



HAL
open science

Configurable EBEN: Extreme Bandwidth Extension Network to enhance body-conducted speech capture

Julien Hauret, Thomas Joubaud, Éric Bavu, Véronique Zimpfer

► To cite this version:

Julien Hauret, Thomas Joubaud, Éric Bavu, Véronique Zimpfer. Configurable EBEN: Extreme Bandwidth Extension Network to enhance body-conducted speech capture. *IEEE/ACM Transactions on Audio, Speech and Language Processing*, 2023, 31, pp.3499-3512. 10.1109/TASLP.2023.3313433 . hal-04037315v2

HAL Id: hal-04037315

<https://hal.science/hal-04037315v2>

Submitted on 12 Sep 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Configurable EBEN: Extreme Bandwidth Extension Network to enhance body-conducted speech capture

1st Hauret Julien

Laboratoire de Mécanique des Structures et des Systèmes Couplés
Conservatoire national des arts et métiers, HESAM Université
Paris, France
ORCID : 0000-0002-1512-2487
julien.hauret@lecnam.net

2nd Joubaud Thomas

Department of Acoustics and Soldier Protection
French-German Research Institute of Saint-Louis (ISL)
Saint-Louis, France
ORCID : 0000-0002-5324-8785
thomas.joubaud@isl.eu

3rd Zimpfer Véronique

Department of Acoustics and Soldier Protection
French-German Research Institute of Saint-Louis (ISL)
Saint-Louis, France
ORCID : 0000-0002-7852-1928
veronique.zimpfer@isl.eu

4th Bavu Éric

Laboratoire de Mécanique des Structures et des Systèmes Couplés
Conservatoire national des arts et métiers, HESAM Université
Paris, France
ORCID : 0000-0001-6395-634X
eric.bavu@lecnam.net

Abstract—This paper presents a configurable version of Extreme Bandwidth Extension Network (EBEN), a Generative Adversarial Network (GAN) designed to improve audio captured with body-conduction microphones. We show that although these microphones significantly reduce environmental noise, this insensitivity to ambient noise happens at the expense of the bandwidth of the speech signal acquired by the wearer of the devices. The obtained captured signals therefore require the use of signal enhancement techniques to recover the full-bandwidth speech. EBEN leverages a configurable multiband decomposition of the raw captured signal. This decomposition allows the data time domain dimensions to be reduced and the full band signal to be better controlled. The multiband representation of the captured signal is processed through a U-Net-like model, which combines feature and adversarial losses to generate an enhanced speech signal. We also benefit from this original representation in the proposed configurable discriminators architecture. The configurable EBEN approach can achieve state-of-the-art enhancement results on synthetic data with a lightweight generator that allows real-time processing.

Index Terms—Speech enhancement, PQMF-banks, Bandwidth extension, Frugal AI, Body-Conduction Microphones

I. INTRODUCTION

Capturing speech involves the use of microphones to transform mechanical vibrations into an electrical signal, later digitalized and eventually used for radio communications. Under quiet conditions, using an airborne-sound microphone near the speaker’s lips is the most appropriate way to capture clean speech. Nevertheless, in presence of ambient noise generated by sources contaminating the sound scene, the speech signal of interest is altered by the acoustic environment, which also contributes to air molecules vibration. This situation – which reduces the intelligibility of communications – is frequently encountered in industry, on the battlefield or in strong winds. In extreme cases, operators are even unable to communicate.

Before using any speech enhancement technique, it is worth pondering the best mechanical signal to rely on in noisy conditions. There are other choices besides recording airborne sound pressure, such as the body-conducted inner vibrations of the speaker. The human body is not as easily moved by environmental noise as ambient air, due to the high damping of the transmitted sound wave in the tissues. Therefore, capturing inner tissues’ vibrations caused by the vocal tract near the speaker’s head has great potential for improving the signal to noise ratio when recording speech in noisy environments. This can be performed with noise-resilient body-conduction microphones (BCMs), which allow sensing the internal vibrations of the equipped person. This family of unconventional voice pickup systems includes bone conduction transducers [1]–[5], throat microphones [6], [7] and in-ear microphones mounted in occlusive earplugs offering hearing protection [8]–[11]. Studies including [2], [9] and [12] demonstrated that they offer higher quality and intelligibility in noise than conventional capture devices. We also conducted our experiment on Sec. III-A to determine when it is preferable to use a BCM over a traditional microphone.

In addition to eliminating external noise pollution, BCMS are less invasive and compatible with helmets, which are often required in noisy environments. Similarly, they are suitable for wearing gas masks or face masks which is not negligible in times of pandemic. In-ear capture devices are also prone to be integrated into hearing protection devices. The protection will isolate the sensor from the external environment, and the wearer’s speech capture will be improved. Finally, the broad adoption of true wireless stereo earbuds and bone conduction earphones also benefits the development of inner voice capture. Indeed, those systems are reversible and could be used as BCMS.

Despite many advantages, the usage of BCMs has not yet been democratized. This can be marginally explained by the fact that they are not always necessary (e.g., in a quiet, distant meeting), but mainly because recordings suffer from reduced bandwidth. Indeed, mid and high frequencies are missing due to the intrinsic low-pass characteristics of the biological pathway. Further processing is then necessary to optimize the effective bandwidth of the captured speech. Moreover, other physiological sounds, such as swallowing, blood flow - or any other sound produced by the body - are also picked up by BCMs and represent a new form of noise contaminating speech capture.

In simple terms, speech capture in surrounding noise can be achieved either by using airborne speech with a denoising algorithm or by using a noise-proof body-conduction microphone with bandwidth extension techniques. The latter is a viable solution for critical noise levels ($\geq 85\text{dB}$) when differential microphones or directional boom microphones cannot eliminate high-level surrounding noise. Therefore, this article proposes an extreme bandwidth extension deep neural network for speech signals captured with noise-resilient body-conduction microphones.

Since the desirable system could be a two-way communication device for industrial or military usecases, this entails real-time execution constraints, *i.e.* a short processing time to be indistinguishable from the human ear. Moreover, edge computations are required to guarantee low latency, necessitating a light algorithm. These considerations also match frugal AI requirements. Finally, the developed model should be robust to speaker identity, physiological and residues of external noises that would have infiltrated the microphone.

To meet the expectations of extreme bandwidth extension and related requirements, research like [4], [13], [14] suggests that frugal deep learning is an appropriate approach. Indeed, compared to conventional signal processing methods that can only extend existing frequency content, deep learning can regenerate missing components such as fricatives. Additionally, deep learning offers the advantage of simultaneous denoising and signal enhancement, eliminating the need for separate denoising procedures. On the other hand, massive deep learning models are not relevant for real-time execution.

Based on the above observations, we developed configurable EBEN, a new deep learning model inspired by a family of lightweight convolutional-based encoder-decoder architecture [13], [15]–[18] to infer mid and high frequencies from speech containing only low frequencies (extreme bandwidth extension). We use a generator that maps the degraded speech signal to an enhanced version. This task is called blind speech enhancement because we do not use any external modality (contrary to Seanet [16], which takes advantage of both airborne speech and accelerometer data). EBEN’s generator is optimized to produce samples close to the reference while maintaining a certain degree of naturalness at different time scales. We still differ from previous work by using a multiband decomposition performed with Pseudo-Quadrature Mirror Filters (PQMF) [19]. Combined with some hypotheses on ad-

ressed degradations, this decomposition is applied to reduce the dimension of input features, which significantly decreases the latency and computational load of the network. In addition, this alternative representation is useful for focusing signal discrimination solely on high frequency bands.

A preliminary version of our research was presented at ICASSP 2023 [20]. The present paper extends the original study by highlighting the usefulness of BCMs in noise, by addressing more diverse and realistic degradations, by discussing the goals and flexibility offered by the configurable aspect of our approach, and by comparing EBEN’s latency and memory footprint to other previously published networks. An extensive discussion of the usefulness of common objective metrics is also proposed for the specific task of bandwidth extension, along with a correlation analysis of objective and neural distances with subjective evaluations. The present paper also proposes an extensive discussion of related work. Finally, details of the training strategy, architecture, and statistical analysis of the evaluation survey are presented. The website <https://jhauret.github.io/eben> also provides example audio files to listen to and the source code of EBEN.

The body-conduction microphone studied in this paper is an in-ear microphone prototype mounted in an earplug. The few minutes of recordings at our disposal being insufficient to serve for supervised training, we instead analyzed bandwidth loss to simulate in-ear-like degradations on the French LibriSpeech dataset [21]. Triplet train/dev/test sets of reference and corrupted speech pairs were produced to train our model, and several baselines [13], [14], [16], [22]. We plan to later release a publicly available dataset of speech capture with BCMs to circumvent the disadvantages of the use of synthetic data, which are discussed in Sec. VI-D.

It is worth noting that focusing on the capture-induced degradation of a specific BCM does not detract from the generality of the proposed approach. This family of sensors consistently degrades speech in a similar manner, acting as a low-pass filter, and adding physiological and frictional noise that would not be captured by a conventional microphone. Variations occur mainly in cut-off frequency, attenuation, and lack of coherence at certain frequencies. Therefore, a suitable dataset would be sufficient to address any other capture system. It is also noteworthy that preliminary experiments have shown that the EBEN approach performs well at conventional bandwidth extension (*e.g.* upsampling 4kHz to 16kHz), which has been confirmed by listening and metrics.

In Sec. II we review related previous studies, which also serve as a baseline for our comparisons. In Sec. III, we show that BCMs are more suitable for recording speech in noise than traditional microphones, present the observed degradation with our in-ear prototype, and describe our protocol for generating synthetic data. Sec. IV provides a brief reminder of the PQMF bank and a detailed presentation of EBEN architecture and loss functions. Sec. V describes the training pipeline, experimental results, and compares EBEN to other approaches. The discussion in Sec. VI provides some insights into the EBEN model by discussing why PQMF bands fit well in the context of real-time

BCM speech enhancement and the configurable aspects of the architecture. This section also includes a statistical analysis of the correlation between subjective and objective metrics and discuss the accordance of the synthetic data generation strategy. Finally, Sec. VII concludes the paper.

II. RELATED WORK

The earliest speech bandwidth extension algorithms, usually applied to telephony applications, were performed with pure signal processing algorithms like spectral folding [23], Linear Predictive Coding [24], modulation techniques [25], [26] or non-linear processing [27]. This simple procedure has also been used in the context of in-ear microphones [9] with fair results, yet to be improved. This method creates missing harmonics in the high frequencies but cannot recover missing formants and fricatives. The earliest data-driven approaches have subsequently offered a more realistic extension. Those approaches are composed of several building blocks, including a statistical model that aims to estimate the high band spectral envelope. In many cases, this statistical model is one of the following: codebooks [28], Gaussian Mixture Model [29], Hidden Markov Model [30] and even some neural networks [31]. Although the quality is generally better with those methods, overly smoothed spectra are still produced at the expense of speech naturalness.

Recent advances in neural speech synthesis [14], [32]–[36] have proven that end-to-end deep learning is state-of-the-art in terms of simplicity and sample quality. Therefore, deep learning seems promising to accomplish this extreme bandwidth extension task. Indeed, the ability of neural networks to extract relevant features for the downstream task will allow the matching of high and low frequency contents. Raw waveform input is preferred over handcrafted features like spectrogram, mel-spectrogram [37], or Mel-frequency cepstral coefficients (MFCCs) [38] to minimize human processing and let the network build its representation. This trend is endorsed by several works in the audio field [39]–[42] and especially for bandwidth extension (synonym of audio super-resolution) to avoid rebuilding the phase separately [22], [43]–[45].

The use of raw audio can also be combined with multiband processing to speed up inference, as in *DurIAN* [46], *MB-MelGAN* [47] or *RAVE* [48]. The speech signal is therefore processed at a reduced sampling rate thanks to the decomposition, unlike other super-resolution networks [49], which use an input signal sampled at the target sampling frequency. To pursue the objective of fast inference, a fully convolutional architecture has been preferred like in [22], [32], specifically U-Net-like as other audio-to-audio tasks [50], [51]. The up-sampling layers use transposed convolutions [52] instead of subpixel layers [53]. Transposed convolutions do not produce checkerboard artifacts when kernel size and strides are chosen to avoid overlapping disparities, as explained in [54].

In addition, a simple reconstruction loss may be insufficient for conditional generation, producing unrealistic samples. As shown in [14], [55]–[57], adversarial networks [58] can significantly improve the naturalness of the produced sound.

Multiple discriminators are even used in [16], [45], [56] to focus signal discrimination at different time scales. Moreover, *feature matching* is also encouraged for the reconstruction loss because it allows to enhance the produced sound quality in an end-to-end fashion. Indeed, discriminators’ embeddings are excelling at building a relevant representation for our problem; it is therefore consistent to compute distance based on those features. Alternatives to this approach are either the L_1/L_2 norms in the time domain, which are known to be misaligned with human perception, or complex losses like multiscale Short Time Fourier Transform, which depend on chosen hyperparameters [22], [59], thus increasing tuning efforts.

Regarding the specific literature on blind (or non-multi modal) speech enhancement for BCMs, different approaches adopted classic processing to achieve bandwidth extension [1], [6], [7], [9]. Subjective quality evaluations have proven those approaches to be inadequate for this task. Then, neural networks started to be employed, firstly as a processing block among others [10] to estimate an enhancement function in a fixed feature domain combined with time-domain filtering. Subsequent research then began to carry out the improvement task and ultimately to perform the enhancement as end-to-end tasks. Among published manuscripts on the subject, the works of Yuang Li et al. [4], Hung-Ping Liu et al. [60], and Dongjing Shan et al. [61] applied this approach for bone conduction microphones, and Mattes Ohlenbusch et al. [11] to in-ear microphones. The main drawback of those approaches relies on the fact that they are based on a pure reconstructive loss, eventually with a regularization part. As they expressed in their articles, *an audible difference between the target and the processed signals remains*. This statement may be irrevocable due to the limited information left in the signal captured by BCMs. However, GANs [62] can produce realistic signals by slightly deviating from the reference. The task of speech enhancement for speech capture with body-conduction microphones is thus complicated. Indeed, [3], [16], [63], [64] only used BCMs as a conditional signal for enhancement in a multi-modal framework. Moreover, even if BCMs mainly capture speech, residues of external and physiological noises persist [65] and would necessitate denoising. Hopefully, deep learning models can perform the denoising task simultaneously with the bandwidth extension. [9] has also proven that the contaminating noise knowledge was helpful, although we will not capitalise on this particular knowledge in the present paper.

Lastly, this research area lacks large public corpora that use body-conduction microphones to reliably train deep models. The ABCS [66] and EMSB¹ datasets, which consist of either bone or in-ear and air-conducted Mandarin speech pairs, are currently the largest corpora, consisting of 42 hours and 128 hours of speech, respectively. Another smaller public² dataset is Speech in-EAR (SpEAR) database proposed in [67] with 25 participants, split in French/English speakers. Other private

¹<https://github.com/elevoctech/ESMB-corpus>

²freely available upon request from a research institution

datasets emerged, like [4], which introduced 200 minutes of speech recorded via bone conduction. The dataset was large enough to train on, likely due to their model’s meager number of parameters (4.5k for the lightest model). [11] opted for a different strategy with their overall 30 minutes in-ear captured speech. The limited-size dataset was first used to simulate meaningful degradations, taking into account the body-produced noise used to train their model. Finally, they re-used real data to fine-tune their model’s decoder.

III. IN-EAR MICROPHONE STUDY

The selected BCM is an early prototype based on a MEMS microphone (STM MP34DT01) driven by an STM32 H7 microcontroller [68] developed by the ISL and Cotral Lab. This device takes advantage of the speaker’s hearing protection by being placed inside a custom-molded earplug, which increases communication capabilities in challenging and noisy environments. The reference speech signals are captured by a B&K Type 4192 condenser microphone connected to a TEAC LX10 data recorder. The reference and in-ear signals are recorded at 48 kHz, resampled at 16 kHz and finally synchronized using cross-correlation. We collected utterances of the Combescure’s phonetically balanced sentences [69] from a single speaker.

A. In-noise comparison with traditional microphone

This section aims at justifying that BCMs are more suitable for noisy environments and at establishing a rough estimate of the noise level above which their use should prevail. We conducted subjective A/B preference tests to compare our in-ear microphone with a traditional microphone. A single speaker was recorded simultaneously by both microphones in different acoustic environments, ranging from a quiet acoustically treated room to a reverberating room with several levels of surrounding noise. When surrounding noise is present, the speaker naturally produces Lombard speech [70]. Comparisons are performed using 7 different utterances from the Combescure’s sentences [69], which have been recorded in an audiometric booth (IAC Acoustics and walls covered with acoustic foam), and in a reverberating room with pink noise levels $\{ \emptyset, 55\text{dB}, 65\text{dB}, 75\text{dB}, 85\text{dB}, 95\text{dB} \}$ without any enhancement techniques. We recruited 38 participants and used the GoListen platform [71] to conduct the test. Participants were asked whether they preferred in-ear or classic recordings for all acoustic environments. Participants were divided into two groups to rate either the quality or the ease of understanding of the audio samples. Obtained results are presented in Fig. 1 along with the corresponding p-values for each A/B comparisons.

According to Fisher’s exact test and a significance rate of 1%, the obtained results allow us to conclude that the use of an in-ear microphone is preferred for both ease of speech understanding and sound quality for surrounding noise levels of 75 dB or more. On the other hand, a traditional microphone is endorsed for ease of understanding and quality for noise




Fig. 1: A/B testing results: in-ear vs traditional microphones. The p-values shown at the top of each bar indicate the significance of the preferred microphone.

levels below or equal to 55dB. No statistically significant difference can be drawn for a 65dB noise level.

B. Degradation study

In-ear own voice capture is more adapted for applications in noisy environments because it mainly contains speech without external noise. However, the acoustic wave propagation between the vocal tract and the transducers causes irreversible information loss: almost no relevant speech signal is picked up above a threshold frequency. Complex interactions with tissues are also responsible for phase shifts and anti-resonances.

This phenomenon is further influenced by the occlusion effect [72] due to the fitting of the individual protectors, causing speech to resonate inside the ear canal. This aspect causes an amplification of the remaining signal, leading the wearer to hear an amplified version of their own voice. The occlusion effect is therefore the consequence of wearing an earplug, but it is also necessary in order to obtain an in-ear signal that is not significantly degraded by environmental noise.

A first coarse approximation of those degradations can be modeled by a linear impulse response ψ that allows to estimate the in-ear signal x from the emitted signal y :

$$x(t) = (\psi * y)(t) \quad (1)$$

To evaluate the corresponding transfer function, we simultaneously use the in-ear prototype and a regular microphone placed in front of the speaker’s mouth under noise-free conditions. The absence of noise allows us to consider airborne speech as the emitted signal. The degradation filter estimates $\{\hat{\Psi}_i\}_{i \in [1,53]}$ were obtained with cross power spectral densities $\{P_{yx,i}\}_{i \in [1,53]}$ and $\{P_{yy,i}\}_{i \in [1,53]}$ approximated by Welch’s method [73], Eq. 2. Short Time Fourier Transforms were computed on 512 samples corresponding to 32 ms for the 16 kHz sampling rate used during this analysis. Welch’s method has a temporal horizon of 1.024 second with a recovery rate of 50%. A Voice Activity Detection (VAD) pre-processing based

on a simple reference’s power thresholding was applied to select meaningful segments. The reference and in-ear signals were normalized before calculating the cross power spectral densities because the in-ear microphone is not calibrated. Therefore, the shape of the transfer function is correct as a function of frequency, but the absolute amplitude does not reflect differences in sound pressure.

$$\tilde{\Psi}_i(f) = \frac{P_{yx,i}(f)}{P_{yy,i}(f)}, \forall i \in [1, 53] \quad (2)$$

53 estimates were necessary to produce a robust estimation of the transfer function, noted $\Psi_{raw_median} = \text{median}(\{\tilde{\Psi}_i\}_{i \in [1,53]})$, because speech signals are not stationary. The analysis was performed on a single-person recording of 23 seconds after the VAD processing, corresponding to 10 utterances of Combescure’s sentences [69]. As Ψ_{raw_median} is still noisy we performed a smoothing step to obtain $\Psi_{smoothed_median}$ which is plotted in Fig. 2, surrounded by its 10% and 90% percentiles, illustrated by $IQR_{80\%}$.




Fig. 2: Transfer function of the in-ear transducer

The estimated coherence function \tilde{C}_{yx} , defined on Eq. 3 and represented on Fig. 3, highlights an absence of causality between x and y above 3kHz. Hence, Fig. 2 does not make sense above that frequency.

$$\tilde{C}_{yx}(f) = \frac{|P_{yx}(f)|^2}{P_{xx}(f)P_{yy}(f)} \quad (3)$$

This shows that the in-ear BCM allows to only capture relevant speech content inside the ear canal for frequencies $\{f \mid \Psi_{smoothed_median}(f) > -20\text{dB}, \forall f \in \mathbb{R}^+\}$ *i.e.* in a range below 2 kHz. Indeed, Fig. 2 indicates that the in-ear microphone exhibits a very high attenuation at mid and high frequencies: no relevant signal is present in this frequency range. Interestingly, at very low frequencies *i.e.* below 80 Hz, the coherence function in Fig. 3 is also close to zero, which denotes a poor correspondance between the two signals. The physiological sounds (*e.g.* swallowing, blood flow, tongue clicking, teeth grinding) are responsible for this phenomenon, as they are only sensed by the in-ear transducers. Some additional phenomenons like microphonics and movement




Fig. 3: Coherence function of the in-ear transducer

artifacts may also occur. A time domain representation of the synchronized capture in Fig. 4 highlights this difference in the quiet region, for $t > 0.5$ s. For the [80, 300] Hz range, the occlusion effect and small shifts in formant frequencies may occur.

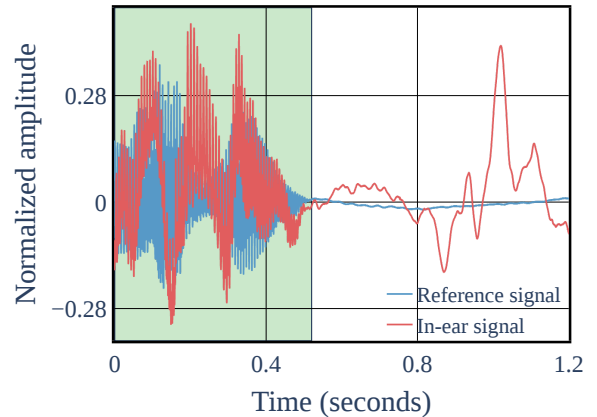


Fig. 4: Time domain representation of speech signals captured in a quiet environment. Active speech is presented in green area.

Finally, two anti-resonances are observed in Fig. 3 at 900 Hz and 1700 Hz, corresponding to vibration nodes of the occluded ear canal and propagation in the bones and tissues of our subject. It is noteworthy that those observations are not universal: acoustic paths differ among speakers because their bone and biological tissue structures are unique - as well as their ear canal geometries and properties - which results in different spectral properties.

C. Simulation of the dataset

Deep learning-based approaches are only efficient in large data regimes; the few minutes of in-ear samples currently available to us are highly insufficient for supervised training. Therefore, we adopted a simulated corrupted wideband speech from the French Librispeech dataset [21] in an in-ear-like

fashion, along with a data augmentation strategy. In the present paper, we simulated two kind of transfer functions to filter the clean speech data: Ψ_{fixed} and Ψ_{random} which are jointly plot on Fig. 2.

In both cases, a gaussian white noise with a power -23 dB below the low-pass filtered signal is added. This noise intends to play the role of physiological noise. It is also masking any high frequency residues.

Ψ_{fixed} : This fixed degradation, used in Sec. V, is obtained using an autoregressive moving-average model. Ψ_{fixed} is a 2nd order low-pass filter with a cutoff frequency of 600 Hz and unitary Q-factor that is applied using a *filtfilt*³ procedure to ensure zero phase shift.

Ψ_{random} : This ever-changing degradation is constructed to fall within the green area of Fig. 2. It is more realistic as it fits better to $\Psi_{smoothed_median}$ and has some randomness. Indeed, Ψ_{random} is sampled from a log-uniform distribution with $IQR_{80\%}$ bounds and brought to a very low gain above 3kHz with an Hann apodization function. We will use this transfer function in Sec. VI-D to discuss the limitations of the linear time-invariant modeling of the in-ear degradation.

Those simulated degradations might lack some realism but still ensure a wide application field for developed algorithms and the ability to focus on the bandwidth extension issue. Naturally, this simulation approach would involve minimizing the disparity between the simulated and real data; however, this process is highly time-consuming and does not guarantee improved performance. Addressing this discrepancy would require incorporating additional speakers, surpassing our current assumption of a linear impulse response, and adopting complex physical models, such as in [74]. Moreover, it would necessitate the realistic blending of pre-recorded physiological noises with speech signals. Instead, we have opted for a simpler yet adequately challenging degradation approach to compare various methods. In parallel, we are focusing on collecting a substantial dataset of speech capture with different BCMS for a subsequent study. Once this dataset is complete, the step of enhancing simulation relevance will be bypassed.

IV. EBEN

A. Theory of Pseudo Quadrature Mirror Filter

The Quadrature Mirror Filter (QMF) banks, introduced in [75], are a set of analysis filters $\{H_i\}_{i \in [0, M-1]}$ used to decompose a signal into several non-overlapping channels of same bandwidth, and synthesis filters $\{G_i\}_{i \in [0, M-1]}$ used to recompose the signal afterward. Fig. 5 shows the entire pipeline. Those filters are obtained from frequency translations of the same low-pass prototype filter $h[n] = \mathcal{Z}^{-1}\{H(z)\}$. A typical frequency response for an M-band Pseudo-QMF (PQMF) bank is given in Fig. 6.

The reconstruction is exact if $\{H_i\}_{i \in [0, M-1]}$ and $\{G_i\}_{i \in [0, M-1]}$ have an infinite support. In practice, this is impossible, but Truong Nguyen proposed a near-perfect reconstruction in [19] by constraining the prototype filter




Fig. 5: PQMF Analysis and Synthesis : block-diagram




Fig. 6: Frequency response of a PQMF filter bank

to be a linear-phase spectral factor of a $2M$ th band filter, significantly reducing aliasing. In other words, the analysis and synthesis impulse responses noted respectively $h_i[n] = \mathcal{Z}^{-1}\{H_i(z)\}$ and $g_i[n] = \mathcal{Z}^{-1}\{G_i(z)\}$, are given by Eq. 4 where N is the filter length.

$$\begin{cases} h_i[n] = 2h[n] \cos\left(\left(2i+1\right)\frac{\pi}{2M}\left(n - \frac{N-1}{2}\right) + (-1)^i \frac{\pi}{4}\right) \\ g_i[n] = 2h[n] \cos\left(\left(2i+1\right)\frac{\pi}{2M}\left(n - \frac{N-1}{2}\right) - (-1)^i \frac{\pi}{4}\right) \end{cases}, \quad 0 \leq n \leq N-1, \quad 0 \leq i \leq M-1 \quad (4)$$

Then, Yuan-Pei Lin and PP Vaidyanathan [76] proposed a more straightforward design methodology by constructing the prototype from a Kaiser window and filling the following conditions:

- Make the prototype filter close to zero out of its passband to minimize the aliasing.

$$|H(e^{j\omega})| \approx 0 \quad \text{for } |\omega| > \frac{\pi}{M} \quad (5)$$

- Make the prototype filter close to one into its passband to minimize the distortion.

$$|H(e^{j\omega})| \approx 1 \quad \text{for } |\omega| < \frac{\pi}{M} \quad (6)$$

Given the desired stopband attenuation and transition bandwidth, these requirements directly translate into a one-degree-of-freedom optimization criterion at the prototype's cutoff frequency. This criterion is minimized to find the optimal cutoff frequency for some M and N . Also note that although minimal band overlap implies very low reconstruction error, it is not equivalent; in fact, phase opposition phenomena between the bands also contribute to the elimination of redundant content in the synthesis phase. In practice, a kernel size of $N = 8M$ is sufficient for pseudo-perfect reconstruction (signal-to-error ratio of 55 dB), and a kernel size of $N = 128M$ is sufficient for pseudo-perfect separation of frequency content between bands. In this article, we have used a convolution kernel of $N = 8M$, which is fast to compute and sufficient to separate frequency content.

³consists in applying a digital filter forward and backward to a signal.

B. Model architecture

1) *Generator*: Unlike frequency approaches [14], [55], [77], which require massive 2D convolutional operations to extract meaningful features from spectrograms or heavy waveform approaches [4], [13], [16], [22], [49], [78] which directly process the audio at the targetted sampling rate, we propose for EBEN to encapsulate a lightweight U-Net-like generator between a PQMF analysis layer and a PQMF synthesis layer. This enclosure reduces the model’s memory footprint by decreasing the first embedding sample rate by a factor of M . It also makes it possible to keep only P subbands with voice content to feed to the first convolution and the last convolution via the most external skip connection. P must lie between 1 and M . Moreover, the number of encoder/decoder blocks is reduced to meet the constraints of real-time applications. Global architecture is exhibited in Fig. 7a and subblocks in Fig. 7b,7c,7d. Convolutions are intertwined with Leaky ReLU activation functions with a negative slope of 0.01. The last non-linearity in the generator is a Hyperbolic tangent placed right before the PQMF synthesis block, in order to bring back values between -1 and 1. Skip connections are additive. We also apply weight normalization [79] on top of every convolution block with trainable weights, in order to ensure a fast convergence during training. Altogether, the EBEN generator is configured by M : the number of PQMF bands and P : the number of bands from which information is extracted.

2) *Discriminators*: EBEN’s discriminators directly exploit the PQMF subbands as inputs without recombining nor upsampling the reconstructed subband signals. We adopt a multiscale ensemble discriminator approach, inspired by the work of Kumar *et al.* in [55], whose inputs are the Q upper bands of the PQMF decomposition, similarly to [80]. Due to the divisibility constraint on the number of input and output channels by the number of groups, Q must be one of $\{1, 2, 3, 5, 6, 10, 15\}$. Like P , it must also satisfy $1 \leq Q \leq M$. The ensemble of discriminators analyzes the generated subband signals at different time scales and helps to improve their quality via the adversarial process, even though each discriminator is relatively simple. The subband discriminators $\{D_k\}_{k \in [1,2,3]}$ exhibit similar receptive fields to the original multiscale MelGAN discriminators [55]. Moreover, we combined our PQMF discriminators with the full scale MelGAN discriminator $D_{k=0}$ to ensure coherence between bands. The exact architecture of discriminators are displayed in Fig. 7e and Fig. 7f together with their positioning in the overall system Fig. 7a. We kept Leaky ReLU as an activation function but used a stronger negative slope of 0.2 to allow for a better gradient transmission to the generator. We also maintained the weight normalization technique. Overall, the EBEN discriminators are configured by M : the number of PQMF bands and Q : the number of enhanced subbands.

C. Loss functions

At each batch, we train alternatively the ensemble of discriminators $\{D_k\}_{k \in [0,1,2,3]}$ to minimize $\mathcal{L}_{\mathcal{D}}$ defined on Eq. 7 and the generator G to minimize $\mathcal{L}_{\mathcal{G}} = \mathcal{L}_{\mathcal{G}}^{adv} + 100 \times \mathcal{L}_{\mathcal{G}}^{rec}$

where $\mathcal{L}_{\mathcal{G}}^{adv}$ and $\mathcal{L}_{\mathcal{G}}^{rec}$ are respectively defined on Eq. 8 and Eq. 9. Our loss setup is inspired by [16]: $\mathcal{L}_{\mathcal{D}}$ and $\mathcal{L}_{\mathcal{G}}^{adv}$ are a classical hinge loss while $\mathcal{L}_{\mathcal{G}}^{rec}$ is a feature matching loss. Using discriminators embeddings for the reconstructive loss allows focusing on the semantic of the signal, which is harder to operate in the time domain because useful information is drowned out amid useless details.

In the underneath definitions, $D_{k,t}^{(l)}$ represents the layer l of the discriminator (among L_k layers) of scale k (among K scales) at time t . $F_{k,l}$ and $T_{k,l}$ are the number of features and temporal length for given indices. We kept x for in-ear signal and y for the reference.

$$\mathcal{L}_{\mathcal{D}} = E_y \left[\frac{1}{K} \sum_{k \in [0,3]} \frac{1}{T_{k,L_k}} \sum_t \max(0, 1 - D_{k,t}(y)) \right] + E_x \left[\frac{1}{K} \sum_{k \in [0,3]} \frac{1}{T_{k,L_k}} \sum_t \max(0, 1 + D_{k,t}(G(x))) \right] \quad (7)$$

$$\mathcal{L}_{\mathcal{G}}^{adv} = E_x \left[\frac{1}{K} \sum_{k \in [0,3]} \frac{1}{T_{k,L_k}} \sum_t \max(0, 1 - D_{k,t}(G(x))) \right] \quad (8)$$

$$\mathcal{L}_{\mathcal{G}}^{rec} = E_x \left[\frac{1}{K} \sum_{k \in [0,3]} \frac{1}{T_{k,l} F_{k,l}} \sum_t \sum_{l \in [1, L_k]} \|D_{k,t}^{(l)}(y) - D_{k,t}^{(l)}(G(x))\|_{L_1} \right] \quad (9)$$

The generator’s loss combination allows to generate audio samples as close as possible to the reference signal thanks to $\mathcal{L}_{\mathcal{G}}^{rec}$, while remaining creative at high frequencies when no information is available in the degraded signal (especially for fricatives) thanks to $\mathcal{L}_{\mathcal{G}}^{adv}$.

V. EXPERIMENTS AND EVALUATION

A. Training strategy

We trained different models [13], [14], [16], [22] and the proposed EBEN model on the French Librispeech [21] dataset resampled uniformly at 16kHz to reverse the Ψ_{fixed} degradation applied on the fly. All the experiments were performed for two days on a single RTX 2080 Ti GPU with a 16 batch size of 2-second randomly sliced audio, corresponding to 13 epochs for the EBEN model. Losses are optimized with Adam [81] using a constant learning rate of $3e-4$ and $\beta = (0.5, 0.9)$ for EBEN and optimizers parameter values found in original papers for the other approaches. No parameter tuning nor early stopping was performed. The EBEN set of hyperparameters is given by $\{M = 4, P = 1, Q = 3\}$. We use $M = 4$ here because this coarse slicing of the spectra is sufficient to separate frequency bands containing valuable cues from non-relevant frequency bands by taking $P = 1$. Such a reduced number of frequency bands also allows to reduce the length

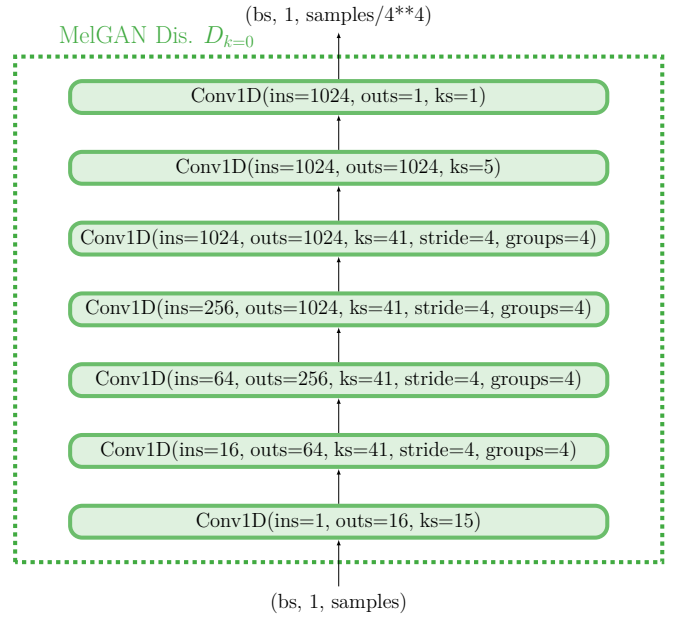

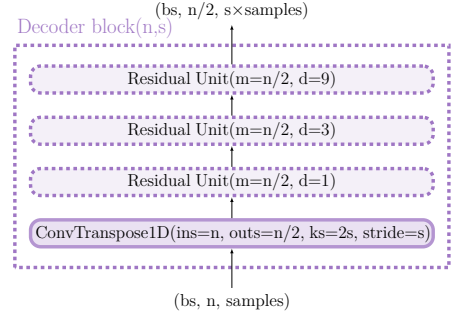
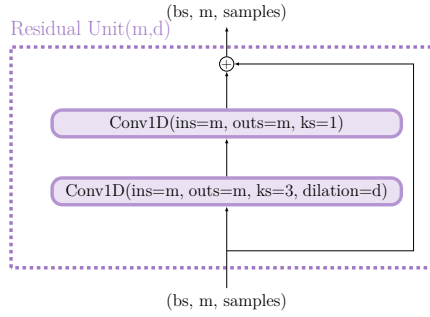
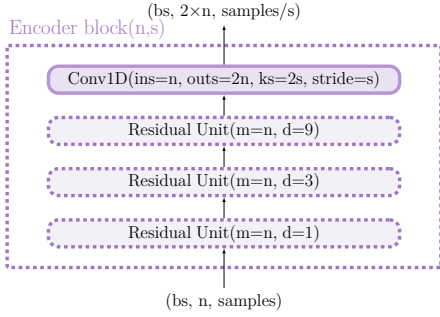
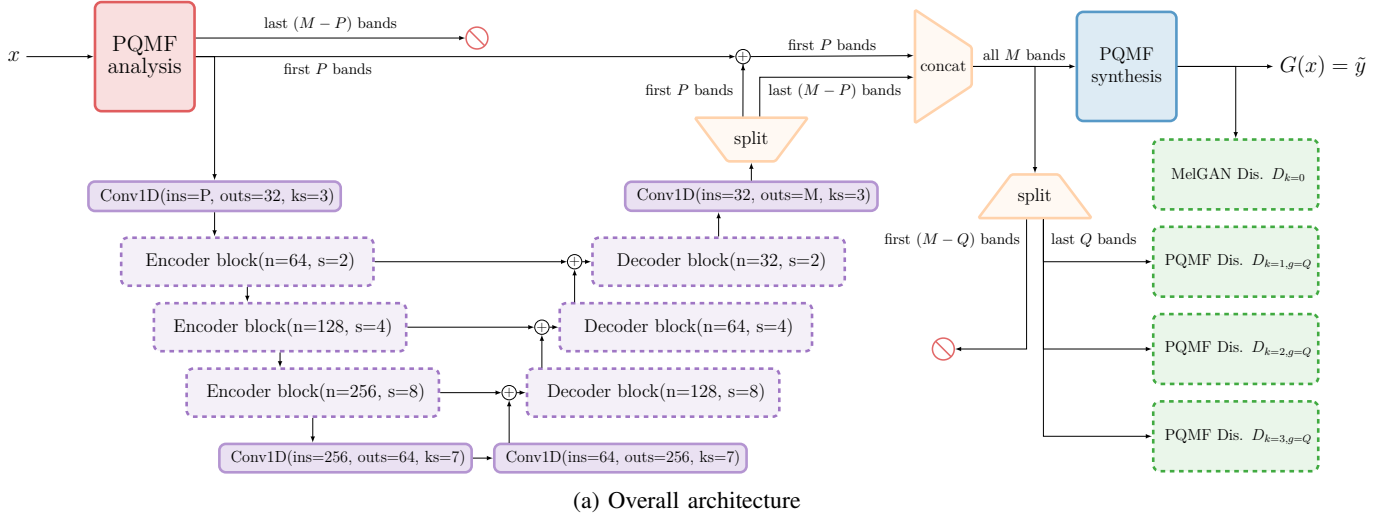


Fig. 7: Architecture of EBEN. *ins*: input channels. *outs*: output channels. *ks*: kernel size

of the PQMF kernel for the analysis and synthesis stages. The value $Q = 3$ was also chosen because we assume that the first frequency band does not require significant enhancement with the proposed degradation.

B. Objective evaluation

1) *Speech quality metrics*: To evaluate the model performances, Tab. I highlights several objective metrics: Perceptual Evaluation of Speech Quality (PESQ) [82], Scale-Invariant Signal-to-Distortion Ratio (SI-SDR) [83], Short-Time Objective Intelligibility (STOI) [84] and Nonesqa-MOS (N-MOS) [85], which was a much better candidate than Nonesqa [86] for this kind of degradation. All the metrics have been computed on the test set for each benchmarked model. Speech enhancement being a one-to-many problem, these results should be analyzed cautiously. Indeed, a plausible signal with perfect intelligibility but still different from reference would be misjudged by the metrics. Note that these metrics are intrusive, since they require groundtruth audio. Generally speaking, speech quality assessment still lacks objective and non-intrusive evaluation metrics, although recent work such as [86] may be part of the solution. This observation is confirmed by [87] which points out that current objective metrics are questionable.

Speech \ Metrics	PESQ	SI-SDR	STOI	N-MOS
Simulated In-ear	2.42 (0.34)	8.4 (3.7)	0.83 (0.05)	2.57 (0.58)
Audio U-net [22]	2.24 (0.49)	11.9 (3.7)	0.87 (0.04)	2.59 (0.44)
Hifi-GAN v3 [14]	1.32 (0.16)	-25.1 (11.4)	0.78 (0.04)	3.70 (0.68)
Seanet [16]	1.92 (0.48)	11.1 (3.0)	0.89 (0.04)	4.25 (0.28)
Streaming-Seanet [13]	2.01 (0.46)	11.2 (3.6)	0.89 (0.04)	3.91 (0.60)
EBEN (ours)	2.08 (0.45)	10.9 (3.3)	0.89 (0.04)	4.02 (0.39)

TABLE I: PESQ/SI-SDR/STOI on test set. Significantly best values (acceptance=0.05) are in **bold**.

Even though purely reconstructive approaches have a clear advantage when evaluated on comparative metrics, Kuleshov’s model [22] does not prevail on STOI, which is the comparative metric that is the most correlated with human evaluation for our specific task, as shown in VI-C. Looking at these results, we could say that best performing models for STOI are either Seanet, Streaming-Seanet or EBEN.

2) *Frugality study*: Enhancing performances need to be qualified by the model’s latency and heaviness to take deep learning from hype to real-world applications. Indeed, the bandwidth extension is applicable for a two-way communication device, if latency is roughly smaller than 20 ms as claimed in [88]. The total number of parameters influencing the memory space should also be reduced. Therefore, we reported on Tab. II :

- P_{gen} : The total number of parameters for the generator, including non-trainable parameters like PQMF-bank for EBEN. For other methods, preprocessing parameters like the mel windows are not counted.
- P_{dis} : The total number of parameters for discriminators.
- τ : The latency corresponding to the generator’s forward pass during inference (no gradients are calculated). We

carefully synchronized GPU to account for any asynchronous execution and chose the fastest kernels by enabling the cudnn benchmark. The reported measures are averaged over 10000 points.

- δ : The maximum memory allocation used during inference measured with `torch.cuda.max_memory_allocated`.

δ and τ are given for a single one-second sample.

Parameters \ Speech	P_{gen}	P_{dis}	τ (ms)	δ (MB)
Audio U-net [22]	71.0 M	\emptyset	37.5	1117.3
Hifi-GAN v3 [14]	1.5 M	70.7 M	3.1	22.2
Seanet [16]	8.3 M	56.6 M	13.1	89.2
Streaming-Seanet [13]	0.7 M	56.6 M	7.5	10.9
EBEN (ours)	1.9 M	27.8 M	4.3	20.0

TABLE II: Parameters, latencies and memory usage of models

Tab. II nuances the simple study of model parameters. Indeed, neither τ nor δ linearly depends on the number of parameters. They are also influenced by models’ depth, embedding width, and hyperparameters that will, for instance, determine the choice of the convolution algorithm (Winograd, FFT, GEMM). Thanks to the reduction operated by PQMF filtering, EBEN is the lightest proportionally to its parameters and one of the fastest networks. It is also more than 3 times faster to infer and 4 times lighter than Seanet [16].

C. Subjective evaluation

1) *Visual inspection of spectrograms*: To visually assess and compare the obtained results with each trained model, Fig. 8 shows some spectrograms obtained from the testing set. It can be observed that a purely reconstructive approach [22] is not sufficient to produce high frequencies. Indeed, when low frequency information is insufficient, the model predicts the mean of speech signals, which is zero. Among generative approaches, our method is competitive. Indeed, EBEN reconstructs a fair amount of formants and minimizes artifacts. As a comparison, Hifi-GAN v3 [14] and Streaming-Seanet [13] are not as efficient for harmonic reconstruction. Seanet [16] seems to be the closest to the reference’s spectrogram. All approaches were able to get rid of the additive Gaussian noise. Some additional zoomable spectrograms confirming these observations are available at <https://jharet.github.io/eben/>.

2) *MUSHRA study*: We conducted a subjective comparative evaluation of the different trained models using the MUltiple Stimuli with Hidden Reference and Anchor [89] (MUSHRA) methodology. According to the MUSHRA specification, a rating scale ranging from 0 to 100 has been used; the higher, the better. A total of 56 samples were rated by 170 participants, corresponding to 7 audios enhanced by five different networks, plus the hidden reference and a hidden low anchor (corresponding to an untrained EBEN network) and the simulated in-ear signal. Participants were recruited by e-mail to complete one of two available tests on the GoListen platform [71] : MUSHRA-Q which allows to rank methods for produced sound quality, and MUSHRA-U, which aims at




Fig. 8: Spectrograms of various bandwidth extension models sandwiched by the simulated in-ear and the reference signals.

ranking methods for ease of speech understanding. Ease of understanding is linked with notions of phonetic confusion and intelligibility (but is not equivalent to standard listening test to assess intelligibility, such as the Modified Rhyme Test [90]), while audio quality reflects the naturalness and listening comfort. For both tests, we recorded participants’ hearing status and type of sound reproduction system to retain 69 participants over 88 for MUSHRA-U and 66 over 82 for MUSHRA-Q. We also followed the two post-screening phases recommended by the International Telecommunication Union (ITU) [89] to retain only participants who provided consistent ratings:

- Stage 1 post-screening: A listener should be excluded from the aggregated responses if he or she rates the hidden reference condition for at least 15% of the test items lower than a score of 90.
- Stage 2 post-screening: Exclude subjects whose individual grades fall outside $1.5 \times$ the upper or lower bound of the IQR of the aggregated listeners for at least 25% of the test items.

After applying those two criteria, we retain 47/88 for MUSHRA-U and 43/82 for MUSHRA-Q. The overall age repartition is as follows: 37% are below 27 years old, 24% are above 50, and 39% between 27 and 50 years old. We found no statistically significant differences in ratings between the age categories. The distribution of obtained gradings are shown Fig. 9 and Fig. 10.

The statistical distributions have been studied using a non-parametric Friedmann Analysis of Variance to confirm the statistical significance of the results. The obtained p-values are lower than $1e-20$, demonstrating that there are significant differences, both in terms of quality and intelligibility among tested approaches. This made it possible to perform a post-hoc Nemenyi-Friedmann analysis, in order to assess the 2-to-2 independence of the distributions. The obtained results show that the EBEN approach ranks first ex aequo with Seanet in terms of quality and ease of understanding (no statistically significant difference between EBEN and Seanet, p-value

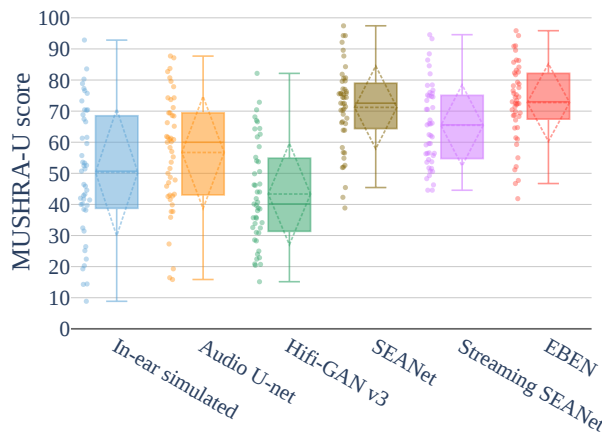


Fig. 9: MUSHRA-U : statistical distributions of scores obtained with a MUSHRA procedure for the ranking of perceived ease of understanding across trained models.

> 0.5) and that these two methods significantly outperform the second best approach Streaming Seanet (p-value < 0.005).

VI. DISCUSSION

A. PQMF insights

PQMF banks are helpful for a wide range of tasks, including audio equalization, noise reduction, or compression, e.g., by reducing the bit rate on sparser bands. This work has used the PQMF analysis outputs to speed up the inference by taking advantage of the decimation operator. The multiband representation has the same dimensionality as the original signal but is condensed along the time axis and extended along channels, allowing parallel computing. Also, by the very nature of our problem, some frequency bands of the input signal do not contain any information, and we can drop them. Furthermore, generating bands reduces redundancy, leading again to a reduction in computational complexity. Finally, it allows the design of discriminators for EBEN that act only where bandwidth extension is needed.

Along with EBEN source code, we also provide a modern and efficient implementation of the PQMF analysis and




Fig. 10: MUSHRA-Q : statistical distributions of scores obtained with a MUSHRA procedure for the ranking of perceived sound quality across trained models.

synthesis with native Pytorch functions, using only strided convolutions and strided transposed convolutions.

B. EBEN’s Configurability

We called our work ”Configurable” because several aspects of EBEN’s architecture can be adapted to address different BCM degradations. The corresponding hyperparameters are:

- M : The number of PQMF bands. It has a direct impact on the width of each band’s frequency bandwidth. Higher values of this parameter enable finer control over the bands, enhancing precision. Furthermore, it influences the downsampling factor. Raising the number of bands reduces the computational burden on the network, but it may result in some loss of temporal resolution. As pointed in [91], downsampling proves to be the most effective option when finetuning models for faster inference, as it offers significant computational gains with minimal performance trade-offs. Moreover, when the number of bands is considerably large, the empirical equation $N = 8M$ may require a convolution in the Fourier domain.
- P : The number of informative PQMF bands sent to the generator. Given the sampling frequency F_s and the cutoff frequency of the low-pass filter F_c , the factor $\frac{P}{M}$ should be just above the reduced cutoff frequency of the degradation $\frac{2F_c}{F_s}$. In this way, all the information remaining in the low-pass filtered signal is captured and the high-frequency noise is discarded. We carried out several ablation studies with EBEN, by varying the number of P bands for the specific degradation described in this article, using a constant number of bands Q sent to the discriminators. By performing these ablation studies, we were able to determine that the inclusion of informationless bands significantly degrades the objective metrics obtained for the enhanced signals when some high-frequency noise is present, but has a minimal effect when the noise is absent. For a pure bandwidth extension

task from $F_c = 4$ kHz to $F_c = 16$ kHz, when there is absolutely no content above the initial half sampling frequency, a value of $P = 1$ is also enough to obtain excellent results (STOI = 0.93 / N-MOS = 4.23 on test set after enhancement). Those conclusions of course heavily depend on the kind of degradation, which further highlights the benefits of a configurable approach.

- Q : The number of PQMF bands sent to the discriminators to be refined. In fact, the very first bands should contain clean low frequencies of the speech signal and would only require a slight modification by the MelGAN discriminator to suppress physiological noise. On the contrary, the last bands, which suffer from information loss, must be filled by the generator network, pushed to do so thanks to the PQMF discriminators. We also carried out several ablation studies with EBEN, by varying Q in [1; 3] for a fixed value of $P = 1$. This study allows to draw a clear relationship between the degree of band refinement and the corresponding enhancement performance: as the number of refined bands decreases, the enhancement performance diminishes accordingly. We believe that discriminator configurability is once again beneficial for specific enhancement objectives in targeted frequency bands. Indeed, some BCMs exhibit a very sharp rolloff. In these cases, P and Q should be chosen to be complementary. On the other hand, other kinds of BCMs exhibit a smaller high-frequency rolloff. In such scenarios, EBEN’s configurability allows Q to be chosen so that the frequency bands input to the PQMF discriminators overlap with the P frequency bands fed to the generator. This design allows to refine the upper bands that are degraded, even if there is still some residual speech content that can be useful to the generator for the bandwidth extension task.

C. Correlation between subjective and objective metrics

General purpose intrusive (or comparative) metrics for assessing speech intelligibility and quality are far from perfect. However, the existing literature still relies on those common set of metrics for evaluation purposes for lack of anything better. In the present study for example, the PESQ metric ranks the Audio U-net [22] approach as the best method, which contradicts the results obtained with subjective evaluation. Having MUSHRA test results and the various metrics at our disposal enables us to perform a correlation analysis between the objective metrics and the subjective tests, in order to provide an overview of which metrics are best/less suited to the bandwidth extension task. The results are shown in Fig .11.

We can deduce from the coefficients shown in Fig .11 that SI-SDR and PESQ are not suitable to evaluate the quality of bandwidth extension methods because of their poor correlation with MUSHRA. On the other hand, this analysis allows to conclude that STOI and Noresqa-MOS are two relevant indicators. Hence, the pseudo-ranking of the relevance level of the metrics for our specific use case is Noresqa-MOS \approx STOI \gg SI-SDR \gg PESQ. It is also noteworthy that

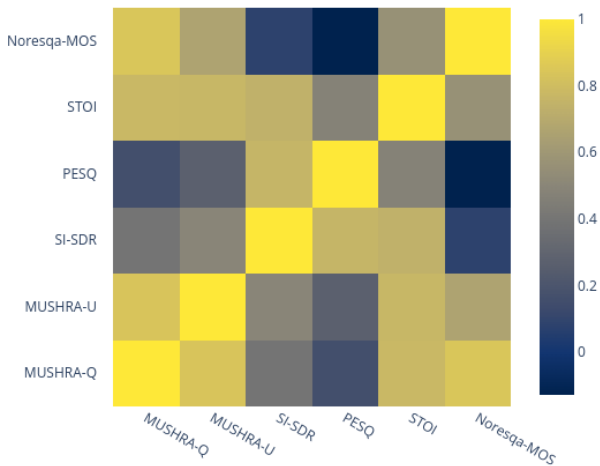


Fig. 11: Pearson product-moment correlation coefficients of objective and subjective metrics

Noresqa-MOS is more correlated with MUSHRA-Q than with MUSHRA-U. This seems logical since Noresqa-MOS is built to predict quality.

D. Accordance of the synthetic data generation

In this part, we would like to reflect on the relevance of the generated synthetic data to tackle some real in-ear captured speech. To do this, we used two EBEN models with the same configuration. One was trained to reverse the Ψ_{fixed} degradation and discussed in V, while the other was trained on Ψ_{random} which is closer to the real degradation by design. We chose EBEN among the other approaches, but this section is independent to the model choice. Although we have selected two models that perform well with in-distribution data, the results obtained on Tab. III show that neither model is able to significantly enhance the raw in-ear speech signal according to objective metrics.

Speech	PESQ	SI-SDR	STOI	N-MOS
In-ear enhanced via EBEN trained on Ψ_{fixed}	1.16	-37.4	0.51	3.82
In-ear enhanced via EBEN trained on Ψ_{random}	1.28	-41.6	0.53	3.80
Raw In-ear	1.5	-37.0	0.56	3.33

TABLE III: EBEN’s ability to enhance real data according to different training sets

Efforts to get closer to $\Psi_{smoothed_median}$ for Ψ_{random} did not pay off because the complex degradation cannot be accurately simulated by a linear transfer function. Rather, the degradation is likely non-linear. Moreover, the additive physiological and frictional noise is time-dependent, making the assumption of a linear time-invariant system untrue in practice. Therefore, instead of investing a significant amount of time and effort to create a suitable simulation model, we will try to use real data in our future works. Indeed, the data-driven nature of deep learning suggests that the training set should be based on real data: we are in the process of building and releasing a complete BCM recording dataset.

We presented Configurable EBEN: a state-of-the-art, real-time compatible, and lightweight neural network architecture to address the problem of unimodal enhancement of speech signals captured with noise-resilient body-conduction microphones. The main challenge encountered with these unconventional microphones is the need to achieve a bandwidth extension of the raw captured signals. We therefore designed EBEN to be fully configurable for the bandwidth enlargement needed. We specifically proposed a multiband approach, where the enhancement is solely conditioned on the first P informative bands, and the adversarial training is mainly targeted to enhance Q bands over a total of M bands through newly designed discriminators. Furthermore, this multiband decomposition – which is using Pseudo Quadrature Mirror Filter bank – enables a reduction of the feature dimensionality from the very first layer of the encoder. This benefits streaming compatibility, because fewer computations are required during the forward pass and reduce data redundancy. Those findings are supported by extensive experimentation and comparisons with existing models. These experiments demonstrate that EBEN is competitive in many aspects, including enhancement performance, latency, and memory footprint. EBEN is therefore a good compromise between frugal AI requirements and speech enhancement performance; ready to be trained on a real BCMs dataset.

Acknowledgements: This work has been partially funded by the French National Research Agency under the ANR Grant No. ANR-20-THIA-0002. This work was also granted access to the HPC/AI resources of [CINES / IDRIS / TGCC] under the allocation 2022-AD011013469 made by GENCI.

REFERENCES

- [1] H. S. Shin, H.-G. Kang, and T. Fingscheidt, “Survey of speech enhancement supported by a bone conduction microphone,” in *Speech Communication; 10. ITG Symposium*. VDE, 2012, pp. 1–4.
- [2] M. McBride, P. Tran, T. Letowski, and R. Patrick, “The effect of bone conduction microphone locations on speech intelligibility and sound quality,” *Applied ergonomics*, vol. 42, no. 3, pp. 495–502, 2011.
- [3] M. Li, I. Cohen, and S. Mousazadeh, “Multisensory speech enhancement in noisy environments using bone-conducted and air-conducted microphones,” in *2014 IEEE China Summit & International Conference on Signal and Information Processing (ChinaSIP)*. IEEE, 2014, pp. 1–5.
- [4] Y. Li, Y. Wang, X. Liu, Y. Shi, S. Patel, and S.-F. Shih, “Enabling real-time on-chip audio super resolution for bone-conduction microphones,” *Sensors*, vol. 23, no. 1, p. 35, 2022.
- [5] B. Acker-Mills, A. Houtsma, and W. Ahroon, “Speech intelligibility with acoustic and contact microphones,” ARMY AEROMEDICAL RESEARCH LAB FORT RUCKER AL, Tech. Rep., 2005.
- [6] A. Shahina and B. Yegnanarayana, “Mapping speech spectra from throat microphone to close-speaking microphone: A neural network approach,” *EURASIP Journal on Advances in Signal Processing*, vol. 2007, pp. 1–10, 2007.
- [7] M. T. Turan and E. Erzin, “Enhancement of throat microphone recordings by learning phone-dependent mappings of speech spectra,” in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2013, pp. 7049–7053.
- [8] J. C. Bos, D. W. Tack, and L. L. Bossi, “Speech input hardware investigation for future dismounted soldier computer systems,” *DRCD Toronto CR*, vol. 64, p. 2005, 2005.

- [9] R. E. Bouserhal, T. H. Falk, and J. Voix, "In-ear microphone speech quality enhancement via adaptive filtering and artificial bandwidth extension," *The Journal of the Acoustical Society of America*, vol. 141, no. 3, pp. 1321–1331, 2017.
- [10] H. Park, Y.-S. Shin, and S.-H. Shin, "Speech quality enhancement for in-ear microphone based on neural network," *IEICE TRANSACTIONS on Information and Systems*, vol. 102, no. 8, pp. 1594–1597, 2019.
- [11] M. Ohlenbusch, C. Rollwage, and S. Doclo, "Training strategies for own voice reconstruction in hearing protection devices using an in-ear microphone," in *2022 International Workshop on Acoustic Signal Enhancement (IWAENC)*, 2022, pp. 1–5.
- [12] J. G. Casali and E. H. Berger, "Technology advancements in hearing protection circa 1995: Active noise reduction, frequency/amplitude-sensitivity, and uniform attenuation," *American Industrial Hygiene Association Journal*, vol. 57, no. 2, pp. 175–185, 1996.
- [13] Y. Li, M. Tagliasacchi, O. Rybakov, V. Ungureanu, and D. Roblek, "Real-time speech frequency bandwidth extension," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 691–695.
- [14] J. Kong, J. Kim, and J. Bae, "Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis," *Advances in Neural Information Processing Systems*, vol. 33, pp. 17 022–17 033, 2020.
- [15] A. Défossez, G. Synnaeve, and Y. Adi, "Real time speech enhancement in the waveform domain," *Proc. Interspeech*, 2020.
- [16] M. Tagliasacchi, Y. Li, K. Misiunas, and D. Roblek, "Seanet: A multi-modal speech enhancement network," *Proc. Interspeech 2020*, pp. 1126–1130, 2020.
- [17] N. Zeghidour, A. Luebs, A. Omran, J. Skoglund, and M. Tagliasacchi, "Soundstream: An end-to-end neural audio codec," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 495–507, 2021.
- [18] A. Défossez, J. Copet, G. Synnaeve, and Y. Adi, "High fidelity neural audio compression," *arXiv preprint arXiv:2210.13438*, 2022.
- [19] T. Q. Nguyen, "Near-perfect-reconstruction pseudo-qmf banks," *IEEE Transactions on signal processing*, vol. 42, no. 1, pp. 65–76, 1994.
- [20] J. Hauret, T. Joubaud, V. Zimpfer, and É. Bavu, "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)*, 2023, pp. 1–5.
- [21] V. Pratap, Q. Xu, A. Sriram, G. Synnaeve, and R. Collobert, "Mls: A large-scale multilingual dataset for speech research," *Proc. Interspeech*, 2020.
- [22] V. Kuleshov, S. Z. Enam, and S. Ermon, "Audio super-resolution using neural nets," in *Proceedings of International Conference on Learning Representations (ICLR)*, 2017.
- [23] J. Makhoul and M. Berouti, "High-frequency regeneration in speech coding systems," in *ICASSP'79. IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 4. IEEE, 1979, pp. 428–431.
- [24] S. Chennoukh, A. Gerrits, G. Miet, and R. Sluijter, "Speech enhancement via frequency bandwidth extension using line spectral frequencies," in *2001 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings (Cat. No. 01CH37221)*, vol. 1. IEEE, 2001, pp. 665–668.
- [25] J. Epps, "Wideband extension of narrowband speech for enhancement and coding," Ph.D. dissertation, UNSW Sydney, 2000.
- [26] A. De Cheveigné and H. Kawahara, "Yin, a fundamental frequency estimator for speech and music," *The Journal of the Acoustical Society of America*, vol. 111, no. 4, pp. 1917–1930, 2002.
- [27] B. Iser and G. Schmidt, "Bandwidth extension of telephony speech," in *Speech and Audio Processing in Adverse Environments*. Springer, 2008, pp. 135–184.
- [28] Y. Yoshida and M. Abe, "An algorithm to reconstruct wideband speech from narrowband speech based on codebook mapping," in *ICSLP*, vol. 94, 1994, pp. 1591–1594.
- [29] K.-Y. Park and H. S. Kim, "Narrowband to wideband conversion of speech using gmm based transformation," in *2000 IEEE international conference on acoustics, speech, and signal processing. Proceedings (Cat. No. 00CH37100)*, vol. 3. IEEE, 2000, pp. 1843–1846.
- [30] P. Jax and P. Vary, "On artificial bandwidth extension of telephone speech," *Signal Processing*, vol. 83, no. 8, pp. 1707–1719, 2003.
- [31] B. Iser and G. Schmidt, "Neural networks versus codebooks in an application for bandwidth extension of speech signals," in *Eighth European Conference on Speech Communication and Technology*, 2003.
- [32] A. v. d. Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, "Wavenet: A generative model for raw audio," *International Speech and Communication Association (ISCA)*, 2016.
- [33] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerrv-Ryan *et al.*, "Natural tts synthesis by conditioning wavenet on mel spectrogram predictions," in *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2018, pp. 4779–4783.
- [34] R. Prenger, R. Valle, and B. Catanzaro, "Waveglow: A flow-based generative network for speech synthesis," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 3617–3621.
- [35] Z. Kong, W. Ping, J. Huang, K. Zhao, and B. Catanzaro, "Diffwave: A versatile diffusion model for audio synthesis," *Proceedings of International Conference on Learning Representations (ICLR)*, 2021.
- [36] R. Yamamoto, E. Song, and J.-M. Kim, "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, 2020, pp. 6199–6203.
- [37] S. Davis and P. Mermelstein, "Comparison of parametric representations for monosyllabic word recognition in continuously spoken sentences," *IEEE transactions on acoustics, speech, and signal processing*, vol. 28, no. 4, pp. 357–366, 1980.
- [38] B. P. Bogert, "The quefrency alalysis of time series for echoes; cepstrum, pseudo-autocovariance, cross-cepstrum and saphe cracking," *Time series analysis*, pp. 209–243, 1963.
- [39] W. Dai, C. Dai, S. Qu, J. Li, and S. Das, "Very deep convolutional neural networks for raw waveforms," in *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2017, pp. 421–425.
- [40] F. G. Germain, Q. Chen, and V. Koltun, "Speech denoising with deep feature losses," *Proc. Interspeech 2019*, pp. 2723–2727, 2019.
- [41] 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.
- [42] K. Goel, A. Gu, C. Donahue, and C. Ré, "It's raw! audio generation with state-space models," *arXiv preprint arXiv:2202.09729*, 2022.
- [43] Z.-H. Ling, Y. Ai, Y. Gu, and L.-R. Dai, "Waveform modeling and generation using hierarchical recurrent neural networks for speech bandwidth extension," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 5, pp. 883–894, 2018.
- [44] S. Birnbaum, V. Kuleshov, Z. Enam, P. W. W. Koh, and S. Ermon, "Temporal film: Capturing long-range sequence dependencies with feature-wise modulations," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [45] X. Hao, C. Xu, N. Hou, L. Xie, E. S. Chng, and H. Li, "Time-domain neural network approach for speech bandwidth extension," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 866–870.
- [46] C. Yu, H. Lu, N. Hu, M. Yu, C. Weng, K. Xu, P. Liu, D. Tuo, S. Kang, G. Lei *et al.*, "Dorian: Duration informed attention network for multimodal synthesis," *Proc. Interspeech*, 2020.
- [47] G. Yang, S. Yang, K. Liu, P. Fang, W. Chen, and L. Xie, "Multi-band melgan: Faster waveform generation for high-quality text-to-speech," in *2021 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2021, pp. 492–498.
- [48] A. Caillon and P. Esling, "Rave: A variational autoencoder for fast and high-quality neural audio synthesis," *arXiv preprint arXiv:2111.05011*, 2021.
- [49] H. Wang and D. Wang, "Towards robust speech super-resolution," *IEEE/ACM transactions on audio, speech, and language processing*, vol. 29, pp. 2058–2066, 2021.
- [50] D. Stoller, S. Ewert, and S. Dixon, "Wave-u-net: A multi-scale neural network for end-to-end audio source separation," *arXiv preprint arXiv:1806.03185*, 2018.
- [51] A. Bosca, A. Guerin, L. Perotin, and S. Kitic, "Dilated u-net based approach for multichannel speech enhancement from first-order ambisonics

- recordings,” in *2020 28th European Signal Processing Conference (EUSIPCO)*. IEEE, 2021, pp. 216–220.
- [52] M. D. Zeiler, D. Krishnan, G. W. Taylor, and R. Fergus, “Deconvolutional networks,” in *2010 IEEE Computer Society Conference on computer vision and pattern recognition*. IEEE, 2010, pp. 2528–2535.
- [53] W. Shi, J. Caballero, F. Huszár, J. Totz, A. P. Aitken, R. Bishop, D. Rueckert, and Z. Wang, “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*, 2016, pp. 1874–1883.
- [54] A. Odena, V. Dumoulin, and C. Olah, “Deconvolution and checkerboard artifacts,” *Distill*, 2016. [Online]. Available: <http://distill.pub/2016/deconv-checkerboard>
- [55] K. Kumar, R. Kumar, T. de Boissiere, L. Gestin, W. Z. Teoh, J. Sotelo, A. de Brébisson, Y. Bengio, and A. C. Courville, “Melgan: Generative adversarial networks for conditional waveform synthesis,” *Advances in neural information processing systems*, vol. 32, 2019.
- [56] S. Kim and V. Sathe, “Bandwidth extension on raw audio via generative adversarial networks,” *arXiv preprint arXiv:1903.09027*, 2019.
- [57] S. E. Eskimez, K. Koishida, and Z. Duan, “Adversarial training for speech super-resolution,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 13, no. 2, pp. 347–358, 2019.
- [58] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative adversarial networks,” *Communications of the ACM*, vol. 63, no. 11, pp. 139–144, 2020.
- [59] B. Feng, Z. Jin, J. Su, and A. Finkelstein, “Learning bandwidth expansion using perceptually-motivated loss,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 606–610.
- [60] H.-P. Liu, Y. Tsao, and C.-S. Fuh, “Bone-conducted speech enhancement using deep denoising autoencoder,” *Speech Communication*, vol. 104, pp. 106–112, 2018.
- [61] D. Shan, X. Zhang, C. Zhang, and L. Li, “A novel encoder-decoder model via ns-lstm used for bone-conducted speech enhancement,” *IEEE Access*, vol. 6, pp. 62 638–62 644, 2018.
- [62] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative adversarial nets,” *Advances in neural information processing systems*, vol. 27, 2014.
- [63] M. Wang, J. Chen, X. Zhang, Z. Huang, and S. Rahardja, “Multi-modal speech enhancement with bone-conducted speech in time domain,” *Applied Acoustics*, vol. 200, p. 109058, 2022.
- [64] H. Wang, X. Zhang, and D. Wang, “Fusing bone-conduction and air-conduction sensors for complex-domain speech enhancement,” *IEEE/ACM transactions on audio, speech, and language processing*, vol. 30, pp. 3134–3143, 2022.
- [65] R. E. Bouserhal, P. Chabot, M. Sarria-Paja, P. Cardinal, and J. Voix, “Classification of nonverbal human produced audio events: a pilot study,” 2018.
- [66] M. Wang, J. Chen, X.-L. Zhang, and S. Rahardja, “End-to-end multi-modal speech recognition on an air and bone conducted speech corpus,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, pp. 1–12, 2022.
- [67] R. E. Bouserhal, A. Bernier, and J. Voix, “An in-ear speech database in varying conditions of the audio-phonation loop,” *The Journal of the Acoustical Society of America*, vol. 145, no. 2, pp. 1069–1077, 2019.
- [68] “BIONEAR outstanding hearing for 4.0 professionals,” <https://www.cotral-communication.com/en/industry.html>, accessed: 2022-10-10.
- [69] P. Combescure, “20 listes de dix phrases phonétiquement équilibrées,” 1981.
- [70] H. Brumm and S. A. Zollinger, “The evolution of the lombard effect: 100 years of psychoacoustic research,” *Behaviour*, vol. 148, no. 11-13, pp. 1173–1198, 2011.
- [71] D. Barry, Q. Zhang, P. W. Sun, and A. Hines, “Go listen: an end-to-end online listening test platform,” *Journal of Open Research Software*, vol. 9, no. 1, 2021.
- [72] M. K. Brummund, F. Sgard, Y. Petit, and F. Laville, “Three-dimensional finite element modeling of the human external ear: Simulation study of the bone conduction occlusion effect,” *The Journal of the Acoustical Society of America*, vol. 135, no. 3, pp. 1433–1444, 2014.
- [73] P. Welch, “The use of fast fourier transform for the estimation of power spectra: a method based on time averaging over short, modified periodograms,” *IEEE Transactions on audio and electroacoustics*, vol. 15, no. 2, pp. 70–73, 1967.
- [74] C. Blondé-Weinmann, T. Joubaud, V. Zimpfer, P. Hamery, and S. Roth, “Numerical and experimental investigation of the sound transmission delay from a skin vibration to the occluded ear canal,” *Journal of Sound and Vibration*, vol. 542, p. 117345, 2023.
- [75] J. Rothweiler, “Polyphase quadrature filters—a new subband coding technique,” in *ICASSP’83. IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 8. IEEE, 1983, pp. 1280–1283.
- [76] Y.-P. Lin and P. Vaidyanathan, “A kaiser window approach for the design of prototype filters of cosine modulated filterbanks,” *IEEE signal processing letters*, vol. 5, no. 6, pp. 132–134, 1998.
- [77] M. Lagrange and F. Gontier, “Bandwidth extension of musical audio signals with no side information using dilated convolutional neural networks,” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 801–805.
- [78] J. Su, Y. Wang, A. Finkelstein, and Z. Jin, “Bandwidth extension is all you need,” in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 696–700.
- [79] T. Salimans and D. P. Kingma, “Weight normalization: A simple reparameterization to accelerate training of deep neural networks,” *Advances in neural information processing systems*, vol. 29, 2016.
- [80] A. Mustafa, N. Pia, and G. Fuchs, “Stylemelgan: An efficient high-fidelity adversarial vocoder with temporal adaptive normalization,” in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 6034–6038.
- [81] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, Y. Bengio and Y. LeCun, Eds., 2015.
- [82] A. W. Rix, J. G. Beerends, M. P. Hollier, and A. P. Hekstra, “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, 2001, pp. 749–752.
- [83] J. Le Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, “Sdr—half-baked or well done?” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 626–630.
- [84] C. H. Taal, R. C. Hendriks, R. Heusdens, and J. Jensen, “A short-time objective intelligibility measure for time-frequency weighted noisy speech,” in *2010 IEEE international conference on acoustics, speech and signal processing*. IEEE, 2010, pp. 4214–4217.
- [85] P. Manocha and A. Kumar, “Speech quality assessment through mos using non-matching references,” in *Interspeech*, 2022. [Online]. Available: <https://arxiv.org/abs/2206.12285>
- [86] P. Manocha, B. Xu, and A. Kumar, “Noresqa: A framework for speech quality assessment using non-matching references,” *Advances in Neural Information Processing Systems*, vol. 34, 2021.
- [87] A. Vinay and A. Lerch, “Evaluating generative audio systems and their metrics,” *arXiv preprint arXiv:2209.00130*, 2022.
- [88] N. Lezzoum, G. Gagnon, and J. Voix, “Echo threshold between passive and electro-acoustic transmission paths in digital hearing protection devices,” *International Journal of Industrial Ergonomics*, vol. 53, pp. 372–379, 2016.
- [89] R. BS.1534-3, “Method for the subjective assessment of intermediate quality level of audio systems,” *International Telecommunication Union Radiocommunication Assembly*, 2015.
- [90] A. S. House, C. Williams, M. H. Hecker, and K. D. Kryter, “Psychoacoustic speech tests: A modified rhyme test,” *The Journal of the Acoustical Society of America*, vol. 35, no. 11, pp. 1899–1899, 1963.
- [91] S. Zaiem, R. Algayres, T. Parcollet, S. Essid, and M. Ravanelli, “Fine-tuning strategies for faster inference using speech self-supervised models: a comparative study,” in *ICASSP 2023-International Conference on Acoustics, Speech, and Signal Processing*, 2023.



Julien Hauret is a PhD candidate at Cnam Paris, pursuing research in machine learning applied to acoustics. He holds two MSc from ENS Paris Saclay, one in Electrical Engineering (2020) and the other in Applied Mathematics (2021). He has a strong research training as evidenced by his internships at Columbia University and the French Ministry of the Armed Forces. He also lectures on algorithms and data structures at the École des Ponts ParisTech. His research focuses on the use of deep learning for speech enhancement applied to body-conducted

speech. With a passion for interdisciplinary collaboration, Julien aims to improve human communication through technology.



Thomas Joubaud is a Research Associate at the Acoustics and Soldier Protection department within the French-German Research Institute of Saint-Louis (ISL), France, since 2019. In 2013, he received the graduate degree from Ecole Centrale Marseille, France, as well as the master's degree in Mechanics, Physics and Engineering, specialized in Acoustical Research, of the Aix-Marseille University, France. He earned the Ph.D. degree in Mechanics, specialized in Acoustics, of the Conservatoire National des Arts et Métiers (Cnam), Paris, France, in 2017. The

thesis was carried out in collaboration with and within the ISL. From 2017 to 2019, he worked as a post-doctorate research engineer with Orange SA company in Cesson-Sévigné, France. His research interests include audio signal processing, hearing protection, psychoacoustics, especially speech intelligibility and sound localization, and high-level continuous and impulse noise measurement.



Véronique Zimpfer is a Scientific Researcher at the Acoustics and Soldier Protection department within the French-German Research Institute of Saint-Louis (ISL), Saint-Louis, France, since 1997. She holds a M.Sc in Signal Processing from the Grenoble INP, France and obtained a PhD in Acoustics from INSA Lyon, France, in 2000. Her expertise lies at the intersection of communication in noisy environments and auditory protection. Her research focuses on improving adaptive auditory protectors, refining radio communication strategies through unconventional microphone methods, and enhancing auditory perception while utilizing protective gear.

radio communication strategies through unconventional microphone methods, and enhancing auditory perception while utilizing protective gear.



Éric Bavu is a Full Professor of Acoustics and Signal Processing at the Laboratoire de Mécanique des Structures et des Systèmes Couplés (LMSSC) within the Conservatoire National des Arts et Métiers (Cnam), Paris, France. He completed his undergraduate studies at École Normale Supérieure de Cachan, France, from 2001 to 2005. In 2005, he earned an M.Sc in Acoustics, Signal Processing, and Computer Science Applied to Music from Université Pierre et Marie Curie Sorbonne University (UPMC), followed by a Ph.D. in Acoustics jointly awarded

by Université de Sherbrooke, Canada, and UPMC, France, in 2008. He also conducted post-doctoral research on biological soft tissues imaging at the Langevin Institute at École Supérieure de Physique et Chimie ParisTech (ESPCI), France. Since 2009, he has supervised six Ph.D. students at LMSSC, focusing on time domain audio signal processing for inverse problems, 3D audio, and deep learning for audio. His current research interests encompass deep learning methods applied to inverse problems in acoustics, moving sound source localization and tracking, speech enhancement, and speech recognition.