



AALBORG UNIVERSITY
DENMARK

Aalborg Universitet

Minimum Processing Near-End Listening Enhancement

Fuglsig, Andreas Jonas; Jensen, Jesper; Tan, Zheng Hua; Bertelsen, Lars Sondergaard; Lindof, Jens Christian; Ostergaard, Jan

Published in:

IEEE/ACM Transactions on Audio, Speech, and Language Processing

DOI (link to publication from Publisher):

[10.48550/arXiv.2210.17154](https://doi.org/10.48550/arXiv.2210.17154)

[10.1109/TASLP.2023.3282094](https://doi.org/10.1109/TASLP.2023.3282094)

Creative Commons License

CC BY-NC-ND 4.0

Publication date:

2023

Document Version

Publisher's PDF, also known as Version of record

[Link to publication from Aalborg University](#)

Citation for published version (APA):

Fuglsig, A. J., Jensen, J., Tan, Z. H., Bertelsen, L. S., Lindof, J. C., & Ostergaard, J. (2023). Minimum Processing Near-End Listening Enhancement. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 31, 2233-2245. <https://doi.org/10.48550/arXiv.2210.17154>, <https://doi.org/10.1109/TASLP.2023.3282094>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

Minimum Processing Near-End Listening Enhancement

Andreas Jonas Fuglsig , *Student Member, IEEE*, Jesper Jensen , Zheng-Hua Tan , *Senior Member, IEEE*, Lars Søndergaard Bertelsen , Jens Christian Lindof, *Member, IEEE*, and Jan Østergaard , *Senior Member, IEEE*

Abstract—The intelligibility and quality of speech from a mobile phone or public announcement system are often affected by background noise in the listening environment. By pre-processing the speech signal it is possible to improve the speech intelligibility and quality — this is known as near-end listening enhancement (NLE). Although, existing NLE techniques are able to greatly increase intelligibility in harsh noise environments, in favorable noise conditions the intelligibility of speech reaches a ceiling where it cannot be further enhanced. Actually, the focus of existing methods solely on improving the intelligibility causes unnecessary processing of the speech signal and leads to speech distortions and quality degradations. In this article, we provide a new rationale for NLE, where the target speech is minimally processed in terms of a processing penalty, provided that a certain performance constraint, e.g., intelligibility, is satisfied. We present a closed-form solution for the case where the performance criterion is an intelligibility estimator based on the approximated speech intelligibility index and the processing penalty is the mean-square error between the processed and the clean speech. This produces an NLE method that adapts to changing noise conditions via a simple gain rule by limiting the processing to the minimum necessary to achieve a desired intelligibility, while at the same time focusing on quality in favorable noise situations by minimizing the amount of speech distortions. Through simulation studies, we show the proposed method attains speech quality on par or better than existing methods in both objective measurements and subjective listening tests, whilst still sustaining objective speech intelligibility performance on par with existing methods.

Index Terms—Minimum processing, adaptive, near-end listening enhancement, speech intelligibility, speech quality, approximated speech intelligibility index, optimization.

I. INTRODUCTION

REAL-LIFE speech communication, e.g., with mobile phones or public announcements, takes place in a large variety of often noisy places. Here, environmental noises such as cars, trains, construction work and other people talking may

Manuscript received 31 October 2022; revised 20 April 2023; accepted 25 May 2023. Date of publication 5 June 2023; date of current version 14 June 2023. This work was supported by Innovation Fund Denmark under Grant 9065-00204B. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Yu Tsao. (*Corresponding author: Andreas Jonas Fuglsig.*)

Andreas Jonas Fuglsig is with the RTX A/S, 9400 Nørresundby, Denmark, and also with the Department of Electronic Systems, Aalborg University, 9000 Aalborg, Denmark (e-mail: ajf@rtx.dk).

Jesper Jensen, Zheng-Hua Tan, and Jan Østergaard are with the Department of Electronic Systems, Aalborg University, 9000 Aalborg, Denmark (e-mail: jje@es.aau.dk; zt@es.aau.dk; jo@es.aau.dk).

Lars Søndergaard Bertelsen and Jens Christian Lindof are with the RTX A/S, 9400 Nørresundby, Denmark (e-mail: lsb@rtx.dk; jcl@rtx.dk).

Digital Object Identifier 10.1109/TASLP.2023.3282094

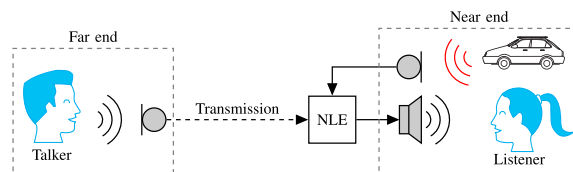


Fig. 1. Basic principle of near-end listening enhancement (NLE), where the received far-end signal is processed prior to playback in a noisy environment.

interfere with speech perception and degrade both the speech intelligibility (SI) and the speech quality (SQ). In near-end listening scenarios, the noise sources are physically present in the environment where the listener is located, cf. the right-hand part of Fig. 1. Therefore, in the near-end scenario, we cannot extract the clean speech signal by removing noise from a noisy input signal as done in far-end speech enhancement, when there is noise in the left-hand side of Fig. 1, e.g., [1], [2]. Instead, several other techniques exist for increasing the SI and SQ in noise by modifying the speech signal received from the far-end prior to playback in the noisy environment, and are known as Near-end Listening Enhancement (NLE).

Both SI and SQ are important factors for the listening experience, but the importance of SI and SQ changes depending on the acoustic situation [3], [4], [5]. Sometimes the requirements to SI and SQ may even be at a conflict and a trade-off must be made [5]. Increased intelligibility can lead to an increase in experienced SQ in noisy conditions [4], [5]. That is, when listening in noise, the SI is an important contributing factor to the experienced SQ [5]. Even though NLE processing may introduce distortion to the target speech signal, which could lead to decreases in SQ, these speech distortions may be masked by the harsher environmental noise in these noisy situations [5], [6]. However, speech becomes naturally more intelligible as noise conditions improve [3], [4], [5]. Hence, in favorable noise conditions, the intelligibility of unprocessed speech approaches 100% so that it cannot be further enhanced [3]. In fact, further excessive or unnecessary processing of the speech signals may lead to speech distortions and quality degradations because the noise can no longer mask these processing distortions [5]. Therefore, it may be useful to adapt the NLE processing to the noise situations, such that it can be activated and deactivated in the best possible way [5].

Traditionally, the goal of NLE algorithms is solely to maximize the SI for the listener in the noisy environment by modifying the time-frequency characteristics of the input speech

such that it is not masked by the environmental noise, cf. the reviews in [7], [8], [9] and the contributions to the Hurricane challenges [10], [11]. Existing NLE techniques can be categorized into two main classes. The first one is the heuristic or expert driven approaches, e.g., [12], [13], [14], [15], [16], [17], [18], [19], [20], [21], [22], [23], [24]. These consider various approaches regarding, e.g., formants, transients, Lombard speech, spectral shaping and dynamic range compression, or plain audibility. Although the class of heuristically based NLE techniques provide significant improvements in SI, they are often not derived according to specific objective optimality criteria. Instead, they are non-parametric and based on subjective expert experiences and knowledge. Hence, they cannot claim optimality and maybe are not adapting well to changing environments.

The second class of NLE algorithms is based on the main idea of manipulating the input speech such that a target intelligibility metric is maximized when the noise conditions are known. One of the most widely used optimization targets for SI enhancement is the Speech Intelligibility Index (SII) [25] or variations thereof, which have been used for NLE in numerous studies, e.g., [26], [27], [28], [29], [30], [31], [32]. The SII based approaches [26], [27] show good performance but fall behind the state-of-the-art heuristic [19] in subjective tests [10], [11], because the optimization targets do not correlate well with subjective intelligibility across varying noises and degradations [33]. Furthermore, some methods, e.g., [26], [27], require solving an optimization problem in real-time with varying execution time [23]. The SII based solutions also rely on simplifying assumptions about frequency gains being constant across frequency subbands [32]. Recently, deep neural network (DNN) based approaches [34], [35] have been able to optimize more advanced measures such as the extended short-time objective intelligibility (ESTOI) measure [36] that correlate more with subjective tests than the simpler metrics such as SII [33]. However, DNN methods come at the cost of significant memory usage, black-box solutions, and possible problems in generalizing to new acoustic scenarios.

We note for near-end listening with headphones, adaptive noise cancellation (ANC) techniques [37], [38] may be employed, which — instead of processing the target speech signal — aim at adaptively cancelling the noise by adding an anti-phase noise component to the speech signal before playout in the noisy environment [39], [40]. However, ANC with classic adaptive filtering has generally been insufficient for improving intelligibility outside headphone use until the use of DNNs [39]. Hence, NLE based on speech modification is still the predominant approach.

Common for the objective SI metrics such as SII and the Glimpse model [41] is that audibility is the decisive factor of intelligibility. Therefore, the principle for all NLE algorithms is to adjust the SNR for the perceptually important parts of the speech [39]. Hence, the simplest solution to the NLE problem is to increase the power of the clean speech signal until the noise is sufficiently masked, i.e., increasing the SNR (almost indefinitely). However, such an increase in speech power may lead to various problems, most significantly possible hearing damage to the near-end listener, but also problems with loudspeaker overload and unpleasant playback levels [28]. Therefore, many

NLE approaches take a power constraint into account, such that an equal power constraint is maintained between the unprocessed and processed clean speech signal. However, this also considerably limits the potential for increasing the SI [24], [42].

Existing methods are designed to always achieve the desired output power level. Thus, the power is always increased (or decreased) to the maximum allowed level even though this may not be necessary and may lead to excessive processing and a decrease in SQ or SI [5], [8]. Thus, increasing the allowed output power too much, in order to overpower the noise, may cause problems for SQ. Hence, we need to control the amount of processing in a different way, such that we can achieve both good SI and SQ depending on the noise situation.

In this work, we take inspiration from the area of far-end noise reduction. Particularly, the work of Zahedi et al. [43], where the concept of *minimum processing beamforming* is proposed. In this new rationale, the goal is to ensure a minimum level of SI performance, while guiding the noise reduction capabilities of the beamforming towards a particular performance; ambient-preserving, aggressive noise reduction or standard Wiener filtering.

In this article, we propose the concept of *minimum processing near-end listening enhancement*. The proposed concept provides a more general formulation of the NLE problem that focuses on both SI and SQ in an adaptive manner. The concept is based on minimizing a speech processing penalty subject to a certain intelligibility performance constraint. The point is no longer only to maximize SI but instead process the target speech signal just enough to achieve a minimum desired intelligibility in the given noise conditions, while minimizing the amount of speech distortions in favorable noise situations. Thus, the concept is adapting to the noise conditions by increasing SI when needed and ensuring an automatic focus on SQ, when the SI is already sufficient by automatically reducing the amount of processing when the desired SI is achieved.

We present an exemplary case study where the processing penalty is the mean-square error (MSE) between the processed near-end signal and the clean speech, and the performance criterion used in the optimization is an intelligibility estimator based on the approximated SII (ASII) [27]. While it is common in SII based works to assume that gains in the same subband are the same, we do not assume this upfront in this work. Instead, the optimization problem in this work is specifically built as a function of the gains in STFT domain to formulate a more general optimization problem without this common assumption. The case study shows we are able to adapt to changing noise conditions via a simple gain rule that is computationally efficient and which does not require online optimization. Furthermore, it is a result of our work that for the given case study the gains within a subband are the same, thus the derived solution mathematically validates the simplifying assumption of constant gains made in existing SII based works.

Finally, the proposed NLE method achieves SQ performance on par or better than existing methods in both objective measurements and subjective listening tests, while maintaining objective SI performance on par with existing methods for a wide range of SNRs.

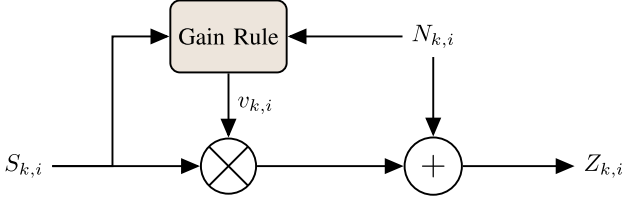


Fig. 2. Block diagram representation of frequency domain NLE signal model.

The rest of the article is organized as follows. In Section II we introduce our signal model. Section III introduces the proposed minimum processing NLE and we make our case study. Section IV introduces practical considerations and summarizes the algorithm. Section V presents the experimental setup and objective SI and SQ performance. Section VI presents a subjective SQ listening test and discusses the results. We conclude the article in Section VII.

II. SIGNAL MODEL

In this work we consider a signal model where speech and noise are represented in the complex short-time discrete Fourier transform (STFT) domain, cf. Fig. 2.

Here $S_{k,i}$ is the clean speech signal received from the far-end in frequency-bin k and time-frame i . In most general settings the far-end signal is noisy, however, we assume an adequately clean version is available or can be achieved through proper noise reduction at the far-end. $N_{k,i}$, is the time-frequency representation of the additive near-end noise. As is common in the literature, we assume the speech, S , and noise, N , are uncorrelated of each other. To increase the SI (and SQ) of the signal received by the near-end listener, the speech signal, S , is linearly pre-processed prior to play out in the near-end environment. Thus, we assume that the signal presented to the near-end listener in STFT domain, $Z_{k,i}$, follows

$$Z_{k,i} = v_{k,i}S_{k,i} + N_{k,i}, \quad (1)$$

where $v_{k,i}$ denotes the linear pre-processing gains. The NLE gains constitute the core of NLE, and determining a proper NLE gain rule is the core task of NLE.

A. Subband Model

We focus on NLE based on (the perceptually driven) ASII for speech intelligibility prediction. ASII is part of a family of well-known speech intelligibility and quality predictors such as ESTOI [36], STOI [44], SII [25], extended SII (ESII) [45], the hearing-aid speech perception index (HASPI) [46], the hearing-aid speech quality index (HASQI) [47] and perceptual evaluation of speech quality (PESQ) [48], that all mimic aspects of human speech perception. That is, signals are analyzed in, e.g., octave bands, fractional octave bands or critical bands similar to the human ear. In general, these bands are referred to as subbands. Perceptually motivated subbands may be defined such that multiple frequency bins contribute to the same and/or multiple subbands, with an individual weight for each frequency bin-subband pair. We index subbands with index j and frequencies with index k .

To illustrate the use of subbands in this work, let $\omega_{j,k}$ denote the non-negative filter weights that implement the j 'th subband filter, \mathbb{B}_j denote the set of frequency bins that contribute to the j 'th subband with $j \in \{1, \dots, J\}$ and J the total number of subbands. Then, the clean speech spectrum level within one subband, j , and time-frame, i , is defined as

$$\sigma_{S_j,i}^2 \triangleq \sum_{k \in \mathbb{B}_j} \omega_{j,k} \sigma_{S_{k,i}}^2, \quad (2)$$

where $\sigma_{S_{k,i}}^2$ is the clean speech spectrum level in time-frequency bin k, i . For the sake of brevity and ease of reading we assume any normalization of the subband filter weights, $\omega_{j,k}$, is already included in the weights. We elaborate further on the definition of $\omega_{j,k}$ and the relation between subbands and frequency bins in Appendix A.

Finally, we assume the speech and noise are processes of complex random vectors comprised of the STFT coefficients. In practice one could estimate the statistics of these processes online, and the mathematical framework can be applied to the time-varying case on a per time-frame basis. Therefore, for brevity of notation we disregard the time-index, i , and assume we are considering a certain time frame i , unless otherwise is stated.

III. MINIMUM PROCESSING NEAR-END LISTENING ENHANCEMENT

A. Concept

For a particular subband j , we denote by $\mathbf{S}_j \in \mathbb{C}^{|\mathbb{B}_j|}$ the vector containing all S_k for $k \in \mathbb{B}_j$, where $|\mathbb{B}_j|$ denotes the number of frequency bins in the j 'th subband. Similarly, we create the vector \mathbf{Z}_j by stacking Z_k for $k \in \mathbb{B}_j$. Furthermore, let $\mathcal{D}_j(\cdot, \cdot)$ and $\mathcal{I}_j(\cdot, \cdot)$ be finite non-negative functionals, where $\mathcal{D}_j(\mathbf{S}_j, \mathbf{Z}_j)$ measures the distortion (processing penalty) between the clean speech signal, S , and the signal presented to the near-end listener, Z , and $\mathcal{I}_j(\mathbf{S}_j, \mathbf{Z}_j)$ is an intelligibility or performance estimator for the NLE in subband j . Then, the minimum processing near-end listening enhancer in subband j is defined as the solution to the following optimization problem:

$$\min_{\{v_k\} \in \mathbb{R}_+, k \in \mathbb{B}_j} \mathcal{D}_j(\mathbf{S}_j, \mathbf{Z}_j) \quad \text{s.t.} \quad \mathcal{I}_j(\mathbf{S}_j, \mathbf{Z}_j) \geq I'_j. \quad (3)$$

The term I'_j is defined as

$$I'_j = \min(I_j, I_j^{\max}), \quad (4)$$

where I_j is a desired minimum requirement on the NLE intelligibility performance $\mathcal{I}_j(\mathbf{S}_j, \mathbf{Z}_j)$ and I_j^{\max} is the maximum achievable performance when disregarding the processing penalty $\mathcal{D}_j(\mathbf{S}_j, \mathbf{Z}_j)$, i.e., when the performance $\mathcal{I}_j(\mathbf{S}_j, \mathbf{Z}_j)$ is maximized in an unconstrained manner. We highlight the generality of the problem formulation in (3), in that $\mathcal{I}_j(\mathbf{S}_j, \mathbf{Z}_j)$ is a general ‘‘performance criterion’’ which could reflect any type of performance aspect of interest to the application, e.g., SI, SQ or other measures of interest.

B. Case Study

In this article, we study the case, where the processing penalty, \mathcal{D}_j , is the MSE criterion, and the performance criterion, \mathcal{I}_j , is an intelligibility estimator based on the ASII [27]. We solve the problem analytically for any given minimum performance constraint and subband definition. Furthermore, we show how we can select I_j to achieve a desired total minimum performance across all subbands, A^* , i.e., we guarantee that the optimum solution satisfies

$$\sum_j \mathcal{I}_j(\mathbf{S}_j, \mathbf{Z}_j) \geq A^*. \quad (5)$$

1) *Processing Penalty*: To ensure the processing does not distort speech excessively and the playback volume is not increased infinitely, we introduce a processing penalty [43]. We consider a MSE processing penalty, where the MSE is evaluated in subband domain instead of DFT domain to ensure compatibility with the minimum processing NLE formulation in (3). That is, the j 'th subband MSE processing penalty is

$$\mathcal{D}_j(\mathbf{S}_j, \mathbf{Z}_j) = \sum_{k \in \mathbb{B}_j} \omega_{j,k} (1 - v_k)^2 \sigma_{S_k}^2. \quad (6)$$

See Appendix B for the derivation.

2) *Performance Criterion*: We consider, as an example of the proposed frame work, a performance criterion based on the ASII [27]. The original work on minimum processing in [43] considered a performance criterion based on the SII [25]. Both SII and ASII consider intelligibility to be a weighted sum of intermediate measures of the audibility of speech in a subband as a function of the subband SNR. In SII, the band audibility is determined by a function where the long-term SNR is log-transformed and clipped between -15 dB and $+15$ dB and normalized to range between zero and one. In ASII, the audibility is determined by a non-linear approximation of the log-transform and clipping via a sigmoidal function of the subband SNR. The advantage of using the ASII is that its nonlinear approximation of the SII provides a nice mathematical tractability circumventing the need for dealing with clipping of SNRs. Based on the ASII, have the following performance criterion for the j 'th subband,

$$\sum_{k \in \mathbb{B}_j} \omega_{j,k} v_k^2 \sigma_{S_k}^2 \geq \sigma_{N_j}^2 I_j^\xi, \quad (7)$$

where $I_j^\xi \triangleq \frac{I_j}{1-I_j}$, and I_j is the desired subband audibility performance. The details are shown in Appendix C.

C. Optimization Problem and Solution

Joining the results of (6) and (7), the minimum processing NLE problem (3), which we consider, and its solution is given in the following theorem.

Theorem 1: The minimum processing NLE problem (3) with MSE processing penalty (6) and ASII performance criterion (7) is

$$\begin{aligned} & \min_{\{v_k\} \in \mathbb{R}_+, k \in \mathbb{B}_j} \sum_{k \in \mathbb{B}_j} \omega_{j,k} (1 - v_k)^2 \sigma_{S_k}^2, \\ & \text{subject to} \quad \sum_{k \in \mathbb{B}_j} \omega_{j,k} v_k^2 \sigma_{S_k}^2 \geq \sigma_{N_j}^2 I_j^\xi. \end{aligned} \quad (8)$$

The optimal *minimum processing NLE* gains are

$$v_{k,j}^{\text{MP}} = \begin{cases} 1 & \text{if } \sigma_{S_j}^2 \geq \sigma_{N_j}^2 I_j^\xi \\ \sqrt{\frac{\sigma_{N_j}^2 I_j^\xi}{\sigma_{S_j}^2}} & \text{otherwise} \end{cases}, \quad \forall k \in \mathbb{B}_j. \quad (9)$$

In Appendix D we derive the *minimum processing NLE*, i.e., the optimal gains in (8).

We immediately have the following result regarding the optimal minimum processing NLE subband gains (9).

Corollary 1: The optimum gains for the minimum processing NLE (8) are equal for all frequencies within a subband j .

This is an important result because most work on subband SNR based enhancement make this assumption for convenience in the joint optimization across all subbands [9], [27], [30], [32], [49]. However, as we show, it can be deduced from (9) that the optimal solution does not depend upon k , and it is optimal to have the same gain across the entire subband when optimizing for each subband individually.

For the optimal gain (9), the first case takes effect when the unprocessed speech signal already satisfies the performance constraint, i.e., the speech has an intelligibility at or above the desired level. Hence, there is no reason to process the speech signal and the optimal solution is to do nothing, i.e., $v_{k,j}^{\text{MP}} = 1$ for $k \in \mathbb{B}_j$, and thus maximize SQ by minimizing the MSE.

In the second case, the unprocessed speech is below the desired audibility level or equivalently the desired SNR. The optimal gain then increases the speech power such that the subband SNR is exactly at the desired level. That is, the speech power in the j 'th subband is increased just enough to satisfy the subband audibility constraint. This illustrates how the optimal solution provides the minimum required processing necessary to achieve the desired intelligibility. Thus, the two cases of (9) illustrate how our approach has the advantage of achieving a desired intelligibility level, while automatically adjusting the processing to minimize processing artifacts.

Squaring the gains of (9), $(v_{k,j}^{\text{MP}})^2$, we can see that the optimal processed speech power (32) is increased proportionally to the near-end noise power. This is also the naturally expected solution to ASII and SII based NLE approaches given sufficient ability to increase speech power [24], [26], [27]. Thus, we provide a mathematical justification for the heuristic approach of increasing speech power proportionally to the noise in [24].

Depending on the choice of subband definition, the subbands may overlap such that multiple frequencies contribute to multiple subbands indexed by \mathbb{F}_k . Therefore, the optimum gain, $v_{j,k}^*$ may also contribute to multiple subbands as indicated by the dependence on both j and k . Using the weights of these contributions, $\omega_{j,k}$, we can weigh the filter gains through the subband filters. Thus, we have the following corollary.

Corollary 2: The NLE processor in frequency bin k is given as

$$v_k^{\text{MP}} = \sqrt{\sum_{j \in \mathbb{F}_k} \omega_{j,k} (v_{k,j}^{\text{MP}})^2}. \quad (10)$$

From this corollary, we see for overlapping subbands, that gains that were previously equal do not remain the same after being projected back to STFT domain.

IV. PRACTICAL CONSIDERATIONS

A. Preventing Excessive Sound Levels

The proposed optimal minimum processing increases speech power to the necessary level to achieve the desired intelligibility. However, this may lead to speech levels that are unpleasant for the user. In extreme cases, if not controlled, such high signal levels could even cause hearing damage. Usually in most NLE work, this is prevented by a power equality constraint [24], [27]. To prevent excessive output levels, in our work, we put an upper limit on the maximum processed speech power within each subband as described in [31, Sec. 2.2.5]. The idea is to limit the gain applied to each subband, j , such that the resulting subband power of the enhanced speech signal does not increase beyond a maximum subband power, P_S^{\max} . That is,

$$\bar{v}_{k,j}^{\text{MP}} = \min \{v_{k,j}^{\text{MP}}, v_{j,\max}\} \quad \forall k \in \mathbb{B}_j, \quad (11)$$

with maximum subband gain

$$v_{j,\max} = \sqrt{\frac{P_S^{\max}}{\sigma_{S_j}^2}}. \quad (12)$$

This is done prior to the subband filtering of the gains in (10). Thus, $\bar{v}_{k,j}^{\text{MP}}$ is inserted on the right-hand side of (10).

B. Choosing the Intelligibility Limit Per Band

The proposed minimum processing NLE as well as the work of [43] considers a performance constraint, I_j , for each particular subband $j \in \{1, \dots, J\}$. Thus, the processed speech achieves a desired performance in each particular subband. This target must be decided upon for each of the J subbands before processing. In this section, we show how, when using the ASII criterion, we can select and achieve a single total target intelligibility, A^* , instead of having to select a desired performance constraint for each of the J subbands. This single total target intelligibility is then converted to a per band intelligibility target.

Let I_j be a given minimum requirement on the performance in a particular subband, j , i.e., we require

$$f(\xi_j) \geq I_j. \quad (13)$$

Then the processed ASII score satisfies

$$\text{ASII} = \sum_j \gamma_j f(\xi_j) \geq \sum_j \gamma_j I_j. \quad (14)$$

That is, given that the minimum performance requirement is met for each band, the processed ASII is greater than or equal to a weighted sum of the performance criterions, where the weights are the band importance functions.

The band importance functions, γ_j , in the SII [25] describe how some subbands are more important for intelligibility than others. Therefore, it is natural to require higher audibility limits for the more important subbands and lower audibility limits for

the less important bands. We can achieve this by using the band importance functions.

Let A^* be the required minimum total ASII performance, and let the subband audibility limits be given as

$$I_j = \frac{A^* \gamma_j}{\sum_i \gamma_i^2}. \quad (15)$$

Then by inserting in (14) the resulting ASII score satisfies

$$\text{ASII} \geq \sum_j \gamma_j \frac{A^* \gamma_j}{\sum_i \gamma_i^2} = \frac{A^*}{\sum_i \gamma_i^2} \sum_j \gamma_j^2 = A^*, \quad (16)$$

that is, the total processed ASII score is lower bounded by the required performance. This means we can guarantee a total ASII greater than or equal to a target ASII, A^* .

As a special case, we see that removing the weights and selecting a fixed audibility limit across subbands $I_j = A^*$, $\forall j$ also achieves $\text{ASII} \geq A^*$. In our simulations we use the weighted audibility limits, as these have shown a slightly better performance than having the same fixed limit for each subband.

C. Estimating Statistics

When processing time-varying signals, such as speech and non-stationary noise, the statistics of the signals change over time. Therefore, the statistics must be estimated and updated in time using, e.g., recursive averaging. However, if the statistics are updated too fast or abruptly, the optimal gains in (9) can change suddenly between time-frames, which causes audible distortions to the target speech signal. To circumvent this, it is common in the literature, e.g., [24], [43], to use slowly time-varying processing, where the recursive averaging is across several frames (seconds). Therefore, in this work, we let the average energy per DFT bin and critical band be based on a long-term average over several short-time frames,

$$\sigma_{S_k}^2 \triangleq \frac{1}{I} \sum_i |S_{k,i}|^2, \quad (17)$$

$$\sigma_{S_j}^2 \triangleq \sum_k \omega_{j,k} \sigma_{S_k}^2, \quad (18)$$

where I is the total number of frames. Similar expressions hold for the near-end noise signal $N_{k,i}$. Thus, the estimated statistics in this work do not change over time, and the implemented processing is time-invariant. For time-varying estimates, (17) should be updated to a recursive or moving average.

D. Algorithm Summary

The proposed minimum processing NLE algorithm with ASII intelligibility criterion is summarized with the flowchart in Fig. 3.

V. OBJECTIVE PERFORMANCE EVALUATION

In this section, we evaluate the performance of the proposed minimum processing NLE method (9) in the NLE setup shown in Fig. 2, where the target speech signal is processed prior to play out in the noisy environment, and the gain rule follows the proposed algorithm summarized in Fig. 3. We evaluate objective

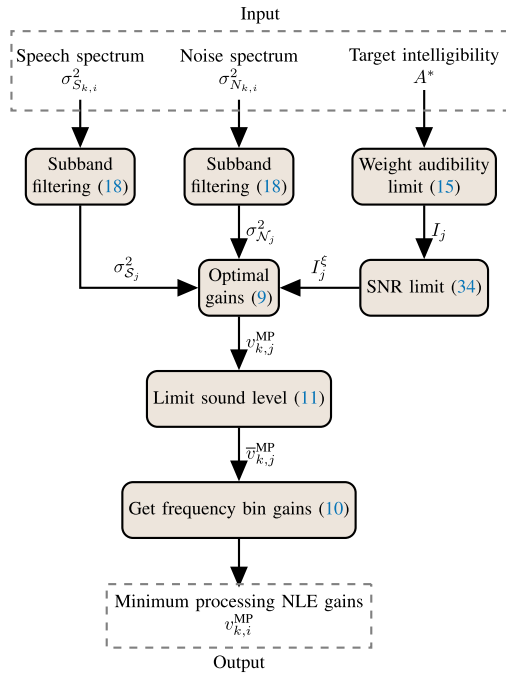


Fig. 3. Flowchart of the proposed minimum processing NLE gain rule.

SI and SQ performance, and compare performance with existing related state-of-the-art NLE methods. For intelligibility scoring, we use the chosen performance criterion ASII [27] and the more general intelligibility predictor ESTOI [36]. The goal of minimum processing is to reproduce the clean unprocessed input speech signal as well as possible. Therefore, quality is estimated using PESQ [50] and Segmental SNR (Seg-SNR) [51].

A. Experimental Setup

The speech material used for the evaluation are English sentences each a few seconds in duration from the test set in the TIMIT-database [52] sampled at 16 kHz. In each trial a speaker is selected randomly without replacement across trials. Afterwards, a random sentence is selected amongst those available for the selected speaker. Finally, the speech excerpt is padded with 0.5 s of silence at the beginning and 0.125 s of silence at the end.

For the noise signals we consider white noise, synthetic speech shaped noise, as well as noise recorded inside a car traveling at 130 km/h. For the recorded noise, a random excerpt of the appropriate length is cut out of the noise recording for each trial.

Speech and noise signals were transformed to time-frequency domain using an STFT with 32 ms Hann windows with 50% overlap at a sampling frequency of 16 kHz. For the subbands, we consider a total of $J = 30$ overlapping auditory filters with linearly spaced center frequencies on the equivalent rectangular bandwidth scale from 150 Hz to 8000 Hz [53], see Appendix A for further details on the subband definition.

The target intelligibility is set to $A^* = 0.7$, unless otherwise is stated. We choose this target, because an ASII value of 0.7 corresponds to almost full intelligibility [33], and since any

higher value requires significantly more processing (see also Section V-C). The band intelligibility limits, I_j are weighted according to (15) for all trials.

For the maximum subband gain, $v_{j,\max}$, in (12), the maximum allowed power P_S^{\max} is chosen such that

$$10 \log \left(\frac{P_S^{\max}}{P_0} \right) = 100 \text{ dB}_{\text{SPL}}. \quad (19)$$

Here P_0 represents the digital reference power corresponding to a reference sound pressure [31]. We make the convention that a signal power of P_0 corresponds to a signal with an $RMS = 1$, thus

$$P_0 = 10^{-\frac{100 \text{ dB}_{\text{SPL}}}{10}}. \quad (20)$$

Finally, we conduct a total of 10 trials and evaluate the performances in each trial. The average performance across the trials is then taken as the final score for each combination of noise, SNR and gain rule.

B. Reference Methods

We consider three reference methods that are similar to our work: (1) The original ASII optimization of [27] because we also optimize for ASII, (2) the well known SII optimization of [26] that clips SNRs above and below a certain level and (3) the very recent NoiseProp algorithm of [24] because it has a simple gain-rule very similar to ours. The method of [24] has an additional tuning parameter, $\rho_{NV} \in [0, 1]$. It is mentioned in [24], that informal listening has shown $\rho_{NV} \in [0.5, 0.8]$ is a good choice. Therefore, we select the value $\rho_{NV} = 0.7$ for all experiments.

All reference methods are implemented following the procedure described in the previous subsection. The only difference is the applied gain rule. That is, for ease of comparison all methods are implemented as time-invariant even though the original work may have considered slowly time-varying processing, e.g., [24], [26].

All the reference methods are based on a target speech power equality constraint. Hence, they are designed to keep the total power across frequency at a certain level, P_{ref} . However, the proposed method is able to increase the power as much as needed, achieving a total processed power level denoted by P_{MP} . Therefore, for a fair comparison, the reference methods are also allowed to increase the power to that level. That is, the reference methods are implemented such that $P_{\text{ref}} = P_{\text{MP}}$.

C. Minimum Processing Effect

Initially, we illustrate the effects of the proposed minimum processing solution in terms of processing penalty versus intelligibility performance and resulting increases in speech power. The effects are considered using a white noise source to limit the importance of the spectral distribution of noise.

Fig. 4 shows that for increasing values of the target intelligibility, A^* , the achieved ASII (top) along with the corresponding MSE processing penalty (middle) and increase in speech power (bottom). The lower limit case of target ASII $A^* = 0$ corresponds to the case of no processing. The ASII is upper bounded by

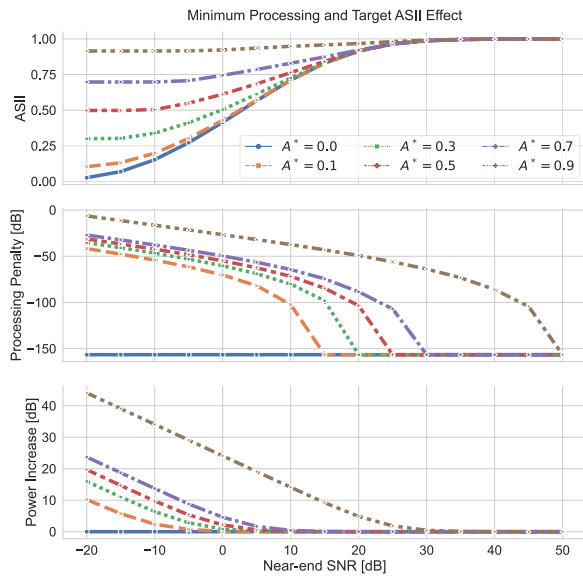


Fig. 4. The achieved ASII (top), and the corresponding MSE processing penalty (middle) and power increase (bottom) versus input near-end SNR for varying choices of target intelligibility. The legend in the top panel applies to all plots.

$A^* \leq 1$ and is only achievable in the limit case of infinite SNR, therefore we only show the trend up to $A^* = 0.9$. The top panel in Fig. 4 shows, that when the target intelligibility $A^* > 0$, we are able to improve the performance over the unprocessed case ($A^* = 0$) for a large range of SNRs. From the middle and bottoms panels in Fig. 4, we see that the greater the target intelligibility and the lower the input SNR, the increases in intelligibility performance come at the cost of a greater need for processing and increased speech power. Finally, as the near-end SNR improves, the target intelligibility is more easily achieved and the processing required to achieve a given target decreases, as can be seen by a transition from a high level of processing to no processing.

D. Objective Performance

Fig. 5 shows the ASII, ESTOI, PESQ and Seg-SNR performance as function of the near-end SNR for the proposed minimum processing NLE and the reference methods in both car noise and speech shaped noise. Further experiments using cafeteria and babble noise sources show performance similar to experiments using speech shaped noise, however these are omitted here do to space limitations.

1) *Estimated Intelligibility Performance:* From the SI metrics ASII and ESTOI, we see that the proposed method improves SI over the unprocessed speech to the desired level at the lower SNRs. The proposed method even achieves the best ESTOI performance in car noise compared to the reference methods. In ASII, the reference methods are able to gain a slightly better SI than the proposed. This is expected, as they try to maximize the SI given the allowed power constraint and not just reach a sufficient level of SI. As the SNR increases and the speech becomes naturally intelligible, the proposed method follows the unprocessed performance above the target intelligibility.

Thus, the proposed method always achieves the desired SI performance or better. Similar performance can be seen with SII maximization by Sauert [26], however this does not happen until the SII clips the subband SNRs above 15 dB and the SII is fully maximized. Additionally, as the unprocessed SI reaches its ceiling the reference methods are no longer able to further enhance SI by maximization.

2) *Estimated Speech Quality Performance:* Generally, from the PESQ and Seg-SNR plots, we see that all Seg-SNR and PESQ scores are at the lowest levels for very low SNRs, and it is difficult to tell any difference in quality performances. This is expected, since the near-end SNR is low and the noise is the dominant part of the signals. As the input SNR increases so does the PESQ and Seg-SNR scores, where the unprocessed clean speech is considered to be the maximum level of SQ. This is due to PESQ and Seg-SNR measuring quality by how similar the signal presented to the listener, Z , is to the clean reference, S . Hence, for high SNR this is the maximum achievable level when the near-end noise is insignificant and $Z \approx S$.

For increasing SNRs in both car noise and speech shaped noise, we see that the proposed method follows the maximum SQ performance of the unprocessed signal, when the intelligibility is at the desired level and the processing switches off. Similarly, the method of Sauert [26] switches off automatically at maximum SII, but does so at a higher level of intelligibility, thereby incurring more speech distortions than the proposed method. On the other hand, the reference methods of Taal [27] and Niermann [24] never switch off and continue processing at the higher SNRs causing large speech distortions.

Comparing between car noise and speech shaped noise we see that the PESQ and Seg-SNR show less speech distortion in speech shaped noise. Further, since car noise is a less severe noise, speech distortions occur at lower SNRs and are more clearly heard in car noise. Thus, excessive processing is more detrimental to SQ in better noise conditions.

E. Gain Dynamics

To show the advantages of minimum processing in comparison with existing methods, we illustrate how the proposed method and [24] affect the speech spectra in different noise conditions. We consider the behavior in car noise for an input SNR of -30 dB and 10 dB. In these SNRs both the proposed method and [24] achieve very high SI, however they have a vastly different SQ at 10 dB as seen in Fig. 5(a).

Fig. 6 shows an example of the clean speech and noise spectrum (top row), as well as the derived optimal gains for both methods (bottom row) for both -30 dB SNR (left column) and 10 dB SNR (right column). Looking at the left-hand column with low SNR, we see that the proposed method provide optimal gains that raise the speech power sufficiently above the near-end noise, and the reference method [24] optimal gains are proportional to the near-end noise. Now considering the difference between the two SNRs (columns), we see how the proposed method automatically minimizes the amount of processing as the SNR increases and the SI is at the ceiling level. This results in high SI at low SNRs, and both high SI and SQ at high SNRs in Fig. 5(a).

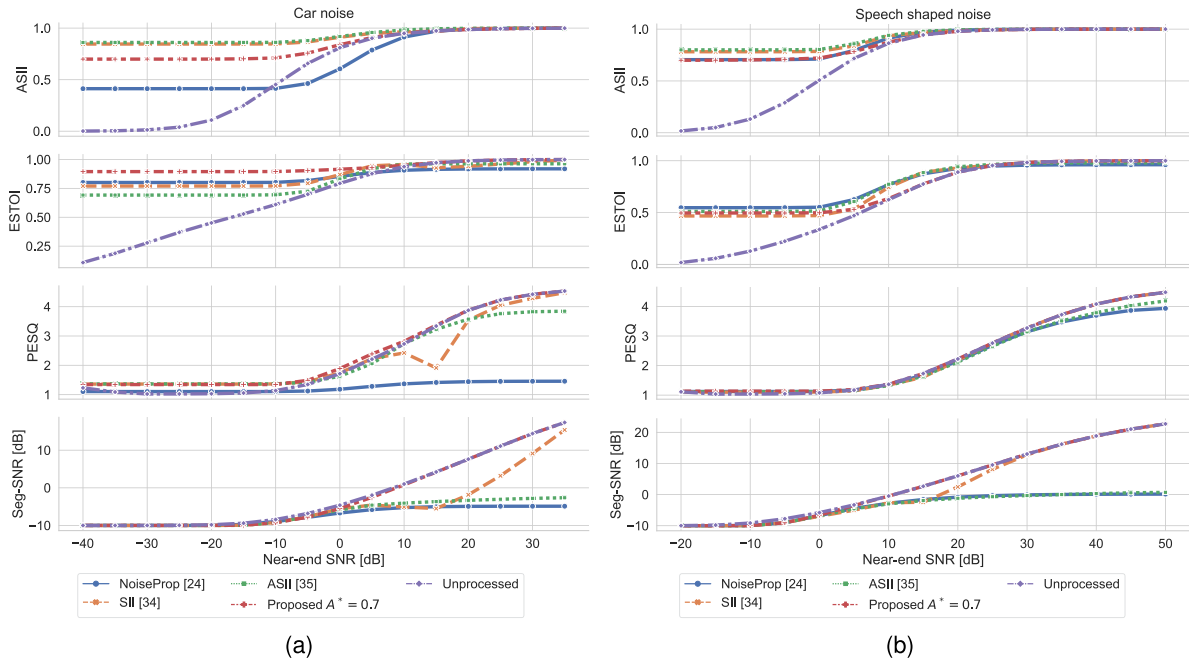


Fig. 5. ASII, ESTOI, PESQ and Seg-SNR scores in (a) car noise and (b) speech shaped noise. Legends apply to all plots within a column.

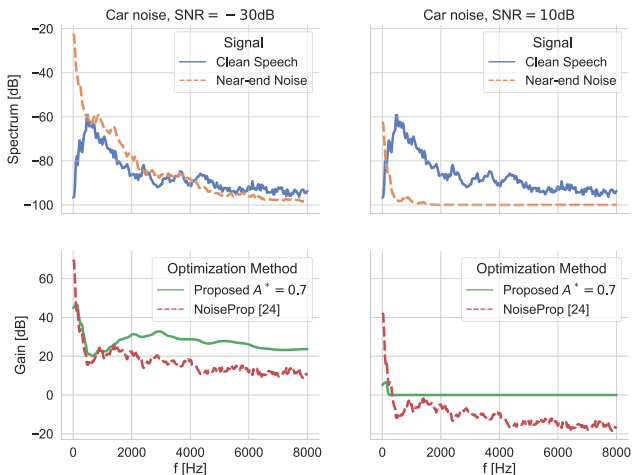


Fig. 6. Clean speech and near-end noise spectra (top row), and optimal gains (bottom row) for car noise in -30 dB SNR (left column) and 10 dB SNR (right column).

On the other hand, the reference [24] method, which always tries to maximize SI, continues to process the speech signal proportionally to the near-end noise even at the high SNR when SI is at the ceiling level. This continued processing introduces audible speech distortions, which result in low SQ scores at the higher SNRs as seen in Fig. 5(a).

VI. SUBJECTIVE SPEECH QUALITY TEST

In this section, we evaluate the performance of the proposed minimum processing NLE method (9) by a subjective listening test for SQ. Speech quality and general audio quality is highly subjective and may not always be represented well by objective measures. For example, grading SQ in the presence of noise using PESQ is difficult due to a high sensitivity to

the environmental noise compared to the speech distortion [43]. Therefore, to further investigate the performance of the proposed procedure we perform a listening test to assess the SQ in the near-end listening scenario, where we cannot process the noise but only the clean speech component. We consider the proposed minimum processing in comparison with the highly similar and state-of-the-art NoiseProp method of [24]. We investigate performance in both car noise and speech shaped noise.

A. Listening Test Setup

The listening test was conducted by 21 (4 female, 17 male) volunteering untrained listeners. The participants had an age span of 22 to 59 years with an average age of 39.6 years. All participants had self-reported normal hearing. The average test time including training was 38 minutes.

The listening test was conducted in a silent room using a Lenovo T570 laptop equipped with a USB sound card (DragonFly Black) and a pair of closed headphones (Beyerdynamic DT-770 Pro 32 O) for reporting and audio playback. The user interface was based on the Web Audio Evaluation Tool [54]. All audio stimuli were normalized to the same perceived loudness of -30 LUFS following [55] and the EBU R128 [56] recommendation for loudness normalization as implemented within the Web Audio Evaluation Tool [57]. Participants were allowed to adjust the general volume to a comfortable level during the training session of the test.

B. Procedure

We conducted a listening test following the MULTi Stimulus with Hidden Reference and Anchor (MUSHRA) [58] paradigm, where the audio quality is assessed on a scale from 0 to 100, divided into five equal intervals labelled as *bad*, *poor*, *fair*, *good* and *excellent*. The participants were instructed to grade the *basic*

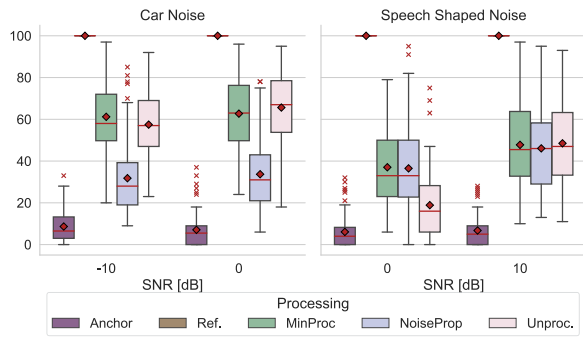


Fig. 7. Boxplot illustration of the MUSHRA listening test results for car noise (left) and speech shaped noise (right). Medians and means are indicated by red horizontal lines and diamonds, respectively. Outliers (according to the 1.5 interquartile range rule) are indicated by red crosses.

audio quality compared to a known reference signal. No other definition was provided of audio quality.

Each test participant was presented with 2 sequences of 8 trials, one sequence for each of the car and speech shaped noise. Each trial consisted of a reference signal (unprocessed signal in high SNR) and five other signals to be rated: 1 hidden reference, 2 systems under test (proposed method and NoiseProp [24]), 1 unprocessed signal, and 1 hidden anchor (unprocessed signal at lower SNR). The SNRs considered for the reference, anchor and test cases vary depending on the noise type, such that the perceptual quality was not affected too severely by a loss in intelligibility. For car noise, the reference and anchor signal SNRs were set to 15 dB and -30 dB respectively. For speech shaped noise, the reference and anchor signal SNRs were set to 25 dB and -5 dB respectively. For both noise types, four trials were used to evaluate the systems at an input SNR of 0 dB. The last 4 trials were used to evaluate the systems at -10 dB and 10 dB for car noise and speech shaped noise, respectively.

Before the actual experiment, each participant went through a short training session, where they were able to get used to the task at hand, the different audio stimuli and the user interface. The sentences in the training were different from those in the actual experiment and did not contribute to the final test score.

C. Listening Test Results and Discussion

The average scores across trials and participants for each noise, processing and SNR condition are shown in Fig. 7 with boxplots.

The speech shaped noise results at an SNR of 0 dB in Fig. 7, show that processing with the proposed method and NoiseProp [24] to increase intelligibility has a positive effect on subjective quality compared to the unprocessed performance, when the noise situation is more severe. As the SNR increases to 10 dB, we cannot detect any effect on the subjective SQ compared to the unprocessed performance. This is in contrast to the PESQ results in Fig. 5(b), where no difference was seen. The listening test results for car noise at both -10 dB and 0 dB SNR show that NoiseProp [24] deteriorates SQ, while the proposed minimum processing does not. Both methods have comparable SI performance at these SNRs with the proposed method slightly outperforming the NoiseProp, as shown by ESTOI in Fig. 5(a).

The SQ results are also confirmed with PESQ at 0 dB SNR in Fig. 5(a), where there is a score difference of almost 2 points. At input SNR of -10 dB PESQ only shows a minor difference between the two methods. This illustrates, when the unprocessed SI is almost maximized due to favorable noise conditions, the excessive processing of NoiseProp, as seen Fig. 6, can harm subjective SQ, even though the SI may be increased slightly.

These observations confirm the results of the objective SQ metrics as well as those in the literature [5], that SI enhancement may lead to both decreases and increases in SQ. Particularly, decreases in SQ are observed when the unprocessed intelligibility is already high and the processing becomes excessive. On the other hand, in harsh noise conditions, the environmental noise masks most speech distortions. This is also in line with the results of [5], where SI is shown to be the dominant factor for SQ in severe noise.

Existing NLE methods focus on maximizing SI. Therefore, given a certain allowed increase in speech power, the reference methods are sometimes able to gain a slightly higher SI than the proposed method as the ASII and ESTOI scores show in Fig 5. However, the natural ceiling effect of SI occurs prior to the maximum value of the metric [33]. Therefore, the objective SQ and listening test results also show the achieved SI score may be unnecessarily high and lead to excessive speech distortion. Instead, by using a single intelligibility target, A^* , our proposed method can be better linked to the onset of the ceiling effect for various noise types. Thus, by varying the target intelligibility, A^* , we see how the proposed method allows controlling the tradeoff between SI and SQ.

VII. CONCLUSION

We proposed a (novel) near-end listening enhancement concept, where the output signal is optimized to have the minimum amount of processing artifacts (for example, distance from the clean speech signal), with the constraint that a certain performance criterion (for example, estimated intelligibility) is satisfied. The proposed concept is adaptive to environmental noise conditions by focusing on estimated intelligibility for high noise levels, and, as a consequence of the minimum processing paradigm, quality for low noise levels. The trade-off between intelligibility and quality can be controlled via the performance criterion. To demonstrate an instance of the proposed framework, we make a thorough investigation of an example case, where the performance criterion is based on the approximated speech intelligibility index and the processing metric is the mean squared error. We show that, this provides a simple gain rule based NLE, that raises the speech spectrum above the noise to exactly the minimum required level for the desired intelligibility. The proposed method, is able to distribute the processing across subbands to maximize or reach a target intelligibility by varying the desired intelligibility target between the subbands. If the noise conditions are favorable, the processing is automatically reduced leading to fewer processing artifacts and hence improved speech quality. Thereby, by allowing an increase in speech power, the proposed NLE technique is able to adaptively achieve any desired intelligibility level with a minimum of

processing artifacts. Experimental studies verify the advantages of the concept in terms of both intelligibility and quality.

Future work will include expanding the proposed concept to jointly consider near-end listening enhancement together with far-end noise reduction for an increased intelligibility and quality performance. Additionally, it is interesting combining the proposed method with acoustic echo cancellation techniques when working with joint far-end and near-end optimization.

APPENDIX A SUBBAND WEIGHTS

Assigning frequency bins to subbands and determining the weight of their contributions can be done in different ways, cf. [43, App. A]. We consider a gammatone filter bank model [53] that we normalize to preserve power between subband and frequency domain.

We focus on subbands in terms of overlapping auditory filters. Let h_j be the impulse response of the j 'th subband auditory filter. Then, the energy of the clean speech signal, $S_{j,i}$, is given as the convolution with between s and h_j , which in in time-subband domain is given as

$$\mathcal{S}_{j,i}^2 \triangleq \sum_{k \in \mathbb{B}_j} |S_{k,i}|^2 |H_j(k)|^2, \quad (21)$$

where $H_j(k)$ represents the DFT of h_j in frequency-bin k . We consider squared magnitude responses, $|H_j(k)|^2$ given by the gammatone filter bank derived in [53]. Furthermore, we want to ensure the total power is the same in both the frequency and subband domain, i.e.,

$$\sum_j \sigma_{\mathcal{S}_{j,i}}^2 = \sum_k \sigma_{S_{k,i}}^2. \quad (22)$$

Now inserting the filtering operation we have that

$$\sum_j \sigma_{\mathcal{S}_{j,i}}^2 = \sum_j \sum_k |H_j(k)|^2 \sigma_{S_{k,i}}^2 \quad (23)$$

$$= \sum_k \sum_j |H_j(k)|^2 \sigma_{S_{k,i}}^2. \quad (24)$$

From (24) we see that (22) is satisfied if $\sum_j |H_j(k)|^2 = 1$. We can achieve this by normalizing $H_j(k)$, i.e.,

$$H'_j(k) = \frac{H_j(k)}{\sqrt{\sum_l |H_l(k)|^2}}, \quad \forall j. \quad (25)$$

Hence, we let the subband filter weights, $\omega_{j,k}$, be the normalized squared magnitude response of h_j , i.e., $\omega_{j,k} = |H'_j(k)|^2$.

APPENDIX B MSE PROCESSING PENALTY

For the j 'th subband we write the MSE processing penalty as

$$\mathcal{D}_j(\mathbf{S}_j, \mathbf{Z}_j) = \sum_{k \in \mathbb{B}_j} \omega_{j,k} \mathbb{E} [|S_k - Z_k|^2] \quad (26)$$

$$= \sum_{k \in \mathbb{B}_j} \omega_{j,k} \left((1 - v_k)^2 \sigma_{S_k}^2 + \sigma_{N_k}^2 \right), \quad (27)$$

where we have assumed the STFT coefficients are uncorrelated across frequency. Thereby, the MSE in a subband is the 'weighted average' of the MSE in the DFT domain where the weights are the subband filter weights. Since the near-end noise power, $\sigma_{N_k}^2$, is unaffected by the processing, v_k , we can disregard it from the processing penalty, and we have

$$\mathcal{D}_j(\mathbf{S}_j, \mathbf{Z}_j) = \sum_{k \in \mathbb{B}_j} \omega_{j,k} (1 - v_k)^2 \sigma_{S_k}^2. \quad (28)$$

APPENDIX C ASII PERFORMANCE CRITERION

The ASII [27] is defined as

$$ASII \triangleq \sum_j \gamma_j f(\xi_j), \quad (29)$$

where the weights γ_j are pre-defined constants (more specifically the band importance functions as defined in [25]),

$$f(\xi_j) \triangleq \frac{\xi_j}{\xi_j + 1}, \quad (30)$$

is the sigmoidal audibility function per subband, and ξ_j is the SNR per subband, i.e.,

$$\xi_j \triangleq \frac{\sigma_{\mathcal{S}_j}^2}{\sigma_{N_j}^2}, \quad (31)$$

where

$$\sigma_{\mathcal{S}_j}^2 = \sum_{k \in \mathbb{B}_j} \omega_{j,k} v_k^2 \sigma_{S_k}^2, \quad (32)$$

is the processed speech power. Hence, the *processed* subband SNR is the ratio of the *processed* speech power to the near-end noise power at the listener.

Let I_j be a given minimum requirement on the audibility performance in a particular subband, j , i.e., we require

$$f(\xi_j) \geq I_j. \quad (33)$$

It is beneficial to interpret the subband audibility limits, I_j , in terms of the subband SNR. Using (30), it follows that (33) can be written as

$$\xi_j \geq \frac{I_j}{1 - I_j} \triangleq I_j^\xi \quad (34)$$

where $I_j^\xi \triangleq \frac{I_j}{1 - I_j}$, may be considered a desired minimum processed SNR. Therefore, instead of focusing on choosing an intelligibility limit, I_j , we can equivalently choose an SNR limit I_j^ξ , or interpret the desired audibility limits as a limit on the subband SNR. Hence, the performance criterion in (3) is chosen as the subband SNR, i.e.,

$$\mathcal{I}_j(\mathbf{S}_j, \mathbf{Z}_j) = \xi_j = \frac{\sigma_{\mathcal{S}_j}^2}{\sigma_{N_j}^2}. \quad (35)$$

Writing out the performance criterion we have that

$$\mathcal{I}_j(\mathbf{S}_j, \mathbf{Z}_j) \geq I_j^\xi \quad (36)$$

$$\sum_{k \in \mathbb{B}_j} \omega_{j,k} v_k^2 \sigma_{S_k}^2 \geq \sigma_{N_j}^2 I_j^\xi. \quad (37)$$

APPENDIX D PROOF OF THEOREM 1

We first see that the optimization problem (3) can be written as (8) by inserting (6) and (7) in (3).

To prove that the solution to (8) is given by the weights in (9), we start by making an observation on the relationship between the MSE cost function and the performance constraint in (8). Clearly, the unconstrained solution to the MSE minimization problem is $v_k = 1$ for $k \in \mathbb{B}_j$, i.e., no processing obviously minimizes the MSE.

The interesting scenario is then that of $\sum_{k \in \mathbb{B}_j} \omega_{j,k} \sigma_{S_k}^2 < \sigma_{N_j}^2 I_j^\xi$. Since we have no energy constraint and we must increase the speech energy to not violate the performance constraint we can lower bound the gains as $v_k \geq 1$ for all $k \in \mathbb{B}_j$, i.e., we only increase the speech power and have no reason to decrease it. Hence, the optimization problem becomes

$$\begin{aligned} \min_{\{v_k\} \in \mathbb{R}_+, k \in \mathbb{B}_j} \quad & \sum_{k \in \mathbb{B}_j} \omega_{j,k} (1 - v_k)^2 \sigma_{S_k}^2, \\ \text{subject to} \quad & \sum_{k \in \mathbb{B}_j} \omega_{j,k} v_k^2 \sigma_{S_k}^2 \geq \sigma_{N_j}^2 I_j^\xi, \\ & v_k \geq 1. \end{aligned} \quad (38)$$

Formulating the Lagrangian we have

$$\begin{aligned} \mathcal{L}(\{v_k\}, \lambda_j, \{\mu_k\}) = & \sum_{k \in \mathbb{B}_j} \omega_{j,k} (1 - v_k)^2 \sigma_{S_k}^2 + \mu_k (1 - v_k) \\ & + \lambda_j \left(\sigma_{N_j}^2 I_j^\xi - \sum_{k \in \mathbb{B}_j} \omega_{j,k} v_k^2 \sigma_{S_k}^2 \right) \end{aligned}$$

Taking the derivative with respect to v_k for a particular $k \in \mathbb{B}_j$, we have

$$\frac{\partial}{\partial v_k} \mathcal{L} = -2\omega_{j,k} (1 - v_k) \sigma_{S_k}^2 - 2\lambda_j \omega_{j,k} v_k \sigma_{S_k}^2 - \mu_k. \quad (39)$$

Since we optimize over a non-convex set the optimization problem is non-convex. Thus, the Karush-Kuhn-Tucker conditions for this problem are not sufficient for a global optimum. However, they are still necessary conditions, thus we use these to determine a solution to the problem.

$$v_k \geq 1, \quad \mu_k \geq 0, \quad \mu_k (1 - v_k) = 0, \quad \lambda_j \geq 0 \quad (40a)$$

$$\sigma_{N_j}^2 I_j^\xi - \sum_{k \in \mathbb{B}_j} \omega_{j,k} v_k^2 \sigma_{S_k}^2 \leq 0 \quad (40b)$$

$$\lambda_j \left[\sigma_{N_j}^2 I_j^\xi - \sum_{k \in \mathbb{B}_j} \omega_{j,k} v_k^2 \sigma_{S_k}^2 \right] = 0 \quad (40c)$$

$$-2\omega_{j,k} (1 - v_k) \sigma_{S_k}^2 - 2\lambda_j \omega_{j,k} v_k \sigma_{S_k}^2 - \mu_k = 0 \quad (40d)$$

Isolating μ_k in the last equation we have,

$$-2\omega_{j,k} (1 - v_k) \sigma_{S_k}^2 - 2\lambda_j \omega_{j,k} v_k \sigma_{S_k}^2 \geq 0. \quad (41)$$

Now solving for v_k ,

$$v_k \geq \frac{1}{1 - \lambda_j} \quad (42)$$

Combining this with $v_k \geq 1$, we have

$$v_k = \max \left(1, \frac{1}{1 - \lambda_j} \right). \quad (43)$$

We observe that $v_k = v_{k'}$ for all $k, k' \in \mathbb{B}_j$. That is, the gains are equal for all frequencies within subband j .

Now to determine λ_j , we first notice that if $\lambda_j = 0$ then $v_k = 1 \forall k \in \mathbb{B}_j$. Then the performance constraint can only be satisfied if $\sigma_{S_j}^2 \geq \sigma_{N_j}^2 I_j^\xi$, i.e., if we are in the scenario that does not require processing. Therefore, we must have $\lambda_j > 0$ if $\sigma_{S_j}^2 < \sigma_{N_j}^2 I_j^\xi$. Furthermore, from (43) we see that the non-trivial solution requires $\lambda_j < 1$. Therefore, (40c) is only satisfied for $\lambda_j \in (0, 1)$ if the optimal λ_j satisfy

$$\sigma_{N_j}^2 I_j^\xi - \sum_{k \in \mathbb{B}_j} \omega_{j,k} \frac{1}{(1 - \lambda_j)^2} \sigma_{S_k}^2 = 0 \quad (44)$$

$$\lambda_j = 1 - \sqrt{\frac{\sigma_{S_j}^2}{\sigma_{N_j}^2 I_j^\xi}}. \quad (45)$$

Inserting this into (43) the optimal gain is

$$v_k^* = \max \left(1, \sqrt{\frac{\sigma_{N_j}^2 I_j^\xi}{\sigma_{S_j}^2}} \right) = \sqrt{\frac{\sigma_{N_j}^2 I_j^\xi}{\sigma_{S_j}^2}}, \quad (46)$$

where the second equality follows since we consider the case of $\sigma_{N_j}^2 I_j^\xi > \sigma_{S_j}^2$. Finally, by inserting the optimum v_k^*, λ_j^* and μ_k^* in (40d), it can be seen that this is a stationary point of the Lagrangian and thus dual feasible. Hence, the solution is

$$v_k^* = \begin{cases} 1 & \text{if } \sigma_{S_j}^2 \geq \sigma_{N_j}^2 I_j^\xi \\ \sqrt{\frac{\sigma_{N_j}^2 I_j^\xi}{\sigma_{S_j}^2}} & \text{otherwise} \end{cases}, \quad \forall k \in \mathbb{B}_j. \quad (47)$$

This completes the proof.

REFERENCES

- [1] P. C. Loizou, *Speech Enhancement: Theory and Practice*, 2nd ed. Boca Raton, FL, USA: CRC, 2013.
- [2] S. Gannot, E. Vincent, S. Markovich-Golan, and A. Ozerov, "A consolidated perspective on multimicrophone speech enhancement and source separation," *IEEE/ACM Trans. Audio, Speech Lang. Process.*, vol. 25, no. 4, pp. 692–730, Apr. 2017.
- [3] J. RENNIES, A. PUSCH, H. SCHEPKER, and S. DOULO, "Evaluation of a near-end listening enhancement algorithm by combined speech intelligibility and listening effort measurements," *J. Acoustical Soc. Amer.*, vol. 144, no. 4, pp. EL315–EL321, Oct. 2018.
- [4] R. PRICKEN, M. WÄLTERMANN, E. PAROTAT, M. SOLODUCHA, and A. RAAKE, "Quality aspects of near-end listening enhancement approaches in telecommunication applications," in *Proc. DAGA*, 2017, pp. 872–875.
- [5] Y. TANG, C. ARNOLD, and T. J. COX, "A study on the relationship between the intelligibility and quality of algorithmically-modified speech for normal hearing listeners," *J. Otorhinolaryngol., Hear. Balance Med.*, vol. 1, no. 1, Jun. 2018, Art. no. 5.

- [6] C. H. Taal, R. C. Hendriks, and R. Heusdens, "Speech energy redistribution for intelligibility improvement in noise based on a perceptual distortion measure," *Comput. Speech Lang.*, vol. 28, no. 4, pp. 858–872, Jul. 2014.
- [7] M. Cooke, C. Mayo, C. Valentini-Botinhao, Y. Stylianou, B. Sauert, and Y. Tang, "Evaluating the intelligibility benefit of speech modifications in known noise conditions," *Speech Commun.*, vol. 55, no. 4, pp. 572–585, May 2013.
- [8] M. Cooke, S. King, M. Garnier, and V. Aubanel, "The listening talker: A review of human and algorithmic context-induced modifications of speech," *Comput. Speech Lang.*, vol. 28, no. 2, pp. 543–571, Mar. 2014.
- [9] W. B. Kleijn, J. B. Crespo, R. C. Hendriks, P. Petkov, B. Sauert, and P. Vary, "Optimizing speech intelligibility in a noisy environment: A unified view," *IEEE Signal Process. Mag.*, vol. 32, no. 2, pp. 43–54, Mar. 2015.
- [10] M. Cooke, C. Mayo, and C. Valentini-Botinhao, "Intelligibility-enhancing speech modifications: The hurricane challenge," in *Proc. Interspeech*, 2013, pp. 3552–3556.
- [11] J. Rennie, H. Schepker, C. Valentini-Botinhao, and M. Cooke, "Intelligibility-enhancing speech modifications – The hurricane challenge 2.0," in *Proc. Interspeech*, Oct. 2020 pp. 1341–1345.
- [12] R. Niederjohn and J. Grotelueschen, "The enhancement of speech intelligibility in high noise levels by high-pass filtering followed by rapid amplitude compression," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 24, no. 4, pp. 277–282, Aug. 1976.
- [13] J. L. Hall and J. L. Flanagan, "Intelligibility and listener preference of telephone speech in the presence of babble noise," *J. Acoustical Soc. Amer.*, vol. 127, no. 1, pp. 280–285, Jan. 2010.
- [14] K. Nathwani, F. Hafiz, A. Swain, and R. Biswas, "Speech intelligibility enhancement using an optimal formant shifting approach," in *Proc. IEEE 12th Int. Symp. Image Signal Process. Anal.*, 2021, pp. 120–125.
- [15] C. Tantubundhit et al., "New signal decomposition method based speech enhancement," *Signal Process.*, vol. 87, no. 11, pp. 2607–2628, Nov. 2007.
- [16] E. Jokinen, U. Remes, P. Alku, E. Jokinen, U. Remes, and P. Alku, "Intelligibility enhancement of telephone speech using Gaussian process regression for normal-to-lombard spectral tilt conversion," *IEEE/ACM Trans. Audio, Speech Lang. Process.*, vol. 25, no. 10, pp. 1985–1996, Oct. 2017.
- [17] R. Zhang, R. Hu, G. Li, and X. Wang, "Spectral Tilt Estimation for Speech Intelligibility Enhancement Using RNN Based on All-Pole Model," in *Proc. Int. Conf. Multimedia Model.*, 2019, pp. 144–156.
- [18] G. Li, H. Ruimin, R. Zhang, and X. Wang, "A mapping model of spectral tilt in normal-to-lombard speech conversion for intelligibility enhancement," *Multimedia Tools Appl.*, vol. 79, no. 27–28, pp. 19471–19491, Jul. 2020.
- [19] T.-C. Zorila, V. Kandia, and Y. Stylianou, "Speech-in-noise intelligibility improvement based on power recovery and dynamic range compression," in *Proc. IEEE 20th Eur. Signal Process. Conf.*, 2012, pp. 2075–2079.
- [20] C. Chermaz and S. King, "A sound engineering approach to near end listening enhancement," in *Proc. Interspeech*, 2020, pp. 1356–1360.
- [21] B. Sauert and P. Vary, "Near end listening enhancement: Speech intelligibility improvement in noisy environments," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process. Proc.*, 2006, pp. 493–496.
- [22] B. Sauert, G. Enzner, and P. Vary, "Near end listening enhancement with strict loudspeaker output power constraining," in *Proc. Int. Workshop Acoust. Echo Noise Control*, Paris, 2006, pp. 1–4.
- [23] M. Niermann, P. Jax, and P. Vary, "Near-end listening enhancement by noise-inverse speech shaping," in *Proc. IEEE 24th Eur. Signal Process. Conf.*, 2016, pp. 2390–2394.
- [24] M. Niermann and P. Vary, "Listening enhancement in noisy environments: Solutions in time and frequency domain," *IEEE/ACM Trans. Audio, Speech Lang. Process.*, vol. 29, pp. 699–709, 2021.
- [25] *Methods for Calculation of the Speech Intelligibility Index*, ANSI/ASA S3.5–1997, Acoustical Society of America, New York, NY, USA, 2017.
- [26] B. Sauert and P. Vary, "Recursive closed-form optimization of spectral audio power allocation for near end listening enhancement," in *Proc. ITG-Fachtagung Sprachkommunikation*, Oct. 2010, pp. 1–4.
- [27] C. H. Taal, J. Jensen, and A. Leijon, "On optimal linear filtering of speech for near-end listening enhancement," *IEEE Signal Process. Lett.*, vol. 20, no. 3, pp. 225–228, Mar. 2013.
- [28] H. Schepker, J. Rennie, and S. Doclo, "Speech-in-noise enhancement using amplification and dynamic range compression controlled by the speech intelligibility index," *J. Acoustical Soc. Amer.*, vol. 138, no. 5, pp. 2692–2706, Nov. 2015.
- [29] R. C. Hendriks, J. B. Crespo, J. Jensen, and C. H. Taal, "Optimal near-end speech intelligibility improvement incorporating additive noise and late reverberation under an approximation of the short-time SII," *IEEE/ACM Trans. Audio, Speech Lang. Process.*, vol. 23, no. 5, pp. 851–862, May 2015.
- [30] M. Niermann, P. Jax, and P. Vary, "Joint near-end listening enhancement and far-end noise reduction," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2017, pp. 4970–4974.
- [31] B. Sauert, *Near-End Listening Enhancement: Theory and Application* (Aachener Beiträge zu digitalen Nachrichtensystemen Series), 1st ed. Aachen, Germany: Wissenschaftsverlag Mainz, 2014.
- [32] A. J. Fuglsig, J. Ostergaard, J. Jensen, L. S. Bertelsen, P. Mariager, and Z.-H. Tan, "Joint far- and near-end speech intelligibility enhancement based on the approximated speech intelligibility index," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2022, pp. 7752–7756.
- [33] S. Van Kuyk, W. B. Kleijn, and R. C. Hendriks, "An evaluation of intrusive instrumental intelligibility metrics," *IEEE/ACM Trans. Audio, Speech Lang. Process.*, vol. 26, no. 11, pp. 2153–2166, Nov. 2018.
- [34] H. Li, S.-W. Fu, Y. Tsao, and J. Yamagishi, "iMetricGAN: Intelligibility enhancement for speech-in-noise using generative adversarial network-based metric learning," in *Proc. Interspeech*, 2020, pp. 1336–1340.
- [35] H. Li and J. Yamagishi, "Multi-metric optimization using generative adversarial networks for near-end speech intelligibility enhancement," *IEEE/ACM Trans. Audio, Speech Lang. Process.*, vol. 29, pp. 3000–3011, 2021.
- [36] J. Jensen and C. H. Taal, "An algorithm for predicting the intelligibility of speech masked by modulated noise maskers," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 24, no. 11, pp. 2009–2022, Nov. 2016.
- [37] S. M. Kuo and D. R. Morgan, "Active noise control: A tutorial review," *Proc. IEEE*, vol. 87, no. 6, pp. 943–973, Jun. 1999.
- [38] N. V. George and G. Panda, "Advances in active noise control: A survey, with emphasis on recent nonlinear techniques," *Signal Process.*, vol. 93, no. 2, pp. 363–377, Feb. 2013.
- [39] G. Li, R. Hu, X. Wang, and R. Zhang, "A near-end listening enhancement system by RNN-based noise cancellation and speech modification," *Multimedia Tools Appl.*, vol. 78, no. 11, pp. 15483–15505, Jun. 2019.
- [40] F. Cheng, X. Wang, L. Gang, W. Tu, and J. Wang, "Speech intelligibility enhancement in strong mechanical noise based on neural networks," in *Proc. Pacific Rim Conf. Multimedia*, 2018, pp. 702–712.
- [41] M. Cooke, "A glimpsing model of speech perception in noise," *J. Acoustical Soc. Amer.*, vol. 119, no. 3, pp. 1562–1573, Feb. 2006.
- [42] B. Sauert and P. Vary, "Near end listening enhancement optimized with respect to speech intelligibility index," in *Proc. IEEE 17th Eur. Signal Process. Conf.*, 2009 pp. 1844–1848.
- [43] A. Zahedi, M. S. Pedersen, J. Østergaard, T. U. Christiansen, L. Bramsløw, and J. Jensen, "Minimum processing beamforming," *IEEE/ACM Trans. Audio, Speech Lang. Process.*, vol. 29, pp. 2710–2724, 2021.
- [44] C. H. Taal, R. C. Hendriks, R. Heusdens, and J. Jensen, "An algorithm for intelligibility prediction of time-frequency weighted noisy speech," *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 19, no. 7, pp. 2125–2136, Sep. 2011.
- [45] K. S. Rhebergen and N. J. Versfeld, "A speech intelligibility index-based approach to predict the speech reception threshold for sentences in fluctuating noise for normal-hearing listeners," *J. Acoustical Soc. Amer.*, vol. 117, no. 4, pp. 2181–2192, Apr. 2005.
- [46] J. M. Kates and K. H. Arehart, "The hearing-aid speech perception index (HASPI)," *Speech Commun.*, vol. 65, pp. 75–93, Nov./Dec. 2014.
- [47] J. M. Kates and K. H. Arehart, "The hearing-aid speech quality index (HASQI) version 2," *J. Audio Eng. Soc.*, vol. 62, no. 3, pp. 99–117, 2014.
- [48] *Perceptual Evaluation of Speech Quality (PESQ)*, Rec. ITU-T P.862, International Telecommunication Union, Geneva, Switzerland, Recommendation, Feb. 2001.
- [49] S. Khademi, R. C. Hendriks, and W. B. Kleijn, "Intelligibility enhancement based on mutual information," *IEEE/ACM Trans. Audio, Speech Lang. Process.*, vol. 25, no. 8, pp. 1694–1708, Aug. 2017.
- [50] *Wideband Extension to Recommendation P.862 for the Assessment of Wideband Telephone Networks and Speech Codecs*, Rec. ITU-T P.862.2, International Telecommunication Union, Geneva, Switzerland, Nov. 2007.
- [51] J. H. L. Hansen and B. L. Pellom, "An effective quality evaluation protocol for speech enhancement algorithms," in *Proc. Int. Conf. Speech Lang. Process.*, 1998, pp. 2819–2822.
- [52] J. S. Garofolo et al., "TIMIT acoustic-phonetic continuous speech corpus LDC93S1," Web Download. Philadelphia: Linguistic Data Consortium, 1993.
- [53] S. van de Par, A. Kohlrausch, R. Heusdens, J. Jensen, and S. H. Jensen, "A perceptual model for sinusoidal audio coding based on spectral integration," *EURASIP J. Adv. Signal Process.*, vol. 2005, no. 9, pp. 1292–1304, Jun. 2005.
- [54] N. Jillings, D. Moffat, B. D. Man, and J. D. Reiss, "Web audio evaluation Tool: A browser-based listening test environment," pp. 147–152, Jul. 2015.

- [55] *Algorithms to Measure Audio Programme Loudness and True-Peak Audio Level*, Rec. ITU-R BS.1770-4, International Telecommunication Union, Geneva, Switzerland, Oct. 2015.
- [56] *Loudness Normalisation and Permitted Maximum Level of Audio Signals*, Rec. R-128, European Broadcasting Union, Geneva, Switzerland, 2014.
- [57] B. D. Man, "Web audio evaluation tool," May 2022. [Online]. Available: <https://github.com/BrechtDeMan/WebAudioEvaluationTool>
- [58] *Method for the Subjective Assessment of Intermediate Quality Level of Audio Systems*, Rec. ITU-R BS.1534-3, International Telecommunication Union, Geneva, Switzerland, Oct. 2015.



Andreas Jonas Fuglsig (Student Member, IEEE) received the B.Sc. and M.Sc. degrees in mathematical engineering from Aalborg University, Aalborg, Denmark, in 2017 and 2019, respectively. Since 2020, he has been working toward the Ph.D. degree with the Centre on Acoustic Signal Processing Research, Aalborg University and RTX A/S, Nørresundby, Denmark. In 2019, he was an R&D Engineer with RTX A/S. His research interests include acoustic signal processing, speech enhancement, and information theory.



Jesper Jensen received the M.Sc. degree in electrical engineering and the Ph.D. degree in signal processing from Aalborg University, Aalborg, Denmark, in 1996 and 2000, respectively. From 1996 to 2000, he was a Ph.D. Student and an Assistant Research Professor with the Center for Person Kommunikation, Aalborg University. From 2000 to 2007, he was a Postdoctoral Researcher and an Assistant Professor with the Delft University of Technology, Delft, The Netherlands, and an External Associate Professor with Aalborg University. He is currently a Fellow with Oticon A/S,

Denmark, where his main responsibility is scouting and development of new signal processing concepts for hearing aid applications. He is a Professor with the Section for Signal and Information Processing, Department of Electronic Systems, Aalborg University. He is also a Co-Founder of the Centre for Acoustic Signal Processing Research, Aalborg University. His research interests include the area of acoustic signal processing, including signal retrieval from noisy observations, coding, speech and audio modification and synthesis, intelligibility enhancement of speech signals, signal processing for hearing aid applications, and perceptual aspects of signal processing.



Zheng-Hua Tan (Senior Member, IEEE) received the B.Sc. and M.Sc. degrees in electrical engineering from Hunan University, Changsha, China, in 1990 and 1996, respectively, and the Ph.D. degree in electronic engineering from Shanghai Jiao Tong University, Shanghai, China, in 1999. He is currently a Professor with the Department of Electronic Systems and a Co-Head of the Centre for Acoustic Signal Processing Research, Aalborg University, Aalborg, Denmark. He is also a Co-Lead of the Pioneer Centre for AI, Copenhagen, Denmark. He was a Visiting

Scientist with the Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA, USA, an Associate Professor with the Department of Electronic Engineering, Shanghai Jiao Tong University, Shanghai, China, and a Postdoctoral Fellow with AI Laboratory, Korea Advanced Institute of Science and Technology, Daejeon, Korea. He has authored or coauthored more than 200 refereed publications. His research interests include machine learning, deep learning, pattern recognition, speech and speaker recognition, noise-robust speech processing, multimodal signal processing, and social robotics. He is the Chair of IEEE Signal Processing Society Machine Learning for Signal Processing Technical Committee. He is an Associate Editor for the IEEE/ACM TRANSACTIONS ON AUDIO, SPEECH AND LANGUAGE PROCESSING. He was an Associate/Guest Editor for several other journals. He was the General Chair of IEEE MLSP 2018 and a TPC Co-Chair of IEEE SLT 2016.

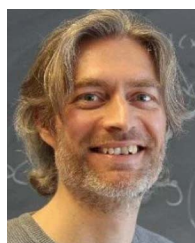


Lars Søndergaard Bertelsen received the M.S.E.E. degree from Aalborg University, Aalborg, Denmark, in 2000. He was with the Telecommunication Industry for several years doing both software, ASIC/FPGA, and system design. Recently, he was also with DSP audio processing in wireless audio products.



Jens Christian Lindof (Member, IEEE) received the M.Sc.EE degree in telecommunication from Aalborg University, Aalborg, Denmark, in 1991. He is currently a Chief Technology Officer of RTX A/S, Nørresundby, Denmark, where he is heading the company's technology road-mapping activities. He was with Texas Instruments Denmark as a Senior Member Technical Staff, Site Manager, and responsible for advanced technology development. He was a Company Supervisor for several Industrial Ph.D. students ranging from novel power amplifier designs to sound

processing algorithms and importance of cooperative technical clusters. He was a co-author of several papers and patents within the area of RF, cellular, and audio. He was keen on entrepreneurship and thus involved with many start-ups including a handful cofounded over the last 25 years.



Jan Østergaard (Senior Member, IEEE) received the M.Sc. degree in electrical engineering from Aalborg University, Aalborg, Denmark, in 1999 and the Ph.D. degree (*cum laude*) in electrical engineering from Delft University of Technology, Delft, The Netherlands, in 2007. He was an R&D Engineer with ETI A/S, Aalborg, from 1999 to 2002, and ETI Inc., Virginia Beach, VA, USA, from 2002 to 2003. From 2007 to 2008, he was a Postdoctoral Researcher with the University of Newcastle, Newcastle, NSW, Australia. He has been a Visiting Researcher with Tel Aviv

University, Tel Aviv, Israel, and Universidad Técnica Federico Santa María, Valparaíso, Chile. He is currently a Full Professor of information theory and signal processing and the Head of the Section on AI and Sound and Centre on Acoustic Signal Processing Research, Aalborg University. His research interests include acoustic signal processing, statistical signal processing, information theory, joint source-channel coding, and networked control theory. Dr. Østergaard was the recipient of the Danish Independent Research Council's Young Researcher's Award, Best Ph.D. Thesis Award by the European Association for Signal Processing, and fellowships from the Danish Independent Research Council and Villum Foundations Young Investigator Programme. He is an Associate Editor for the IEEE TRANSACTIONS ON INFORMATION THEORY.