

Phase-Aware Deep Speech Enhancement: It's All About The Frame Length

Tal Peer^{a)} and Timo Gerkmann

Signal Processing (SP), Universität Hamburg, Germany
tal.peer@uni-hamburg.de, timo.gerkmann@uni-hamburg.de

Abstract: Algorithmic latency in speech processing is dominated by the frame length used for Fourier analysis, which in turn limits the achievable performance of magnitude-centric approaches. As previous studies suggest the importance of phase grows with decreasing frame length, this work presents a systematical study on the contribution of phase and magnitude in modern Deep Neural Network (DNN)-based speech enhancement at different frame lengths. Results indicate that DNNs can successfully estimate phase when using short frames, with similar or better overall performance compared to using longer frames. Thus, interestingly, modern phase-aware DNNs allow for low-latency speech enhancement at high quality.

Date: 5 October 2022

1. Introduction

Single-channel speech enhancement is often carried out in the time-frequency domain where the signals are represented by their time-varying frequency content. To obtain the time-frequency representation, one applies a transformation such as the short-time Fourier transform (STFT) which has a number of free parameters. These parameters (namely frame length, frame shift and window function, see also Section 2) must be chosen appropriately, e.g. based on physical characteristics of the speech signal. However, not only the signal should be considered, but also the algorithm which will be applied on the time-frequency representation; the choice of STFT parameters should result in a representation that is the most *useful* for the algorithm at hand.¹

In this letter, we focus expressly on the choice of frame length for speech enhancement algorithms based on deep neural networks (DNNs) and specifically consider phase-aware approaches. Choosing an adequate frame length is an important decision in the design of STFT-based systems. On one hand, it largely determines the overall algorithmic latency of the system, hence real-time systems such as hearing aids would tend to use short frames. On the other hand, short frames lead to limited spectral resolution which hinders many algorithms and also result in a larger number of frames for a given signal, potentially increasing computational complexity, e.g. when using temporal convolutions. The STFT representation is complex-valued and commonly separated into a *magnitude spectrogram* and *phase spectrogram*. The relevance of the phase spectrogram to the speech enhancement task has been a topic of debate. Traditionally it has been considered to be of little to no importance due to empirical studies² as well as theoretical results.³ However, more recent studies have shown that phase does carry speech-relevant information.^{4,5} Motivated by these findings, phase-aware speech processing has been enjoying a certain renaissance and several phase-aware methods have been proposed.^{6–10}

In recent years DNNs have rapidly become the tool of choice in many fields, including audio and speech processing. Consequently, many recent phase-aware speech enhancement and source separation methods use a DNN to either directly estimate the phase spectrogram^{11–13} or estimate phase derivatives and reconstruct the phase from them.^{14,15} Other DNN-based approaches include directly operating on complex spectrograms without separating into magnitude and phase^{16–18} or simply taking phase into consideration for improved magnitude estimation.¹⁹

Some authors have taken a different path and altogether replaced the STFT-based representation with a learned encoder-decoder mechanism, which usually results in a real-valued representation.^{20–22} An interesting aspect of these learned encoder-decoder approaches is that they show very good performance when using very short frames of about 2 ms, even going as short as 0.125 ms.²¹ This stands in sharp contrast to traditional STFT-based approaches which generally use frame lengths of about 20 ms to 60 ms. Note that while learned encoder-decoder approaches have been originally proposed for source separation, they also show good performance on the speech enhancement task.^{23,24}

Following the publication of the pioneering learned encoder-decoder Conv-TasNet model,²⁰ several authors have proposed extensions and analyses. Among other results, it has been shown that the main contributing factors to the performance of Conv-TasNet are the use of short frames and time-domain loss function, not the learned encoder-decoder.^{25,26} It has also been shown that when replacing the learned encoder with the STFT, the optimal set of input features depends on the chosen frame length;^{26,27} for longer frames (25 ms to 64 ms) the magnitude spectrum works well, while shorter frames (2 ms to 4 ms) show better performance only with the full complex spectrogram as input (in form of concatenated real and imaginary parts). This observation is especially important, since it means that phase-aware speech processing (with either implicit or explicit phase estimation) should possibly employ different frame lengths than magnitude-only processing.

While the choice of loss function in phase-aware speech enhancement DNNs has been studied with respect to perceptual measures,²⁴ and the effect of STFT parameters on magnitude-only DNNs has also been analyzed,²⁸ we are not aware of an analysis regarding the choice of frame length in the phase-aware setting. Previous studies unrelated to DNNs have shown that the importance of phase to speech-related tasks varies with the choice of STFT parameters. In particular, it has been shown that departing from the typical frame lengths used in speech processing (corresponding to about 20 ms to 40 ms) and either using shorter frames^{29,30} or using a window shape that effectively shortens the frame⁴ can result in very good signal reconstruction from the phase spectrogram alone. A similar result has been observed for longer-than-typical frames.^{29,31} However, as long frames yield algorithmic latencies which may be prohibitive for many real-time

^{a)} Author to whom correspondence should be addressed.

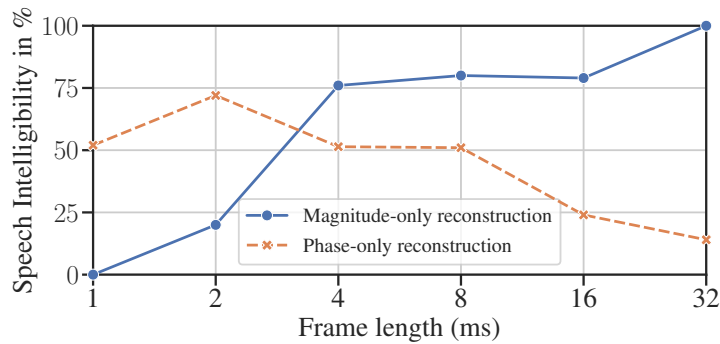


Fig. 1. Redrawn excerpt of experimental results by Kazama et al.²⁹ (shown here for illustration purposes). When reconstructing speech either from the magnitude spectrogram or phase spectrogram (replacing the other component by noise), the intelligibility of the resulting signal strongly depends on the choice of frame length.

speech processing devices, here we choose to focus on short frames for their potential benefit regarding latency. Figure 1, based on results from Ref. 29, shows how the contribution of phase and magnitude to the intelligibility of reconstructed clean speech changes with varying frame length. One observes that as the frames become shorter, phase becomes more important while magnitude gradually loses relevance. Note, however, that these findings are based on somewhat artificial signal reconstruction experiments on oracle data — whether or not they also apply to actual speech enhancement, source separation, etc. remains unclear.

Typical frame lengths in speech processing — around 32 ms — correspond to an interval that is short enough to be considered quasi-stationary but long enough to cover multiple fundamental periods of voiced speech (whose fundamental period lies between 2 ms and 12.5 ms).³² These considerations apply to the magnitude spectrogram but not necessarily to the phase spectrogram. Indeed, it seems that the irrelevance often attributed to the phase spectrogram is in part due to the choice of frame length in experiments. Note that existing model-based phase estimation methods do not typically operate on short frames^{9,10,33,34} and thus do not attempt to take advantage of the findings regarding phase importance and frame length in perceptual studies.

Based on these previous results and observations, we seek to answer two questions in this letter: a) How does the choice of frame length affect magnitude and phase estimation in a phase-aware DNN? b) Does DNN-based phase estimation allow the use of shorter frames, thus reducing algorithmic latency? Using an example DNN with explicit phase estimation, we analyze and compare the performance under different frame lengths. In order to gain further insight, we also attempt to characterize the relative contribution of the magnitude and phase spectrograms at each frame length and show that the aforementioned observations on the importance of phase in short frames also carry over to the context of speech enhancement.

2. Preliminaries

The STFT of a discrete time-domain signal $x(n)$ is computed by segmenting the signal into overlapping frames of length M and shift H . A real-valued multiplicative window function $w(n)$ is applied to each frame, which is then transformed to the frequency domain with the discrete Fourier transform (DFT). Assuming the one-sided M -point DFT is used, we obtain the *complex spectrogram* $X \in \mathbb{C}^{K \times L}$, defined as

$$X_{k,\ell} = \sum_{n=0}^{M-1} x(\ell H + n)w(n)e^{-j2\pi \frac{kn}{M}}, \quad (1)$$

where k is the frequency index, ℓ is the frame index, $K = \frac{M}{2} + 1$ is the number of frequency bins and L is the number of time frames. Unless otherwise noted, we always consider the whole spectrogram and thus omit the indices in the following. We also define an *overlap ratio* $R = \frac{M-H}{M}$ for convenience. As M is the number of samples in a single frame, we define $M_t = \frac{M}{f_s}$ (where f_s is the sampling frequency) as the physical frame length, measured in seconds. The term *frame length* will refer to M_t from this point onwards.

In the context of speech enhancement we consider an additive noise model, which is expressed in polar coordinates in terms of the magnitude spectrogram $|X|$ and phase spectrogram ϕ_X :

$$|X|e^{j\phi_X} = |S|e^{j\phi_S} + |V|e^{j\phi_V}, \quad (2)$$

where S and V are the clean speech signal and an additive noise component, respectively. Given the noisy signal X , the task is to compute an estimate $\hat{S} = |\hat{S}|e^{j\hat{\phi}_S}$ which is subsequently transformed back into the time-domain, yielding the estimated clean signal $\hat{s}(n)$.

3. Neural network architecture

The DNN architecture proposed here is an adaptation of a previously proposed model¹² for audio-visual speech separation and enhancement, consisting of loosely coupled magnitude and phase sub-networks. Although we do not consider an audio-visual input here, this model is relatively simple and performs explicit estimation of both magnitude and phase, which is essential for our experiments. The parts pertaining to the video stream are simply omitted and the model is adapted accordingly. The resulting network is depicted in Fig. 2 and described below.

Both sub-networks are realized as convolutional neural networks (CNNs) using one-dimensional depthwise separable convolution layers³⁵ along the time axis (the different frequency bins at the input are considered as channels in this setup). Both networks consist of multiple identical residual blocks; The basic building block is composed of a pre-activation (ReLU), a batch normalization layer and a convolutional layer whose output is added to the block's input.

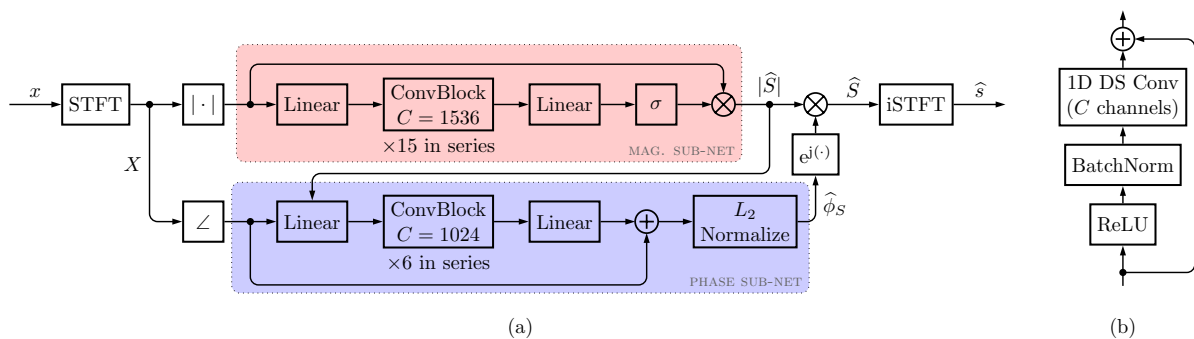


Fig. 2. (a) Overview of the proposed network architecture, based on the audio-visual model by Afouras et al.,¹² albeit only using the noisy speech signal as input. Note that the input to the phase sub-network consists of the estimated magnitude along with the cosine and sine of the noisy phase, which are shown here as a single input for simplicity. (b) Detail of the basic convolutional block.

3.1 Magnitude sub-network

The magnitude sub-network takes the noisy magnitude spectrogram $|X|$ outputs a real mask which is applied to the noisy STFT magnitude spectrogram to produce a magnitude estimate (note that the original network¹² used the mel-scale spectrogram along with video features as input). The noisy magnitude spectrogram is fed through a chain of 15 convolutional blocks with 1536 input/output channels each. Linear layers at the input and output help to model the inter-frequency relationships and project the data into the correct dimensions. A sigmoid activation function is applied to the output, resulting in a real mask with values in $[0, 1]$. The real mask is multiplied with the input, resulting in a magnitude estimate $|\hat{S}|$.

3.2 Phase sub-network

The input to the phase sub-network is a concatenation of $|\hat{S}|$, $\cos(\phi_X)$ and $\sin(\phi_X)$ along the frequency axis. This is fed into a linear input layer and subsequently through 6 convolutional blocks with 1024 channels and a linear output layer. The output of the linear layer is treated as the concatenated cosine and sine of the phase residual, which are added to the respective inputs. The resulting estimate is L_2 -normalized to ensure that the cosine and sine outputs are consistent with each other (i.e. that they represent a unit vector on the complex plane). By having the estimated magnitude as an additional input, the phase sub-network can learn to focus on high-energy spectrogram regions, resulting in an overall better phase estimate.

3.3 Training procedure

The magnitude and phase sub-networks are trained jointly using a training set consisting of pairs of noisy and clean speech samples at different signal-to-noise ratios (SNR), see Section 4.4.1 for further details. We employ a time-domain loss, namely the negative scale invariant signal to distortion ratio (SI-SDR).³⁶ The SI-SDR loss has been shown to produce superior results over frequency-domain loss functions when both magnitude and phase spectrograms are estimated.²⁶

4. Experiments

The main experiment we conduct is a comparison of the model’s performance under variation of the STFT frame length M_t , in terms of perceptual measures. Since the model we consider includes explicit estimation of phase and magnitude, we are able to also analyze and quantify the relative contribution of magnitude and phase estimation, again as a function of frame length. This analysis is conducted in a manner comparable with the perceptual experiments in previous works,^{4,29,30} although here we use estimates of the clean magnitude and phase, rather than the clean or noisy signals. For each frame length, we produce three estimates of the clean speech signal — the actual output of the network as well as two synthetic signals composed of the estimated magnitude and noisy phase or vice-versa:

$$\hat{s} = \text{iSTFT}\{|\hat{S}|e^{j\hat{\phi}_s}\}, \quad \hat{s}_{\text{mag}} = \text{iSTFT}\{|\hat{S}|e^{j\hat{\phi}_x}\}, \quad \hat{s}_{\text{ph}} = \text{iSTFT}\{|X|e^{j\hat{\phi}_s}\}. \quad (3)$$

To allow for a fair comparison we must keep the number of DNN parameters constant. In the case of the network architecture we consider, the number of parameters depends on the number of frequency bins K . Hence, we zero-pad the frames prior to applying the DFT, resulting in a constant number of bins $K = 257$, which corresponds to the longest frames we consider ($M_t = 32$ ms) at $f_s = 16$ kHz. All experiments employ a square-root Hann window with an overlap ratio $R = \frac{1}{2}$. The same window is used for the forward and inverse STFT operations.

4.1 Data and training details

For training we use clean and noisy excerpts from the 2020 Deep Noise Suppression (DNS) dataset³⁷ with SNR $\in \{-5, 0, \dots, 10\}$ dB. Each excerpt is 2 s long and the data set contains in total 100 h of speech, from which 80 % are used for training and the remaining 20 % for validation. We train all models using the Adam optimizer, a batch size of 32 and a learning rate of 10^{-4} . Training is stopped if the validation loss has not decreased for 10 epochs.

Evaluation is performed on a custom test set composed of clean speech from the WSJ corpus³⁸ and noise from the CHiME3 dataset,³⁹ mixed with SNR $\in \{-10, -5, \dots, 20\}$ dB. This test set contains 672 excerpts in total. All training and evaluation data is sampled at $f_s = 16$ kHz.

4.2 Evaluation details

In addition to evaluation in terms of instrumental quality and intelligibility measures (POLQA, ESTOI) using the entire evaluation dataset, we also report the results of a small-scale listening experiment in which participants were asked to rate the overall quality of the estimated signals introduced in Eq. (3). The listening experiment consisted of 12 trials. In each trial the participants were presented with a reference clean signal from the WSJ corpus and then asked to rate the quality of eight different signals: A noisy version from the evaluation dataset (at 0 dB SNR), the reference itself, the model output \hat{s} , and the phase/magnitude-based reconstructions (\hat{s}_{ph} and \hat{s}_{mag}) for two representative frame lengths — 32 ms and 4 ms. Ten normal-hearing individuals aged 25 to 43 participated in the experiment and rated each signal on a continuous quality scale (CQS) from 0 to 100.

5. Results and discussion

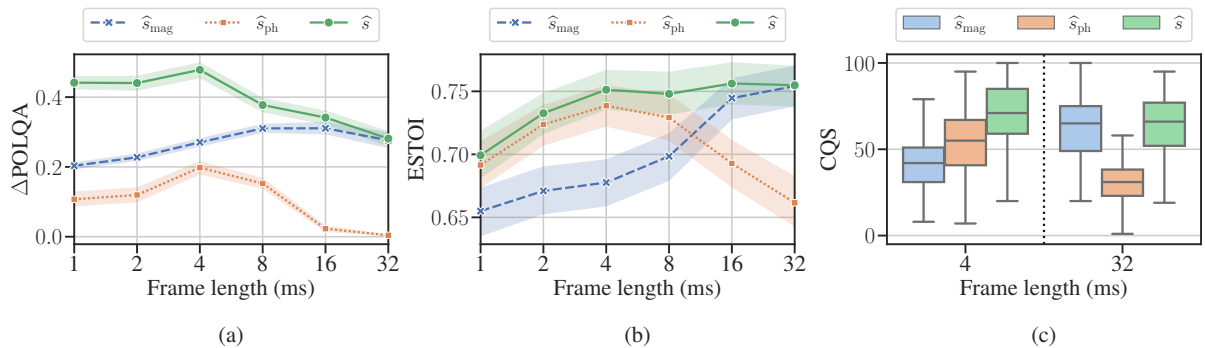


Fig. 3. (a),(b) Evaluation results on the WSJ/CHiME test set (see Section 4.4.1). The plots show the mean improvement of POLQA score and the mean ESTOI score over all input SNRs as function of the frame length. Error bands depict the 95 % confidence interval. (c) Results of listening experiment on two frame lengths, where participants were asked to rate the different signals on a continuous quality scale (CQS) in the range 0–100.

Evaluation results are depicted in Fig. 3. We first consider the effect of frame length on the overall estimate \hat{s} , which uses both the estimated magnitude and estimated phase for iSTFT reconstruction. Intelligibility (in terms of ESTOI) is not affected much by the choice of frame length except for very short frames, which cause rapid degradation.

In terms of speech quality (ΔPOLQA), however, we observe a consistent improvement with decreasing frame length. The improvement in speech quality reaches a maximum at $M_t = 4$ ms, after which it starts to decline, while still reaching relatively high values for very short frames of 1 ms to 2 ms.

For both POLQA and ESTOI, the magnitude-based and phase-based estimates (\hat{s}_{mag} and \hat{s}_{ph} , respectively) show an interesting picture: At $M_t = 32$ ms, \hat{s}_{mag} reaches similar values to \hat{s} and the phase-based estimate \hat{s}_{ph} shows virtually no improvement over the noisy input. As the frame length decreases this gradually changes: The magnitude-based estimate loses quality and intelligibility, while the opposite holds for the phase-based estimate. Note that although the magnitude-based estimate shows degraded quality for short frames, the phase spectrum’s contribution is still sufficient to maintain and even boost overall performance of the joint estimate. We attribute this behavior to the relative contribution of phase and magnitude spectra and the interplay between them.

Results of the listening experiment largely agree with the instrumental measures. The joint estimate \hat{s} received very similar quality scores for short and long frames (however without the rising trend shown by POLQA). This indicates, again, that the use of shorter frames (e.g. for reduced latency) does not imply a sacrifice in terms of enhancement performance. The scores given to the phase and magnitude-based estimates are consistent with POLQA and ESTOI scores and support the proposition, that the importance of phase and magnitude estimation is dependent on the chosen frame length. In particular, at the short frame regime, phase estimation plays a larger role than magnitude estimation.

Since the results in Fig. 3 are averaged across all SNRs, we provide further insight in Fig. 4 by showing the POLQA improvement at different SNRs for two selected frame lengths (4 ms, 16 ms). Besides the overall better performance of the shorter frame length and the unimportance of phase estimation with long frames (cf. Fig. 3), the quality of magnitude-based and phase-based estimates shows much stronger dependence on SNR. In particular, the phase-based estimate actually surpasses the magnitude-based estimate in terms of POLQA at low SNRs (≤ 0 dB), suggesting that phase estimation is especially beneficial at difficult noise conditions, in accordance with previous perceptual studies.⁴⁰ This also translates into the joint estimation case (i.e. \hat{s}), where the difference in ΔPOLQA between the frame lengths is more pronounced at low SNRs.

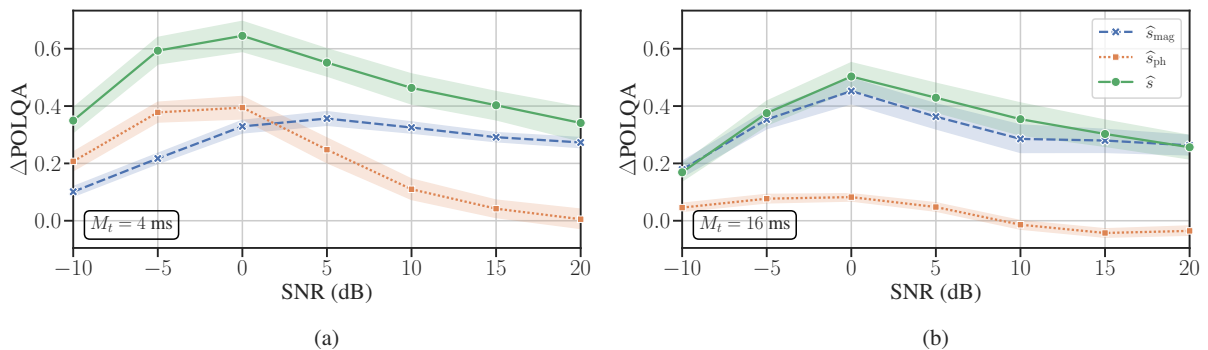


Fig. 4. Mean POLQA improvement on the WSJ/ChiME test set for $M_t \in \{4, 16\}$ ms, shown as a function of input SNR.

6. Conclusion

In this work we have presented a study on the effect of frame length in STFT-based phase-aware speech enhancement with DNNs. Results indicate that the use of short frames in this context (down to 4 ms) does yield similar or better performance in terms of instrumental and subjective quality measures compared to long frames (32 ms). Furthermore, by explicitly estimating phase and magnitude we are able to show that varying the frame length affects the individual contribution of magnitude and phase estimation to the quality of the combined output — at short frames, it is dominated by phase estimation. These findings suggest that by employing explicit phase estimation, speech enhancement DNNs can achieve reduced latency, which is of great importance for speech communication devices such as hearing aids or virtual assistants.

Acknowledgments

This work has been funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – project number 247465126. We would like to thank J. Berger and Rohde&Schwarz SwissQual AG for their support with POLQA.

References and links

- ¹ T. Virtanen, E. Vincent, and S. Gannot, "Time-Frequency Processing: Spectral Properties," in *Audio Source Separation and Speech Enhancement* (John Wiley & Sons, Ltd, 2018), Chap. 2, pp. 15–29.
- ² D. Wang and J. Lim, "The unimportance of phase in speech enhancement," *IEEE Transactions on Acoustics, Speech, and Signal Processing* **30**(4), 679–681 (1982).
- ³ Y. Ephraim and D. Malah, "Speech enhancement using a minimum-mean square error short-time spectral amplitude estimator," *IEEE Transactions on Acoustics, Speech, and Signal Processing* **32**(6), 1109–1121 (1984).
- ⁴ K. Paliwal, K. Wójcicki, and B. Shannon, "The importance of phase in speech enhancement," *Speech Communication* **53**(4), 465–494 (2011).
- ⁵ T. Gerkmann, M. Krawczyk-Becker, and J. Le Roux, "Phase Processing for Single-Channel Speech Enhancement: History and recent advances," *IEEE Signal Processing Magazine* **32**(2), 55–66 (2015).
- ⁶ J. Le Roux and E. Vincent, "Consistent Wiener Filtering for Audio Source Separation," *IEEE Signal Processing Letters* **20**(3), 217–220 (2013).
- ⁷ T. Gerkmann, "MMSE-optimal enhancement of complex speech coefficients with uncertain prior knowledge of the clean speech phase," in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (2014).
- ⁸ T. Gerkmann, "Bayesian Estimation of Clean Speech Spectral Coefficients Given a Priori Knowledge of the Phase," *IEEE Transactions on Signal Processing* **62**(16), 4199–4208 (2014).
- ⁹ M. Krawczyk and T. Gerkmann, "STFT Phase Reconstruction in Voiced Speech for an Improved Single-Channel Speech Enhancement," *IEEE/ACM Transactions on Audio, Speech, and Language Processing* **22**(12), 1931–1940 (2014).
- ¹⁰ P. Mowlae and J. Kulmer, "Harmonic Phase Estimation in Single-Channel Speech Enhancement Using Phase Decomposition and SNR Information," *IEEE/ACM Transactions on Audio, Speech, and Language Processing* **23**(9), 1521–1532 (2015).
- ¹¹ N. Takahashi, P. Agrawal, N. Goswami, and Y. Mitsufuji, "PhaseNet: Discretized Phase Modeling with Deep Neural Networks for Audio Source Separation," in *Interspeech 2018* (2018).
- ¹² T. Afouras, J. S. Chung, and A. Zisserman, "The Conversation: Deep Audio-Visual Speech Enhancement," in *Interspeech 2018* (2018).
- ¹³ J. Le Roux, G. Wichern, S. Watanabe, A. Sarroff, and J. R. Hershey, "Phasebook and Friends: Leveraging Discrete Representations for Source Separation," *IEEE Journal of Selected Topics in Signal Processing* **13**(2), 370–382 (2019).
- ¹⁴ N. Zheng and X.-L. Zhang, "Phase-Aware Speech Enhancement Based on Deep Neural Networks," *IEEE/ACM Transactions on Audio, Speech, and Language Processing* **27**(1), 63–76 (2019).
- ¹⁵ Y. Masuyama, K. Yatabe, Y. Koizumi, Y. Oikawa, and N. Harada, "Phase Reconstruction Based On Recurrent Phase Unwrapping With Deep Neural Networks," in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (2020).
- ¹⁶ D. S. Williamson, Y. Wang, and D. Wang, "Complex ratio masking for joint enhancement of magnitude and phase," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (2016).
- ¹⁷ K. Tan and D. Wang, "Learning Complex Spectral Mapping With Gated Convolutional Recurrent Networks for Monaural Speech Enhancement," *IEEE/ACM Transactions on Audio, Speech, and Language Processing* **28**, 380–390 (2020).
- ¹⁸ Y. Hu, Y. Liu, S. Lv, M. Xing, S. Zhang, Y. Fu, J. Wu, B. Zhang, and L. Xie, "DCCRN: Deep Complex Convolution Recurrent Network for Phase-Aware Speech Enhancement," in *Interspeech 2020* (2020).
- ¹⁹ H. Erdogan, J. R. Hershey, S. Watanabe, and J. Le Roux, "Phase-sensitive and recognition-boosted speech separation using deep recurrent neural networks," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (2015).
- ²⁰ Y. Luo and N. Mesgarani, "Conv-TasNet: Surpassing Ideal Time-Frequency Magnitude Masking for Speech Separation," *IEEE/ACM Transactions on Audio, Speech, and Language Processing* **27**(8), 1256–1266 (2019).
- ²¹ Y. Luo, Z. Chen, and T. Yoshioka, "Dual-Path Rnn: Efficient Long Sequence Modeling for Time-Domain Single-Channel Speech Separation," in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (2020).
- ²² C. Subakan, M. Ravanelli, S. Cornell, M. Bronzi, and J. Zhong, "Attention Is All You Need In Speech Separation," in *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (2021).
- ²³ Y. Koyama, T. Vuong, S. Uhlich, and B. Raj, "Exploring the Best Loss Function for DNN-Based Low-latency Speech Enhancement with Temporal Convolutional Networks," arXiv:2005.11611 [cs, eess] (2020).
- ²⁴ Z.-Q. Wang, G. Wichern, and J. Le Roux, "On the Compensation Between Magnitude and Phase in Speech Separation," *IEEE Signal Processing Letters* **28**, 2018–2022 (2021).

- ²⁵ D. Ditter and T. Gerkmann, "A Multi-Phase Gammatone Filterbank for Speech Separation Via Tasnet," in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (2020).
- ²⁶ J. Heitkaemper, D. Jakobeit, C. Boeddeker, L. Drude, and R. Haeb-Umbach, "Demystifying TasNet: A Dissecting Approach," in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (2020).
- ²⁷ M. Pariente, S. Cornell, A. Deleforge, and E. Vincent, "Filterbank Design for End-to-end Speech Separation," in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (2020).
- ²⁸ D. Takeuchi, K. Yatabe, Y. Koizumi, Y. Oikawa, and N. Harada, "Effect of spectrogram resolution on deep-neural-network-based speech enhancement," *Acoustical Science and Technology* **41**(5), 769–775 (2020).
- ²⁹ M. Kazama, S. Gotoh, M. Tohyama, and T. Houtgast, "On the significance of phase in the short term Fourier spectrum for speech intelligibility," *The Journal of the Acoustical Society of America* **127**(3), 1432–1439 (2010).
- ³⁰ T. Peer and T. Gerkmann, "Intelligibility Prediction of Speech Reconstructed From Its Magnitude or Phase," in *Speech Communication: 14th ITG Conference*, Kiel (online) (2021).
- ³¹ L. Alsteris and K. Paliwal, "Importance of window shape for phase-only reconstruction of speech," in *2004 IEEE International Conference on Acoustics, Speech, and Signal Processing* (2004).
- ³² K. K. Paliwal, J. G. Lyons, and K. K. Wójcicki, "Preference for 20-40 ms window duration in speech analysis," in *2010 4th International Conference on Signal Processing and Communication Systems* (2010).
- ³³ P. Magron, R. Badeau, and B. David, "Phase reconstruction of spectrograms with linear unwrapping: Application to audio signal restoration," in *2015 23rd European Signal Processing Conference (EUSIPCO)* (2015).
- ³⁴ T. Peer, K.-J. Ziegert, and T. Gerkmann, "Plosive Enhancement Using Phase Linearization and Smoothing," in *Speech Communication: 14th ITG Conference*, Kiel (online) (2021).
- ³⁵ F. Chollet, "Xception: Deep Learning with Depthwise Separable Convolutions," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Honolulu, HI (2017).
- ³⁶ 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)* (2019).
- ³⁷ C. K. Reddy, V. Gopal, R. Cutler, E. Beyrami, R. Cheng, H. Dubey, S. Matuselych, R. Aichner, A. Aazami, S. Braun, P. Rana, S. Srinivasan, and J. Gehrke, "The INTERSPEECH 2020 Deep Noise Suppression Challenge: Datasets, Subjective Testing Framework, and Challenge Results," in *Interspeech 2020* (2020).
- ³⁸ D. B. Paul and J. M. Baker, "The design for the wall street journal-based CSR corpus," in *Proceedings of the Workshop on Speech and Natural Language - HLT '91* (1992).
- ³⁹ J. Barker, R. Marxer, E. Vincent, and S. Watanabe, "The third 'CHiME' speech separation and recognition challenge: Dataset, task and baselines," in *2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)* (2015).
- ⁴⁰ M. Krawczyk-Becker and T. Gerkmann, "An evaluation of the perceptual quality of phase-aware single-channel speech enhancement," *The Journal of the Acoustical Society of America* **140**(4), EL364–EL369 (2016).