



HAL
open science

Learning the Proximity Operator in Unfolded ADMM for Phase Retrieval

Pierre-Hugo Vial, Paul Magron, Thomas Oberlin, Cédric Févotte

► **To cite this version:**

Pierre-Hugo Vial, Paul Magron, Thomas Oberlin, Cédric Févotte. Learning the Proximity Operator in Unfolded ADMM for Phase Retrieval. *IEEE Signal Processing Letters*, 2022, 29, pp.1619-1623. 10.1109/LSP.2022.3189275 . hal-03790178

HAL Id: hal-03790178

<https://hal.science/hal-03790178>

Submitted on 28 Sep 2022

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

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

Learning the Proximity Operator in Unfolded ADMM for Phase Retrieval*

Pierre-Hugo Vial[†], Paul Magron[‡], Thomas Oberlin[§], Cédric Févotte[†]

Abstract

This paper considers the phase retrieval (PR) problem, which aims to reconstruct a signal from phaseless measurements such as magnitude or power spectrograms. PR is generally handled as a minimization problem involving a quadratic loss. Recent works have considered alternative discrepancy measures, such as the Bregman divergences, but it is still challenging to tailor the optimal loss for a given setting. In this paper we propose a novel strategy to automatically learn the optimal metric for PR. We unfold a recently introduced ADMM algorithm into a neural network, and we emphasize that the information about the loss used to formulate the PR problem is conveyed by the proximity operator involved in the ADMM updates. Therefore, we replace this proximity operator with trainable activation functions: learning these in a supervised setting is then equivalent to learning an optimal metric for PR. Experiments conducted with speech signals show that our approach outperforms the baseline ADMM, using a light and interpretable neural architecture.

Keywords— Phase retrieval, audio, proximity operator learning, deep unfolding, ADMM.

1 Introduction

Phase retrieval (PR) consists in reconstructing data from phaseless nonnegative measurements. This problem finds practical applications in various areas such as optical imaging [1, 2] and audio signal processing [3, 4], which is the domain of interest of this paper. Traditionally, PR is formulated as the following optimization problem:

$$\min_{\mathbf{x} \in \mathbb{R}^L} \|\mathbf{A}\mathbf{x}\|^d - \mathbf{r}\|^2, \quad (1)$$

where $\mathbf{A} \in \mathbb{C}^{K \times L}$ is the measurement operator, $\mathbf{r} \in \mathbb{R}_+^K$ are the phaseless measurements, and $\|\cdot\|$ denotes the Euclidean norm. In audio signal processing, \mathbf{A} is often the short-time Fourier transform (STFT), and the measurements are either magnitude ($d = 1$) or power ($d = 2$) spectrograms. In the seminal work [5], the authors consider STFT magnitude measurements and propose an iterative procedure, known as the Griffin-Lim algorithm (GLA) which is proved to converge to a critical point of (1). Other optimization algorithms were proposed to tackle Problem (1), such as majorization-minimization [6], gradient descent [7] or alternating direction method of multipliers (ADMM) [8]. PR has also been addressed by replacing the quadratic loss in (1) with alternative discrepancy measures such as Bregman divergences [9], which have been shown to lead to superior results in many audio applications [10, 11]. However, prescribing a loss that is optimal for all signal processing problems and classes of audio signals remains challenging [9].

On the other hand, recent PR approaches have leveraged deep neural networks (DNNs) [12, 13, 14, 15]. Despite their successful performances in a large number of tasks, the enthusiasm for DNNs can be tempered by a general lack of explainability due to their black box structure, and by their limited ability to generalize to unseen data or experimental conditions. Deep unfolding (or unrolling) [16, 17] is a promising attempt to alleviate these limitations with model-based architectures derived from iterative algorithms. This strategy consists in considering each iteration of an optimization algorithm as a (trainable) layer of a DNN. It has proved very promising when applied to many algorithms in signal processing [17, 18, 19], including audio PR [20, 21, 22]. However, these approaches treat the linear layers as trainable parameters while activation functions remain fixed, although recent studies show that trainable activation functions can upgrade performance [23, 24, 25]. Besides, they do not address the problem of learning the metric involved in the PR formulation problem.

In this paper, we propose to unfold the ADMM algorithm for PR proposed in [9]. Our method builds upon observing that the choice of the discrepancy measure only affects the computation of a proximity operator in the

*This work is supported by the European Research Council (ERC FACTORY-CoG-6681839) and by ANITI under grant agreement ANR-19-PI3A-0004.

[†]IRIT, Université de Toulouse, CNRS, Toulouse, France (e-mail: firstname.lastname@irit.fr).

[‡]Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France (e-mail: firstname.lastname@inria.fr).

[§]ISAE-SUPAERO, Université de Toulouse, France (e-mail: firstname.lastname@isae-supaero.fr).

ADMM updates. Therefore, we can recast the problem of metric learning as a problem of proximity operator learning in the unfolded ADMM. To that end, we replace this proximity operator with a trainable activation function. We show that the proposed parametrization of the network is connected to the metric involved in the original optimization problem, which yields an interpretable architecture. Experiments performed on speech signals demonstrate the efficiency of our method, which outperforms a baseline ADMM [8] with a number of iterations equal to the number of layers in the unfolded ADMM.

The rest of this paper is structured as follows. Section 2 presents the related work. Then, the proposed method is introduced in Section 3 and tested experimentally in Section 4. Finally, Section 5 draws some concluding remarks.

Mathematical notation:

- \mathbf{A} (capital, bold font): matrix.
- \mathbf{x} (lower case, bold font): vector.
- z (regular): scalar.
- $|\cdot|, \angle(\cdot)$: magnitude and complex angle, respectively.
- \mathbf{x}^H : Hermitian transpose.
- $\odot, (\cdot)^d$, fraction bar: element-wise matrix or vector multiplication, power, and division, respectively.

2 Related work

2.1 PR with Bregman divergences

In [9], we proposed to reformulate the PR problem by substituting the quadratic loss in (1) with a Bregman divergence. Bregman divergences encompass beta-divergences [26], which include the Kullback-Leibler and Itakura-Saito divergences, as well as the quadratic loss. They are defined as follows:

$$\mathcal{D}_\psi(\mathbf{p} | \mathbf{q}) = \sum_{k=1}^K [\psi(p_k) - \psi(q_k) - \psi'(q_k)(p_k - q_k)], \quad (2)$$

where $\mathbf{p}, \mathbf{q} \in \mathbb{R}^K$ and $\psi : \mathbb{R} \rightarrow \mathbb{R}$ is a strictly-convex, continuously-differentiable generating function (with derivative ψ'). Since Bregman divergences are not symmetric in general (i.e., $\mathcal{D}_\psi(\mathbf{p} | \mathbf{q}) \neq \mathcal{D}_\psi(\mathbf{q} | \mathbf{p})$), we considered the two following problems, respectively termed “right” and “left”:

$$\min_{\mathbf{x} \in \mathbb{R}^L} \mathcal{D}_\psi(\mathbf{r} | |\mathbf{A}\mathbf{x}|^d) \text{ and } \min_{\mathbf{x} \in \mathbb{R}^L} \mathcal{D}_\psi(|\mathbf{A}\mathbf{x}|^d | \mathbf{r}). \quad (3)$$

Two algorithms were derived for solving these problems, based on gradient descent and ADMM [9].

2.2 PR via ADMM

In this work we focus on the “left” PR problem using ADMM, since it produced the best reconstruction results in [9]. Drawing on [8], this approach consists in reformulating (3) into the following constrained optimization problem:

$$\min_{\mathbf{x} \in \mathbb{R}^L, \mathbf{u} \in \mathbb{R}_+^K, \theta \in [0; 2\pi[K} \mathcal{D}_\psi(\mathbf{u} | \mathbf{r}) \quad \text{s.t.} \quad (\mathbf{A}\mathbf{x})^d = \mathbf{u} \odot e^{i\theta}, \quad (4)$$

where \mathbf{u} and θ are auxiliary variables for the magnitude and phase of $(\mathbf{A}\mathbf{x})^d$. Then, the augmented Lagrangian is:

$$\begin{aligned} \mathcal{L}(\mathbf{x}, \mathbf{u}, \theta, \boldsymbol{\lambda}) &= \mathcal{D}_\psi(\mathbf{u} | \mathbf{r}) + \Re \left(\boldsymbol{\lambda}^H ((\mathbf{A}\mathbf{x})^d - \mathbf{u} \odot e^{i\theta}) \right) \\ &\quad + \frac{\rho}{2} \left\| (\mathbf{A}\mathbf{x})^d - \mathbf{u} \odot e^{i\theta} \right\|^2, \end{aligned} \quad (5)$$

where \Re is the real part function, $\boldsymbol{\lambda} \in \mathbb{C}^K$ are the Lagrange multipliers and $\rho > 0$ is the augmentation parameter. Alternate minimization of \mathcal{L} leads to the following update rules [9]:

$$\mathbf{h}_{t+1} = (\mathbf{A}\mathbf{x}_t)^d + \frac{\boldsymbol{\lambda}_t}{\rho} \quad (6)$$

$$\mathbf{u}_{t+1} = \text{prox}_{\rho^{-1}\mathcal{D}_\psi(\cdot | \mathbf{r})}(|\mathbf{h}_{t+1}|) \quad (7)$$

$$\boldsymbol{\theta}_{t+1} = \angle \mathbf{h}_{t+1} \quad (8)$$

$$\mathbf{x}_{t+1} = \mathbf{A}^H \left(\mathbf{u}_{t+1} \odot e^{i\boldsymbol{\theta}_{t+1}} - \frac{\boldsymbol{\lambda}_t}{\rho} \right)^{1/d} \quad (9)$$

$$\boldsymbol{\lambda}_{t+1} = \boldsymbol{\lambda}_t + \rho(\mathbf{A}\mathbf{x}_{t+1} - \mathbf{u}_{t+1} \odot e^{i\boldsymbol{\theta}_{t+1}}), \quad (10)$$

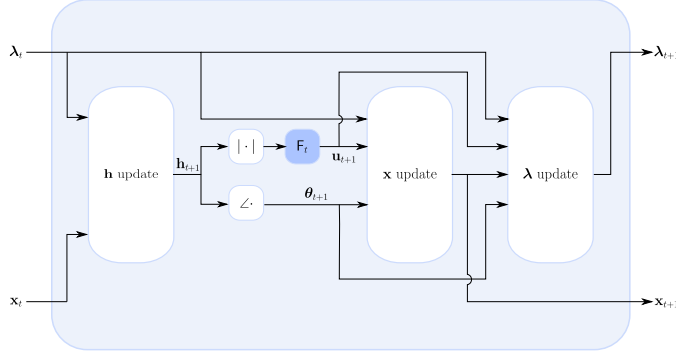


Figure 1: One layer of the proposed unfolded architecture.

where prox_f denotes the proximity operator of a function f , defined as the mapping of a vector \mathbf{y} to the solution of the minimization problem:

$$\text{prox}_{\rho^{-1}f}(\mathbf{y}) := \arg \min_{\mathbf{z} \in \mathbb{R}^K} f(\mathbf{z}) + \frac{\rho}{2} \|\mathbf{z} - \mathbf{y}\|^2. \quad (11)$$

However, a closed-form expression of this operator is not available for every Bregman divergence. To alleviate this issue, we propose in this paper to replace this proximity operator with a learnable neural activation function, as detailed in the next section.

3 Proposed method

3.1 General architecture

The ADMM updates detailed in 2.2 consist in successive linear and nonlinear computations. As such, this algorithm can be viewed as a neural network \mathbf{U} via unfolding:

$$(\mathbf{x}_T, \boldsymbol{\lambda}_T) = \mathbf{U}(\mathbf{x}_0, \boldsymbol{\lambda}_0) = \mathbf{U}_1 \circ \dots \circ \mathbf{U}_T(\mathbf{x}_0, \boldsymbol{\lambda}_0), \quad (12)$$

where \mathbf{U}_t denotes the t -th layer of the network, mimicking the t -th iteration of the ADMM algorithm, as illustrated in Fig. 1. The layer \mathbf{U}_t can be decomposed into two linear parts denoted by $\mathbf{L}_t^{(1)}$ and $\mathbf{L}_t^{(2)}$, and a nonlinear part \mathbf{NL}_t as follows:

$$\mathbf{L}_t^{(1)} : (\mathbf{x}_{t-1}, \boldsymbol{\lambda}_{t-1}) \mapsto \mathbf{h}_t \quad (13)$$

$$\mathbf{NL}_t : \mathbf{h}_t \mapsto (\mathbf{u}_t, \boldsymbol{\theta}_t) = (\mathbf{F}_t(|\mathbf{h}_t|, \mathbf{r}), \angle \mathbf{h}_t) \quad (14)$$

$$\mathbf{L}_t^{(2)} : (\mathbf{x}_{t-1}, \boldsymbol{\lambda}_{t-1}, \mathbf{u}_t, \boldsymbol{\theta}_t) \mapsto (\mathbf{x}_t, \boldsymbol{\lambda}_t), \quad (15)$$

with $\mathbf{h}_t, \mathbf{x}_t, \boldsymbol{\lambda}_t$ respectively defined as in (6), (9) and (10). \mathbf{F}_t denotes a parameterized sublayer modeling the proximity operator of equation (7). Since the choice of the discrepancy measure \mathcal{D}_ψ only affects the proximity operator (7) in the updates, we can recast the problem of metric learning as the problem of proximity operator learning. We propose to leverage a trainable activation function in order to model this layer and learn the proximity operator.

3.2 Proposed parameterization

To build the non-linear sublayers \mathbf{F}_t that model $\text{prox}_{\rho^{-1}\mathcal{D}_\psi(\cdot|\mathbf{r})}$, we first reformulate this operator as follows. Let $\mathbf{v} \in \mathbb{R}^K$ and $f(\mathbf{z}) = \sum_{k=1}^K [\psi(z_k) + v_k z_k]$. We have [27]:

$$\text{prox}_{\rho^{-1}f}(\mathbf{y}) = \text{prox}_{\rho^{-1}\tilde{\psi}}(\mathbf{y} - \rho^{-1}\mathbf{v}), \quad (16)$$

where $\tilde{\psi}(\mathbf{z}) = \sum_k \psi(z_k)$. Setting $\mathbf{v} = -\psi'(\mathbf{r})$ in (16), with ψ' applied entrywise, it is straightforward to see that:

$$\text{prox}_{\rho^{-1}\mathcal{D}_\psi(\cdot|\mathbf{r})}(\mathbf{y}) = \text{prox}_{\rho^{-1}\tilde{\psi}}(\mathbf{y} + \rho^{-1}\psi'(\mathbf{r})). \quad (17)$$

This formulation of the proximity operator is more convenient than (7) since the measurements \mathbf{r} no longer appear in the input function of the proximity operator, but instead in the argument of the latter (with \mathbf{y}). This leads to a more natural parametrization for unrolling.

Let us first derive the proximity operator (17) in a simple scenario, namely the quadratic loss ($\tilde{\psi} = \frac{1}{2} \|\cdot\|^2$, $\psi'(\mathbf{r}) = \mathbf{r}$). In this case we have [8]:

$$\text{prox}_{\rho^{-1}\frac{1}{2}\|\cdot - \mathbf{r}\|^2}(\mathbf{y}) = \frac{\mathbf{y} + \rho^{-1}\mathbf{r}}{1 + \rho^{-1}}. \quad (18)$$

As a result, a first simple approach for proximity operator learning would consist in treating ρ as a learnable parameter. However, early experiments have shown poor performance with this approach, which is due to the very low expressive power of such a model (only one scalar value). More generally, one can consider a beta-divergence with shape parameter β , for which $\psi'(\mathbf{r}) = \frac{\mathbf{r}^{\beta-1}}{\beta-1}$ [26]. However, the proximity operator of $\tilde{\psi}$ is not available for every beta-divergence.

To alleviate this issue, we model this unavailable proximity operator using Adaptive Piecewise Linear (APL) activations [28]. They are defined by:

$$\text{APL}(\mathbf{y}) := \max(\mathbf{y}, 0) + \sum_{c=1}^C w_c \max(-\mathbf{y} + b_c, 0), \quad (19)$$

where w_c and b_c are learnable parameters controlling the slopes and biases of the linear segments, and the max is applied entry-wise. Then, we propose the following parametrization of the nonlinear layer \mathbf{F}_t :

$$\mathbf{F}_t(\mathbf{y}, \mathbf{r}) = \text{APL}_t \left(\gamma_t^{(1)} \mathbf{y} + \gamma_t^{(2)} \frac{\mathbf{r}^{\beta_t-1}}{\beta_t-1} \right), \quad (20)$$

with learnable parameters $w_{c,t}, b_{c,t}, \gamma_t^{(1)}, \gamma_t^{(2)}$, and β_t . Even though *ad hoc*, this parametrization is motivated by the following considerations:

- APL can represent any continuous piecewise linear function over a subset of real numbers. As such, it generalizes the proximity operator obtained in the quadratic case [8].
- The term in the form of $\frac{\mathbf{r}^{\beta-1}}{\beta-1}$ in (20) is reminiscent of ψ' for beta-divergences, as mentioned above.
- Introducing learnable weights $\gamma_t^{(1)}$ and $\gamma_t^{(2)}$ allows to increase the model capacity, as it was shown beneficial in our preliminary experiments.

Note that when $w_c = 0$, $\gamma_t^{(1)} = \frac{1}{1+\rho-1}$, $\gamma_t^{(2)} = \frac{\rho^{-1}}{1+\rho-1}$ and $\beta_t = 2$, \mathbf{F}_t is equal to the proximity operator for the quadratic loss (18). Overall, our parametrization (20) is a good trade-off between tractability, interpretability, and expressiveness.

Two variants of the proposed architecture will be considered in our experiments. In the “untied” variant, each layer uses different parameters, and the global set of parameters is $\left\{ \{w_{c,t}, b_{c,t}\}_{c=1}^C, \gamma_t^{(1)}, \gamma_t^{(2)}, \beta_t \right\}_{t=1}^T$, while in the “tied” variant, the parameters are shared among layers, i.e., constant with t .

In the end, after learning these parameters (see Section 4.1.2), the proposed method, termed unfolded ADMM (UADMM), estimates a signal \mathbf{x}_T via $(\mathbf{x}_T, \boldsymbol{\lambda}_T) = \mathbf{U}(\mathbf{x}_0, \boldsymbol{\lambda}_0)$, where \mathbf{x}_0 is some initial estimate.

3.3 Discussion about interpretability

Under mild assumptions and as detailed in the supplementary material,¹ we can prove that there exists a close-form function $f_{\mathbf{r},t} : \mathbb{R}^K \rightarrow \mathbb{R} \cup \{+\infty\}$ such that $\mathbf{F}_t(\mathbf{y}, \mathbf{r}) = \text{prox}_{f_{\mathbf{r},t}}(\mathbf{y})$, (see (32) in the supplementary material). In the “tied” variant, where $f_{\mathbf{r},t} = f_{\mathbf{r}}$, reconstructing $f_{\mathbf{r}}$ from the learned parameters is analogous to identifying the metric $\mathcal{D}_\psi(\cdot | \mathbf{r})$ involved in the PR optimization problem. With the relaxation proposed in the “untied” case, this interpretation is more limited as $f_{\mathbf{r},t}$ is different in each layer of the network.

Note that when replacing (17) with (20), we have disentangled the proximity operator of $\tilde{\psi}$ and the derivative ψ' , in addition to introducing weights $\gamma_t^{(1)}$ and $\gamma_t^{(2)}$. As a result, the function $f_{\mathbf{r}}$ is no longer guaranteed to be a Bregman divergence, strictly speaking. Nevertheless, we can still interpret it as a measure of discrepancy between \mathbf{y} and \mathbf{r} .

4 Experiments

In this section, we assess the potential of UADMM for PR of speech signals. Our code is implemented using the PyTorch framework [29] and is available online for reproducibility.²

4.1 Protocol

4.1.1 Data

We build a set of speech signals from the TIMIT dataset [30]. The dataset is split into training, validation, and test subsets containing 1000, 10, and 50 utterances, respectively (note that we did not observe a significant performance improvement when using a larger training set). The signals are mono, sampled at 16 kHz and cropped to 2 seconds. The STFT is computed with a 1024 samples-long (46 ms) self-dual sine window [9] and 50% overlap. STFT magnitudes ($d = 1$) are considered as nonnegative observations \mathbf{r} .

¹In particular, the APL unit must be strictly increasing, something we ensure by imposing the weights to be negative, using $w_{c,t} = -\tilde{w}_{c,t}^2$.

²<https://github.com/phvial/LearningProxPR>

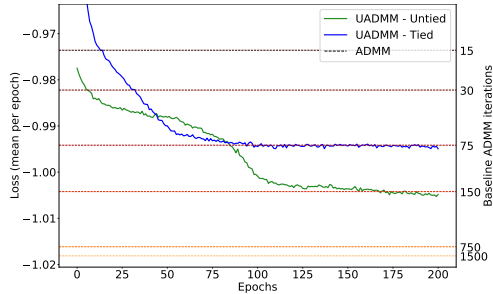


Figure 2: Training loss (negative STOI) over epochs. Note that `pytorch-stoi` implementation does not exactly replicate the original metric and consequently yields values lower than -1 .

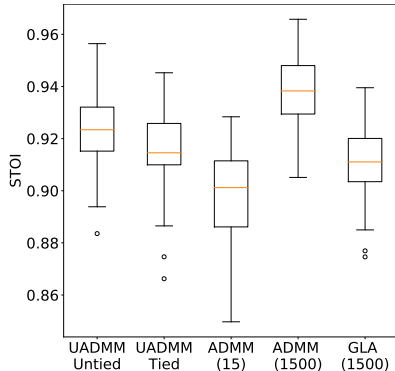


Figure 3: Performance on the test set. Each box-plot is made up of a central line indicating the median, box edges indicating the 1st and 3rd quartiles, whiskers indicating the extremal values, and circles representing the outliers.

4.1.2 Training

The network is trained with the ADAM algorithm [31] using a learning rate of 10^{-4} . We use a structure with $T = 15$ layers and $C = 3$, as these values have shown to be a good trade-off between performance and number of parameters in preliminary experiments. Batches of 10 signals with a maximum of 200 epochs are used for training. Training is stopped when the loss function on the validation subset starts increasing. Given that we consider speech signals, we train the network by minimizing the negative short-term objective intelligibility measure (STOI) [32] between the estimated and ground truth signals. Indeed, this strategy was shown to be efficient for speech enhancement applications [33, 34]. The negative STOI metric used for training the network is implemented in PyTorch via the `pytorch-stoi` library [35].

4.1.3 Methods

As baselines, we consider GLA [5] (run for 1500 iterations), and ADMM using a quadratic loss and $\rho = 10^{-3}$, since this setup has exhibited good performance in our previous study under similar conditions [9]. ADMM is run for a variable number of iterations, with a maximum of 1500 (performance does not further improve beyond). For fairness, the linear parts of UADMM use the same value for ρ , and it is initialized such that F_t replicates the quadratic proximity operator (*cf.* Section 3.2). All methods use the same initial signal estimate \mathbf{x}_0 computed using the ground truth magnitudes \mathbf{r} and a random uniform phase, and $\lambda_0 = 0$.

4.1.4 Evaluation

Reconstruction performance is assessed with the STOI metric (which ranges between 0 and 1, higher is better) computed on the test set with the `pystoi` library [36].

4.2 Results

First, we display the training loss over epochs in Fig. 2. Both UADMM variants outperform the baseline ADMM with 15 iterations on the training set. UADMM-untied reaches a lower loss value than its tied counterpart, which was expected since this variant contains more learnable parameters. They reach a performance comparable to that of ADMM using 150 and 75 iterations, respectively. The results on the test set presented in Fig. 3 confirm that the proposed UADMM approach significantly outperforms the classical ADMM using the same number of

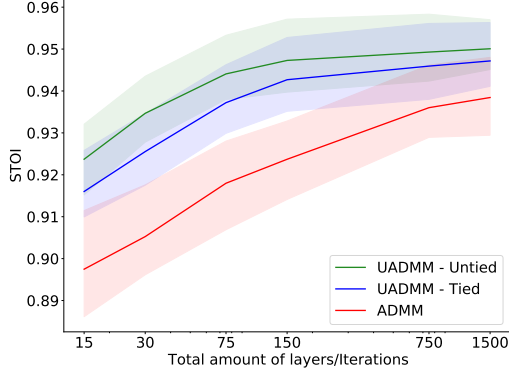


Figure 4: Evaluation with STOI over test dataset with iterated model. The solid lines denote the mean STOI and the light colored areas the values between the first and the third quartile.

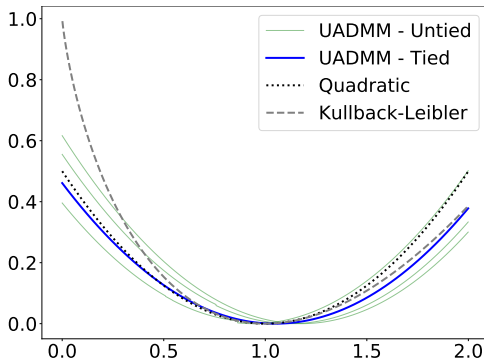


Figure 5: Learned metrics $f_{r,t}(y)$ with $r = 1$. The quadratic loss and Kullback-Leibler divergence $\mathcal{D}_{KL}(y|r)$ are also displayed for the sake of comparison. In the “tied” case, f_r in analogous to $\mathcal{D}_\psi(\cdot|r)$ involved in the PR optimization problem. For clarity, only 3 of the 15 trained layers $f_{r,t}$ are displayed for the “untied” case.

iterations, as well as the GLA baseline. A fully-converged ADMM algorithm (using 1500 iterations) exhibits a higher STOI than our 15 layers-based approach. Nonetheless, a more fair comparison would involve that both approaches use the same total number of iterations/layers.

To that end, we consider an *ad hoc* extension of our method, where we duplicate the 15-layer trained UADMM network in order to increase the total amount of layers without additional training. The results presented in Fig. 4 show that this method consistently and significantly outperforms ADMM for any number of iterations. In particular, the performance of the fully-converged ADMM (after 1500 iterations) is reached at only 30 “iterations” for UADMM-untied (i.e., twice the number of trained layers), which exhibits the computational advantage of the proposed approach.

Finally, let us point out that UADMM-tied with T layers is equivalent to applying T iterations of a standard ADMM algorithm using a learned metric f_r (note however that it differs from the ADMM baseline used in these experiments, which uses a quadratic loss). Following the derivations in the Supplementary, we compute these metrics (f_r in the tied case and $f_{r,t}$ in the untied case) from the trained activation functions, and display them in Fig. 5. These resemble beta-divergences with $\beta \in [1.5, 2.5]$. This is consistent with previous results from the literature [37], where this range of values has shown good performance for audio spectral decomposition.

5 Conclusion

In this paper, we have addressed the problem of metric learning for phase retrieval by unfolding the recently proposed ADMM algorithm [9] into a neural network. We proposed to replace the proximity operator involved in this algorithm with learnable activation functions, since this operator conveys the information about the discrepancy measure used in formulating the PR problem. Experiments conducted on speech signals show that this approach outperforms the ADMM algorithm while keeping a light and interpretable structure.

References

- [1] A. Walther, “The question of phase retrieval in optics,” *Optica Acta: International Journal of Optics*, vol. 10, no. 1, pp. 41–49, 1963.
- [2] R. W. Harrison, “Phase problem in crystallography,” *Journal of the Optical Society of America A*, vol. 10, no. 5, pp. 1046–1055, 1993.
- [3] T. Gerkmann, M. Krawczyk-Becker, and J. Le Roux, “Phase processing for single-channel speech enhancement: History and recent advances,” *IEEE Signal Processing Magazine*, vol. 32, no. 2, pp. 55–66, March 2015.
- [4] P. Mowlae, R. Saeidi, and Y. Stylianou, “Advances in phase-aware signal processing in speech communication,” *Speech Communication*, vol. 81, pp. 1–29, July 2016.
- [5] D. Griffin and J. Lim, “Signal estimation from modified short-time Fourier transform,” *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 32, no. 2, pp. 236–243, April 1984.
- [6] T. Qiu, P. Babu, and D. P. Palomar, “PRIME: Phase retrieval via majorization-minimization,” *IEEE Transactions on Signal Processing*, vol. 64, no. 19, pp. 5174–5186, October 2016.
- [7] E. J Candès, X. Li, and M. Soltanolkotabi, “Phase retrieval via Wirtinger flow: Theory and algorithms,” *IEEE Transactions on Information Theory*, vol. 61, no. 4, pp. 1985–2007, April 2015.
- [8] J. Liang, P. Stoica, Y. Jing, and J. Li, “Phase retrieval via the alternating direction method of multipliers,” *IEEE Signal Processing Letters*, vol. 25, no. 1, pp. 5–9, January 2018.
- [9] P.-H. Vial, P. Magron, T. Oberlin, and C. Févotte, “Phase retrieval with Bregman divergences and application to audio signal recovery,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 15, no. 1, pp. 51–64, January 2021.
- [10] C. Févotte, N. Bertin, and J.-L. Durrieu, “Nonnegative matrix factorization with the Itakura-Saito divergence: With application to music analysis,” *Neural computation*, vol. 21, no. 3, pp. 793–830, March 2009.
- [11] J. Le Roux, H. Kameoka, N. Ono, A. de Cheveigné, and S. Sagayama, “Computational auditory induction as a missing-data model-fitting problem with Bregman divergence,” *Speech Communication*, vol. 53, no. 5, pp. 658 – 676, May-June 2011.
- [12] 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 *Proc. International Workshop on Acoustic Signal Enhancement (IWAENC)*, September 2018, pp. 286 – 290.
- [13] L. Thieling, D. Wilhelm, and P. Jax, “Recurrent phase reconstruction using estimated phase derivatives from deep neural networks,” in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, June 2021, pp. 7088–7092.
- [14] N. B. Thien, Y. Wakabayashi, K. Iwai, and T. Nishiura, “Two-stage phase reconstruction using DNN and von Mises distribution-based maximum likelihood,” in *Proc. Annual Summit and Conference of the Asia-Pacific Signal and Information Processing Association (APSIPA ASC)*, December 2021, pp. 995–999.
- [15] S. Ö. Arik, H. Jun, and G. Diamos, “Fast spectrogram inversion using multi-head convolutional neural networks,” *IEEE Signal Processing Letters*, vol. 26, no. 1, pp. 94–98, January 2018.
- [16] K. Gregor and Y. LeCun, “Learning fast approximations of sparse coding,” in *Proc. International Conference on Machine Learning (ICML)*, June 2010, pp. 399–406.
- [17] J. R. Hershey, J. Le Roux, and F. Weninger, “Deep unfolding: Model-based inspiration of novel deep architectures,” *arXiv preprint arXiv:1409.2574*, 2014.
- [18] C. Bertocchi, E. Chouzenoux, M.-C. Corbineau, J.-C. Pesquet, and M. Prato, “Deep unfolding of a proximal interior point method for image restoration,” *Inverse Problems*, vol. 36, no. 3, pp. 034005, February 2020.
- [19] V. Monga, Y. Li, and Y. C. Eldar, “Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing,” *IEEE Signal Processing Magazine*, vol. 38, no. 2, pp. 18–44, March 2021.
- [20] Y. Masuyama, K. Yatabe, Y. Koizumi, Y. Oikawa, and N. Harada, “Deep Griffin–Lim iteration,” in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2019, pp. 61–65.
- [21] Z.-Q. Wang, J. Le Roux, D. Wang, and J. R. Hershey, “End-to-end speech separation with unfolded iterative phase reconstruction,” in *Proc. Interspeech*, September 2018.
- [22] G. Wichern and J. Le Roux, “Phase reconstruction with learned time-frequency representations for single-channel speech separation,” in *Proc. International Workshop on Acoustic Signal Enhancement (IWAENC)*, September 2018, pp. 396–400.
- [23] S. Qian, H. Liu, C. Liu, S. Wu, and H. San Wong, “Adaptive activation functions in convolutional neural networks,” *Neurocomputing*, vol. 272, pp. 204–212, January 2018.
- [24] A. Apicella, F. Isgrò, and R. Prevede, “A simple and efficient architecture for trainable activation functions,” *Neurocomputing*, vol. 370, pp. 1–15, December 2019.

- [25] S. Scardapane, S. Van Vaerenbergh, S. Totaro, and A. Uncini, “Kafnets: Kernel-based non-parametric activation functions for neural networks,” *Neural Networks*, vol. 110, pp. 19–32, February 2019.
- [26] R. Hennequin, B. David, and R. Badeau, “Beta-divergence as a subclass of Bregman divergence,” *IEEE Signal Processing Letters*, vol. 18, no. 2, pp. 83–86, February 2011.
- [27] P. L. Combettes and V. R. Wajs, “Signal recovery by proximal forward-backward splitting,” *Multiscale Modeling & Simulation*, vol. 4, no. 4, pp. 1168–1200, November 2005.
- [28] F. Agostinelli, M. Hoffman, P. Sadowski, and P. Baldi, “Learning activation functions to improve deep neural networks,” in *Proc. International Conference on Learning Representations (ICLR) workshop*, May 2015.
- [29] “PyTorch,” <https://pytorch.org/>.
- [30] J. S. Garofolo, L. F. Lamel, W. M. Fisher, J. G. Fiscus, D. S. Pallett, N. L. Dahlgren, and V. Zue, “TIMIT acoustic-phonetic continuous speech corpus,” *Linguistic data consortium*, 1993.
- [31] D. P. Kingma and J. L. Ba, “Adam: a method for stochastic optimization,” in *Proc. International Conference on Learning Representations (ICLR)*, May 2015.
- [32] C. H. Taal, R. C. Hendriks, R. Heusdens, and J. Jensen, “An algorithm for intelligibility prediction of time–frequency weighted noisy speech,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 7, pp. 2125–2136, February 2011.
- [33] Y. Zhao, B. Xu, R. Giri, and T. Zhang, “Perceptually guided speech enhancement using deep neural networks,” in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, April 2018, pp. 5074–5078.
- [34] G. Naithani, J. Nikunen, L. Bramslo, and T. Virtanen, “Deep neural network based speech separation optimizing an objective estimator of intelligibility for low latency applications,” in *Proc. International Workshop on Acoustic Signal Enhancement (IWAENC)*, September 2018, pp. 386–390.
- [35] M. Pariente, “PyTorch implementation of STOI,” https://github.com/mpariente/pytorch_stoi, 2021, [Online; accessed 13-February-2022].
- [36] M. Pariente, “PySTOI,” <https://github.com/mpariente/pystoi>, 2018.
- [37] D. Fitzgerald, M. Cranitch, and E. Coyle, “On the use of the beta divergence for musical source separation,” in *Proc. of IET Irish Signals and Systems Conference (ISSC)*, June 2008, pp. 1–6.
- [38] J.-J. Moreau, “Proximité et dualité dans un espace hilbertien,” *Bulletin de la Société mathématique de France*, vol. 93, pp. 273–299, 1965.
- [39] P. L. Combettes and J.-C. Pesquet, “Proximal thresholding algorithm for minimization over orthonormal bases,” *SIAM Journal on Optimization*, vol. 18, no. 4, pp. 1351–1376, 2008.
- [40] R. Gribonval and M. Nikolova, “A characterization of proximity operators,” *Journal of Mathematical Imaging and Vision*, vol. 62, no. 6, pp. 773–789, July 2020.

Supplementary: Characterization of $F(\mathbf{y}, \mathbf{r})$ as a proximity operator.

In this supplementary, we address the problem of identifying a function $f_{\mathbf{r}} : \mathbb{R}^K \rightarrow \mathbb{R} \cup \{+\infty\}$ such that $F(\mathbf{y}, \mathbf{r}) = \text{prox}_{f_{\mathbf{r}}}(\mathbf{y})$, with F defined in (20) of the main paper. Note that we ignore here the layer index t for simplicity.

First, we present in Section A sufficient conditions for the existence of such a function in a general setting. Then, in Section B we consider more particularly the case of APL units. Finally, in Section C we recover the function $f_{\mathbf{r}}$ corresponding to our sublayer F .

5.1 Proximity operators

First, let us recall that the proximity operator of a real-valued convex function $\sigma : \mathbb{R}^K \rightarrow \mathbb{R} \cup \{+\infty\}$ is defined as the mapping of a vector \mathbf{z} to the solution of the minimization problem:

$$\text{prox}_{\rho^{-1}\sigma}(\mathbf{z}) := \arg \min_{\mathbf{p} \in \mathbb{R}^K} \sigma(\mathbf{p}) + \frac{\rho}{2} \|\mathbf{p} - \mathbf{z}\|^2. \quad (21)$$

This definition can be extended to nonconvex functions, resulting in a possibly set-valued operator. Following [38, 39, 40], a function $g : \mathcal{Z} \subset \mathbb{R}^K \rightarrow \mathbb{R}^K$ can be characterized as the proximity operator of a convex function when it is strictly increasing over each coordinate. More precisely, the authors in [40] show that if there exists a convex, lower semi-continuous function $\tilde{g} : \mathbb{R}^K \rightarrow \mathbb{R}$ such that $g(\mathbf{z})$ is a subgradient of $\tilde{g}(\mathbf{z})$ for all $\mathbf{z} \in \mathcal{Z}$, i.e.,

$$g(\mathbf{z}) \in \{\mathbf{v} \in \mathbb{R}^K \mid \forall \mathbf{p} \in \mathbb{R}^K, \langle \mathbf{p} - \mathbf{z}, \mathbf{v} \rangle + \tilde{g}(\mathbf{z}) \leq \tilde{g}(\mathbf{p})\}, \quad (22)$$

then there exists σ such that $g(\mathbf{z}) \in \text{prox}_{\sigma}(\mathbf{z})$ and the following relation stands:

$$\forall \mathbf{z} \in \mathcal{Z}, \sigma(g(\mathbf{z})) = \langle \mathbf{z}, g(\mathbf{z}) \rangle - \frac{1}{2} \|g(\mathbf{z})\|^2 - \tilde{g}(\mathbf{z}). \quad (23)$$

Finally, when g is invertible and with $\mathbf{y} = g(\mathbf{z})$, we have:

$$\sigma(\mathbf{y}) = \langle g^{-1}(\mathbf{y}), \mathbf{y} \rangle - \frac{1}{2} \|\mathbf{y}\|^2 - \tilde{g}(g^{-1}(\mathbf{y})). \quad (24)$$

5.B Characterization with APL

We now address the case of (strictly increasing) APL functions as defined in (19), with negative weights (cf. 3.3) and at least one nonnegative bias b_c . Let us consider the following convex, lower semi-continuous function $\overline{\text{APL}}$ such that $\forall z \in \mathbb{R}$:

$$\overline{\text{APL}}(z) = \frac{z^2}{2} \chi_{[0; +\infty)}(z) + \sum_{c=1}^C w_c \left(\frac{-z^2}{2} + b_c z \right) \chi_{]-\infty; b_c]}(z), \quad (25)$$

where χ_{Π} is the indicator function of set Π . Since for any $z \in \mathbb{R}$, $\overline{\text{APL}}(z)$ is a subgradient of $\overline{\text{APL}}(z)$, and denoting $\widetilde{\text{APL}}(\mathbf{z}) = \sum_{k=1}^K \overline{\text{APL}}(z_k)$, it is straightforward to show that the relation (22) stands for $g(\mathbf{z}) = \overline{\text{APL}}(\mathbf{z})$ and $\tilde{g}(\mathbf{z}) = \widetilde{\text{APL}}(\mathbf{z})$. Besides, since $\overline{\text{APL}}$ is invertible we can use the relation (24) to identify σ :

$$\sigma(\mathbf{y}) = \langle \overline{\text{APL}}^{-1}(\mathbf{y}), \mathbf{y} \rangle - \frac{1}{2} \|\mathbf{y}\|^2 - \widetilde{\text{APL}}(\overline{\text{APL}}^{-1}(\mathbf{y})), \quad (26)$$

with:

$$\overline{\text{APL}}^{-1}(\mathbf{y}) = \frac{\mathbf{y} - \sum_{c=1}^C w_c b_c \chi_{]-\infty, \text{APL}(b_c)]}(\mathbf{y})}{\chi_{[\overline{\text{APL}}(0), +\infty)}(\mathbf{y}) - \sum_{c=1}^C w_c \chi_{]-\infty, \text{APL}(b_c)]}(\mathbf{y})}. \quad (27)$$

5.C Characterization with F

Finally, let us retrieve $f_{\mathbf{r}} : \mathbb{R}^K \rightarrow \mathbb{R} \cup \{+\infty\}$ such that $F(\mathbf{y}, \mathbf{r}) = \text{prox}_{f_{\mathbf{r}}}(\mathbf{y})$. Drawing on the previous section and using the definition of F from (20), we have:

$$F(\mathbf{y}, \mathbf{r}) = \text{prox}_{\sigma} \left(\gamma^{(1)} \mathbf{y} + \gamma^{(2)} \frac{\mathbf{r}^{\beta-1}}{\beta-1} \right). \quad (28)$$

To fully identify $f_{\mathbf{r}}$, we first need to reformulate (28) so that the argument of the right hand side term simply becomes \mathbf{y} . To that end, we leverage a property from [27], which consists in first rewriting (28) as follows:

$$F(\mathbf{y}, \mathbf{r}) = \text{prox}_{\sigma} \left(\frac{\mathbf{y} - \mathbf{q}}{2\alpha + 1} \right), \quad (29)$$

with $\alpha = \frac{1 - \gamma^{(1)}}{2\gamma^{(1)}}$ and $\mathbf{q} = -\frac{\gamma^{(2)}}{\gamma^{(1)}} \frac{\mathbf{r}^{\beta-1}}{\beta-1}$. The property from [27] then states that:

$$\text{prox}_{\varphi + \alpha \|\cdot\|^2 + \langle \mathbf{q}, \cdot \rangle}(\mathbf{y}) = \text{prox}_{\varphi/(2\alpha+1)}\left(\frac{\mathbf{y} - \mathbf{q}}{2\alpha + 1}\right). \quad (30)$$

Let $\varphi = (2\alpha + 1)\sigma$. Combining (29) and (30) yields:

$$\mathbf{F}(\mathbf{y}, \mathbf{r}) = \text{prox}_{(2\alpha+1)\sigma + \alpha \|\cdot\|^2 + \langle \mathbf{q}, \cdot \rangle}(\mathbf{y}). \quad (31)$$

As a result, from (31) we can identify $f_{\mathbf{r}}$ such that its proximity operator is \mathbf{F} . If we further exploit the definition of σ from (26), we finally have:

$$f_{\mathbf{r}}(\mathbf{y}) = \frac{1}{\gamma^{(1)}} \left\langle \text{APL}^{-1}(\mathbf{y}) - \gamma^{(2)} \frac{\mathbf{r}^{\beta-1}}{\beta-1}, \mathbf{y} \right\rangle - \frac{1}{2} \|\mathbf{y}\|^2 - \frac{1}{\gamma^{(1)}} \widetilde{\text{APL}}(\text{APL}^{-1}(\mathbf{y})). \quad (32)$$

Therefore, using (32) one can recover the function associated with the learned proximity operator, and consequently identify the metric involved in the formulation of the PR problem.