DEEP GRIFFIN-LIM ITERATION

Yoshiki Masuyama[†], Kohei Yatabe[†], Yuma Koizumi[‡], Yasuhiro Oikawa[†], Noboru Harada[‡]

[†]Department of Intermedia Art and Science, Waseda University, Tokyo, Japan [‡]NTT Media Intelligence Laboratories, Tokyo, Japan

ABSTRACT

This paper presents a novel phase reconstruction method (only from a given amplitude spectrogram) by combining a signal-processingbased approach and a deep neural network (DNN). To retrieve a time-domain signal from its amplitude spectrogram, the corresponding phase is required. One of the popular phase reconstruction methods is the Griffin–Lim algorithm (GLA), which is based on the redundancy of the short-time Fourier transform. However, GLA often involves many iterations and produces low-quality signals owing to the lack of prior knowledge of the target signal. In order to address these issues, in this study, we propose an architecture which stacks a sub-block including two GLA-inspired fixed layers and a DNN. The number of stacked sub-blocks is adjustable, and we can trade the performance and computational load based on requirements of applications. The effectiveness of the proposed method is investigated by reconstructing phases from amplitude spectrograms of speeches.

Index Terms— Phase reconstruction, spectrogram consistency, deep neural network, residual learning.

1. INTRODUCTION

In recent years, phase reconstruction has gained much attention in the signal processing community [1, 2]. Many ordinary speech processing methods defined in the time-frequency domain have considered only amplitude spectrograms and utilized the phase of the observed signal without modifying it. Meanwhile, recent studies have proven that phase reconstruction can improve the quality of the reconstructed signal [3], and thus several methods have been proposed for that [4–8]. Phase reconstruction solely from an amplitude spectrogram has also received increasing attention along the development of the short-time Fourier transform (STFT)-based speech synthesis [9, 10] which generates an amplitude spectrogram and requires phase reconstruction for generating a time-domain signal. This paper focuses on such a situation where only an amplitude spectrogram is available for reconstructing the phase.

When only an amplitude spectrogram is available and no explicit information is given for the phase, such as in STFT-based speech synthesis, the Griffin–Lim algorithm (GLA) is one of the popular methods for phase reconstruction [11]. GLA promotes the consistency of a spectrogram by iterating two projections (see Section 2.1), where a spectrogram is said to be consistent when its inter-bin dependency owing to the redundancy of STFT is retained [12]. GLA is based only on the consistency and does not take any prior knowledge about the target signal into account. Consequently, GLA often requires many iterations and results in low-quality signals.

For incorporating prior knowledge of target signals into phase reconstruction, deep neural networks (DNNs) have been applied recently [13–16]. There exist a number of approaches to reconstruct phase using DNNs. One approach is to treat it as a classification problem by discretizing the candidates of phase [13, 14], which is



Fig. 1. A block diagram of the proposed architecture for reconstructing phase from a given amplitude spectrogram (top), which stacks a common sub-block (bottom). The sub-block consists of two fixed GLA-inspired layers (red, blue) and a trainable DNN (green).

effectively utilized in speech separation. Other approaches handle phase as a continuous periodic variable [15] or treat complex-valued spectrogram [16]. While these DNN-based phase reconstruction methods have obtained successful results, the number of layers is determined when they are trained. That is, their performance and computational load are fixed at the training. It should be beneficial if one can easily trade the performance and computational load at the time of inference depending on requirements of applications.

In this study, we propose a phase reconstruction method which incorporates a DNN into GLA. The proposed method stacks a common sub-block motivated by the iterative procedure of GLA, which constructs a deep architecture, named deep Griffin-Lim iteration (DeGLI), as illustrated in Fig. 1. In the proposed architecture, the number of total layers corresponds to the number of stacking, and its depth can be adjusted afterward based on the allowable computational load in applications. Its training procedure is also proposed to effectively train the DNN within the sub-block. Our main contributions are twofold: (1) proposing a deep architecture whose sub-block contains the fixed GLA-inspired layers which enable reduction of the amount of trainable parameters (Section 3.1); and (2) proposing its training procedure which instructs the sub-block to be a denoiser, instead of requiring it to reconstruct the phase (Section 3.2). Thanks to this training procedure, the difficulty of training a DNN in phase reconstruction arisen from the periodic nature of phase is circumvented. To evaluate the effectiveness of the proposed method, the quality of the reconstructed signal by GLA and the proposed method is compared.

2. RELATED WORKS

2.1. Griffin–Lim Algorithm (GLA)

GLA is a popular phase recovery algorithm based on the consistency of a spectrogram [11]. This algorithm expects to recover a complexvalued spectrogram, which is consistent and maintains the given amplitude A, by the following alternative projection procedure:

$$\mathbf{X}^{[m+1]} = P_{\mathcal{C}} \left(P_{\mathcal{A}}(\mathbf{X}^{[m]}) \right), \tag{1}$$

where **X** is a complex-valued spectrogram updated through the iteration, P_S is the metric projection onto a set S, and m is the iteration index. Here, C is the set of consistent spectrograms, and A is the set of spectrograms whose amplitude is the same as the given one. The metric projections onto these sets C and A are given by,

$$P_{\mathcal{C}}(\mathbf{X}) = \mathcal{G}\mathcal{G}^{\dagger}\mathbf{X},\tag{2}$$

$$P_{\mathcal{A}}(\mathbf{X}) = \mathbf{A} \odot \mathbf{X} \oslash |\mathbf{X}|, \tag{3}$$

where \mathcal{G} represents STFT, \mathcal{G}^{\dagger} is the pseudo inverse of STFT (iSTFT), \odot and \oslash are element-wise multiplication and division, respectively, and division by zero is replaced by zero. GLA is obtained as an algorithm for the following optimization problem [12]:

$$\min_{\mathbf{X}} \|\mathbf{X} - P_{\mathcal{C}}(\mathbf{X})\|_{\text{Fro}}^2 \quad \text{s.t.} \quad \mathbf{X} \in \mathcal{A},$$
(4)

where $\|\cdot\|_{\text{Fro}}$ is the Frobenius norm. This equation minimizes the energy of the inconsistent components under the constraint on amplitude which must be equal to the given one. Although GLA has been widely utilized because of its simplicity, GLA often involves many iterations until it converges to a certain spectrogram and results in low reconstruction quality. This is because the cost function in Eq. (4) only requires the consistency, and the characteristics of the target signal are not taken into account. Introducing prior knowledge of the target signal into the algorithm can improve the quality of reconstructed signals as discussed in [7, 17].

2.2. DNN-based phase reconstruction with fixed STFT layers

Recently, DNNs including fixed STFT (and iSTFT) layers were considered for treating phase information within the networks. A generative adversarial network (GAN)-based approach to reconstruct a complex-valued spectrogram solely from a given amplitude spectrogram was presented in [16]. The output of the generator (a complexvalued spectrogram) is converted back to the time domain by iSTFT layer and inputted to the discriminator, where this iSTFT layer is essential for its training as discussed in [16]. As another example, a DNN for speech separation [18] employed the multiple input spectrogram inverse (MISI) layer which consists of the pair of STFT and iSTFT as in GLA. The MISI layer is applied to the output of the DNN for speech separation to improve its performance by considering the effect of the phase reconstruction together with the separation. In addition, in [19], the time-frequency representation was also trained with the DNN for speech separation. The success of these DNNs indicates that considering STFT (and iSTFT) together with a DNN is important for treating phase.

The common strategy for these DNNs is that fixed STFT-related layers are placed after a rich DNN. Their loss functions are evaluated after going through such STFT-related layers, and their effect is propagated for updating the parameters of DNNs. Based on this observation, loss functions tied with STFT (and iSTFT) seem important in phase reconstruction because such loss functions are related to the concept of the consistency. At the same time, fixed STFT-related layers have several benefits for training. Since they do not contain trainable parameters, adding STFT-related layers does not increase the number of trainable parameters while they capture the structure of complex-valued spectrograms efficiently. Therefore, use of the STFT-related layers within DNNs may be recommended for treating phase information. However, there are little research on such DNN containing STFT within the network.

3. PROPOSED DEEP ARCHITECTURE

Based on the above discussions, we propose an architecture for phase reconstruction, named deep Griffin–Lim iteration (DeGLI), which is a unification of GLA and a DNN. As illustrated in Fig. 1, the proposed architecture consists of a common sub-block, and it is stacked to form the whole deep architecture based on the iterative procedure of GLA. The architecture of DeGLI is introduced in Section 3.1, while its training procedure is described in Section 3.2.

3.1. Deep Griffin–Lim Iteration (DeGLI)

One interesting trend of research in deep learning is to interpret an optimization algorithm as a recurrent neural network (RNN) and construct a DNN architecture following that [20–22]. The DNN introduced in the previous section [18, 19] was also obtained by a similar approach called deep unfolding [23, 24]. In this context, the iterative procedure of GLA in Eq. (1) is interpreted as an RNN which stacks the fixed linear layer $P_{\mathcal{C}}$ and target-dependent nonlinear layer $P_{\mathcal{A}}$. By looking close at Eq. (1), it can be seen that the complex-valued spectrogram at *m*th iteration $\mathbf{X}^{[m]}$ is inputted into the nonlinear layer $P_{\mathcal{A}}$, and then its output passes through the fixed linear layer $P_{\mathcal{C}}$ consisting of STFT \mathcal{G} and iSTFT \mathcal{G}^{\dagger} as in Eq. (2). That is, GLA is a parameter-fixed RNN consisting of STFT and iSTFT layers within the network. Inspired from the above observations, the proposed deep architecture for phase reconstruction, or DeGLI, is defined through a sub-block based on GLA.

Let us consider the intermediate representation of GLA,

$$\mathbf{Y}^{[m]} = P_{\mathcal{A}}(\mathbf{X}^{[m]}), \tag{5}$$

$$\mathbf{Z}^{[m]} = P_{\mathcal{C}}(\mathbf{Y}^{[m]}), \tag{6}$$

where the combination of these equations recovers Eq. (1). Since $\mathbf{Y}^{[m]}$ is the amplitude-replaced version of $\mathbf{X}^{[m]}$, their difference indicates the amount of mismatch between the amplitude of current spectrogram $|\mathbf{X}^{[m]}|$ and the desired amplitude \mathbf{A} . Similarly, since $\mathbf{Z}^{[m]}$ is the closest consistent spectrogram to $\mathbf{Y}^{[m]}$ (in the Euclidean sense), the difference between them indicates the amount of inconsistent components [12]. Such differences should be quite informative for phase reconstruction because the aim of GLA is to reduce them as much as possible. However, such intermediate information is not considered in the original GLA in Eq. (1).

To effectively use those intermediate information in a learning scheme, we propose DeGLI as the following architecture:

$$\mathbf{X}^{[m+1]} = \mathscr{B}(\mathbf{X}^{[m]}),\tag{7}$$

$$= \mathbf{Z}^{[m]} - \mathscr{F}(\mathbf{X}^{[m]}, \mathbf{Y}^{[m]}, \mathbf{Z}^{[m]}),$$
(8)

where \mathscr{B} is the proposed DeGLI-block inspired by GLA as in Fig. 1, and \mathscr{F} is a DNN. The whole architecture can also be viewed as an RNN or a feed-forward network in which the weights are shared. By stacking M DeGLI-blocks (which is equivalent to iterate Eq. (7) M times), the whole DeGLI architecture becomes M-times deeper without increasing the number of trainable parameters. That is, the total depth of the DeGLI architecture can be adjusted afterward, which enables one to easily trade its performance and computational load for adapting the allowable computational time of various applications. Note that, as a specific case, DeGLI reduces to the ordinary GLA when $\mathscr{F}(\mathbf{X}^{[m]}, \mathbf{Y}^{[m]}, \mathbf{Z}^{[m]}) = \mathbf{O}$, where \mathbf{O} is the zero matrix. A variant of GLA in [25] can also be obtained by setting $\mathscr{F}(\mathbf{X}^{[m]}, \mathbf{Y}^{[m]}, \mathbf{Z}^{[m]}) = \gamma(\mathbf{X}^{[m]} - \mathbf{Z}^{[m]})$ ($0 < \gamma < 1$), which indicates that DeGLI is a general architecture including several GLAtype algorithms as spacial cases.



Fig. 2. The block diagram for training the sub-block.

One of the key points of the DeGLI architecture is that $\mathbf{Z}^{[m]}$ (= $P_{\mathcal{C}}(P_{\mathcal{A}}(\mathbf{X}^{[m]})))$ is the output of GLA-inspired layers at *m*th iteration, and the proposed DeGLI-block \mathscr{B} is defined as the subtraction of the DNN output $\mathscr{F}(\mathbf{X}^{[m]}, \mathbf{Y}^{[m]}, \mathbf{Z}^{[m]})$ from the output of GLA-inspired layers $\mathbf{Z}^{[m]}$. Defining the DeGLI-block in this way is based on two reasons: (1) differences between each intermediate representation indicates the *residual* to the desired ones as discussed in the above paragraph; and (2) it is known that treating a residual is easier than directly estimating the target according to the literature on the residual learning [26,27].

3.2. Training procedure for DeGLI-block

Since the proposed DeGLI architecture can be interpreted as a large DNN, a simple strategy for training the DeGLI-block is directly minimizing the loss of phase reconstruction measured by the output:

$$\min D(\mathscr{G}_{\theta}(\mathbf{A}), P_{\mathcal{C}}(\mathscr{G}_{\theta}(\mathbf{A}))), \qquad (9)$$

where $\mathscr{G}(\cdot) = P_{\mathcal{A}}(\mathscr{B}(\cdots \mathscr{B}(\mathscr{B}(\cdot))))$ represents the whole DeGLI architecture, θ represents all trainable parameters in \mathscr{G} (i.e., the parameters in \mathscr{F}), $D(\cdot, \cdot)$ is a measure of mismatch such as a norm of difference, and the minimization is considered for all **A**. This problem is related to the optimization problem for GLA in Eq. (4) when D is the squared Frobenius norm of the difference of the variables (note that, since $P_{\mathcal{A}}$ is applied at the last of \mathscr{G} , the constraint in Eq. (4) is always satisfied). Although the above training strategy is straightforward, the number of the blocks should be defined in advance for applying it. In addition, it did not work well in our preliminary experiments.

In order to tackle this issue, we train the DeGLI-block \mathscr{B} to be a denoiser by the training procedure illustrated in Fig. 2. Let \mathbf{X}^* be a complex-valued spectrogram of a target signal, and $\mathbf{\tilde{X}} = \mathbf{X}^* + \mathbf{N}$ be its noisy counterpart degraded by complex-valued noise \mathbf{N} . Then, the DeGLI-block \mathscr{B} is trained so that $\mathscr{B}(\mathbf{\tilde{X}}) \approx \mathbf{X}^*$, i.e.,

$$\mathscr{B}(\widetilde{\mathbf{X}}) = \widetilde{\mathbf{Z}} - \mathscr{F}(\widetilde{\mathbf{X}}, \widetilde{\mathbf{Y}}, \widetilde{\mathbf{Z}}) \approx \mathbf{X}^{\star}, \tag{10}$$

based on the definition in Eq. (8). Since $\tilde{\mathbf{Z}}$ is obtained only from the fixed layers $P_{\mathcal{A}}$ and $P_{\mathcal{C}}$, the optimization problem for training the DNN \mathscr{F} is given by

$$\min_{\mathbf{A}} D(\widetilde{\mathbf{Z}} - \mathbf{X}^{\star}, \mathscr{F}_{\theta}(\widetilde{\mathbf{X}}, \widetilde{\mathbf{Y}}, \widetilde{\mathbf{Z}})).$$
(11)

In such denoising, the DNN \mathscr{F} estimates the residual components $\widetilde{\mathbf{Z}} - \mathbf{X}^*$ which should not be contained in the GLA output $\widetilde{\mathbf{Z}} = P_{\mathcal{C}}(P_{\mathcal{A}}(\widetilde{\mathbf{X}}))$. To be specific, the DNN takes the mismatch to the consistency and amplitude into account by inputting $\widetilde{\mathbf{Y}}$ and $\widetilde{\mathbf{Z}}$, and it implicitly eliminates the latent target signal (such as clean speech) through hidden layers in \mathscr{F} . This training strategy is closely related to the residual learning strategy. It has been shown that a denoising sub-block with the residual learning strategy is robust to the type and level of noise, and it can be applied to a variety of tasks as discussed



Fig. 3. The illustration of the DNN used in the experiment. It maps real and imaginary parts of three complex-valued spectrograms (\mathbf{X} , \mathbf{Y} , and \mathbf{Z}) to those of the residual. Here, "Conv" indicates a convolutional layer with the zero padding for keeping the input size, where k, s, and c are the kernel size, stride size, and the number of channels, respectively. "GLU" represents the gated linear unit.

in [27, 28]. The idea of applying a denoising DNN for general tasks can also be found in [29, 30].

Note that, after passing through the fixed nonlinear layer P_A , the amplitude of the complex-valued spectrogram is always replaced by the desired one. That is, the difference between $\widetilde{\mathbf{Y}}$ and the target \mathbf{X}^* is only phase, and thus denoising of $\widetilde{\mathbf{Y}} (= P_A(\widetilde{\mathbf{X}}))$ corresponds to phase reconstruction. It can be expected that the denoising sub-block including GLA-inspired layers also works well in phase reconstruction. In any case, the trained DNN \mathscr{F} (and thus \mathscr{B}) only affects the phase of the final output because the amplitude is always set to the given one by P_A after the last DeGLI block.

4. EXPERIMENT

In order to validate the effectiveness of DeGLI, the quality of reconstructed speeches was evaluated by objective measures. The proposed method was compared with GLA as a baseline method.

4.1. Experimental settings

A DNN \mathscr{F} used in the DeGLI-block \mathscr{B} for the experiment is illustrated in Fig. 3. The 2D Convolutional layers (Conv) and the gated linear units (GLU) [31] are stacked with the skip connections. In the Conv layers, the complex-valued spectrograms are treated as images, where the real and imaginary parts are concatenated along the channel direction. Note that the input of the DNN is three complex-valued spectrograms as in Figs. 1 and 2, which results in six channels as each of the three consists of the real and imaginary parts.

As the training dataset for denoising, the Wall Street Journal (WSJ-0) corpus recorded at the sampling rate of 16 kHz was utilized. 14 250 speech files were randomly selected from the database to form a training set, and the rest of the data was used as a validation set. During the mini-batch training, the utterances were divided into about 2-second-long segments (32768 samples), and the Adam optimizer was utilized as the optimization solver. The network was trained for 50 epochs with a learning rate control, where the learning rate was decayed by multiplying $10^{-0.5}$ if the loss function on the validation set did not decrease for 2 consecutive epochs, and the initial learning rate was set to 10^{-3} . As the noise utilized for training in the time-frequency domain (described in Section 3.2), the complex Gaussian noise was added so that the signal-to-noise ratio was randomly selected from -6 to 0 dB, and the measure of mismatch D as the loss function in Eq. (11) was set to the ℓ_1 -norm of difference. STFT was implemented with the Hann window, whose duration was 64 ms, with 32 ms shifting. As the test dataset, randomly selected 500 utterances from the TIMIT dataset were utilized for obtaining amplitude spectrograms for phase reconstruction, where the initial phases were set to zero in the time-frequency domain (i.e., the amplitude spectrogram was directly inputted as the initial value).



Fig. 4. An example of the spectrograms within the proposed DeGLIblock. "GLA" represents the output of the GLA-inspired layers $\mathbf{Z} = P_{\mathcal{C}}(P_{\mathcal{A}}(\mathbf{X}))$, and "Original" is the clean speech signal \mathbf{X}^{\star} to be recovered. The difference between them $\mathbf{Z} - \mathbf{X}^{\star}$ is shown as "Residual", while its estimation $\mathscr{F}(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ is denoted by "DNN output". The DNN was able to accurately estimate the Residual.

4.2. Experimental results

An example of the results of the residual learning is shown in Fig. 4 for illustrating how the DNN in the proposed DeGLI-block works. As shown in the figure, the DNN appropriately estimated the residual which is the difference between the output of the current GLA-inspired layers and the target spectrogram. Such estimated residual is subtracted so that the difference to the ideal spectrogram is reduced. We expect that the estimation by the DNN is reasonably accurate to improve the output of the DeGLI-block.

The performance of phase reconstruction was evaluated by STOI [32] and PESQ [33]. The score per iteration averaged among the test set is shown in the upper row of Fig. 5. Both STOI and PESQ of the proposed method were always higher than those of GLA at each iteration, and it improved the performance as the number of iteration increased. Since the iteration corresponds to the depth of the whole architecture of DeGLI, this result indicates that one can iterate the DeGLI-block until the quality of the reconstructed signal become satisfactory. Namely, one can eliminate unnecessary computation, or decide the depth based on the available computational resource at that time. We stress that this unique feature of the proposed method cannot be achieved by a single rich DNN directly mapping an inputted amplitude spectrogram into the final reconstructed signal.

Since the computational time per iteration is different between GLA and the proposed DeGLI-block, the performance was also investigated in terms of computational time for fair comparison. In this experiment, "Intel Core i9-7980XE (2.60 GHz)" and "NVIDIA GeForce GTX 1080 Ti" were employed for the CPU and GPU, respectively. For both methods, STFT and iSTFT were implemented by TensorFlow. The scores per computational time are illustrated in the bottom row of Fig. 5. Since the computational time per iteration of the proposed method was about 2.4 and 9.6 times slower than GLA by using GPU and CPU, respectively, the difference of the scores between the methods is closer than in the top row. Nevertheless, the proposed method notably outperformed GLA especially for PESO. To see the scores at some specific iterations, box plots of the scores are also shown in Fig. 6. The results were obtained from the 100th iteration for GLA and the 10th iteration for the proposed method because the computational times of these methods are roughly the same at those iteration numbers. It can be seen that the tendencies of the scores are the same as the averaged values in Fig. 5, and the effectiveness of the proposed DeGLI architecture was confirmed by a paired one-side *t*-test (p < 0.01).

In summary, it was confirmed that the proposed DeGLI archi-



Fig. 5. Average scores of STOI and PESQ per iteration (top) and per computational time (bottom) for GLA (blue, circles) and the proposed method (red, cross marks). The yellow dashed line indicates that the real time factor is 1. For measuring the computational time, both methods were implemented by using CPU and GPU.



Fig. 6. Box plots of the scores of STOI and PESQ, where the results of GLA were evaluated at the 100th iteration, while those of DeGLI were obtained from the output of the 10th stacks. The red lines are the median, and the boxes indicates the first and third quartiles.

tecture can be trained so that utilizing the common block for every iteration improves the performance, which should be because of training as a denoiser and the residual learning strategy. Note that the trainable DNN used in this experiment was merely an example, and it must be possible to improve the performance by considering a DNN more suitable for phase reconstruction.

5. CONCLUSION

In this study, we proposed a deep architecture, named DeGLI, which combines a DNN with the iterative procedure of GLA. The key idea was to stack the same sub-block, so that the depth of whole architecture can be adjusted without increasing the number of trainable parameters. This feature enables one to trade the quality of the reconstructed signal and computational load depending on applications. The residual learning strategy was applied to train the sub-block as a denoiser, where the DNN removes the undesired components introduced by GLA. Experimental results confirmed that a denoising sub-block is applicable to phase reconstruction, which indicates that the task of training can be different from the phase reconstruction which is not an easy task for a DNN owing to the periodic nature of phase. Investigation of a DNN suitable for the proposed DeGLI remains as a future work.

6. REFERENCES

- T. Gerkmann, M. Krawczyk-Becker, and J. Le Roux, "Phase processing for single-channel speech enhancement: History and recent advances," *IEEE Signal Process. Mag.*, vol. 32, no. 2, pp. 55–66, Mar. 2015.
- [2] P. Mowlaee, R. Saeidi, and Y. Stylianou, "Advances in phaseaware signal processing in speech communication," *Speech Commun.*, vol. 81, pp. 1–29, July 2016.
- [3] K. Paliwal, K. Wójcicki, and B. Shannon, "The importance of phase in speech enhancement," *Speech Commun.*, vol. 53, no. 4, pp. 465–494, Apr. 2011.
- [4] M. Krawczyk and T. Gerkmann, "STFT phase reconstruction in voiced speech for an improved single-channel speech enhancement," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 22, no. 12, pp. 1931–1940, Dec. 2014.
- [5] Y. Wakabayashi, T. Fukumori, M. Nakayama, T. Nishiura, and Y. Yamashita, "Single-channel speech enhancement with phase reconstruction based on phase distortion averaging," *IEEE/ACM Trans. Audio, Speech, Lang. Proc.*, vol. 26, no. 9, pp. 1559–1569, Sept. 2018.
- [6] Y. Masuyama, K. Yatabe, and Y. Oikawa, "Model-based phase recovery of spectrograms via optimization on Riemannian manifolds," in *Int. Workshop Acoust. Signal Enhance.* (*IWAENC*), Sept. 2018, pp. 126–130.
- [7] K. Yatabe, Y. Masuyama, and Y. Oikawa, "Rectified linear unit can assist Griffin–Lim phase recovery," in *Int. Workshop Acoust. Signal Enhance. (IWAENC)*, Sept. 2018, pp. 555–559.
- [8] Y. Masuyama, K. Yatabe, and Y. Oikawa, "Griffin–Lim like phase recovery via alternating direction method of multipliers," *IEEE Signal Process. Lett.*, vol. 26, no. 1, pp. 184–188, Jan. 2019.
- [9] S. Takaki, H. Kameoka, and J. Yamagishi, "Direct modeling of frequency spectra and waveform generation based on phase recovery for DNN-based speech synthesis," in *INTERSPEECH*, Aug. 2017, pp. 1128–1132.
- [10] Y. Saito, S. Takamichi, and H. Saruwatari, "Text-to-speech synthesis using STFT spectra based on low-/multi-resolution generative adversarial networks," in *IEEE Int. Conf. Acoust.*, *Speech, Signal Process. (ICASSP)*, Apr. 2018, pp. 5299–5303.
- [11] D. Griffin and J. Lim, "Signal estimation from modified shorttime Fourier transform," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 32, no. 2, pp. 236–243, Apr. 1984.
- [12] J. Le Roux, H. Kameoka, N. Ono, and S. Sagayama, "Fast signal reconstruction from magnitude STFT spectrogram based on spectrogram consistency," in *Proc. 13th Int. Conf. Digit. Audio Eff. (DAFx-10).* Sept. 2010, 397–403.
- [13] N. Takahashi, P. Agrawal, N. Goswami, and Y. Mitsufuji, "PhaseNet: Discretized phase modeling with deep neural networks for audio source separation," in *INTERSPEECH*, Sept. 2018, pp. 2713–2717.
- [14] J. Le Roux, G. Wichern, S. Watanabe, A. Sarroff, and J. R. Hershey, "Phasebook and friends: Leveraging discrete representations for source separation," *arXiv preprint arXiv:1810.01395*, 2018.
- [15] S. Takamichi, Y. Saito, N. Takamune, D. Kitamura, and H.Saruwatari, "Phase reconstruction from amplitude spectrograms based on von–Mises-distribution deep neural network," in *Int. Workshop Acoust. Signal Enhance. (IWAENC)*, Sept. 2018, pp. 286–290.

- [16] K. Oyamada, H. Kameoka, K. Tanaka T. Kaneko, N. Hojo, and H. Ando, "Generative adversarial network-based approach to signal reconstruction from magnitude spectrograms," in *Eur. Signal Process. Conf. (EUSIPCO)*, Sept. 2018.
- [17] P. Magron, J. Le Roux, and T. Virtanen, "Consistent anisotropic wiener filtering for audio source separation," in *IEEE Workshop Appl. Signal Process. Audio Acoust. (WAS-PAA)*, Oct. 2017, pp. 269–273.
- [18] Z.-Q. Wang, J. Le Roux, D. Wang, and J. R. Hershey, "Endto-end speech separation with unfolded iterative phase reconstruction," in *INTERSPEECH*, Sept. 2018, pp. 2708–2712.
- [19] G. Wichern and J. Le Roux, "Phase reconstruction with learned time-frequency representations for single-channel speech separation," in *Int. Workshop Acoust. Signal Enhance. (IWAENC)*, Sept. 2018, pp. 396–400.
- [20] K. Greger and Y. LeCun, "Learning fast approximations of sparse coding," in *Int. Conf. March. learn.*, 2010, pp. 399–406.
- [21] M. Borgerding, P. Schniter, and S. Rangan, "AMP-inspired deep networks for sparse linear inverse problems," *IEEE Trans. Signal Process.*, vol. 65, no. 16, pp. 4294–4308, Aug. 2017.
- [22] Y. Yang, J. Sun, H. Li, and Z. Xu, "Deep ADMM-Net for compressive sensing MRI," in *Adv. Neural Process. Syst. (NIPS)*, pp. 10–18. Curran Associates, Inc., 2016.
- [23] J. R. Hershey, J. Le Roux, and F. Weninger, "Deep unfolding: Model-based inspiration of novel deep architectures," *arXiv* preprint arXiv:1409.2574, 2014.
- [24] S. Wisdom, J. Hershey, and and S. Watanabe J. Le Roux, "Deep unfolding for multichannel source separation," in *IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP)*, Mar. 2016, pp. 121–125.
- [25] N. Perraudin, P. Balazs, and P. L. Søndergaard, "A fast Griffin– Lim algorithm," in *IEEE Workshop Appl. Signal Process. Audio Acoust. (WASPAA)*, Oct. 2013, pp. 1–4.
- [26] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE Conf. Comp. Vision Pattern Recognit. (CVPR)*, June 2016.
- [27] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, "Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising," *IEEE Trans. Image Process.*, vol. 26, no. 7, pp. 3142–3155, July 2017.
- [28] K. Zhang, W. Zuo, S. Gu, and L. Zhang, "Learning deep CNN denoiser prior for image restoration," in *IEEE Conf. Comp. Vision Pattern Recognit. (CVPR)*, July 2017, pp. 2808–2817.
- [29] S. V. Venkatakrishnan, C. A. Bouman, and B. Wohlberg, "Plug-and-Play priors for model based reconstruction," in *IEEE Glob. Conf. Signal, Inf. Process.*, Dec. 2013, pp. 945– 948.
- [30] Y. Romano, M. Elad, and P. Milanfar, "The little engine that could: Regularization by denoising (RED)," *SIAM J. Imaging Sciences*, vol. 10, no. 4, pp. 1804–1844, 2017.
- [31] Y. N. Dauphin, A. Fan, M. Auli, and D. Grangier, "Language modeling with gated convolutional networks," *arXiv preprint* arXiv:1612.08083, 2016.
- [32] 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. 2155–2136, 2011.
- [33] A. W. Rix, J. G. Beerends, M. P. Hollier, and A. P. Hekstra, "Perceptual evaluation of speech quality (PESQ) — A new method for speech quality assessment of telephone networks and codecs," in *IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP)*, May 2001, vol. 19, pp. 2125–2136.