



HAL
open science

Signal Inpainting from Fourier Magnitudes

Louis Bahrman, Marina Krémé, Paul Magron, Antoine Deleforge

► **To cite this version:**

Louis Bahrman, Marina Krémé, Paul Magron, Antoine Deleforge. Signal Inpainting from Fourier Magnitudes. 2022. hal-03832480v1

HAL Id: hal-03832480

<https://hal.science/hal-03832480v1>

Preprint submitted on 27 Oct 2022 (v1), last revised 22 Jun 2023 (v3)

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.

SIGNAL INPAINTING FROM FOURIER MAGNITUDES

Louis Bahrman^{1,2}, Marina Krémé¹, Paul Magron¹, Antoine Deleforge¹

¹Université de Lorraine, CNRS, Inria, Loria, F-54000 Nancy, France

²LTCI, Télécom Paris, Institut Polytechnique de Paris, France

ABSTRACT

Signal inpainting is the task of restoring degraded or missing samples in a signal. In this paper we address signal inpainting when Fourier magnitudes are observed. We propose a mathematical formulation of the problem that highlights its connection with phase retrieval, and we introduce two methods for solving it. First, we derive an alternating minimization scheme, which shares similarities with the Gerchberg-Saxton algorithm, a classical phase retrieval method. Second, we propose a convex relaxation of the problem, which is inspired by recent approaches that reformulate phase retrieval into a semi-definite program. We assess the potential of these methods for the task of inpainting gaps in speech signals. Our methods exhibit both a high probability of recovering the original signals and robustness to magnitude noise.

Index Terms— Signal inpainting, phase retrieval, audio restoration, convex relaxation, alternating minimization.

1. INTRODUCTION

Signal inpainting [1] is an inverse problem that consists in restoring signals degraded by sample loss. Such a problem typically arises as a result of degradation during signal transmission (packet loss concealment [2]) or in digitization of physically degraded media. Inpainting can also be used to restore signal samples subject to a degradation so heavy that the information about the samples can be considered lost (e.g., signal clipping [3] or impulsive noises [4]). More specifically, in this paper we focus on inpainting compact gaps, which occurs, e.g., when an audio signal is corrupted with clicks [5].

Approaches that tackle this issue can be divided into two categories depending on the number of missing samples or duration of the gaps. When considering short gaps (less than 100 ms), approaches based on autoregressive modeling [6], convex optimization [7], sparse modeling [8], or Bayesian estimation [9] have shown promising results. Conversely, approaches based on sinusoidal modeling [10] or graphs [11] are more suitable for longer gaps (more than 100 ms). However, signal inpainting remains a challenging problem, and

these approaches are not adapted to scenarios where some additional information about the signal is available.

We focus here on a setting where the Fourier magnitudes of the signal are observed. Such setting was notably studied in [12] in the different context of source separation. Its interest lies in the observation that the Fourier magnitudes of natural signals often exhibit smoother and hence more predictable evolution than their respective Fourier phases. The problem then shares a close connection with phase retrieval [13], the task of retrieving a signal from nonnegative measurements (usually magnitudes of a set of inner products). From the seminal work of Gerchberg and Saxton [14], this task has been revived over the last decade with the development of novel optimization approaches based on gradient descent [15] or convex relaxations [16, 17]. While phase retrieval has experienced considerable progress in recent years, its connection with signal inpainting is left to explore, both from the methodological and the theoretical standpoint.

In this paper, we propose to bridge the gap between signal inpainting and phase retrieval, specifically for restoring the missing samples in a signal whose Fourier magnitudes are observed. We formulate it as a constrained phase retrieval problem. Inspired by phase retrieval algorithms, we derive two methods to solve this problem, based on alternating minimization (AM) and convex relaxation (CR). Experiments conducted on speech signals reveal the potential of these techniques for audio inpainting, since they exhibit a high probability of recovering the original signal, as well as a property of robustness to the magnitude noise.

The rest of this paper is organized as follows. Section 2 formulates the problem and introduces two methods for solving it. Experiments on speech signals are conducted in Section 3. Finally, Section 4 draws some concluding remarks.

2. METHODS

2.1. From inpainting to phase retrieval

Let $\mathbf{x}^{\natural} \in \mathbb{R}^L$ denote a signal. We partition its support $\{0, \dots, L-1\}$ into two sets v and \bar{v} such that $\mathbf{x}_v^{\natural} \in \mathbb{R}^{L-d}$ and $\mathbf{x}_{\bar{v}}^{\natural} \in \mathbb{R}^d$ denote the observed and missing samples, respectively, and where $d \leq L$ denotes the number of missing samples, whose location \bar{v} is assumed to be known. We also

This work was made with the support of the French National Research Agency through project DENISE (ANR-20-CE48-0013), and was conducted while L. Bahrman was an intern with Inria in Nancy, France.

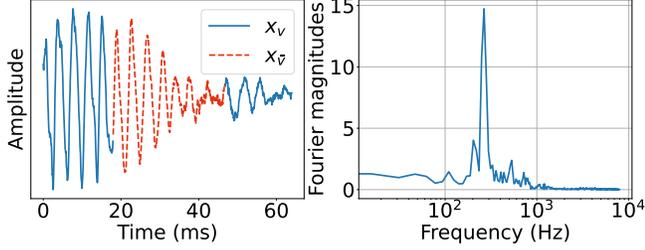


Fig. 1: Illustration of restoring the missing samples of a signal (left), assuming its Fourier magnitudes (right) are known.

assume that the magnitudes of its discrete Fourier transform (DFT), denoted by $\mathbf{b} \in \mathbb{R}_+^L$ are known: $\mathbf{b} = |\Phi \mathbf{x}^h|$, where $\Phi \in \mathbb{C}^{L \times L}$ is the DFT matrix (note that we consider a *complete* transform, i.e., the DFT does not use zero-padding). The task of signal inpainting from Fourier magnitudes is illustrated in Figure 1 and can be stated as:

$$\text{Find } \mathbf{x} \in \mathbb{R}^L \text{ such that } |\Phi \mathbf{x}| = \mathbf{b} \text{ and } \mathbf{x}_v = \mathbf{x}_v^h. \quad (1)$$

We theoretically show in a supplementary material that for almost all $\mathbf{x}^h \in \mathbb{R}^L$, this problem admits a unique solution if at most 25% of the signal is missing [18]. Nonetheless, the problem is non-convex and is an instance of quadratically constrained quadratic program, which are known to be NP-hard in general [19]. There is hence a need for efficient algorithms that provide a solution. Let us formulate (1) as the following optimization problem:

$$\min_{\mathbf{x} \in \mathbb{R}^L} \|\Phi \mathbf{x} - \mathbf{b}\|^2 \quad \text{s. t.} \quad \mathbf{x}_v = \mathbf{x}_v^h, \quad (2)$$

where $\|\cdot\|$ denotes the Euclidean norm.¹ Problem (2) reads as a constrained phase retrieval problem. If the whole signal is missing ($d = L$), the constraint vanishes and it reduces to Fourier phase retrieval. The cornerstone of our approaches lies in introducing an auxiliary phase variable $\mathbf{u} \in \mathbb{C}^L$ such that $|\mathbf{u}| = 1$. Then, $\mathbf{b} = |\text{diag}(\mathbf{b})\mathbf{u}|$, where $\text{diag}(\mathbf{b})$ is the diagonal matrix whose entries are given by the vector \mathbf{b} . We then turn our attention to minimizing the following auxiliary function, which is *exact* in the sense of [22]:

$$\min_{\mathbf{x} \in \mathbb{R}^L, \mathbf{u} \in \mathbb{C}^L} \|\Phi \mathbf{x} - \text{diag}(\mathbf{b})\mathbf{u}\|^2 \quad \text{s. t.} \quad \mathbf{x}_v = \mathbf{x}_v^h, |\mathbf{u}| = 1. \quad (3)$$

2.2. Alternating minimization

The first approach we propose to solve (3) is an AM scheme. Let us first fix \mathbf{u} and derive the update on \mathbf{x} , for which we propose to incorporate the constraint $\mathbf{x}_v = \mathbf{x}_v^h$ within the loss function. To that end, let us reorder \mathbf{x} as $\mathbf{x} = \begin{bmatrix} \mathbf{x}_{\bar{v}} \\ \mathbf{x}_v \end{bmatrix}$ (and similarly for \mathbf{x}^h), and split Φ accordingly as $\Phi = \begin{bmatrix} \Phi_{\bar{v}} & \Phi_v \end{bmatrix}$, with $\Phi_{\bar{v}} \in \mathbb{C}^{L \times d}$ and $\Phi_v \in \mathbb{C}^{L \times (L-d)}$. Using these notations, we

¹Recent works such as [20, 21] have investigated alternative discrepancy measures for phase retrieval. We focus on the Euclidean norm in this study.

Algorithm 1 AM for signal inpainting

Input: $\begin{cases} \mathbf{b} \in \mathbb{R}_+^L : \text{observations, } \mathbf{x}_v^h : \text{known signal} \\ \Phi : \text{Fourier matrix,} \end{cases}$

- 1: Initialize $\mathbf{x}_{\bar{v}}^{(0)}$ and $\mathbf{x}^{(0)} \leftarrow \begin{bmatrix} \mathbf{x}_{\bar{v}}^{(0)} \\ \mathbf{x}_v^h \end{bmatrix}$
- 2: $i \leftarrow 0$
- 3: **while** convergence not reached **do**
- 4: $\mathbf{u}^{(i+1)} \leftarrow \frac{\Phi \mathbf{x}^{(i)}}{|\Phi \mathbf{x}^{(i)}|}$
- 5: $\mathbf{x}_{\bar{v}}^{(i+1)} \leftarrow \Re(\Phi_{\bar{v}}^H \text{diag}(\mathbf{b})\mathbf{u}^{(i+1)})$
- 6: $\mathbf{x}^{(i+1)} \leftarrow \begin{bmatrix} \mathbf{x}_{\bar{v}}^{(i+1)} \\ \mathbf{x}_v^h \end{bmatrix}$
- 7: $i \leftarrow i + 1$
- 8: **end while**

Output: Reconstructed signal $\mathbf{x}^{(i)}$

have $\Phi \mathbf{x} = \Phi_{\bar{v}} \mathbf{x}_{\bar{v}} + \Phi_v \mathbf{x}_v$, and (3) rewrites:

$$\min_{\mathbf{x}_{\bar{v}} \in \mathbb{R}^d} \|\Phi_{\bar{v}} \mathbf{x}_{\bar{v}} + \Phi_v \mathbf{x}_v^h - \text{diag}(\mathbf{b})\mathbf{u}\|^2. \quad (4)$$

Since $\Phi_{\bar{v}}$ is full-rank, it has a left inverse which is its Hermitian transpose $\Phi_{\bar{v}}^H$. Besides, recall that since the DFT is an orthogonal transform, then $\Phi_{\bar{v}}^H \Phi_v = 0$. Altogether this yields the following solution to (4):

$$\mathbf{x}_{\bar{v}} = \Phi_{\bar{v}}^H \text{diag}(\mathbf{b})\mathbf{u}. \quad (5)$$

Let us now derive the update for \mathbf{u} when \mathbf{x} is fixed, for which (3) rewrites:

$$\min_{\mathbf{u} \in \mathbb{C}^L} \|\Phi \mathbf{x} - \text{diag}(\mathbf{b})\mathbf{u}\|^2 \quad \text{s. t.} \quad |\mathbf{u}| = 1, \quad (6)$$

which is straightforward to solve:

$$\mathbf{u} = \frac{\Phi \mathbf{x}}{|\Phi \mathbf{x}|}. \quad (7)$$

Alternating (5) and (7) yields a procedure summarized in Algorithm 1. We discuss the initialization strategy in Section 3.1. Note that at line 5 we apply the real part function to ensure a real-valued signal estimate.²

Remark: Algorithm 1 consists in computing the DFT of a signal, setting its magnitude at a target value, inverting the DFT, and finally keeping the observed samples unchanged. As such, it is similar to the Gerchberg-Saxton algorithm [14], except that the signal-domain magnitude constraint is here replaced with a support constraint on the observed samples.

2.3. Convex relaxation

Let us now derive a method inspired by the PhaseCut algorithm [17], which is based on a CR of the problem. This approach consists in reformulating phase retrieval as a constrained trace minimization problem by lifting it to a higher

²It can be proven rigorously that doing so does not affect convergence guarantees [21], but we do not detail it here for brevity.

dimensional space and relaxing the rank-one constraint. The resulting problem can then be efficiently solved via semi-definite programming. We consider the formulation (3) in which we inject the expression of \mathbf{x} given by (5):

$$\min_{\mathbf{u} \in \mathbb{C}^L} \left\| (\Phi_{\bar{v}} \Phi_{\bar{v}}^H - I) \text{diag}(\mathbf{b}) \mathbf{u} + \Phi_{\bar{v}} \mathbf{x}_{\bar{v}} \right\|^2 \text{ s.t. } |\mathbf{u}| = 1. \quad (8)$$

Now, let us introduce the following:

$$\tilde{\mathbf{m}} := [(\Phi_{\bar{v}} \Phi_{\bar{v}}^H - I) \text{diag}(\mathbf{b}), \Phi_{\bar{v}} \mathbf{x}_{\bar{v}}] \text{ and } \tilde{\mathbf{u}} = \begin{bmatrix} \mathbf{u} \\ 1 \end{bmatrix}, \quad (9)$$

from which we can rewrite (8) as:

$$\min_{\tilde{\mathbf{u}} \in \mathbb{C}^{L+1}} \|\tilde{\mathbf{m}} \tilde{\mathbf{u}}\|^2 \text{ s.t. } |\tilde{\mathbf{u}}| = 1 \text{ and } \tilde{\mathbf{u}}[L] = 1. \quad (10)$$

Drawing on [17], we lift and relax (10) to the following convex problem:

$$\min_{\tilde{\mathbf{U}} \in \mathbb{C}^{(L+1) \times (L+1)}} \text{trace}(\tilde{\mathbf{M}} \tilde{\mathbf{U}}) \text{ s.t. } \text{diag}(\tilde{\mathbf{U}}) = \mathbf{1}, \tilde{\mathbf{U}} \succeq 0, \quad (11)$$

where $\tilde{\mathbf{M}} = \tilde{\mathbf{m}}^H \tilde{\mathbf{m}} \in \mathbb{C}^{(L+1) \times (L+1)}$ and $\tilde{\mathbf{U}} = \tilde{\mathbf{u}} \tilde{\mathbf{u}}^H$. As in [17, 12], problem (11) can be solved by means of a block coordinate descent algorithm, which we summarize in Algorithm 2 (lines 1 to 11). It consists of a nested loop where at each iteration i , the columns of $\tilde{\mathbf{U}}$ are updated sequentially using the notation:

$$k^c = \{0, \dots, k-1, k+1, \dots, L-1\}. \quad (12)$$

This yields a global solution $\tilde{\mathbf{U}}$ to problem (11). If this solution is of rank 1, i.e., $\tilde{\mathbf{U}} = \tilde{\mathbf{u}} \tilde{\mathbf{u}}^H$, then $\tilde{\mathbf{u}}[0, \dots, L-1] / \tilde{\mathbf{u}}[L]$ is guaranteed to globally solve (8). We then obtain a global solution of (3) via $\mathbf{x}_{\bar{v}} = \Phi_{\bar{v}}^H \text{diag}(\mathbf{b}) \tilde{\mathbf{u}}[0, \dots, L-1] / \tilde{\mathbf{u}}[L]$. However, $\tilde{\mathbf{U}}$ needs not be rank-1 in general. Hence, as commonly done in semi-definite relaxations, its closest rank-1 approximation is used, which by the Eckart-Young-Mirsky theorem amounts to setting $\tilde{\mathbf{u}}$ to the eigenvector associated to the largest eigenvalue of $\tilde{\mathbf{U}}$.

3. EXPERIMENTS

In this Section we assess the potential of our methods for inpainting gaps in audio signals. Our code is available online.³

3.1. Experimental setting

Data We consider 100 speech signals from the Librispeech dataset [23]. Signals are sampled at 16 kHz. For each signal, we extract a non-silent sub-signal of variable length L at a random location, where we create a gap of d samples.

Methods State-of-the-art audio restoration methods such as [6, 8] are not appropriate comparison baselines since they do not allow to incorporate magnitude constraints. To avoid unfair comparison, we focus our experiments on evaluating our proposed algorithms. For the AM algorithm, iterations stop when the the maximum loss variation over the previous

³<https://github.com/Louis-Bahrman/Inpainting-Fourier>

Algorithm 2 CR for signal inpainting

Input: $\begin{cases} \mathbf{b} \in \mathbb{R}_+^L : \text{observations, } \mathbf{x}_v^h : \text{known signal} \\ \Phi : \text{Fourier matrix, } \nu \geq 0 : \text{barrier parameter} \\ \tilde{\mathbf{m}} \in \mathbb{C}^{L \times (L+1)} : \text{matrix defined by (9)} \\ n_{\text{iter}} : \text{number of iterations} \end{cases}$

- 1: $\tilde{\mathbf{M}} \leftarrow \tilde{\mathbf{m}}^H \tilde{\mathbf{m}}$ and $\tilde{\mathbf{U}}^{(0)} \leftarrow I$
- 2: **for** $i = 1, \dots, n_{\text{iter}}$ **do**
- 3: **for** $k = 0, \dots, L-1$ **do**
- 4: $\mathbf{z} \leftarrow \tilde{\mathbf{U}}_{k^c, k^c}^{(i)} \tilde{\mathbf{M}}_{k^c, k^c}$ and $\gamma \leftarrow \mathbf{z}^H \tilde{\mathbf{M}}_{k^c, k^c}$
- 5: **if** $\gamma > 0$ **then**
- 6: $\tilde{\mathbf{U}}_{k^c, k^c}^{(i+1)}$ and $(\tilde{\mathbf{U}}_{k^c, k^c}^{(i+1)})^H \leftarrow -\sqrt{\frac{1-\nu}{\gamma}} \mathbf{z}$
- 7: **else**
- 8: $\tilde{\mathbf{U}}_{k^c, k^c}^{(i+1)}$ and $(\tilde{\mathbf{U}}_{k^c, k^c}^{(i+1)})^H \leftarrow 0$
- 9: **end if**
- 10: **end for**
- 11: **end for**
- 12: $\tilde{\mathbf{u}} =$ eigenvector associated to the largest eigenvalue of $\tilde{\mathbf{U}}^{(i+1)}$, and $\tilde{\mathbf{u}} \leftarrow \frac{\tilde{\mathbf{u}}}{\tilde{\mathbf{u}}[L-1]}$
- 13: $\mathbf{x}_{\bar{v}} \leftarrow \Re(\Phi_{\bar{v}}^H \text{diag}(\mathbf{b}) \tilde{\mathbf{u}}[0:L-1])$ and $\mathbf{x} = \begin{bmatrix} \mathbf{x}_{\bar{v}} \\ \mathbf{x}_v^h \end{bmatrix}$

Output: Reconstructed signal \mathbf{x}

5 iterations does not exceed 10^{-10} , or when a maximum number of 1000 iterations is reached. Preliminary experiments have revealed no significant difference between basic initialization schemes (either using a random or zero phase in the Fourier domain, or using a random or zero missing signal in the time domain). Besides, let us outline that we derived an initialization scheme inspired by the spectral initialization method [24], but it did not yield any significant performance improvement. Therefore, we display hereafter the results obtained with an initialization that consists in replacing the missing signal values with zeros. The CR algorithm uses $n_{\text{iter}} = 10$, since the performance does not further improved beyond, while it significantly increases the computational burden. We set the barrier parameter ν to 0 as in [12]. Finally, we also consider a combined CR+AM algorithm. This technique consists in first estimating the signal with CR, and then using this estimate as initialization for AM, which uses the same stopping criterion as above.

Evaluation metric To assess the quality of the reconstruction, we resort to the signal-to-error ratio (SER) expressed in dB:

$$\text{SER}(\mathbf{x}_{\bar{v}}, \mathbf{x}_{\bar{v}}^h) = 10 \log_{10} \left(\frac{\|\mathbf{x}_{\bar{v}}^h\|^2}{\|\mathbf{x}_{\bar{v}} - \mathbf{x}_{\bar{v}}^h\|^2} \right), \quad (13)$$

where \mathbf{x} is the signal estimated with one of our methods. Note that we only compute the SER on the set of missing samples \bar{v} since $\mathbf{x}_{\bar{v}} = \mathbf{x}_{\bar{v}}^h$ elsewhere. We consider that perfect reconstruction is achieved when the SER is greater than 20 dB.

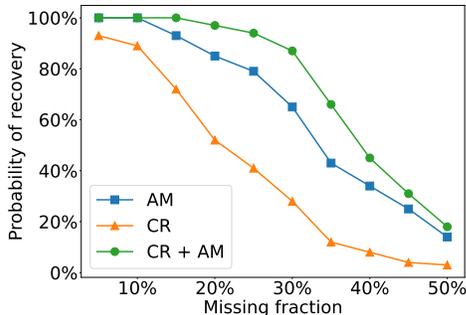


Fig. 2: Influence of the missing fraction onto performance.

Note that AM’s initialization (replacing the missing signal with zeros) corresponds to an SER of 0.

3.2. Results

First, we compare the recovery rate of the three proposed methods with respect to the missing signal fraction. The results corresponding to $L = 1024$ (64 ms) are displayed in Figure 2. CR+AM achieves the best results and consistently outperforms the other approaches. This is explained by the ability of CR to provide a solution that is more likely to converge to a global minimum than AM’s basic initialization. While all methods exhibit a performance drop when the missing fraction increases, CR+AM and AM still yield perfect reconstruction in at least 80 % of the cases when less than 30 % of the signal is missing. Overall, a general guideline might be to resort to the CR+AM method, except when relatively few samples are missing (5 to 10 %), since in this case AM yields perfect reconstruction at a lower computational cost.

Then, Figure 3 shows the recovery rate achieved by the best performing method CR+AM as a function of the signal length L (which ranges from 128 to 4096 samples) and the fraction of missing samples d/L (which ranges from 5 % to 50 %). We observe that perfect reconstruction is achieved in most cases when the fraction is lower than 25 % and $L \leq 1024$, suggesting that the approach performs near the theoretical optimum of [18] in this regime. While for short signals the method even yields satisfactory results under slightly larger missing fractions, its performance eventually drops for larger signals. This may be explained by the increased dimensionality of the problem which may increase the risk of getting trapped in local minima.

Finally, let us evaluate the robustness of our methods to magnitude noise. Indeed, in the previous experiments we considered oracle magnitudes, but in practical applications these are estimated beforehand and therefore contaminated by estimation errors. To simulate such a scenario, we consider noisy magnitudes $\mathbf{b} = \max(0, |\Phi \mathbf{x}^{\mathbf{h}}| + \mathbf{n})$, where \mathbf{n} is a white Gaussian noise whose variance is adjusted to fit a given magnitude signal-to-noise (SNR) value. The results are presented in Figure 4, where the signal length is fixed at $L = 1024$ and the fraction of missing signal at 25%. Over-

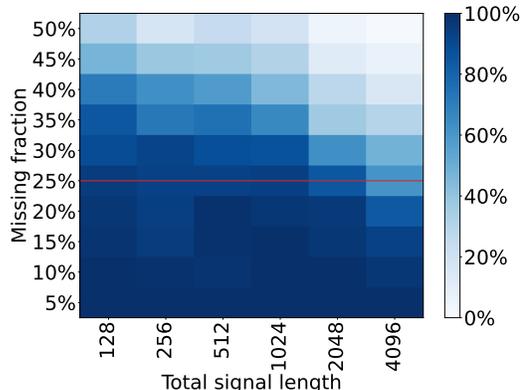


Fig. 3: Probability of achieving perfect reconstruction for CR+AM, as a function of the signal length (in samples) and the missing signal fraction. The red line represents the 25 % theoretical limit under which perfect reconstruction is achievable for almost all signals [18].

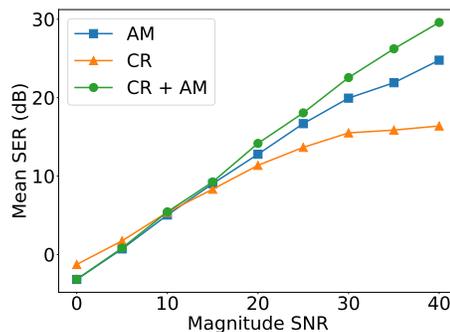


Fig. 4: Influence of the magnitude noise onto performance.

all, we observe that all methods exhibit a similar linear decay in performance in the log-log plot when the SNR falls below 20 dB, suggesting an encouraging robustness to errors. Interestingly, while CR performs worse than the other methods at higher SNRs, it still yields solutions that provide a better initialization than zero to AM.

4. CONCLUSION

We have investigated the problem of signal inpainting from known Fourier magnitude measurements. After formulating the associated optimization problem, we proposed two estimation methods based on alternating minimization and convex relaxation. Experiments in the noiseless regime highlighted that combining the two methods yields near-exact reconstruction most of the time when less than 25% of the signal is missing, closely following our theoretical limit. Further experiments demonstrated a relative robustness of the methods to magnitude errors. Future work will therefore focus on applying such techniques to more realistic settings where magnitudes are estimated, e.g., via light magnitude interpolation methods in the short-time Fourier domain or via data-driven models such as deep neural networks.

5. REFERENCES

- [1] A. Adler, V. Emiya, M. G. Jafari, M. Elad, R. Gribonval, and M. D. Plumbley, "Audio inpainting," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 20, no. 3, pp. 922–932, 2012.
- [2] C.A. Rodbro, M.N. Murthi, S.V. Andersen, and S.H. Jensen, "Hidden Markov model-based packet loss concealment for voice over IP," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 14, no. 5, pp. 1609–1623, September 2006.
- [3] G. Chantas, S. Nikolopoulos, and I. Kompatsiaris, "Sparse audio inpainting with variational bayesian inference," in *Proc. IEEE International Conference on Consumer Electronics (ICCE)*, October 2018.
- [4] B. Derebssa Dufera, K. Eneman, and T. Van Waterschoot, "Missing sample estimation based on high-order sparse linear prediction for audio signals," in *Proc. European Signal Processing Conference (EUSIPCO)*, September 2018.
- [5] P. Magron, R. Badeau, and B. David, "Phase reconstruction of spectrograms with linear unwrapping: Application to audio signal restoration," in *Proc. European Signal Processing Conference (EUSIPCO)*, August 2015.
- [6] A. Janssen, R. Veldhuis, and L. Vries, "Adaptive interpolation of discrete-time signals that can be modeled as autoregressive processes," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 34, no. 2, pp. 317–330, 1986.
- [7] G. Taubock, S. Rajbanshi, and P. Balazs, "Dictionary learning for sparse audio inpainting," *IEEE Journal of Selected Topics in Signal Processing*, vol. 15, no. 1, pp. 104–119, January 2021.
- [8] O. Mokry, P. Zaviska, P. Rajmic, and V. Vesely, "Introducing SPAIN (SParse audio INpainter)," in *Proc. European Signal Processing Conference (EUSIPCO)*, September 2019.
- [9] S.J. Godsill and P.J.W. Rayner, "A bayesian approach to the restoration of degraded audio signals," *IEEE Transactions on Speech and Audio Processing*, vol. 3, no. 4, pp. 267–278, 1995.
- [10] M. Lagrange, S. Marchand, and J.-B. Rault, "Long interpolation of audio signals using linear prediction in sinusoidal modeling," *Journal of the Audio Engineering Society*, vol. 53, no. 10, pp. 891–905, 2005.
- [11] N. Perraudin, N. Holighaus, P. Majdak, and P. Balazs, "Inpainting of long audio segments with similarity graphs," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 6, pp. 1083–1094, 2018.
- [12] A. Deleforge and Y. Traonmilin, "Phase unmixing: Multichannel source separation with magnitude constraints," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017.
- [13] A. Walther, "The question of phase retrieval in optics," *Optica Acta: International Journal of Optics*, vol. 10, no. 1, pp. 41–49, 1963.
- [14] R. W. Gerchberg and O. W. Saxton, "A practical algorithm for the determination of phase from image and diffraction plane pictures," *Optik*, vol. 35, pp. 237–246, 1972.
- [15] 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.
- [16] E. J. Candès, T. Strohmer, and V. Voroninski, "Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming," *Communications on Pure and Applied Mathematics*, vol. 66, no. 8, pp. 1241–1274, 2013.
- [17] I. Waldspurger, A. D'Aspremont, and S. Mallat, "Phase recovery, maxcut and complex semidefinite programming," *Mathematical Programming*, vol. 149, no. 1, pp. 47–81, 2015.
- [18] M. Krémé, A. Deleforge, P. Magron, and L. Bahrman, "Signal inpainting from Fourier magnitudes: A uniqueness result," <https://magronp.github.io/files/ICASSP23sup.pdf>, October 2022.
- [19] C. Audet, P. Hansen, B. Jaumard, and G. Savard, "A branch and cut algorithm for nonconvex quadratically constrained quadratic programming," *Mathematical Programming*, vol. 87, no. 1, pp. 131–152, 2000.
- [20] Z. Li, K. Lange, and J. A. Fessler, "Poisson phase retrieval with Wirtinger flow," in *Proc. IEEE International Conference on Image Processing (ICIP)*, September 2021.
- [21] 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.
- [22] Y. G. Yevtushenko and V. G. Zhadan, "Exact auxiliary functions in optimization problems," *USSR computational mathematics and mathematical physics*, vol. 30, no. 1, pp. 31–42, 1990.
- [23] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An ASR corpus based on public domain audio books," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, April 2015.
- [24] 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.