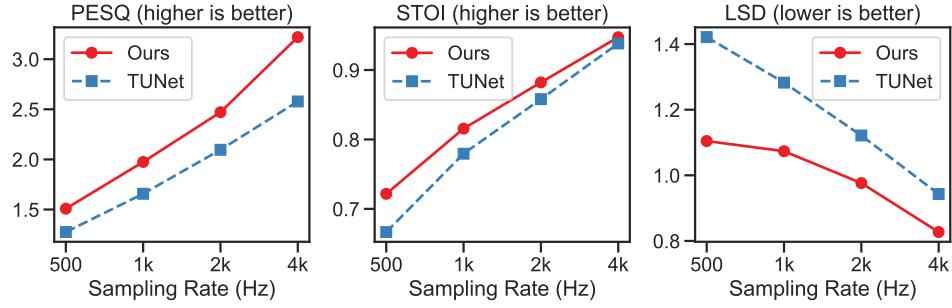


Fig. 6. Effects of the sampling rate on speech super resolution.

Table 1 summarizes the test performance across all methods. We see that although some methods did not show improvements in SNR, they demonstrated significant improvements in audio quality in metrics such as PESQ, STOI, and LSD, reflecting that the upsampled audio showed better quality compared to the original low-frequency audio. TRAMBA (highlighted) has the best performance across all metrics and sampling rates, while requiring similar processing times and memory as U-Net models. The Aero GAN method slightly outperforms TRAMBA in the LSD metric. LSD measures signal similarity more than perceptual quality and TRAMBA outperforms Aero in all other perceptual and noise metrics (SNR, PESQ, STOI). This demonstrates the addition of transformers and Mamba improved the modeling and generation of local speech formants than traditional U-Net architectures and encoder-decoder architectures. This also demonstrates that transformer and Mamba-based architectures can outperform state-of-art GANs with only a fraction of the memory and inference time. Both GANs require significantly more time to perform inference on an L40 GPU. Aero requires more time than the length of the window size it processes (1080ms vs. 512ms), meaning it cannot run in real-time. For these reasons, we compare with the best performing U-Net architecture (TUNet) for the rest of this section.

4.4 Reducing Sampling Rate

Figure 6 shows the super resolution performance of TRAMBA to 16kHz as the sampling rate of the input signal varies from 500Hz to 4kHz. Even as the sampling rate decreases down to 500Hz, TRAMBA outperforms TUNet. Even if the sampling rate decreases down to 2kHz or 1kHz, TRAMBA achieves comparable performance to other models at higher sampling rates.

4.5 Super Resolution Architecture Ablation Study

	PESQ	STOI	LSD	Parameter	Train Time Per Epoch
Replace Mamba with Performer	3.1208	0.9474	0.8042	6.5M	243.93s
Remove SAFiLM	2.8056	0.9428	0.8268	2.6M	114.04s
TRAMBA	3.2344	0.9478	0.8271	5.2M	191.61s

Table 2. Ablation study.

Table 2 shows the results of our ablation study, which explores the impact of combining self-attention, transformers, and Mamba on speech enhancement. Replacing Mamba in our bottleneck layer with Performer, a variant of transformers, yielded similar performance metrics (PESQ, STOI, LSD). However, the training time increased significantly from 191.61 seconds to 243.93 seconds per epoch and the number of parameters increased

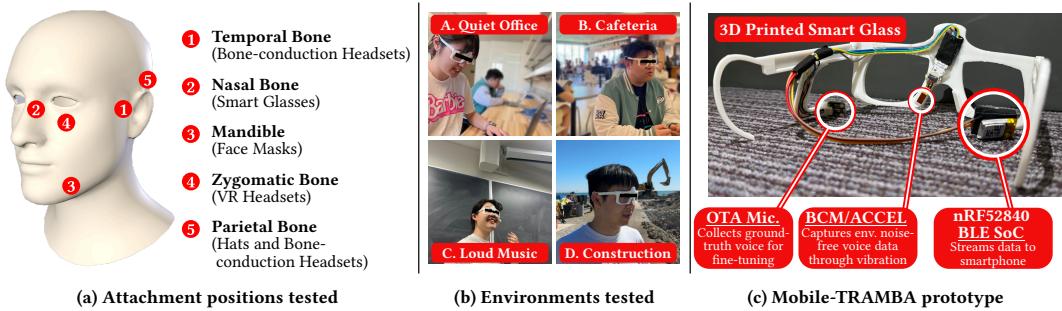


Fig. 7. Data collection setup, environments deployed, and wearable prototype.

to 6.5M. Mamba not only maintains performance but also significantly increases model training speed and reduces the required number of model parameters.

Next, we removed attention (SAFiLM) from the contracting and expansion layers and observed that the training time per epoch decreased by more than 50%. However, the LSD, STOI, and particularly PESQ dropped significantly. This performance loss highlights the importance of transformers and self-attention in enabling models to learn and generate complex patterns of speech. The results of our ablation study demonstrates the importance of both self-attention and Mamba in creating a robust model for generating high quality speech, while being efficient in time and memory footprint.

5 EVALUATION: FINETUNING FROM AUDIO SUPER RESOLUTION TO ACCEL AND BCM SPEECH ENHANCEMENT

In this section, we discuss the performance of TRAMBA on enhancing speech captured by accelerometer readings and bone conduction microphones on areas of the face where (future) wearables are commonly worn. As discussed in Section 3.1 collecting ground truth pairs (ACCEL/BCM, audio) is extremely labor intensive. As such, we only fine-tune our over-the-air audio speech super resolution model (Section 4) with a small amount of paired data collected from an individual, which is similar in procedure to many existing AI voice generation applications [4].

5.1 Data Collection

We recruit 7 volunteers (4 males, 3 females) between 18-30 years old. All studies are approved by the Human Research Protection Office and IRB at Columbia University. We collect data from each volunteer, placing the ACCEL or BCM at each location marked in Figure 7a. At each location, we collect a total of approximately 20 minutes of data. We instruct each volunteer to read a set of sentences from an e-book on basic spoken English [6]. We vary the amount of data used to fine-tune TRAMBA to analyze the effects of the quantity of data on performance.

Figure 7c shows the glasses wearable we created to both collect data and deploy in various environments. There is a holder for the ACCEL/BCM to collect inputs and a holder for an over-the-air microphone to capture high resolution ground truth speech. For the BCM, we use the V2S200D voice vibration sensor [26] and the IM-42352 for the MEMS accelerometer [24].

5.2 Effects of Preprocessing: Downsampling

As discussed in Section 3.1, most works that develop super resolution methods for both over-the-air and vibration-based speech sensing 1) still sample at the non-decimated sampling rate (e.g., 16kHz), or 2) perform low pass

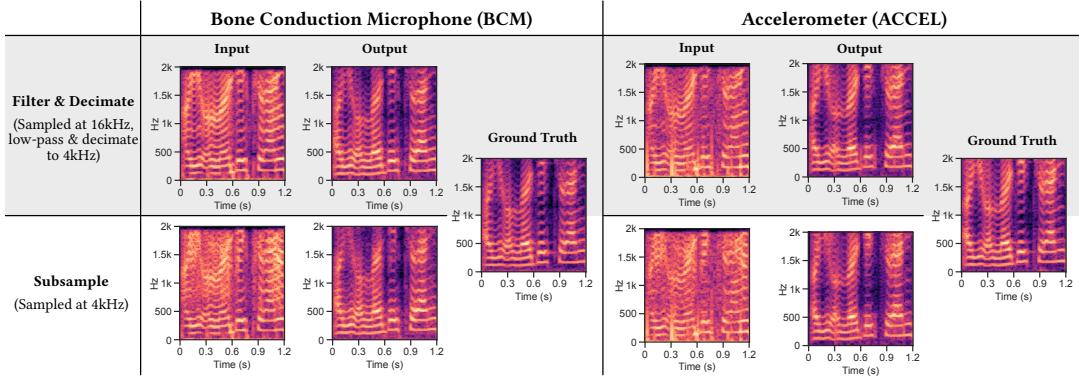


Fig. 8. Effects of downsampling.

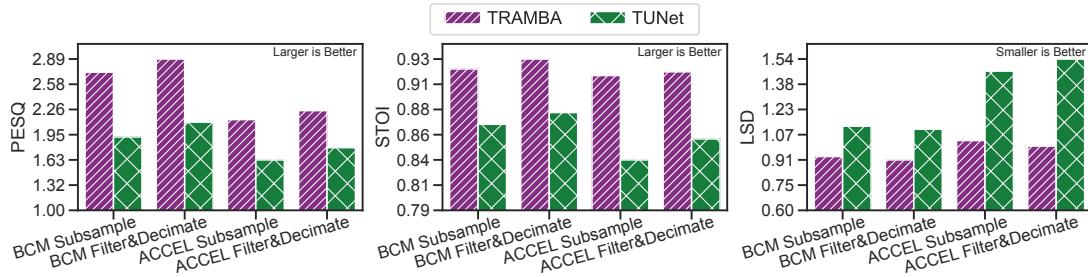


Fig. 9. Downsampling effects on performance.

filtering before subsampling. In both scenarios, aliasing is avoided or significantly reduced, but we cannot reduce power consumption by sampling at a lower rate. As such, we directly 3) decimate speech signals, without any filtering, for an accurate representation of a slower sampling rate.

Figure 8 and Figure 9 compare the performance of TRAMBA after downsampling from 16kHz to 4kHz with the two downsampling approaches, 2) and 3). For method 2), we apply a low pass filter with a cut-off frequency of 2kHz before subsampling. We see that TRAMBA outperforms other methods across all metrics. Although there is a drop in performance for TRAMBA if we remove filtering before decimation (3), TRAMBA still outperforms most other methods even if the other methods filter before decimating (2). Additionally, the performance drop-off is significantly steeper for other methods. For the rest of this section, we directly decimate audio signals, 3).

5.3 BCM and ACCEL-based Speech Enhancement

Table 3 summarizes the performance of TRAMBA for speech super resolution on BCMs and ACCELS, subsampled at 4kHz to 16kHz sampling rate that most works benchmark. The reported values represent the average performance across all participants. Each model was fine-tuned for individual volunteers using 15 of data collected at Location 2 (Nasal Bone) as shown in Figure 7 until convergence. Compared to most other methods, TRAMBA achieved the best performance while requiring orders of magnitude less fine-tuning (30-40 minutes vs. hours and days). TUNet required only 5 minutes to converge but performed significantly lower across all other metrics.

Many of these models require hours to converge, which is impractical in more realistic scenarios. As such, we explored how performance would be impacted by limiting the fine-tuning time to less than an hour. We chose

Method	BCM				ACCEL			
	PESQ	STOI	LSD	Convergence Time	PESQ	STOI	LSD	Convergence Time
TFiLM	2.5800	0.9117	1.3737	674min (1643epoch)	1.8565	0.9146	1.5197	1219min (2971epoch)
AFiLM	2.5112	0.9012	1.4832	129min (272epoch)	1.7593	0.8767	1.6788	1082min (2281epoch)
TUNet	1.9195	0.8701	1.1220	5min (113epoch)	1.6337	0.8366	1.4649	4min (90epoch)
ATS-UNet	1.4387	0.7424	1.5321	472min (2194epoch)	1.4877	0.8019	1.4535	608min (2826epoch)
Aero	2.6762	0.9036	0.8987	1370min (2768epoch)	2.0934	0.9086	1.0727	1400min (2828epoch)
TRAMBA	2.7273	0.9221	0.9329	34min (1291epoch)	2.1365	0.9159	1.0326	43min (1633epoch)

Table 3. Performance summary of speech enhancement from 4kHz to 16kHz. Models were fine-tuned using 15 minutes of data collected at Location 2 (Nasal Bone) until convergence. Models were fine-tuned on an NVIDIA L40 GPU.

Method	Fine-tuning Time Per Epoch	BCM				ACCEL			
		PESQ	STOI	LSD	WER	PESQ	STOI	LSD	WER
TFiLM	24.62s	2.1390	0.8769	1.3342	8.73%	1.6577	0.8486	1.4760	12.89%
AFiLM	28.45s	2.2618	0.8784	1.4657	6.48%	1.6220	0.8422	1.7090	15.66%
TUNet	2.65s	1.9195	0.8701	1.1220	6.21%	1.6337	0.8366	1.4649	13.29%
ATS-UNet	12.91s	1.4155	0.7398	1.5370	29.86%	1.2981	0.6356	1.5979	71.80%
Aero	29.70s	2.6669	0.9072	0.8917	4.48%	2.0172	0.8993	1.1009	14.04%
TRAMBA	1.58s	2.7273	0.9221	0.9329	2.65%	2.1365	0.9159	1.0326	7.76%

Table 4. Performance summary of speech enhancement from 4kHz to 16kHz. Models were fine-tuned using 15 minutes of data collected at Location 2 (Nasal Bone) for 52 minutes. Models were fine-tuned on an NVIDIA L40 GPU.

to fine-tune each model for 52 minutes because this was the time needed to train TRAMBA for 2000 epochs, significantly past the point of convergence. Table 4 summarizes these results. By limiting the fine-tuning time, we see that the performance difference between TRAMBA and existing methods only grows across all metrics.

One additional metric we compare is the word error rate (WER), which is the ratio between the number of incorrect words perceived and the total number of spoken words. To compute this metric, we input all generated speech signals into Whisper [47], a state-of-the-art automatic speech recognition (ASR) model and compare with ground truth transcript.

Compared to the larger GAN and U-Net methods, TRAMBA requires only a fraction of the time to fine-tune (1.6 seconds vs 30 seconds per epoch). Figure 10 visualizes the BCM and accelerometer reconstruction result of one speech segment compared to the best performing GAN (Aero) and U-Net (TUNet). Both TRAMBA and Aero clearly generate higher quality results than TUNet. TRAMBA rectifies aliased speech better than the GAN, as evidenced at the end of the segment. The formants in the original clean speech stays level or trends upwards. However, the BCM and ACCEL are both aliased when sampled at 4kHz. Aero reconstructs the end of the segment with the formants trending downward in the higher frequencies due to aliasing. However, TRAMBA reconstructs speech much more closely to the original speech pattern.

Next, we explore how TRAMBA performs depending on sensor placement, across different volunteers, and availability of labeled data. For the rest of the section, we compare TRAMBA against method TUNet, as it was the second best performing method overall outside of the GANs. While the GAN architectures generally outperformed the U-Net architectures, they take orders of magnitude longer to train, fine-tune, and perform inference, making them impractical to deploy.

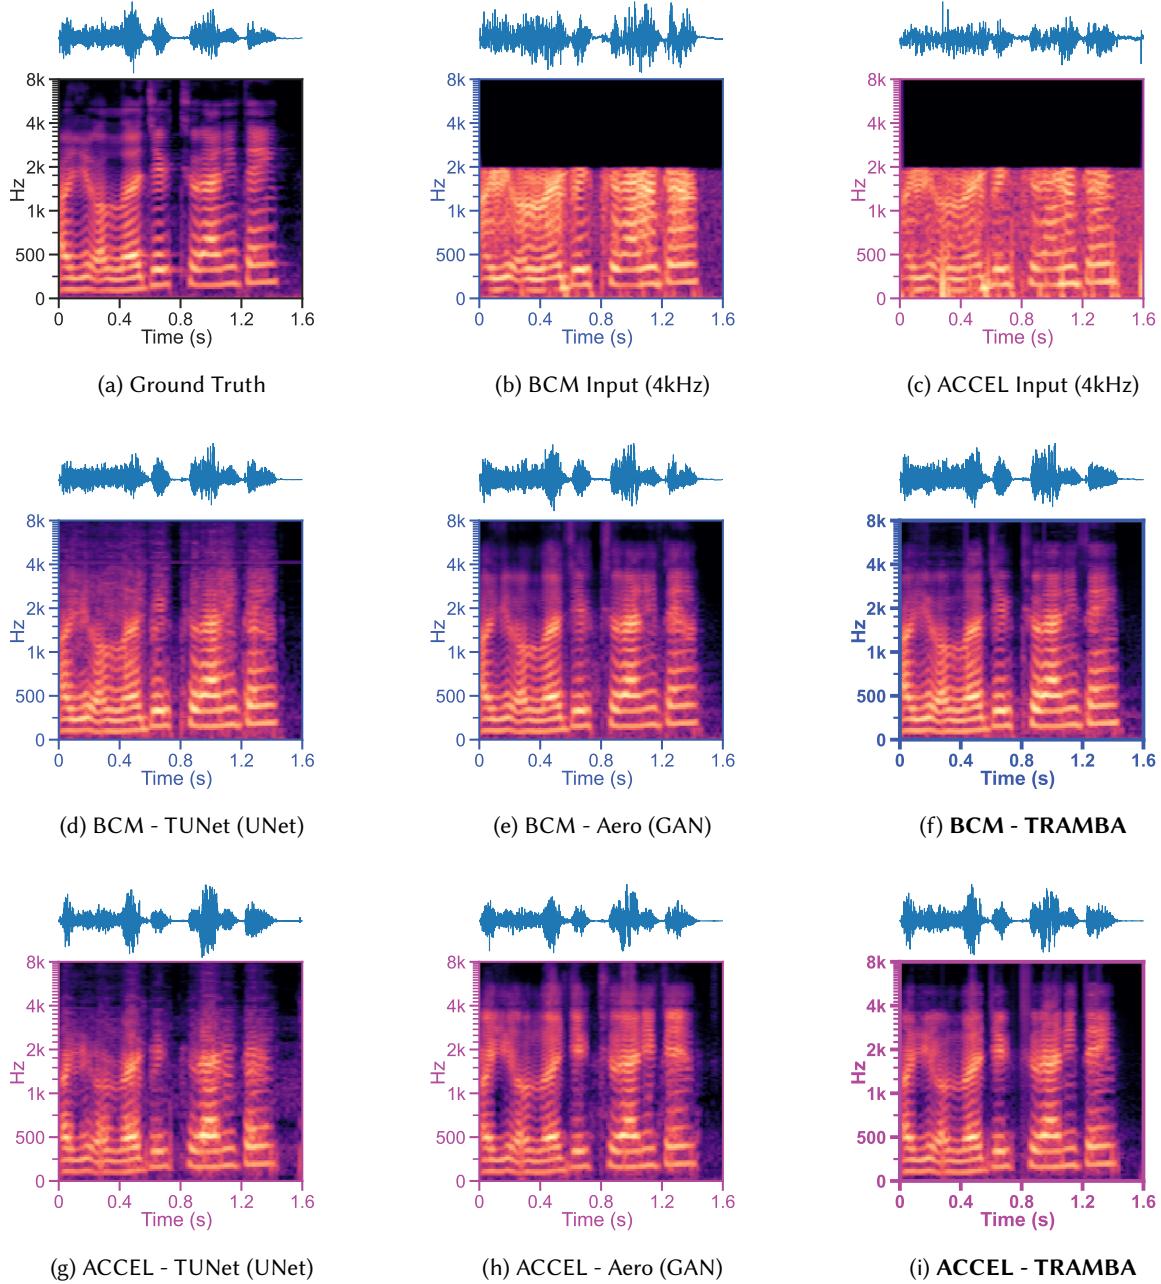


Fig. 10. Visualizing reconstructed audio from ACCEL and BCM after finetuning BCM/ACCEL + audio data pairs. The BCM/ACCEL is placed at location 2 (Nasal Bone). (b) and (c) show no energy at frequencies higher than 2kHz because the BCM and ACCEL were sampled at 4kHz.

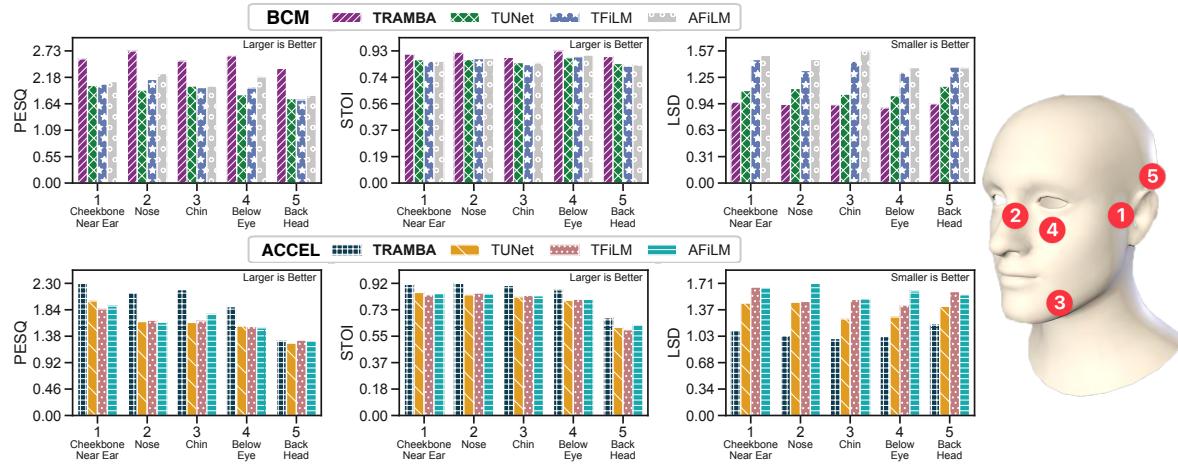


Fig. 11. Performance of BCM/ACCEL speech enhancement at different locations.

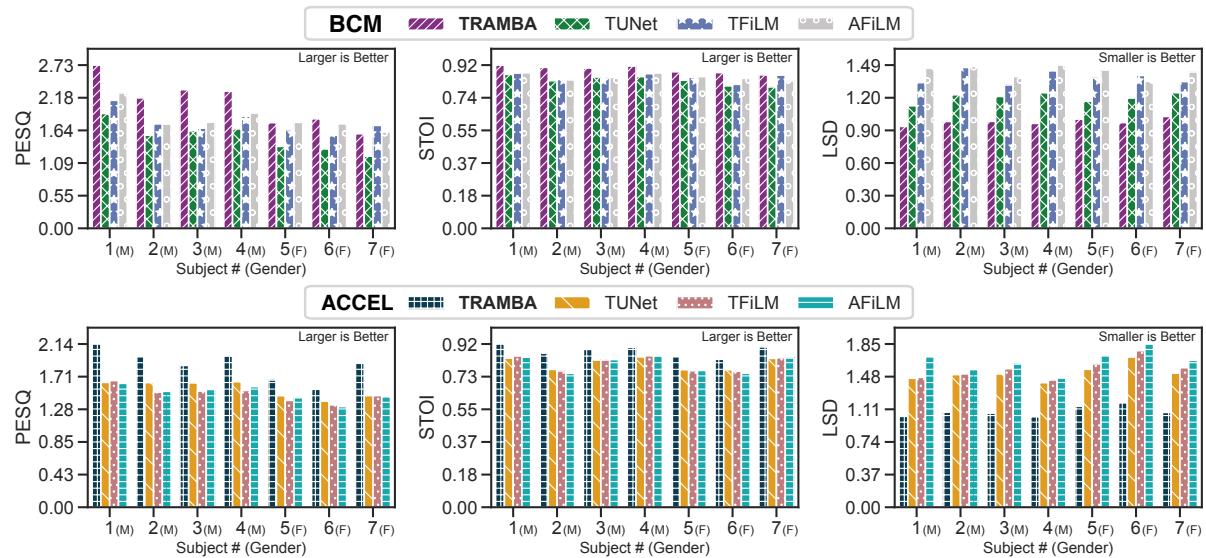


Fig. 12. Performance of BCM/ACCEL speech enhancement breakdown by individuals.

5.3.1 BCM and ACCEL Placement. Figure 11 illustrates the performance breakdown of TRAMBA by BCM and ACCEL placement location. TRAMBA has consistent and better performance across all common and practical locations for a BCM or ACCEL.

5.3.2 Breakdown Across Individuals. Figure 12 shows the performance of TRAMBA broken down by each volunteer. The BCM and ACCEL was placed at location 2 (Nasal Bone). Again, TRAMBA consistently yielded the best performance metrics across all participants. One interesting trend among all methods is that performance

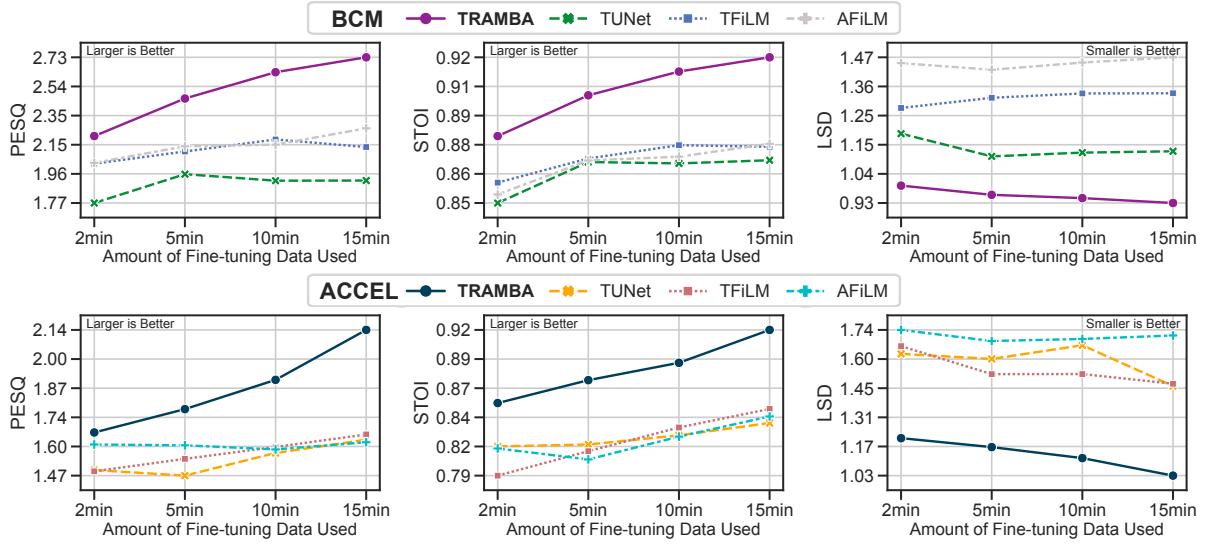


Fig. 13. Performance of BCM/ACCEL speech enhancement varying the amount of data available to fine-tune.

is lower among females than males. We suspect that this is due to their higher frequency voices experiencing greater attenuation through bone and skin. In future work, we plan to explore causes and solutions in more depth.

5.3.3 Data Availability. Figure 13 varies the amount of data we collect from an individual. TRAMBA outperforms all other methods with all fine-tuning dataset size. With only 2 minutes of data, TRAMBA outperforms TUNet across all metrics even if TUNet was fine-tuned with 15 minutes of data collected from the user.

5.3.4 Fine-tuning Time. Figure 14 varies the number of epochs used for fine-tuning up to 2000 epochs, which corresponded to 52 minutes of training TRAMBA. We compare TRAMBA to all other methods trained for the same amount of time. Again, TRAMBA outperforms TUNet at all training times. With only 2.6 minutes (100 epochs) of training, TRAMBA outperforms TUNet trained for 52 minutes across all metrics, except for PESQ, which is only slightly lower.

5.3.5 Sampling Rate. Figure 15 varies the starting sampling rate to perform super resolution to 16kHz. TRAMBA outperforms TUNet across all starting sampling rates. We observe that TRAMBA achieves acceptable speech quality when upsampling from around 2kHz or higher for both BCM and accelerometer. TUNet only achieves similar performance as TRAMBA when the sampling rate is at least double, meaning TRAMBA can enhance and generate higher quality and resolution speech with less information, sampling, and transmission bandwidth.

5.3.6 Unseen Individuals and the Need for Fine-tuning. For completeness, we explore the need for fine-tuning by looking at two cases: (i) directly applying the model pre-trained on audio speech (Section 4) and (ii) creating a model for BCM and accelerometer speech that is generalizable to unseen individuals.

(i) Applying audio speech pre-trained model to BCM and ACCEL. Table 5 shows the performance of the direct application of TRAMBA that is only pre-trained on downsampled audio speech (Section 4) to enhancing BCM and accelerometer speech. PESQ and STOI see very low values, suggesting that training TRAMBA using only OTA audio is not sufficient.

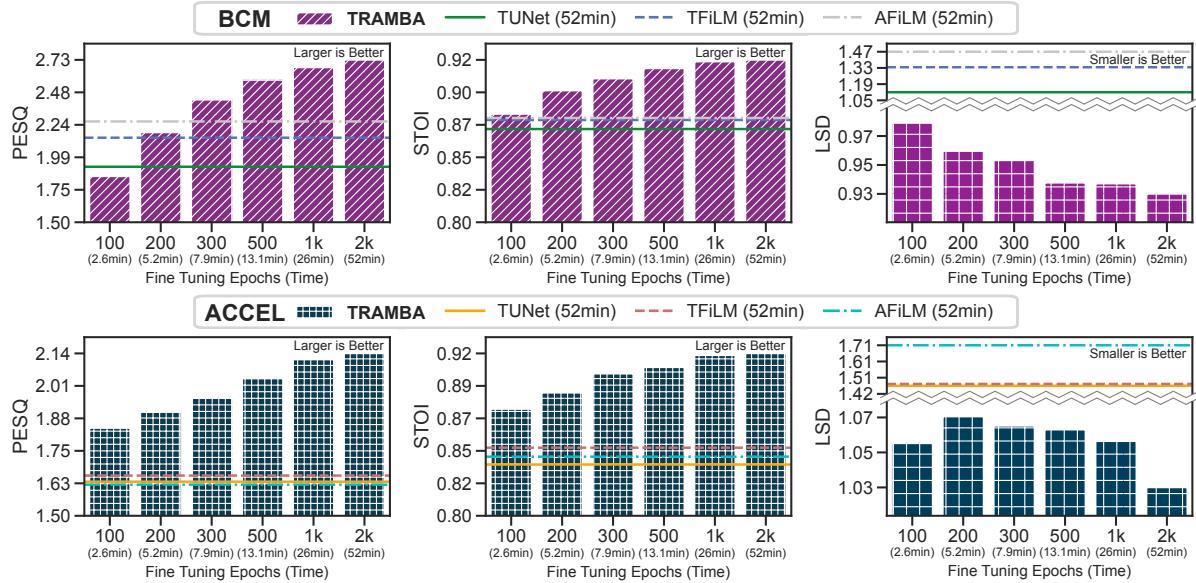


Fig. 14. Performance of BCM/ACCEL speech enhancement varying the number of epochs used to fine-tune.

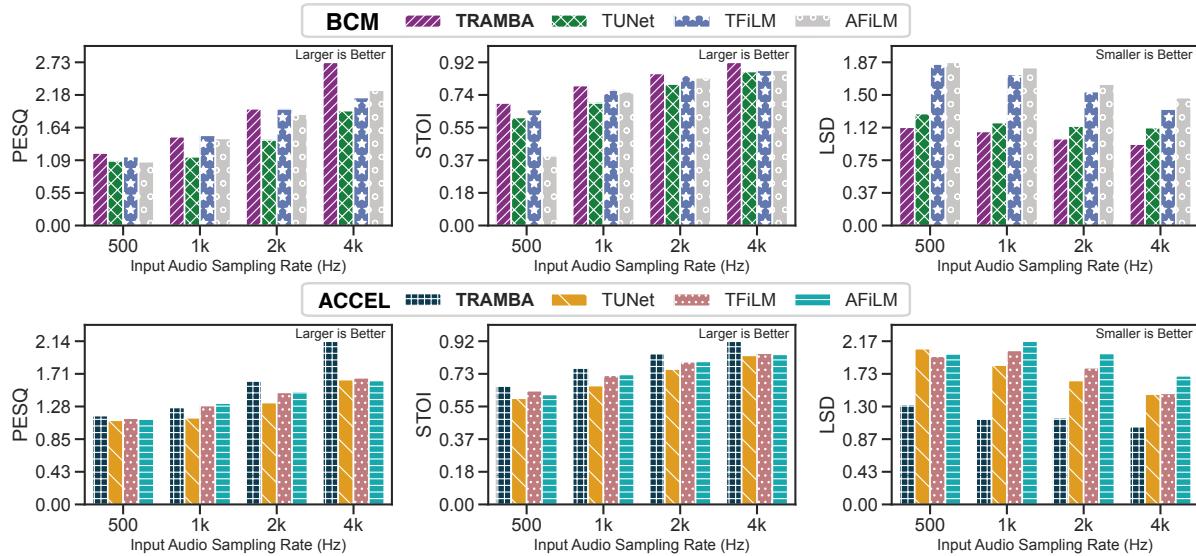


Fig. 15. Performance of speech enhancement, varying the the input sampling rate.

(ii) Creating a generalizable model. Going one step further, we attempt to create a generalizable model by fine-tuning TRAMBA with as big of a dataset as possible. Table 6 compares TRAMBA with baseline models after training the models with most of our volunteers (6 volunteers) and testing on the unseen (1 volunteer). All methods see equally poor performance, suggesting the need for more data from more individuals to create a

	PESQ	STOI	LSD	WER	PESQ	STOI	LSD	WER
BCM - TRAMBA								
(i) Pre-train on audio only	1.0859	0.6515	1.8639	65.51%	1.1141	0.6223	2.2536	74.57%
(ii) Creating generalizable model	1.1614	0.7544	1.1713	34.67%	1.1130	0.7239	1.3392	44.52%
(iii) Fine-tuning to the individual	2.7273	0.9221	0.9329	2.65%	1.9195	0.8701	1.1220	6.21%
ACC - TRAMBA								
(i) Pre-train on audio only	1.1893	0.7095	2.1775	33.33%	1.1590	0.6714	2.1154	18.69%
(ii) Creating generalizable model	1.2926	0.7709	1.3226	37.50%	1.3073	0.7515	1.7991	38.92%
(iii) Fine-tuning to the individual	2.1365	0.9195	1.0326	7.76%	1.6337	0.8366	1.4649	13.29%

Table 5. Performance with fine-tuning and without fine-tuning.

Method	BCM				ACCEL			
	PESQ	STOI	LSD	WER	PESQ	STOI	LSD	WER
TFiLM	1.0827	0.6875	1.5167	73.36%	1.3512	0.7378	1.9940	34.31%
AFiLM	1.0979	0.6903	1.5710	72.99%	1.2900	0.7433	1.9806	39.73%
TUNet	1.1130	0.7239	1.3392	44.52%	1.3073	0.7515	1.7991	38.92%
ATS-UNet	1.0967	0.5717	2.1571	81.15%	1.1371	0.6583	1.7805	76.52%
Aero	1.1557	0.7084	1.2218	54.42%	1.1944	0.7661	1.4720	23.45%
TRAMBA	1.1614	0.7544	1.1713	34.67%	1.2926	0.7709	1.3226	37.50%

Table 6. Performance of BCM/ACCEL speech enhancement on unseen individuals after training with 6 volunteers and testing on an unseen volunteer. The poor performance across all methods suggests that creating a generalizable model with a limited amount of data, due in part to the labor intensive data collection process, is not feasible. Rather, fine-tuning to specific users is currently more practical.

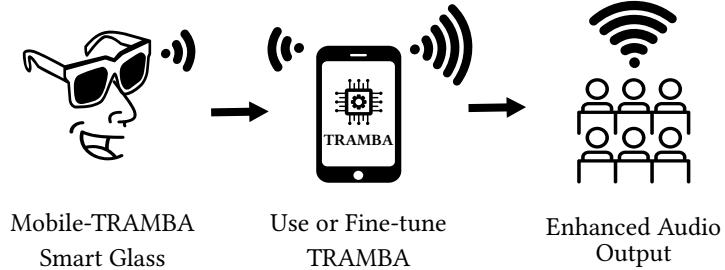


Fig. 16. Architecture of the mobile and wearable platform (mobile-TRAMBA) that leverages TRAMBA to enhance vibration-based speech.

general model. Table 5 further highlights the disparity in performance between simply pre-training on audio (i), attempting to create a generalizable model (ii), and fine-tuning to the individual (iii). Until tools and platforms exist for collecting large amounts of paired audio and vibration data, creating and incorporating a general model is not practically feasible.

6 SYSTEM DESIGN AND IMPLEMENTATION WITH TRAMBA

We incorporate TRAMBA into a vibration-based mobile and wearable system, which we call mobile-TRAMBA for the rest of the paper, shown in Figure 16. Audio is sampled by the head-worn wearable that is transmitted via Bluetooth Low Energy to the smartphone platform or computer, depending on which device the user is currently

Method	Cafeteria (SNR: -0.355 dB)	Loud Music (SNR: -3.742 dB)	Construction Site (SNR: -10.458 dB)
Raw Audio	15.85%	45.12%	86.59%
Conv-TasNet [37]	79.88%	20.12%	85.37%
DCUNet [8]	17.07%	83.54%	86.59%
TRAMBA- BCM	7.93%	6.10%	14.63%
TRAMBA- ACCEL	4.27%	16.46%	10.37%

Table 7. Word error rate in different environments.

using. The smartphone or computer takes incoming data and runs TRAMBA. To implement TRAMBA on a smartphone, we convert our PyTorch model to a TensorFlow Lite (tflite) model by converting to an ONNX model and a standard TensorFlow model as intermediate steps.

Wearable Platform: The wearable prototype we built uses the Seeed Studio XIAO nRF52840 platform, which contains a Nordic nRF52840 Bluetooth Low Energy (BLE) SoC, to transmit speech that is sensed from an accelerometer or a BCM to the smartphone. For our user studies, we built a case for the wearable such that the accelerometer or BCM is situated on the bridge of the nose, which could be possible if the ACCEL or BCM is placed in the nose pads of a glasses wearable, as demonstrated here [41]. We chose location 2 (Nasal Bone) as shown in Figure 7 for our user studies and end-to-end evaluation because we saw the most consistent performance for both the BCM and ACCEL across all methods at this location. However we believe this platform can be adapted to a multitude of locations, as evidenced in Section 5.

Fine-tuning: The first time a user puts on the wearable, they collect up to 15 minutes of their own voice for fine-tuning. The fine-tuning process collects paired OTA speech and vibration-based speech from the smartphone’s microphone and the BCM or accelerometer on the wearable. For the computer version of mobile-TRAMBA, the fine-tuning occurs directly onboard. The computer we used has access to an NVIDIA L40 GPU [45]. Tools for fine-tuning deep learning models on a mobile phone are currently extremely limited. As such, for the smartphone version of mobile-TRAMBA, we transmit the paired data to our GPU-equipped computer to fine-tune and update the model. Once the model is fine-tuned, it is transmitted back to the smartphone, where inference occurs locally. In future work, we plan to explore methods for fine-tuning on-device.

7 END-TO-END EVALUATION

Our recruited volunteers (Section 5.1) fine-tuned the mobile-TRAMBA system with 15 minutes of their own voice before reciting 8 minutes of passages from [6] across various environments with unique noise profiles at different walking speeds. We compare the performance of TRAMBA against over-the-air systems. Namely, we compare against DCUNet [8], a deep learning model for speech denoising, and Conv-TasNet [37], a sound source separation deep neural network that extracts speech from OTA audio and the backbone of the ClearBuds earable platform [7] that leverages deep learning to denoise and enhance OTA speech.

The primary quantitative metric we use for comparison is the word error rate (WER) because it is not possible to obtain ground truth clean speech in naturally noisy environments. However, if speech is properly enhanced by significantly attenuating ambient noises and/or sufficiently generating the higher frequency formants of speech, then a speech-to-text model should naturally translate a larger portion of words correctly. All WER metrics were computed after performing super resolution on BCM or ACCEL sampled speech from 4kHz to 16kHz.

7.1 Performance in Different Environments

We recorded 12 total minutes of data across all modalities (BCM, ACCEL, OTA) in various environments shown in Figure 7b. The environments we chose had varying types and levels of background noise. Table 7 compares the WER between TRAMBA, raw audio, and denoising over-the-air signals. There was almost no background

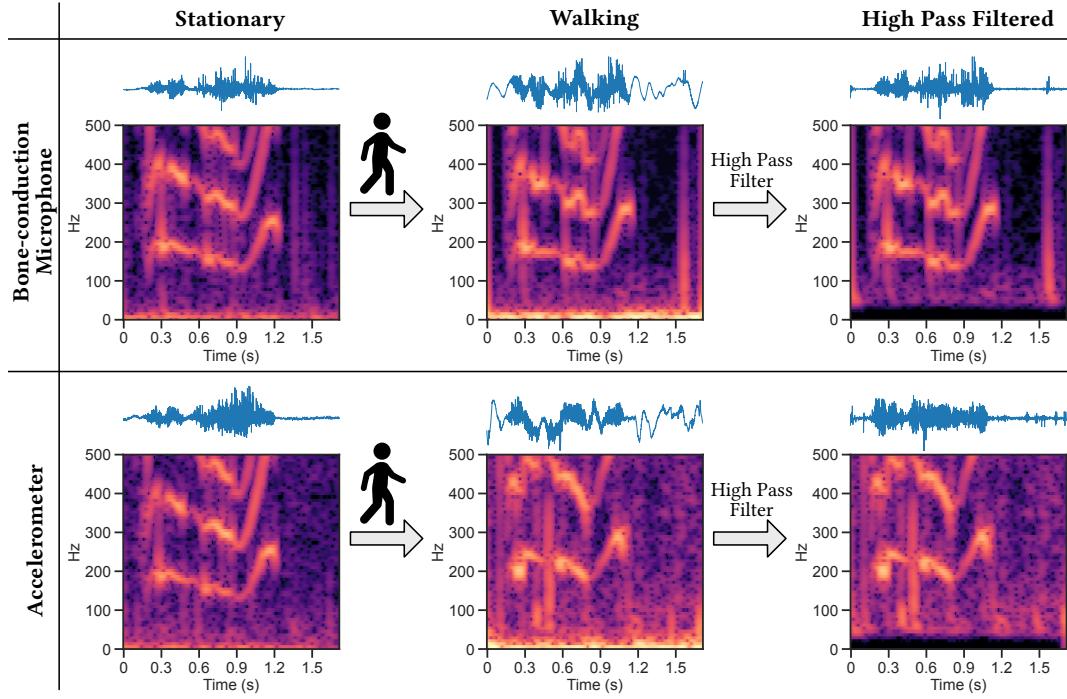


Fig. 17. Spectrogram of ACCEL and BCM while user is moving. There is only a small amount of energy in the extreme low frequencies, well below the lowest frequencies of speech (Typically around 100Hz). This means, movement has little to no effect on performance if a high pass filter is applied (Table 8).

noise in the office setting, yielding a WER of almost 0 across all methods, so we leave out this column. As the noise level increases in the different environments, the WER increases steadily when no processing is performed. Additionally, applying denoising and sound source separation on OTA signals does not always yield a lower WER due to distortions that such models commonly introduce when removing noise [3]. On the other hand, TRAMBA sees much higher and consistent performance across all noise settings, even when the noise is 10 orders of magnitude louder. Bone conduction and vibration-based sensing modalities are naturally not sensitive to ambient noises, unlike OTA microphones.

7.2 Movement

Figure 17 shows the spectrogram of signals collected by the BCM and ACCEL while the user is walking. In this scenario, high energy frequencies appear in the single Hz range or lower, even when the user is running and moving faster, which is significantly lower than the typical lowest frequencies of human voice (approximately 100 Hz). This suggests that applying a high pass filter would be sufficient to remove the effects of movement from speech, as shown in Table 8. Here, we applied a first order Butterworth filter with a 10Hz cut-off frequency. Performance across all metrics remain similar when the user is moving.

	PESQ	STOI	LSD	PESQ	STOI	LSD	
BCM - TRAMBA				BCM - TUNet			
Standing Still	1.8063	0.8114	1.0814	1.5064	0.7772	1.1403	
Walking - 1.2 m/s	1.7639	0.8031	1.0801	1.4956	0.7665	1.1511	
ACC - TRAMBA				ACC - TRAMBA			
Standing Still	1.8300	0.8480	1.2461	1.7077	0.8044	1.4213	
Walking - 1.2 m/s	1.7390	0.8413	1.2722	1.6888	0.8021	1.4323	

Table 8. Performance under movement.

	Data Rate	Power Consumption
TRAMBA- 500 Hz	8 kbps	2.49 mW
TRAMBA- 1 kHz	16 kbps	2.58 mW
TRAMBA- 2 kHz	32 kbps	2.75 mW
TRAMBA- 4 kHz	64 kbps	3.21 mW
TRAMBA- 8 kHz	128 kbps	4.09 mW
No processing (16 kHz)	256 kbps	6.48 mW

Table 9. Total power consumption of TRAMBA’s wearable, while sampling and streaming BCM/ACCEL data at different frequencies to the smartphone. Sampling at 4kHz more than halves the power consumption from sampling and transmitting at full resolution (16kHz).

Device	iPhone 15 Pro	iPhone 14 Pro	iPhone 13 Pro	iPhone 12
Inference Time (ms)	19.584	20.343	22.992	27.931

Table 10. Inference time (latency) on various phone models. TRAMBA processes windows of 512ms at a time and can perform inference in real-time on most modern smartphones.

7.3 Power and Latency

Table 9 shows the power consumption of sampling and transmitting data from the wearable to the smartphone platform at different sampling rates. Performing speech enhancement with TRAMBA on 4kHz signals us to sample and transmit 75% less data than sampling high resolution, 16kHz, OTA speech and reduce power consumption by more than 50%. We achieve even more aggressive savings if the resolution is further decreased.

Table 10 shows the average inference time for one window (512ms) on various smartphones. On all tested devices, the inference time is less than 30ms. Since we process 512ms windows, TRAMBA can perform speech enhancement on many modern smartphones in real-time.

7.4 User Study

To evaluate the sound quality in the noisy environments, we asked our recruited volunteers to rate the enhanced audio on a 5 point Likert scale [25], where a score of 1 is very poor and 5 is very high quality. Figure 18 shows these results across all environments compared to raw audio and OTA source separation methods. The over-the-air methods were perceived to have higher quality in the office environment because were no loud ambient noises interfering with speech. However, environments B, C, D contained increasingly louder ambient backgrounds.

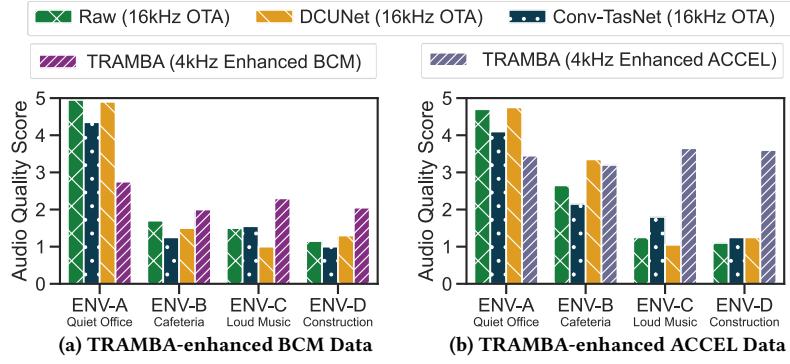


Fig. 18. Perceived quality evaluated by users. Higher number denotes better audio quality.

We see that the perceived quality of TRAMBA remained fairly constant in all scenarios, demonstrating its lack of sensitivity to ambient noises. However, the quality of speech generated by over-the-air methods drastically decreases, as ambient noise overpowers OTA speech.

One interesting observation is that our participants consistently scored ACCEL higher than BCM. While BCMs are less sensitive to external noises than OTA microphones, our participants still reported the presence of external non-speech sound sources in particularly noisy scenarios, degrading the quality of the output. While ACCELS are less sensitive to bone-conduction voice, they are also even less sensitive than BCMs to external noises. Our participants did not report hearing external sound sources on the ACCEL recordings, which contributed to the higher perception scores. Since our work focuses more on enhancing speech components rather than removing external sound sources, we leave this problem to future work.

8 DISCUSSION

We have demonstrated TRAMBA, a novel hybrid transformer and Mamba-based architecture for speech super resolution that outperforms existing state-of-art speech enhancement methods across multiple acoustic modalities (over-the-air, bone-conduction microphones, and accelerometers). Through integrating and deploying TRAMBA into real systems and environments, we show that TRAMBA enables practical deployment of acoustic in mobile and wearable platforms by surpassing the performance gap with a memory footprint that is orders of magnitude less than the state-of-art, reducing reliance on paired data that is labor-intensive to collect by pre-training on easily obtainable downsampled clean speech audio, and leveraging the reduced data rate required to lower power consumption and improve battery life. While our results show great promise, there are several lines of future work that we believe will further strengthen the practical adoption of acoustic speech enhancement in mobile and wearable platforms.

On-Device Fine-tuning. In this work, we did not fine-tune TRAMBA directly on the smartphone due to a lack of frameworks that allow developers to fine-tune directly on the phone. We plan to explore efficient on-device fine-tuning and create open-source, easy-to-use tools to enable developers, engineers, and the greater community to fine-tune large deep learning models on smartphone and other edge platforms.

Generalizability. While fine-tuning the pre-trained model with the acoustic modality of choice at box opening allows us to adapt TRAMBA to a specific user, generalizing the model to unseen individuals without fine-tuning is still unsolved. We plan to build upon the open-source platforms and techniques created in this work to further reduce the barrier for collecting paired vibration and OTA data. In doing so, we hope to enable community-driven activity that will increase the amount of publicly available data needed to train and create generalizable

vibration-based speech enhancement models. Moreover, we observed a degradation in performance in females compared to males, which we hypothesize is due to their naturally higher pitched voices experiencing more severe attenuation through skin and bone. In future work, we plan to explore causes and solutions in more depth. **Security and Privacy.** In this work, we transmit sensitive speech signals between multiple devices. In future work, we plan to explore privacy-aware transmission schemes that would enable safe and reliable communication of sensitive speech between multiple devices. Moreover, we plan to explore other methods of data reduction, beyond sampling at a lower rate (e.g., compression and learning latent embeddings), that would maintain or further reduce power consumption while improving privacy guarantees.

9 CONCLUSION

We present TRAMBA, a novel hybrid transformer and Mamba-based architecture for speech super resolution and enhancement for mobile and wearable platforms. One of the biggest challenges in creating speech enhancement methods for vibration-based sensing modalities is the lack of publicly available datasets arising from the labor-intensive data collection process required. To overcome this challenge, we pre-train TRAMBA on widely available, downsampled, clean speech audio before fine-tuning on a small amount of data collected from the user upon a one-time initial setup. We demonstrate that TRAMBA outperforms existing state-of-art speech super resolution methods across multiple acoustic modalities (over-the-air, BCM, and accelerometer) with a memory footprint of only 19.7 MBs, compared to GANs that require at least hundreds of MBs or more. We integrate TRAMBA into a real mobile and head-worn wearable system and show, through real deployments and user studies, that TRAMBA can generate higher quality speech, in noisy environments, compared to systems that denoise audio collected by microphones that often hear ambient background noises. On the systems side we demonstrate that TRAMBA can improve the battery life of wearable systems, by up to 160%, by reducing the resolution of audio that needs to be sampled and transmitted. TRAMBA is a critical step towards practical integration of speech enhancement into mobile and wearable platforms.

REFERENCES

- [1] [n. d.]. MAX78000. <https://www.analog.com/en/products/max78000.html>. Accessed: 2024-05-01.
- [2] [n. d.]. STM32 32-bit ARM Cortex-MCUs. <https://www.st.com/en/microcontrollers-microprocessors/stm32-32-bit-arm-cortex-mcus.html>. Accessed: 2024-05-01.
- [3] Anis Ben Aicha and Sofia Ben Jebara. 2012. Perceptual speech quality measures separating speech distortion and additive noise degradations. *Speech Communication* 54, 4 (2012), 517–528.
- [4] Apple Inc. 2023. Create a Personal Voice on your Phone, iPad, or Mac. <https://support.apple.com/en-us/104993>, Last accessed on 2024-04-27.
- [5] Sawyer Birnbaum, Volodymyr Kuleshov, Zayd Enam, Pang Wei W Koh, and Stefano Ermon. 2019. Temporal FiLM: Capturing Long-Range Sequence Dependencies with Feature-Wise Modulations. *Advances in Neural Information Processing Systems* 32 (2019).
- [6] Carl Gene Fordham. 2015. Basic Spoken English. https://www.carlgene.com/Basic_Spoken_English_carlgene.com.pdf, Last accessed on 2024-04-30.
- [7] Ishan Chatterjee, Maruchi Kim, Vivek Jayaram, Shyamnath Gollakota, Ira Kemelmacher, Shwetak Patel, and Steven M Seitz. 2022. ClearBuds: wireless binaural earbuds for learning-based speech enhancement. In *Proceedings of the 20th Annual International Conference on Mobile Systems, Applications and Services*. 384–396.
- [8] Hyeong-Seok Choi, Jang-Hyun Kim, Jaesung Huh, Adrian Kim, Jung-Woo Ha, and Kyogu Lee. 2018. Phase-aware speech enhancement with deep complex u-net. In *International Conference on Learning Representations*.
- [9] Krzysztof Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarlos, Peter Hawkins, Jared Davis, Afroz Mohiuddin, Lukasz Kaiser, et al. 2020. Rethinking attention with performers. *arXiv preprint arXiv:2009.14794* (2020).
- [10] Alexandre Defossez, Gabriel Synnaeve, and Yossi Adi. 2020. Real time speech enhancement in the waveform domain. *arXiv preprint arXiv:2006.12847* (2020).
- [11] Di Duan, Yongliang Chen, Weitao Xu, and Tianxing Li. 2024. EarSE: Bringing Robust Speech Enhancement to COTS Headphones. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 7, 4 (2024), 1–33.

- [12] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. *Advances in neural information processing systems* 27 (2014).
- [13] Grand View Research, Inc. 2023. Smart Glasses Market Size & Share Analysis Report, 2030. <https://www.grandviewresearch.com/industry-analysis/smart-glasses-market-report>, Last accessed on 2024-04-25.
- [14] Grand View Research, Inc. 2024. Earphones And Headphones Market Size Report 2030. <https://www.grandviewresearch.com/industry-analysis/earphone-and-headphone-market>.
- [15] Augustine Gray and John Markel. 1976. Distance measures for speech processing. *IEEE Transactions on Acoustics, Speech, and Signal Processing* 24, 5 (1976), 380–391.
- [16] Albert Gu and Tri Dao. 2023. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752* (2023).
- [17] Albert Gu, Karan Goel, and Christopher Ré. 2021. Efficiently modeling long sequences with structured state spaces. *arXiv preprint arXiv:2111.00396* (2021).
- [18] Albert Gu, Isys Johnson, Karan Goel, Khaled Saab, Tri Dao, Atri Rudra, and Christopher Ré. 2021. Combining recurrent, convolutional, and continuous-time models with linear state space layers. *Advances in neural information processing systems* 34 (2021), 572–585.
- [19] Julien Hauret, Thomas Joubaud, Véronique Zimpfer, and Éric Bavu. 2023. EBEN: Extreme bandwidth extension network applied to speech signals captured with noise-resilient body-conduction microphones. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 1–5.
- [20] Lixing He, Haozheng Hou, Shuyao Shi, Xian Shuai, and Zhenyu Yan. 2023. Towards Bone-Conducted Vibration Speech Enhancement on Head-Mounted Wearables. In *Proceedings of the 21st Annual International Conference on Mobile Systems, Applications and Services*. 14–27.
- [21] Jen-Cheng Hou, Syu-Siang Wang, Ying-Hui Lai, Yu Tsao, Hsiu-Wen Chang, and Hsin-Min Wang. 2018. Audio-visual speech enhancement using multimodal deep convolutional neural networks. *IEEE Transactions on Emerging Topics in Computational Intelligence* 2, 2 (2018), 117–128.
- [22] Kaiyuan Hou, Stephen Xia, and Xiaofan Jiang. 2022. BuMA: Non-Intrusive Breathing Detection using Microphone Array. In *Proceedings of the 1st ACM International Workshop on Intelligent Acoustic Systems and Applications*. 1–6.
- [23] Zili Huang, Shinji Watanabe, Shu-wen Yang, Paola García, and Sanjeev Khudanpur. 2022. Investigating self-supervised learning for speech enhancement and separation. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 6837–6841.
- [24] Invensense. 2021. IIM-42352. <https://invensense.tdk.com/products/smartindustrial/iim-42352/>, Last accessed on 2024-04-27.
- [25] Ankur Joshi, Saket Kale, Satish Chandel, and D Kumar Pal. 2015. Likert scale: Explored and explained. *British journal of applied science & technology* 7, 4 (2015), 396–403.
- [26] Knowles Electronics, LLC. 2024. Knowles V2S. <https://www.knowles.com/V2S>, Last accessed on 2024-04-27.
- [27] Volodymyr Kuleshov, S Zayd Enam, and Stefano Ermon. 2017. Audio super-resolution using neural nets. In *ICLR (Workshop Track)*.
- [28] Kehuang Li and Chin-Hui Lee. 2015. A deep neural network approach to speech bandwidth expansion. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 4395–4399.
- [29] Yunpeng Li, Marco Tagliasacchi, Oleg Rybakov, Victor Ungureanu, and Dominik Roblek. 2021. Real-time speech frequency bandwidth extension. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 691–695.
- [30] Yuang Li, Yuntao Wang, Xin Liu, Yuanchun Shi, Shwetak Patel, and Shao-Fu Shih. 2022. Enabling Real-Time On-Chip Audio Super Resolution for Bone-Conduction Microphones. *Sensors* 23, 1 (2022), 35.
- [31] Yunji Liang, Yuchen Qin, Qi Li, Xiaokai Yan, Zhiwen Yu, Bin Guo, Sagar Samtani, and Yanyong Zhang. 2022. AccMyrinx: Speech Synthesis with Non-Acoustic Sensor. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 6, 3 (2022), 1–24.
- [32] Teck Yian Lim, Raymond A Yeh, Yijia Xu, Minh N Do, and Mark Hasegawa-Johnson. 2018. Time-frequency networks for audio super-resolution. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 646–650.
- [33] Hung-Ping Liu, Yu Tsao, and Chiou-Shann Fuh. 2018. Bone-conducted speech enhancement using deep denoising autoencoder. *Speech Communication* 104 (2018), 106–112.
- [34] Jonathan Long, Evan Shelhamer, and Trevor Darrell. 2015. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 3431–3440.
- [35] Xugang Lu, Yu Tsao, Shigeki Matsuda, and Chiori Hori. 2013. Speech enhancement based on deep denoising autoencoder.. In *Interspeech*, Vol. 2013. 436–440.
- [36] Yi Luo and Nima Mesgarani. 2018. Tasnet: time-domain audio separation network for real-time, single-channel speech separation. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 696–700.
- [37] Yi Luo and Nima Mesgarani. 2019. Conv-tasnet: Surpassing ideal time–frequency magnitude masking for speech separation. *IEEE/ACM transactions on audio, speech, and language processing* 27, 8 (2019), 1256–1266.
- [38] Moshe Mandel, Or Tal, and Yossi Adi. 2023. Aero: Audio super resolution in the spectral domain. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 1–5.

[39] Market.us. 2024. Wearable Technology Market Size, Share. <https://market.us/report/wearable-technology-market/>, Last accessed on 2024-04-25.

[40] Juan Manuel Martin-Donas, Angel Manuel Gomez, Jose A Gonzalez, and Antonio M Peinado. 2018. A deep learning loss function based on the perceptual evaluation of the speech quality. *IEEE Signal processing letters* 25, 11 (2018), 1680–1684.

[41] Héctor A Cordourier Maruri, Paulo Lopez-Meyer, Jonathan Huang, Willem Marco Beltman, Lama Nachman, and Hong Lu. 2018. V-speech: Noise-robust speech capturing glasses using vibration sensors. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 4 (2018), 1–23.

[42] Daniel Michelsanti, Zheng-Hua Tan, Shi-Xiong Zhang, Yong Xu, Meng Yu, Dong Yu, and Jesper Jensen. 2021. An overview of deep-learning-based audio-visual speech enhancement and separation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 29 (2021), 1368–1396.

[43] Viet-Anh Nguyen, Anh HT Nguyen, and Andy WH Khong. 2022. Tunet: A block-online bandwidth extension model based on transformers and self-supervised pretraining. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 161–165.

[44] Aditya Arie Nugraha, Antoine Liutkus, and Emmanuel Vincent. 2016. Multichannel audio source separation with deep neural networks. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 24, 9 (2016), 1652–1664.

[45] NVIDIA Corporation. 2023. NVIDIA L40 GPU for Data Center. <https://www.nvidia.com/en-us/data-center/l40/>, Last accessed on 2024-04-27.

[46] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. 2018. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 32.

[47] Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2023. Robust speech recognition via large-scale weak supervision. In *International Conference on Machine Learning*. PMLR, 28492–28518.

[48] Nathanaël Carraz Rakotonirina. 2021. Self-attention for audio super-resolution. In *2021 IEEE 31st International Workshop on Machine Learning for Signal Processing (MLSP)*. IEEE, 1–6.

[49] Dario Rethage, Jordi Pons, and Xavier Serra. 2018. A wavenet for speech denoising. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 5069–5073.

[50] Antony W Rix, John G Beerends, Michael P Hollier, and Andries P Hekstra. 2001. Perceptual evaluation of speech quality (PESQ)-a new method for speech quality assessment of telephone networks and codecs. In *2001 IEEE international conference on acoustics, speech, and signal processing. Proceedings (Cat. No. 01CH37221)*, Vol. 2. IEEE, 749–752.

[51] Wenzhe Shi, Jose Caballero, Ferenc Huszár, Johannes Totz, Andrew P Aitken, Rob Bishop, Daniel Rueckert, and Zehan Wang. 2016. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1874–1883.

[52] Tanmay Srivastava, Prerna Khanna, Shijia Pan, Phuc Nguyen, and Shubham Jain. 2022. Muteit: Jaw motion based unvoiced command recognition using earable. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 6, 3 (2022), 1–26.

[53] Cees H Taal, Richard C Hendriks, Richard Heusdens, and Jesper Jensen. 2010. A short-time objective intelligibility measure for time-frequency weighted noisy speech. In *2010 IEEE international conference on acoustics, speech and signal processing*. IEEE, 4214–4217.

[54] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* 30 (2017).

[55] Emmanuel Vincent, Tuomas Virtanen, and Sharon Gannot. 2018. *Audio source separation and speech enhancement*. John Wiley & Sons.

[56] Heming Wang, Xueliang Zhang, and DeLiang Wang. 2022. Fusing bone-conduction and air-conduction sensors for complex-domain speech enhancement. *IEEE/ACM transactions on audio, speech, and language processing* 30 (2022), 3134–3143.

[57] Stephen Xia and Xiaofan Jiang. 2020. Pams: Improving privacy in audio-based mobile systems. In *Proceedings of the 2nd International Workshop on Challenges in Artificial Intelligence and Machine Learning for Internet of Things*. 41–47.

[58] Stephen Xia and Xiaofan Jiang. 2021. Improving acoustic detection and classification in mobile and embedded platforms. In *Proceedings of the 20th International Conference on Information Processing in Sensor Networks (Co-Located with CPS-IoT Week 2021)*. 402–403.

[59] Stephen Xia and Xiaofan Jiang. 2022. Ava: An adaptive audio filtering architecture for enhancing mobile, embedded, and cyber-physical systems. In *2022 21st ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*. IEEE, 118–131.

[60] Stephen Xia, Jingping Nie, and Xiaofan Jiang. 2021. Csafe: An intelligent audio wearable platform for improving construction worker safety in urban environments. In *Proceedings of the 20th International Conference on Information Processing in Sensor Networks (Co-Located with CPS-IoT Week 2021)*. 207–221.

[61] Junichi Yamagishi, Christophe Veaux, and Kirsten MacDonald. 2019. CSTR VCTK Corpus: English Multi-speaker Corpus for CSTR Voice Cloning Toolkit. <https://doi.org/10.7488/ds/2645>.

[62] Ryuichi Yamamoto, Eunwoo Song, and Jae-Min Kim. 2020. Parallel WaveGAN: A fast waveform generation model based on generative adversarial networks with multi-resolution spectrogram. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 6199–6203.

[63] Karren Yang, Dejan Marković, Steven Krenn, Vasu Agrawal, and Alexander Richard. 2022. Audio-visual speech codecs: Rethinking audio-visual speech enhancement by re-synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 8227–8237.