



HAL
open science

Spurious minimizers in non uniform Fourier sampling optimization

Frédéric de Gournay, Alban Gossard, Pierre Weiss

► **To cite this version:**

Frédéric de Gournay, Alban Gossard, Pierre Weiss. Spurious minimizers in non uniform Fourier sampling optimization. *Inverse Problems*, inPress. hal-03212145v3

HAL Id: hal-03212145

<https://hal.science/hal-03212145v3>

Submitted on 20 Jul 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.



Distributed under a Creative Commons Attribution 4.0 International License

Spurious minimizers in non uniform Fourier sampling optimization

Alban Gossard^{1,2}
alban.paul.gossard@gmail.com

Frédéric de Gournay^{1,3}
degourna@insa-toulouse.fr

Pierre Weiss^{1,2}
pierre.armand.weiss@gmail.com

¹ Institut de Mathématiques de Toulouse; UMR5219; Université de Toulouse; CNRS

² Université de Toulouse, F-31062 Toulouse Cedex 9, France

³ INSA, F-31077 Toulouse, France

July 20, 2022

Abstract— A recent trend in the signal/image processing literature is the optimization of Fourier sampling schemes for specific datasets of signals. In this paper, we explain why choosing optimal non Cartesian Fourier sampling patterns is a difficult nonconvex problem by bringing to light two optimization issues. The first one is the existence of a combinatorial number of spurious minimizers for a generic class of signals. The second one is a vanishing gradient effect for the high frequencies. We conclude the paper by showing how using large datasets can mitigate the first effect and illustrate experimentally the benefits of using stochastic gradient algorithms with a variable metric.

1 Introduction

Finding efficient Fourier sampling schemes is a critical issue in communications and imaging. This led to various theories including the celebrated Shannon-Nyquist theorems for bandlimited signals and compressed sensing for sparse signals. Unfortunately - in most practical cases - the signals to reconstruct are quite loosely described by these generic classes. For instance, magnetic resonance images of brains or knees have a rich structure due to the underlying object. It is therefore tempting to optimize a sampling scheme directly for a given dataset rather than relying on a rough mathematical model. The recent progresses in Graphical Processing Units (GPU) programming, automatic differentiation and machine learning make this idea even more tantalizing. In the sole field of Magnetic Resonance Imaging (MRI), the following list of references [9, 10, 15, 3, 19, 16, 17, 16, 8, 11, 14, 1] illustrates this novel trend.

Unfortunately, most of the above works report (more or less explicitly) optimization issues. Fig. 1 illustrates one of them. In this example, we tried to optimize a sampling scheme for a single image

from the fastMRI challenge [18]. To this end, we minimize the ℓ^2 reconstruction error using a simple back-projection reconstructor with a subsampling factor of 2. The trajectory of a gradient descent is displayed in Fig. 1b. As can be seen, the final sampling set covers approximately uniformly the Fourier domain, while we would expect the low frequencies to be sampled more densely. This likely highlights the presence of a spurious minimizer.

The aim of this paper is to explain this phenomenon from a mathematical perspective and to bring some solutions to mitigate the difficulties. We focus on linear reconstruction methods, which simplifies the analysis and we highlight the critical role of the non uniform Fourier transform as an oscillation generator. We expect that some of the arguments can be reused for more complex nonlinear reconstruction methods, which suffer from the same experimental issues. We also focus on optimization schemes that continuously optimize the positions of some sampling locations. These techniques have the advantage of not relying on a grid, which is an essential feature for various applications such as magnetic resonance imaging or radio-interferometry. In addition, they spark the hope of avoiding the curse of dimensionality encountered in combinatorial problems. We show that this dream is not realistic, but that the situation improves by considering large signals datasets and specific variable metric techniques. We conclude the paper by illustrating our findings on 1D experiments.

2 Notation

In this paper, we will focus on discrete 1D signals, for the ease of exposition. However, the main arguments apply to arbitrary dimensions and continuous signals as well.

We consider a signal u as a vector of \mathbb{C}^N with $N \in 2\mathbb{N}$. We let $\mathcal{N} = \llbracket -\frac{N}{2}, \frac{N}{2} - 1 \rrbracket$. An alternative way to represent a signal $u \in \mathbb{C}^N$ is to use a discrete

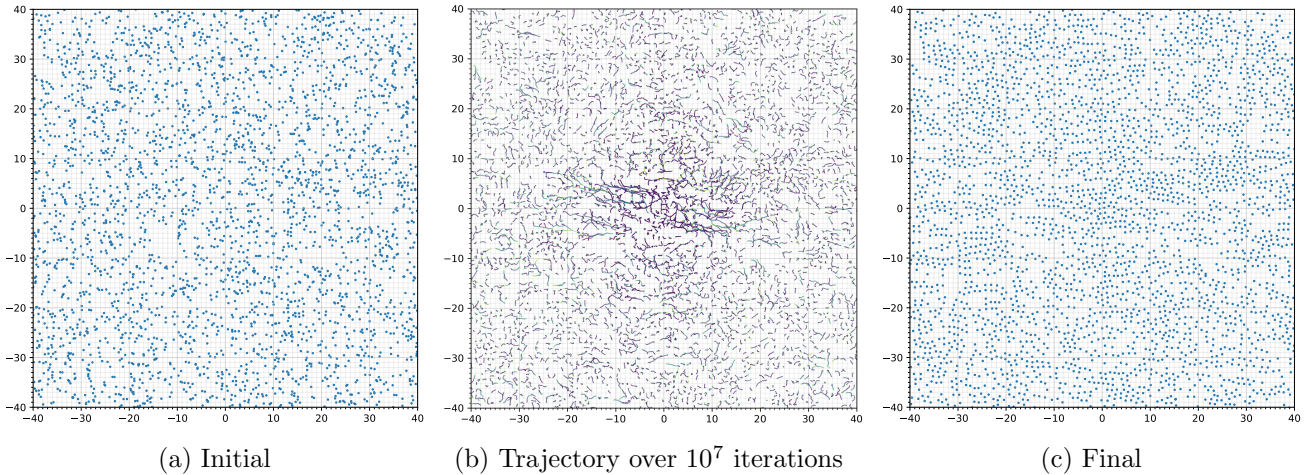


Figure 1: A typical sampling optimization trajectory. Starting from the sampling configuration on the left (uniform point process), we obtain the sampling scheme on the right after 10^7 iterations. The trajectory in the center corresponds to the 10^7 iterations of a gradient descent with fixed step size. Notice that the points clusters have disappeared, but that the scheme is still essentially uniform, while we would expect the low frequencies to be sampled more densely.

measure μ of the form:

$$\mu = \sum_{n \in \mathcal{N}} u_n \delta_{\frac{n}{N}}. \quad (1)$$

Given a location $\xi \in \mathbb{R}$, we define:

$$\hat{u}(\xi) \stackrel{\text{def}}{=} \frac{1}{\sqrt{N}} \sum_{n \in \mathcal{N}} u_n e^{-2i\pi \langle \xi, \frac{n}{N} \rangle}, \quad (2)$$

which can be seen as the continuous Fourier transform of the measure μ . We consider $\Xi = [\xi_1, \dots, \xi_M] \in \mathbb{R}^M$ a set of M locations. The Fourier transform $\hat{u}(\Xi) \in \mathbb{C}^M$ at the locations Ξ can be written as a matrix-vector product of the form $\hat{u}(\Xi) = A(\Xi)^* u$ with the normalized Vandermonde matrix $A(\Xi) \in \mathbb{C}^{N \times M}$ defined by

$$A(\Xi)_{n,m} \stackrel{\text{def}}{=} \frac{1}{\sqrt{N}} e^{2i\pi \langle \xi_m, \frac{n}{N} \rangle}.$$

In what follows, we let $a(\xi) \in \mathbb{C}^N$ denote the vector defined for all $n \in \mathcal{N}$ by

$$a(\xi)[n] \stackrel{\text{def}}{=} \frac{1}{\sqrt{N}} e^{2i\pi \langle \xi, \frac{n}{N} \rangle},$$

so that

$$A(\Xi) = [a(\xi_1), \dots, a(\xi_M)].$$

The matrix $A(\Xi)^*$ can be seen as the nonuniform Fourier transform [13] from the grid to the set of sampling locations Ξ . We let $(A(\Xi)^*)^+$ denote the pseudo-inverse of $A(\Xi)^*$.

3 Preliminaries

Below, we first describe the precise mathematical setting and then turn to some preliminary results.

3.1 The setting

Let $u \in \mathbb{C}^N$ denote a signal. We assume that a sampling device allows to pick M frequencies ξ_1, \dots, ξ_M in \mathbb{R} , yielding the set of measurements $y = A(\Xi)^* u + w$ with $w \sim \mathcal{N}(0, \sigma^2 \text{Id})$ a white Gaussian noise. A vast amount of reconstruction techniques have been designed in the literature to reconstruct u from y . A generic reconstructor can be defined as a mapping $\mathcal{R} : (\mathbb{C}^M \times \mathbb{R}^M) \rightarrow \mathbb{C}^N$ that takes as an input a measurement $y \in \mathbb{C}^M$ and a sampling scheme $\Xi \in \mathbb{R}^M$ and outputs a reconstructed signal $\mathcal{R}(y, \Xi)$. Given a collection of signals u_1, \dots, u_P and a reconstructor \mathcal{R} , a natural framework to find the best sampling scheme Ξ is to solve the following optimization problem:

$$\inf_{\Xi \in \mathbb{R}^M} \frac{1}{2P} \sum_{p=1}^P \mathbb{E}_w (\| \mathcal{R}(A(\Xi)^* u_p + w, \Xi) - u_p \|_2^2). \quad (3)$$

This problem can be attacked with first order methods that continuously optimize the sampling locations ξ_m , see for instance [17, 8, 16]. In this work, we will concentrate on three simple linear reconstruction methods of the form $\mathcal{R}(y, \Xi) = R(\Xi)y$:

The back-projection method which consists in defining the reconstructor as $R_1(\Xi) = A(\Xi)$ or

$$\mathcal{R}_1(y, \Xi) \stackrel{\text{def}}{=} A(\Xi)y. \quad (4)$$

The pseudo-inverse method where the reconstructor is defined with $R_2(\Xi) = (A(\Xi)^*)^+$ or

$$\mathcal{R}_2(y, \Xi) \stackrel{\text{def}}{=} (A(\Xi)^*)^+ y. \quad (5)$$

The Tikhonov method (or regularized inverse) which consists in solving the following quadratic problem:

$$\mathcal{R}_3(y, \Xi) \stackrel{\text{def}}{=} (1 + \lambda) \arg \min_{f \in \mathbb{C}^N} \frac{1}{2} \|A(\Xi)^* f - y\|_2^2 + \frac{\lambda}{2} \|f\|_2^2 \quad (6)$$

for $\lambda > 0$. Hence

$$R_3(\Xi) = (1 + \lambda) (A(\Xi)A(\Xi)^* + \lambda \text{Id})^{-1} A(\Xi). \quad (7)$$

The multiplication by $(1 + \lambda)$ is there to compensate the bias introduced by the regularization and will later simplify the expressions. A similar analysis can be carried out for the more standard solver $R_3(\Xi) = (A(\Xi)A(\Xi)^* + \lambda \text{Id})^{-1} A(\Xi)$, but it leads to significantly more complicated formulas, which we prefer avoiding for the sake of readability.

These techniques are quite popular in the actual practice. We restrict our analysis to linear reconstructors of the type (4), (5) and (6) for simplicity reasons. Note that (5) corresponds to the limit case of (6) when λ tends to zero. Numerical experiments reveal that the optimization issues raised in Theorems 1 and 2 also apply to nonlinear reconstructors such as sparsity promoting convex penalties. However, the techniques used in the proofs do not directly extend to this framework.

We first analyze the problem with a single image u in the dataset, i.e. $P = 1$. Let us define three cost functions J_1 , J_2 and J_3 which respectively correspond to the back-projection, the pseudo-inverse and the regularized inverse.

Definition 1 (Cost function). *Given a signal u , a sampling scheme Ξ and a reconstruction method $R(\Xi)$, the cost function reads*

$$J(\Xi) \stackrel{\text{def}}{=} \mathbb{E}_w \left(\frac{1}{2} \|R(\Xi)(A(\Xi)^* u + w) - u\|_2^2 \right) \quad (8)$$

where $w \sim \mathcal{N}(0, \sigma^2 \text{Id})$ is white Gaussian noise.

3.2 Elementary observations

We will make use of the following definitions.

Definition 2 (The min distance). *Given a set of sampling points Ξ , the min distance $\text{md}(\Xi)$ is defined by*

$$\text{md}(\Xi) \stackrel{\text{def}}{=} \min_{m \neq m'} \text{dist}(\xi_m, \xi_{m'})$$

where dist is the distance on the torus defined for $(\xi_1, \xi_2) \in \mathbb{R}^2$ as

$$\text{dist}(\xi_1, \xi_2) \stackrel{\text{def}}{=} \inf_{k \in \mathbb{Z}} \|\xi_1 - \xi_2 - kN\|_\infty. \quad (9)$$

Definition 3 (Subgrid). *Throughout the paper, we say that $\Xi \in [-N/2, N/2]^M$ is a subgrid if $\xi_m - \xi_{m'} \in \mathbb{Z}^*$ for all $m \neq m'$.*

Proposition 1 (J is N -periodic). *We have*

$$J(\Xi \bmod N) = J(\Xi). \quad (10)$$

Proof. Let $n = kN$ with $k \in \mathbb{N}$. The proof simply stems from the fact that $a(\xi + n) = a(\xi)$. \square

The previous proposition shows that we can restrict our attention to frequencies ξ belonging to the set $[-N/2, N/2[$.

Proposition 2 (Existence of minimizers). *For any $M \in \mathbb{N}$ and any $u \in \mathbb{C}^N$, there exists at least one minimizer of J on $[-N/2, N/2]^M$.*

Proof. We start by noticing that J is a C^∞ function since it is defined as a composition of C^∞ functions. Hence it is also continuous on $[-N/2, N/2]^M$. This yields the existence of at least one minimizer. \square

Now we proceed to a reformulation of the problem by rearranging the terms involved in the definition of J .

Proposition 3. *The reconstructors associated to J_1 , J_2 and J_3 defined in (8) can be expressed as*

$$R(\Xi) = A(\Xi)Q(\Xi) \quad (11)$$

(i.e. the solution lives in $\text{ran}(A)$) with:

- $Q_1(\Xi) = \text{Id}$
- $Q_2(\Xi) = (A(\Xi)^* A(\Xi))^+$
- $Q_3(\Xi) = (1 + \lambda)(A(\Xi)^* A(\Xi) + \lambda \text{Id})^{-1}$.

Proof. For Q_1 , there is nothing to prove. For Q_2 , we use one of the standard properties of the pseudo-inverse. For Q_3 , we use the equality $A(A^* A + \lambda \text{Id}) = (AA^* + \lambda \text{Id})A$ and then left multiply by $(AA^* + \lambda \text{Id})^{-1}$ and right multiply by $(A^* A + \lambda \text{Id})^{-1}$. \square

Proposition 4. *Letting $\hat{u}(\Xi) = A(\Xi)^* u$, we have*

$$J(\Xi) = \frac{1}{2} \|u\|_2^2 - \langle Q(\Xi)\hat{u}(\Xi), \hat{u}(\Xi) \rangle + \frac{1}{2} \|R(\Xi)\hat{u}(\Xi)\|_2^2 + \frac{1}{2} \mathbb{E}_w (\|R(\Xi)w\|_2^2) \quad (12)$$

Proof. We drop the dependency in Ξ to simplify the notation.

$$\begin{aligned} 2J &= \|u\|_2^2 + \|RA^* u\|_2^2 + \mathbb{E}_w (\|Rw\|_2^2) \\ &\quad + 2\mathbb{E}_w (\text{Re}\langle RA^* u - u, w \rangle) - 2\text{Re}\langle RA^* u, u \rangle \\ &= \|u\|_2^2 + \|RA^* u\|_2^2 + \mathbb{E}_w (\|Rw\|_2^2) - 2\langle Q\hat{u}, \hat{u} \rangle \end{aligned}$$

where we used $Q(\Xi)^* = Q(\Xi)$ and $\mathbb{E}_w(w) = 0$. \square

Equation (12) greatly simplifies when Ξ is a subgrid. Let us define the following function

$$\tilde{J}(\Xi) \stackrel{\text{def}}{=} \frac{1}{2}\|u\|_2^2 - \frac{1}{2}\|\hat{u}(\Xi)\|_2^2 + \frac{\sigma^2 M}{2}. \quad (13)$$

Proposition 5. *When Ξ is a subgrid, $J(\Xi) = \tilde{J}(\Xi)$.*

Proof. When Ξ is a subgrid, we have $A(\Xi)^*A(\Xi) = \text{Id}$ and we use the decomposition of Proposition 4 with $Q(\Xi) = \text{Id}$, $R(\Xi)^*R(\Xi) = \text{Id}$. \square

4 Theoretical issues

In this section, we give the main theoretical results of the paper.

4.1 Spurious minimizers

The aim of this section is to illustrate a common situation where the function J possesses a combinatorial number of minimizers. We construct examples where the function \tilde{J} defined in (13) is very oscillatory, while $J - \tilde{J}$ is of small amplitude. The function J is close to \tilde{J} not only for subgrids as in Proposition 5 but also for well-spread schemes. Following the proof of Proposition 5, and the decomposition of Proposition 4, it is sufficient to control how close $Q(\Xi)$ and $R(\Xi)^*R(\Xi)$ are to Id . This is the aim of the following proposition.

Proposition 6 (Bound on Q and R^*R). *Consider a sampling pattern Ξ such that $\text{md}(\Xi) > 1$ and set $\epsilon = 1/\text{md}(\Xi)$. Then*

$$-a_i \text{Id} \preceq Q_i - \text{Id} \preceq a_i \text{Id} \quad (14)$$

$$-b_i \text{Id} \preceq R_i^* R_i - \text{Id} \preceq b_i \text{Id}, \quad (15)$$

with

$$a_1 = 0, \quad a_2 = \frac{\epsilon}{1-\epsilon}, \quad a_3 = \frac{\epsilon}{1-\epsilon}$$

$$b_1 = \epsilon, \quad b_2 = \frac{\epsilon}{1-\epsilon}, \quad b_3 = \frac{4\epsilon}{(1-\epsilon)^2}$$

The proof is postponed to Section 7.1.

Theorem 1 (A combinatorial number of minimizers). *Set a number of samples $M \in \mathbb{N}$ and consider a vector $u \in \mathbb{C}^N$ such that the following properties are verified*

- i) *The modulus $|\hat{u}|^2$ possesses a subset of $K \geq M$ local maximizers $Z = \{\zeta_1, \dots, \zeta_K\}$ separated by a distance at least $\delta = \text{md}(Z)$ with $\delta > 1 + 2r$ for some $r > 0$.*

- ii) *The modulus $|\hat{u}|^2$ is locally strictly concave for each ζ_k :*

$$|\hat{u}|^2(\zeta_k + h) \leq |\hat{u}|^2(\zeta_k) - \frac{c}{2}h^2, \forall h \in [-r, r]$$

for some $c > 0$.

- iii) *For any subset $\bar{\Xi}$ of M distinct points in Z , we have*

$$\frac{cr^2}{2} > (b + 2a) \|\hat{u}(\bar{\Xi})\|_2^2 + bM\sigma^2 \quad (16)$$

where a and b are given in Proposition 6 with $\epsilon = \frac{1}{\delta - 2r}$.

Then, the function J possesses at least $\binom{K}{M} \cdot M!$ local minimizers.

The proof of Theorem 1 is postponed to Section 7.2. The conditions in Theorem 1 may look cryptic at first sight. We first show a simple example of a function u that verifies the hypotheses and leads to a huge number of critical points.

Corollary 1. *Assume that $N \in 4\mathbb{N}$ and define $u \in \mathbb{C}^N$ as follows*

$$u[n] = \begin{cases} \sqrt{N}/2 & \text{if } n = \pm N/4, \\ 0 & \text{otherwise.} \end{cases} \quad (17)$$

Let $M = \lfloor \eta\sqrt{N} \rfloor$ with $\eta = \frac{\pi^2\sqrt{2}}{256 \cdot (20 + 16\sigma^2)}$ then all the functions J_i possess a number of minimizers larger than $M! \cdot \left(\frac{1}{2\eta}\right)^{\eta\sqrt{N}}$.

For $\sigma \leq 1$, the bound holds for $\eta = 3 \cdot 10^{-3}$. For $\sigma = 0$ and J_1 the bound can be increased to $\eta = 1.09 \cdot 10^{-1}$.

Proof. The choice of u in (17) leads to the oscillatory function $\hat{u}(\xi) = \cos\left(\frac{\pi}{2}\xi\right)$. The modulus $|\hat{u}|$ is maximal at every point $\xi \in 2\mathbb{N}$. Let $\xi_0 \in 2\mathbb{N}$ and set $r = \frac{1}{4}$. For any $\xi \in [\xi_0 - r, \xi_0 + r]$, we have

$$\begin{aligned} (|\hat{u}|^2)''(\xi) &= \frac{\pi^2}{2} \left(\sin^2\left(\frac{\pi}{2}\xi\right) - \cos^2\left(\frac{\pi}{2}\xi\right) \right) \\ &\leq -\frac{\pi^2}{2\sqrt{2}}. \end{aligned}$$

Let $p \in \mathbb{N}$. The conditions i) and ii) of Theorem 1 are satisfied with $Z = 2p\mathbb{N} \cap [-N/2, N/2[$, $K = \lfloor N/(2p) \rfloor$, $r = 1/4$, $c = \frac{\pi^2\sqrt{2}}{8}$, $\delta = 2p$. Further notice that for every set $\bar{\Xi} \in Z^M$, $\|\hat{u}(\bar{\Xi})\|_2^2 = M$.

For this example, the condition (16) therefore reads

$$M < \frac{\pi^2\sqrt{2}}{256} \cdot \left(\frac{1}{b + 2a + b\sigma^2} \right). \quad (18)$$

As long as this condition is satisfied, Theorem 1 allows to conclude on the existence of $\binom{\lfloor N/2p \rfloor}{M} \cdot M!$ maximizers.

Now, if $\delta - 2r \geq 2$, we can coarsely simplify the bounds in Proposition 6 as

$$a \leq \frac{2}{\delta - 2r} \quad \text{and} \quad b \leq \frac{16}{\delta - 2r}.$$

Hence, a combinatorial number of minimizers is granted given that

$$M < \frac{\pi^2 \sqrt{2}}{256} \cdot \left(\frac{\delta - 2r}{20 + 16\sigma^2} \right). \quad (19)$$

Now, take $p = \lfloor \sqrt{N} \rfloor$ and $M = \lfloor \eta \cdot \sqrt{N} \rfloor$ with $\eta = \frac{\pi^2 \sqrt{2}}{128 \cdot (20 + 16\sigma^2)}$. Then Theorem 1 yields a number of minimizers larger than $\binom{\lfloor \sqrt{N}/2 \rfloor}{\lfloor \eta \cdot \sqrt{N} \rfloor} \cdot M!$. Using the standard bound

$$\binom{n}{k} \geq \left(\frac{n}{k} \right)^k \quad (20)$$

yields a number of minimizers larger than $\left(\frac{1}{2\eta} \right)^{\eta \sqrt{N}}$.

In particular for $\sigma < 1$ this yields $\eta = 0.003$. The bound can be increased to $\eta = 0.109$ for $\sigma = 0$ and J_1 . \square

4.2 Numerical illustration of Theorem 1

In this section we illustrate Theorem 1 through numerical examples in Fig. 2. We first consider the noiseless settings $\sigma = 0$ and illustrate the existence of spurious minimizers for the back-projection and the pseudo-inverse methods.

We introduce the following function

$$F(\Xi) \stackrel{\text{def}}{=} \frac{1}{2} \sum_{m=1}^M |\hat{u}(\xi_m)|^2 = \frac{1}{2} \|\hat{u}(\Xi)\|_2^2, \quad (21)$$

which somehow measures the energy captured within a sampling scheme Ξ . We also introduce the functions G_1 and G_2 such that $J_1 = \frac{1}{2} \|u\|_2^2 - F + G_1$ and $J_2 = \frac{1}{2} \|u\|_2^2 - F + G_2$. Using Proposition 4, we have

$$G_1(\Xi) = \frac{1}{2} \langle (A(\Xi)^* A(\Xi) - \text{Id}) \hat{u}(\Xi), \hat{u}(\Xi) \rangle, \quad (22)$$

and

$$G_2(\Xi) = \frac{1}{2} \langle (\text{Id} - (A(\Xi)^* A(\Xi))^+) \hat{u}(\Xi), \hat{u}(\Xi) \rangle. \quad (23)$$

From the left to the right, we used three different 1D signals: a high frequency cosine, a low frequency sine and a Gaussian. We plot the different

energy landscapes, for $M = 2$ measurements at locations $\Xi = \{\xi_1, \xi_2\}$ and $N = 16$. From the top to the bottom, we display the functions J_1 , J_2 , G_1 , G_2 , $-F$ and the modulus of the Fourier transform $\xi \mapsto |\hat{u}(\xi)|$. In order to understand the effect of the signal's structure, the local minima of J_1 , J_2 , G_1 , G_2 and $-F$ are represented with red dots.

First notice that the cost functions are symmetric with respect to the diagonal. This simply reflects the fact that permutation of points lead to the same energy, and this illustrates the factor $M!$ in Theorem 1.

As can be seen in all cases, the functions G_1 and G_2 vanish far away from the diagonal (see Proposition 6). These point configurations correspond to well-spread sampling schemes. On the contrary, the function $-F$ can have a large amplitude even outside the diagonal. These two properties are the main ingredients to prove Theorem 1.

The left column (high frequency cosine), corresponds to the example in Corollary 1. We see a number of minimizers that seems quadratic in N for $M = 2$. The center column (low frequency sine) shows that the number of minimizers decreases with a higher regularity of the signal, by reducing the oscillations in F . On the right (Gaussian function), we illustrate a case where F has only one local maximum. Even in this case, the function J still has valleys with shallow local minima. The same phenomenon appears in the center (low frequency sine). Notice that this phenomenon is not captured by Theorem 1, which only relies on local maximizers of F . In these two examples, the oscillations are induced by the function G , which we do not explore in this paper.

In Fig. 3, the energy profile of J_3 is displayed with the low frequency signal (see Fig. 2 center column) for $M = 2$ and for various noise levels σ and regularization parameters λ .

Neither the noise, nor the regularization parameter λ are able to remove the local minimizers of J_3 which are displayed by red dots. The last column is a critical case where the noise prevails over the signal and the reconstruction error is high (typically ~ 0.18 in the noiseless setting and ~ 0.5 with $\sigma = 5 \times 10^{-1}$).

4.3 Flatness for high frequencies

In this paragraph we show that the partial derivatives of the cost function may vanish, for indexes corresponding to high frequencies. This explains another practical difficulty in Fourier sampling optimization: without using variable metric techniques, the sampling points located in the high frequencies

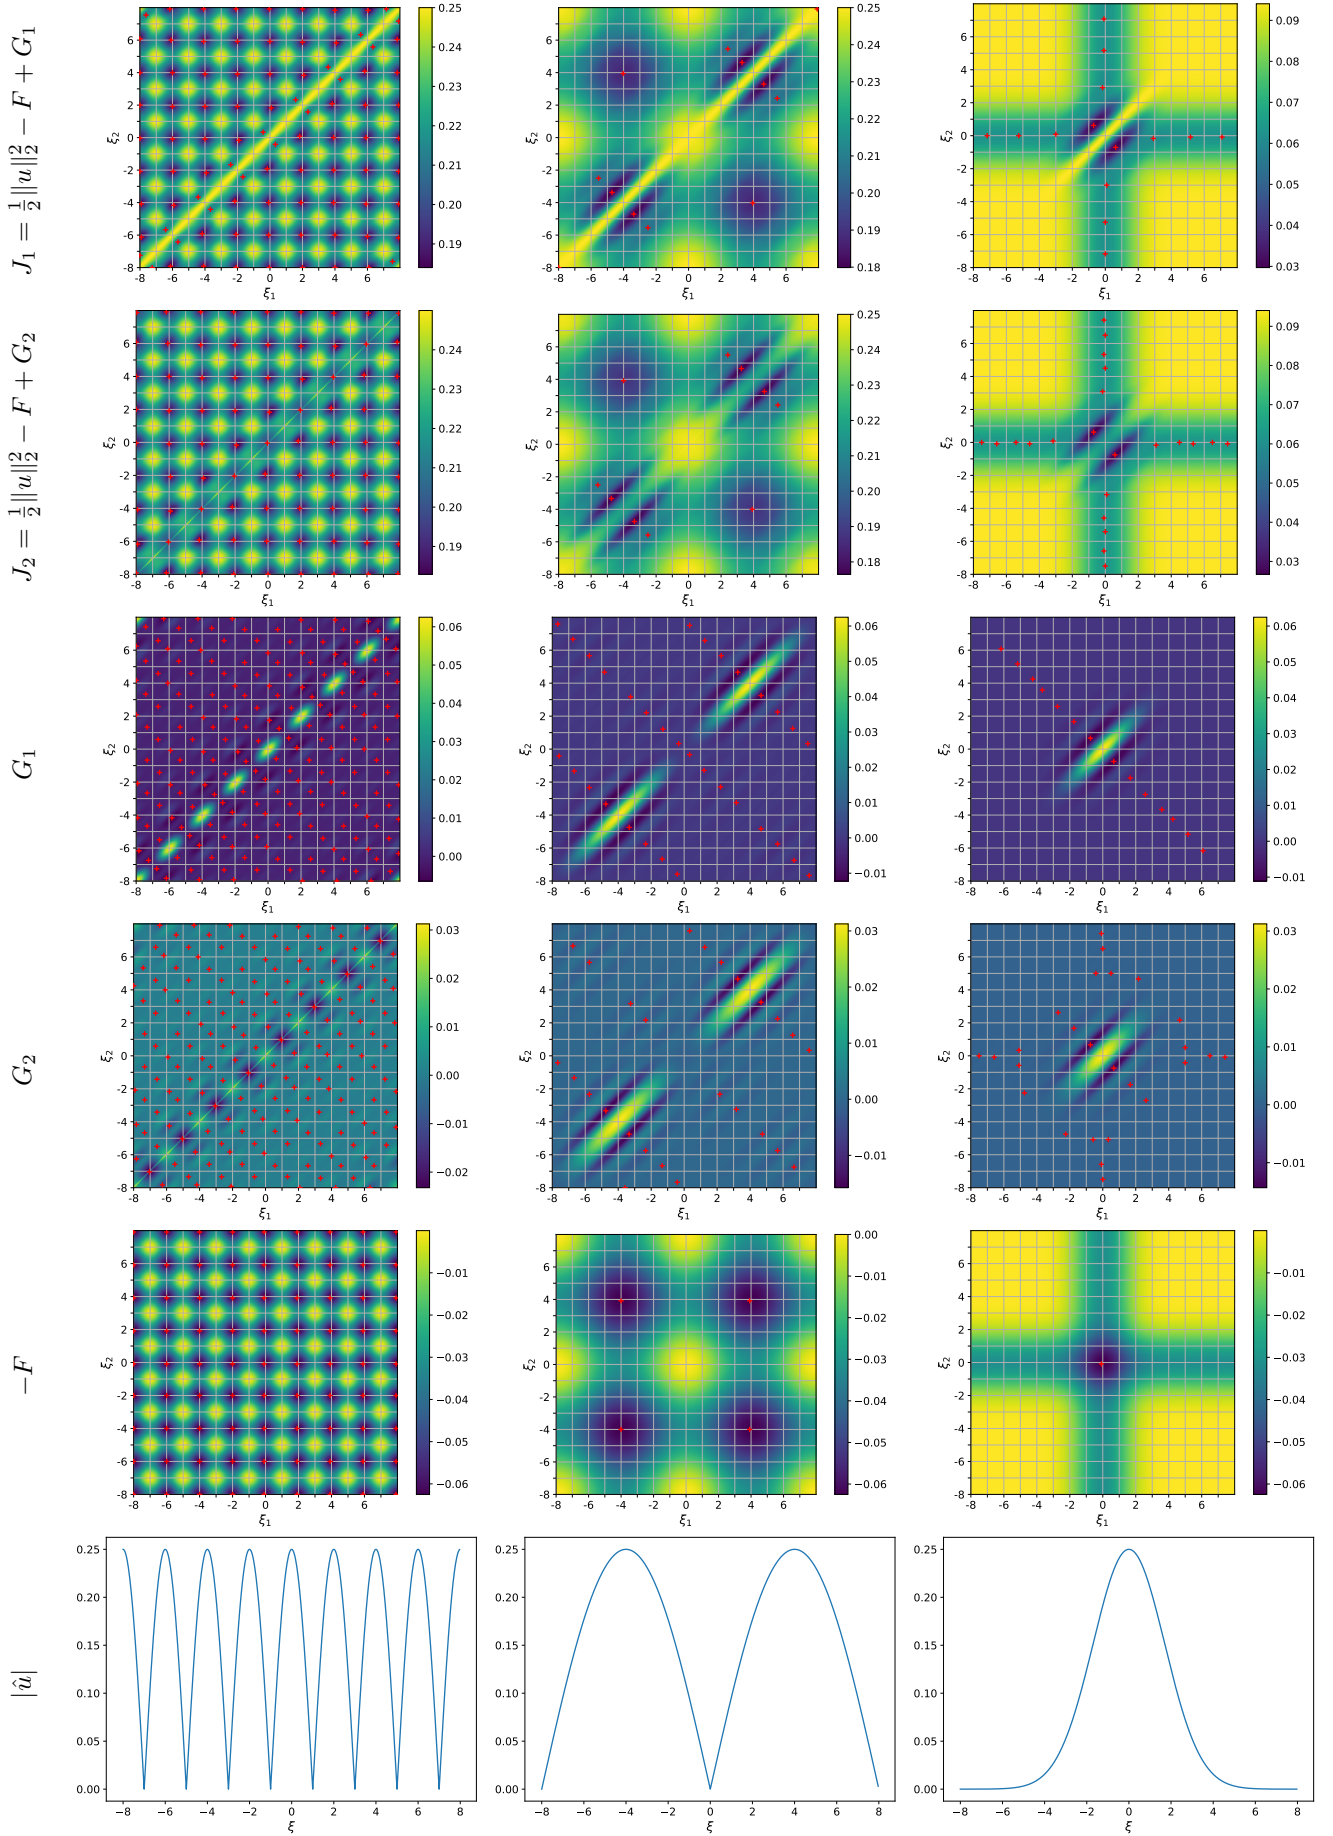


Figure 2: The energy profile for $M = 2$ and three different signals \hat{u} : a high frequency cosine, a low frequency sine and a Gaussian (from left to right). From top to bottom, we represent J_1 , J_2 , G_1 , G_2 , F and $|\hat{u}|$. The red dots represent local minima.

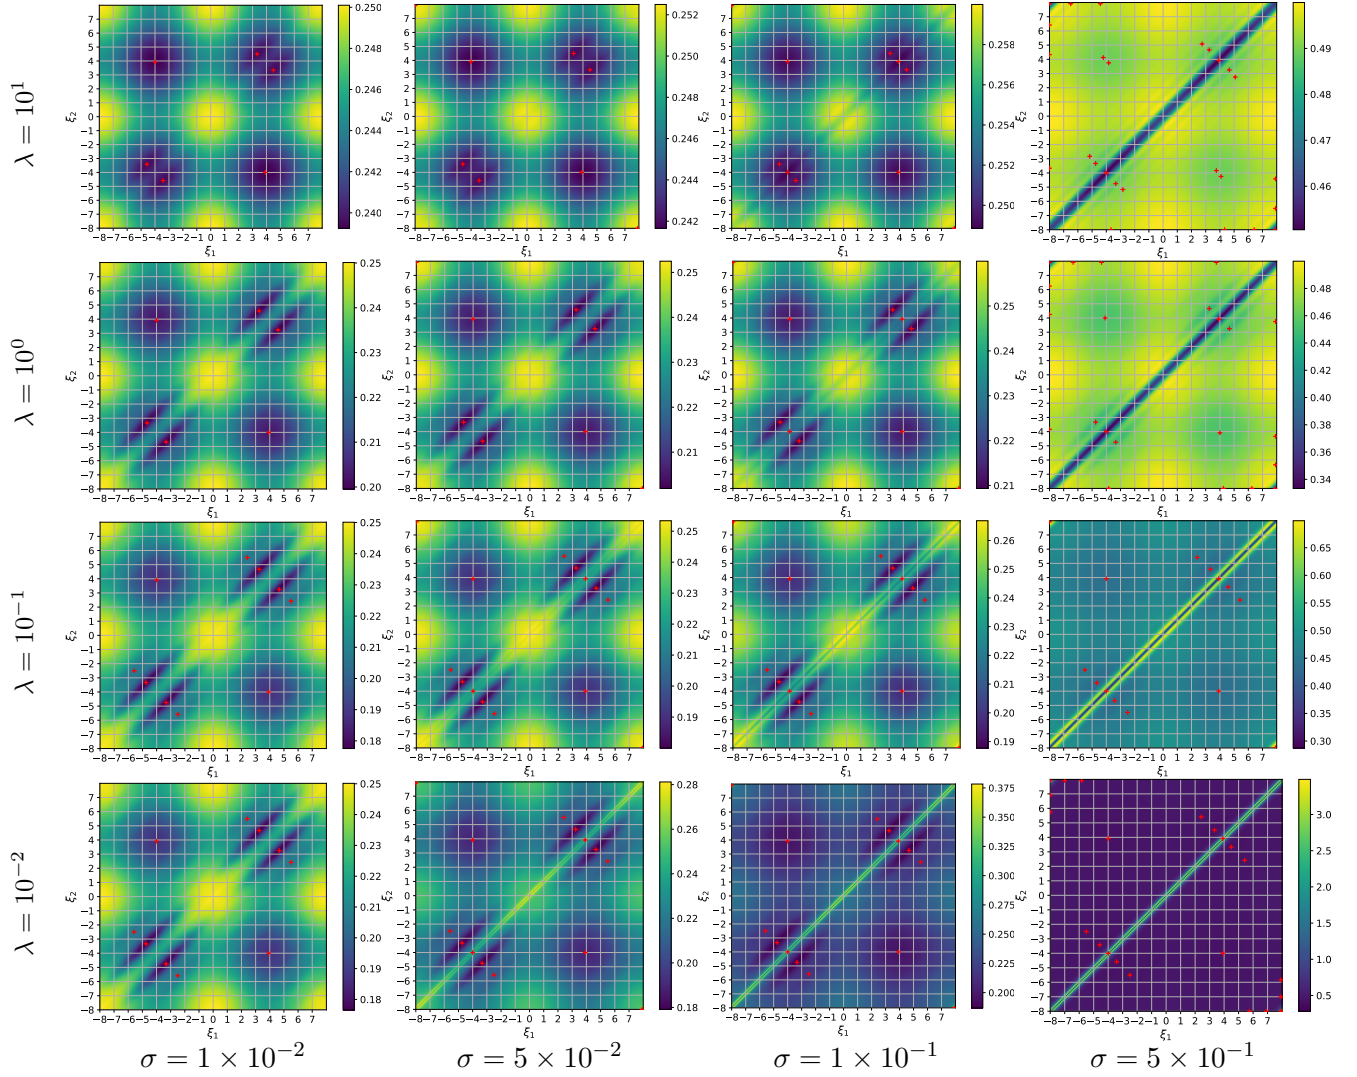


Figure 3: The energy profile of J_3 for $M=2$ with the low frequency signal (see center column of Fig. 2) for different noise levels σ and regularization λ . The red dots represent local minima.

move very slowly. Though our proof only applies to the function J_1 , this effect also seems to occur for J_2 . See for instance the four corners of Fig. 2, right.

Proposition 7. *Letting r denote the residual error function*

$$r(\Xi) = A(\Xi)A(\Xi)^*u - u, \quad (24)$$

the gradient of the cost function J_1 reads:

$$\nabla J_1(\Xi) = \text{Re} \left(\nabla \left(\hat{u}(\Xi) \odot \overline{\hat{r}(\Xi)} \right) \right), \quad (25)$$

where ∇ in the right-hand-side denotes the usual derivative in 1D or the gradient in higher dimension and where \odot is the coordinate-wise (Hadamard) product.

The proof of Proposition 7 is postponed to Section 7.3.

Theorem 2 (Vanishing gradients for high frequencies). *Consider a signal $u \in \mathbb{C}^N$ and a point configuration $\Xi \in \mathbb{R}^M$. Under the decay assumptions*

$$|\hat{u}(\xi)| \lesssim \frac{1}{|\xi|^\alpha} \quad \text{and} \quad |\hat{u}'(\xi)| \lesssim \frac{1}{|\xi|^\alpha}, \quad (26)$$

with $\alpha > 0$, we have

$$\left| \frac{\partial J_1(\Xi)}{\partial \xi_m} \right| \lesssim \frac{\|\hat{u}(\Xi)\|_1}{\text{md}(\Xi)|\xi_m|^\alpha}. \quad (27)$$

The decay assumption appear naturally in the continuous setting, when considering signals u from Sobolev spaces H^k with k derivatives in L^2 .

5 Escaping the minimizers

In this section we propose some solutions to mitigate the issues raised in Section 4 and we illustrate them numerically.

5.1 The effect of using a large dataset

In Theorem 1, we proved existence of many local minimizers in the case $P = 1$, which corresponds to a unique signal. Let us now assume that we have access to P signals u_1, \dots, u_P in \mathbb{C}^N . The analysis carried out to prove Theorem 1 can be replicated verbatim. The only difference being that every occurrence of $|\hat{u}|^2$ must be replaced by $\rho_P \stackrel{\text{def}}{=} \frac{1}{P} \sum_{p=1}^P |\hat{u}_p|^2$. The function ρ_P can be understood as the average power spectral density of the family u_1, \dots, u_P . As highlighted in Theorem 1, two important factors that can create spurious minimizers are i) the number K of strict maximizers of ρ_P and ii) the curvature c at these maximizers.

As P increases, we typically expect the density ρ_P to become smoother. This effect is illustrated for a simple family of shifted and dilated rectangular functions in Fig. 4. As can be seen, both the number of maxima and the curvature c of ρ_P in Theorem 1 decay with P . For $N = 128$, we display the average power spectral density for P ranging from 1 to 10^3 . Each signal is defined by

$$u[n] = \int_{n-\frac{1}{2}}^{n+\frac{1}{2}} \mathbb{1}_{[a,b]}(x) dx, \quad (28)$$

where a and b are drawn uniformly in the range $[-N/2 + 1, N/2 - 1]$. The discrete signals are then renormalized so that $\|u\|_2 = 1$.

The same experiment can be reproduced in a more relevant framework from a practical viewpoint. The average power spectral density for 2D knee images of the fastMRI database [18] are represented in Fig. 5. The image are of size 320×320 . The local maximizers are computed and displayed with red dots in Fig. 5. In that case, increasing the family size P reduces the number of maximizers at a slow rate. Indeed they slightly increase from 13k points in the case $P = 1$ to 14k in the case $P = 100$ and then start to decrease to 11k for $P = 10000$. However, the curvature c decays much faster. As a conclusion, we see that *using large families of signals can reduce asymptotically the number and the size of the basins of attraction of some spurious minimizers.*

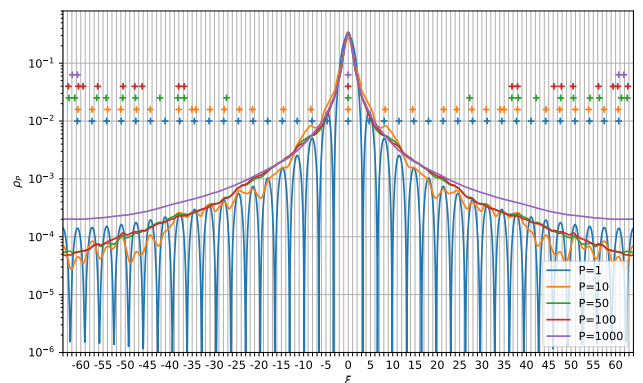


Figure 4: Average power spectral density ρ_P for families of rectangular functions with different sizes P . The dots represent local maxima of ρ_P for different values of P .

5.2 Stochastic gradient descent

When using a large family of signals, the cost function (3) naturally lends itself to the use of stochastic gradient descents (SGD), see [16, 17] that address

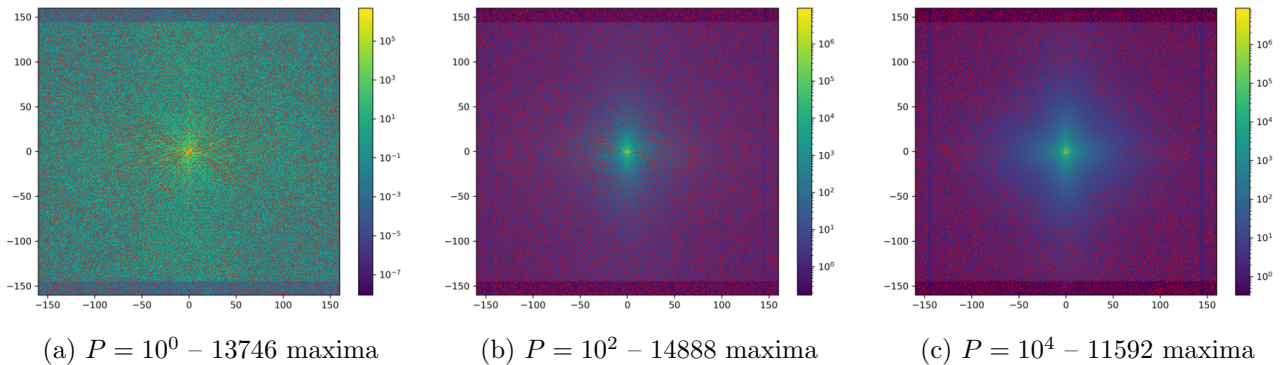


Figure 5: Average power spectral density ρ_P for a subset of images from the knee dataset of fastMRI. The image size is $N = 320$ and the red dots represent local maximizers.

large MRI datasets. Contrarily to a deterministic gradient descent, which is known to converge to critical points under mild regularity conditions, the stochastic gradient with a fixed step size does not converge. The method is known to end up frolicking in the neighborhood of local critical points [6]. The radius of the neighborhood depends on the stochastic gradient variance and on the step-size. Intuitively, *using stochastic gradients algorithms should therefore allow escaping local minimizers*. We will showcase this effect in the forthcoming numerical experiments.

5.3 Variable metric

In Section 4.3, Theorem 2 states that the gradient of J_1 might vanish in the high frequency domain. Using second order information is a well known remedy to mitigate this effect. In this work, we propose a simple method which corresponds to a variable diagonal metric with well-chosen coefficients.

As shown in Theorem 2, the gradient vanishes with a rate depending on the Fourier transform magnitude $|\hat{u}|$. For a dataset, this decay is somewhat captured by the average power spectral density $\rho_P(\xi) \stackrel{\text{def}}{=} \frac{1}{P} \sum_{p=1}^P |\hat{u}_p(\xi)|^2$. Hence, we propose to compute ρ_P once and for all on a fine grid ($20 \times N$ discretization points in our example). The function ρ_P is then linearly interpolated in between the grid points during the gradient descent. At each gradient iteration we replace $\frac{\partial J_1(\Xi)}{\partial \xi_m}$ by

$$\frac{1}{\rho_P(\xi_m)^\beta} \frac{\partial J_1(\Xi)}{\partial \xi_m}, \quad (29)$$

where β is a constant that has to be set empirically. From numerical experiments $\beta \in [1, 2]$ shows good performance. In all the experiments presented hereafter we use $\beta = 1$. We will see later in the numerical experiments, that *this variable metric significantly accelerates the convergence for sampling*

points located in high frequencies.

5.4 Numerical illustrations

In this section, we aim at illustrating numerically the different results established previously. We aim at reconstructing 1D signals of size $N = 128$ from $M = 64$ measurements in the Fourier domain. We suppose that P rectangular signals generated using (28) are given. We illustrate our findings with the back-projection reconstructor associated to the cost function J_1 , but similar results have been obtained with the pseudo-inverse. As we are working in 1D with small dimensions N and M , at each iteration, the whole matrix $A(\Xi)^*$ is evaluated and the gradient ∇J_1 is computed directly from the analytic expression (25). We first use a fixed step gradient descent algorithm in order to showcase the convergence dynamics of the algorithm. The initialization of Ξ is a subgrid with a constant spacing of 2. The following experiments are conducted:

Effect of the dataset size P We first vary the number of signals by taking $P = 1$ and $P = 1000$. The evolution of Ξ is displayed in Fig. 6, respectively top-left and top-center. The history of the cost function is given in Fig. 7. For this experiment, we expect that a good sampling scheme consists of low frequencies sampled at the Shannon-Nyquist rate. In this regard, the sampling scheme obtained in Fig. 6 for $P = 1000$ is more satisfactory than the one obtained for $P = 1$. In the case $P = 1000$, the displacement of Ξ is more important, suggesting that some local minima have been discarded.

Variable metric We then study, for $P = 1000$ the effect of a variable metric gradient descent as described in Section 5.3. We also compare this approach to an L-BFGS algorithm [7] with a line search and with a Hessian estimated using the last

8 gradients. In Fig. 6, the usual gradient algorithm is at the top-center, the variable metric gradient descent is at the bottom-center and the L-BFGS algorithm is at the bottom-left. The cost function evolution is displayed in Fig. 7. Using a variable metric results in a huge speed-up of the algorithm. This is particularly visible for points ξ located at high frequencies, which is another illustration of Theorem 2. For this example, the L-BFGS algorithm converges slightly faster than the variable metric gradient descent in the early iterations. However, its per-iteration cost is much higher since it uses a line search and a non diagonal metric. Since the L-BFGS algorithm can be seen as a state-of-the-art quasi-Newton method, the proposed empirical metric (29) seems remarkably efficient.

Stochastic gradient descent Finally in Fig. 6 right column, we investigate the use of a fixed-step stochastic gradient descent algorithm with a batch size of 1. In that experiment, a new random signal is generated at every iteration using the model (28) and the stochastic gradient is computed with respect to that signal only. The trajectory of the vanilla SGD is comparable with the one obtained using a deterministic gradient descent for $P = 1000$ in Fig. 6 top-center. The variable metric trick significantly improves the convergence speed and more importantly, the final points configuration. As a conclusion, the variable metric SGD algorithm seems to be able to escape spurious minimizers and to take advantage of the averaging effect of the large dataset without the struggle of computing the gradient over a large dataset.

Comparison of the sampling schemes The final sampling schemes are not directly comparable in terms of cost function because the objective function is computed over different datasets. In Table 1, we therefore report the cost function computed on a specific set of signals. This set contains the $P = 1000$ signals that are used in the numerical illustrations of Fig. 6 center column. When tested against a large dataset, the final configuration obtained for $P = 1$ seems highly sub-optimal. This effect is most likely due to a convergence to a local minimizer and also to the fact that the sampling scheme is not adapted to a whole family but only to a single signal. The remarkable observation that can be made from Table 1 is that the optimal configuration obtained with the variable metric SGD performs better on the dataset of $P = 1000$ signals than the experiment conducted in Fig. 6 which is tailored for this dataset. This shows that the

the usual deterministic algorithms are stuck in local minima even with large datasets. On the contrary, the variable metric SGD algorithm seems effective.

These numerical results highlight the effectiveness of the different tricks suggested in this section: the use of a variable metric to handle high frequencies and a stochastic optimization to avoid local minima.

6 Conclusion

We highlighted two obstacles to the convergence of gradient based algorithms for Fourier sampling schemes optimization. The first one is a high number of local minimizers and the second one is a vanishing gradient phenomenon for high frequencies. As far as we know, this is the first theoretical study explaining why optimizing sampling patterns with modern automatic differentiation tools might result in algorithms being stucked at unsatisfactory locations. We also proposed three approaches to mitigate these effects. First, the number of spurious minimizers, the width and the depth of their basins of attraction can be reduced by considering large databases of signals. This acts as a regularization by averaging. Second, the vanishing gradient effect can be attacked with variable metric gradient descents. Finally, the use of a stochastic gradient instead of a deterministic gradient approach seems to allow escaping the narrow basins in a simplified 1D setting. These remarks may help explaining why the recent approaches in the literature based on the Adam optimizer manage to slightly improve the sampling pattern efficiency. Our work suggests that increasing the database sizes may help further easing the numerical resolution of the sampling pattern optimization by further smoothing the energy profiles. Many state-of-the-art reconstructors are based on a quadratic data fidelity term and we expect that some of the techniques used in this paper in the linear case can be reused even in a nonlinear setting. This is left for future research.

Acknowledgment

The authors acknowledge a support from ANR JCJC Optimization on Measures Spaces, ANR-17-CE23-0013-01. P.W and F.G. acknowledge the ANR-3IA Artificial and Natural Intelligence Toulouse Institute. This work was performed using HPC resources from GENCI-IDRIS (Grant 2021-AD011012210R1).

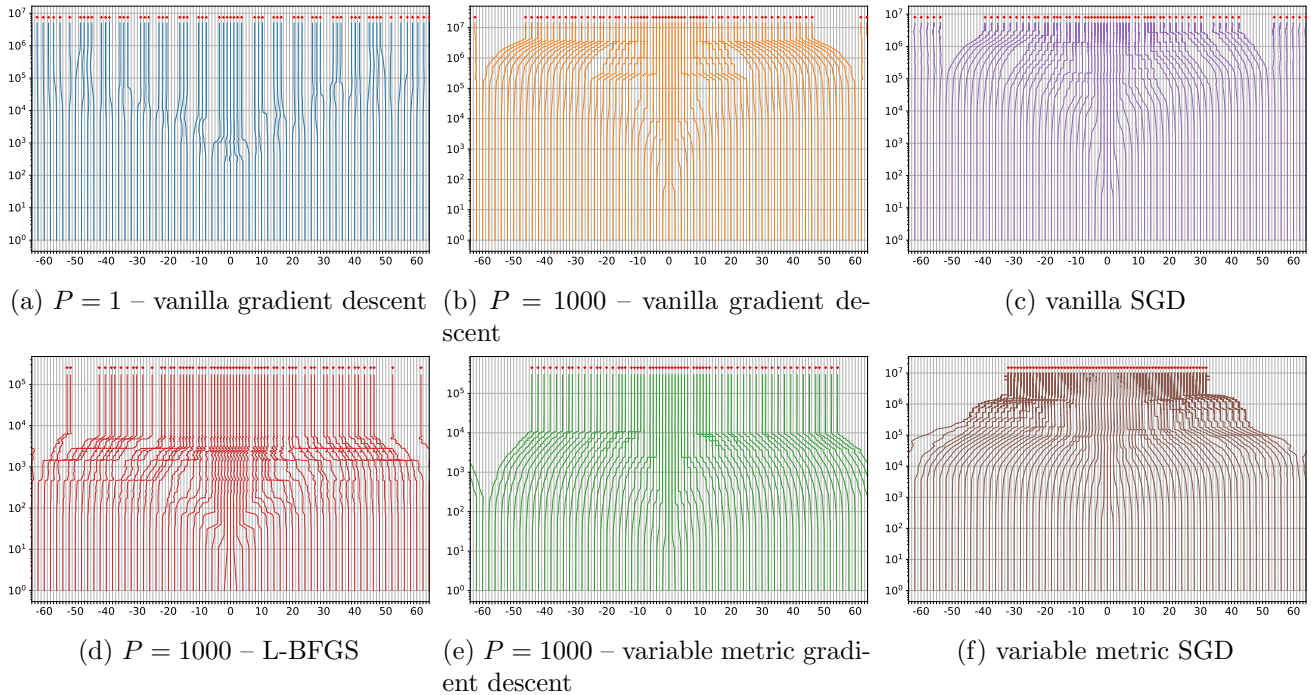


Figure 6: Trajectories of Ξ the back-projection reconstructor J_1 and a fixed-step gradient descent. The iterations are represented on the vertical axis, and the horizontal axis corresponds to ξ and is periodic. The initialization is a uniform subgrid and is seen on the axis $y = 0$ of the top and middle figures. Left and center: trajectories of Ξ for different sizes of signals families. The objective function is given in Fig. 7. The right column represents trajectories of Ξ using a stochastic gradient descent with one signal in the batch that is different at each iteration. The trajectories in the stochastic case have been averaged over the last 10000 iterations.

| Test case | $P = 1$ | $P = 1000$ | $P = 1000$ with var. metric | L-BFGS | SGD | SGD with var. metric |
|-----------|-----------------------|-----------------------|--------------------------------|-----------------------|-----------------------|-------------------------|
| Eff. | 9.07×10^{-2} | 2.68×10^{-2} | 2.38×10^{-2} | 2.41×10^{-2} | 6.63×10^{-2} | 1.00×10^{-2} |

Table 1: Effectiveness of the sampling schemes obtained with different strategies on a dataset of 1000 signals. The table contains the average reconstruction error J_1 over the dataset. This dataset is the one used in the case $P = 1000$, see Fig. 6 center column.

7 Proofs

7.1 Proof of Proposition 6

Significant progress have been made lately in the control of the extreme eigenvalues of Vandermonde matrices, which play a pivotal role in algebraic number theory [5, 12, 4, 2]. The tightest results for well separated schemes was recently obtained in [2]. Rewriting their result in our formalism, we obtain the following inequality.

Proposition 8 (Conditioning of Vandermonde matrices [2]). *Let $\Xi = (\xi_1, \dots, \xi_M)$ denote a set of distinct sampling points. The following inequalities hold*

$$\left(1 - \frac{1}{\text{md}(\Xi)}\right) \text{Id} \preceq A(\Xi)^* A(\Xi) \preceq \left(1 + \frac{1}{\text{md}(\Xi)}\right) \text{Id} \quad (30)$$

Proof. Relation (30) is a direct consequence of [2, eq. (31)] up to renormalizations. \square

Proof. For $i = 1$, recall that $Q_1 = \text{Id}$ and $R_1^* R_1 - \text{Id} = A^* A$. Then (14) is trivial and (15) follows from Proposition 8.

Let τ_m denote the eigenvalues of $A(\Xi)^* A(\Xi)$. By Proposition 8, $|\tau_m - 1| \leq \epsilon < 1$.

For $i = 2$, $R_2^* R_2 = Q_2^* A^* A Q_2 = (A^* A)^+$, $Q_2 = (A^* A)^+$. With $\epsilon < 1$, $A(\Xi)^* A(\Xi)$ is invertible and we have

$$\frac{1}{1 + \epsilon} \text{Id} \preceq (A(\Xi)^* A(\Xi))^{-1} \preceq \frac{1}{1 - \epsilon} \text{Id}. \quad (31)$$

And we finally get

$$\frac{-\epsilon}{1 + \epsilon} \text{Id} \preceq (A(\Xi)^* A(\Xi))^{-1} - \text{Id} \preceq \frac{\epsilon}{1 - \epsilon} \text{Id}. \quad (32)$$

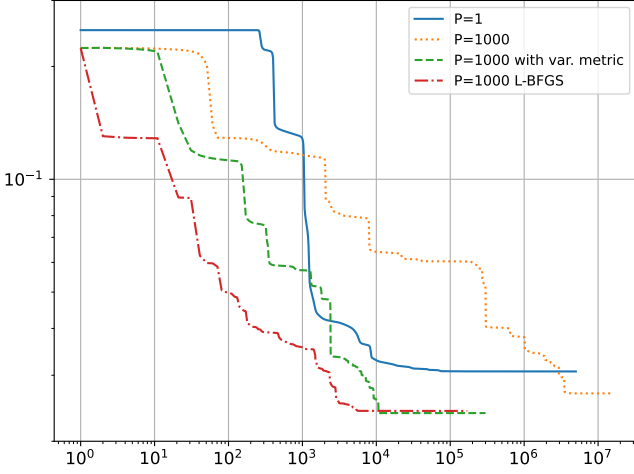


Figure 7: Objective function J_1 (back-projection) for the different experiments in Fig. 6 in the deterministic case.

For $i = 3$, $Q_3 = (1 + \lambda)(A^*A + \lambda \text{Id})^{-1}$, so that

$$\frac{1 + \lambda}{1 + \lambda + \epsilon} \text{Id} \preceq Q_3 \preceq \frac{1 + \lambda}{1 + \lambda - \epsilon} \text{Id}.$$

This gives

$$\frac{-\epsilon}{1 + \lambda + \epsilon} \text{Id} \preceq Q_3 - \text{Id} \preceq \frac{\epsilon}{1 + \lambda - \epsilon} \text{Id}$$

and using that $\epsilon > 0$ and $\lambda \geq 0$ allows to conclude.

In order to prove (15), Proposition 8 yields $1 - \epsilon \leq \tau_m \leq 1 + \epsilon$. In addition,

$$R_3^* R_3 = (1 + \lambda)^2 (A^*A + \lambda \text{Id})^{-1} A^*A (A^*A + \lambda \text{Id})^{-1}$$

can be diagonalized and its eigenvalues are therefore of the form

$$(1 + \lambda)^2 \frac{\tau_m}{(\tau_m + \lambda)^2}.$$

By taking the upper-bound on the numerator and the lower-bound on the denominator, we obtain the following bound:

$$R_3^* R_3 \preceq (1 + \lambda)^2 \frac{1 + \epsilon}{(1 - \epsilon + \lambda)^2} \text{Id}.$$

We can continue as follows:

$$\begin{aligned} & (1 + \lambda)^2 \frac{1 + \epsilon}{(1 - \epsilon + \lambda)^2} \\ &= (1 + \lambda)^2 \frac{1 + \epsilon}{(1 + \lambda)^2 (1 - \epsilon / (1 + \lambda))^2} \\ &\leq \frac{1 + \epsilon}{(1 - \epsilon)^2}. \end{aligned}$$

By a similar reasoning with respect to the smallest eigenvalue of $R_3^* R_3$, we get:

$$\frac{1 - \epsilon}{(1 + \epsilon)^2} \text{Id} \preceq R_3^* R_3 \preceq \frac{1 + \epsilon}{(1 - \epsilon)^2} \text{Id}.$$

Subtracting the identity on both sides and using the fact that $\epsilon^2 < \epsilon$ since $\epsilon < 1$ yields

$$-\frac{4\epsilon}{(1 - \epsilon)^2} \text{Id} \preceq R_3^* R_3 - \text{Id} \preceq \frac{4\epsilon}{(1 - \epsilon)^2} \text{Id}.$$

□

7.2 Proof of Theorem 1

Under the hypotheses of Theorem 1, first notice that any set $\Xi \in Z^M$ is a local maximizer of $\Xi \mapsto \|\hat{u}(\Xi)\|_2^2$. Indeed any perturbation of the individual sampling locations ξ_m results in a decay of the captured energy.

There are $\binom{K}{M}$ possible sampling configurations when all the points belong to Z . Let $\bar{\Xi} = \{\bar{\xi}_1, \dots, \bar{\xi}_M\}$ denote one of them. The idea of the proof is to show that there is a local minimizer of J in the following neighborhood $B = [\bar{\xi}_1 - r, \bar{\xi}_1 + r] \times \dots \times [\bar{\xi}_M - r, \bar{\xi}_M + r]$. A sufficient condition for the set B to contain a local minimizer of J is that $J(\bar{\Xi}) < J(\Xi)$ for all $\Xi \in \partial B$ (the boundary of B) since J is continuous. Throughout this proof, we use $\epsilon = \frac{1}{\delta - 2r}$ since we always have for all Ξ under consideration, $\text{md}(\Xi) \geq \delta - 2r$.

Using the bounds of Proposition 6, we obtain

$$\left| J(\Xi) - \tilde{J}(\Xi) \right| \leq \left(\frac{b}{2} + a \right) \|\hat{u}(\Xi)\|_2^2 + \frac{b}{2} \sigma^2 M. \quad (33)$$

For all $\Xi \in \partial B$, at least one index m must verify $\bar{\xi}_m - \xi_m = r$ and we have by strict concavity of $|\hat{u}|$ around $\bar{\xi}_m$

$$\|\hat{u}(\bar{\Xi})\|_2^2 - \|\hat{u}(\Xi)\|_2^2 \geq \frac{cr^2}{2} \quad (34)$$

$$\|\hat{u}(\bar{\Xi})\|_2^2 + \|\hat{u}(\Xi)\|_2^2 \leq 2\|\hat{u}(\bar{\Xi})\|_2^2. \quad (35)$$

Hence for $\Xi \in \partial B$, using (34) yields

$$\tilde{J}(\Xi) - \tilde{J}(\bar{\Xi}) \geq \frac{cr^2}{2} \quad (36)$$

Combining the previous inequalities yields

$$\begin{aligned} & J(\Xi) - J(\bar{\Xi}) \\ &= J(\Xi) - \tilde{J}(\Xi) + \tilde{J}(\Xi) - \tilde{J}(\bar{\Xi}) + \tilde{J}(\bar{\Xi}) - J(\bar{\Xi}) \\ &\stackrel{(36)}{\geq} \frac{cr^2}{2} + J(\Xi) - \tilde{J}(\Xi) + \tilde{J}(\bar{\Xi}) - J(\bar{\Xi}) \\ &\stackrel{(33)}{\geq} \frac{cr^2}{2} - \left[\left(\frac{b}{2} + a \right) (\|\hat{u}(\Xi)\|_2^2 + \|\hat{u}(\bar{\Xi})\|_2^2) + bM\sigma^2 \right] \\ &\stackrel{(35)}{\geq} \frac{cr^2}{2} - [(b + 2a) \|\hat{u}(\bar{\Xi})\|_2^2 + bM\sigma^2] \end{aligned}$$

Therefore, the condition

$$\frac{cr^2}{2} > (b + 2a) \|\hat{u}(\bar{\Xi})\|_2^2 + bM\sigma^2 \quad (37)$$

suffices to conclude on the existence of a maximizer of J in the interior of B . The multiplicative factor $M!$ is related to the fact that for a given maximizer, all the possible permutations of indices give rise to different maximizers.

7.3 Proof of Proposition 7

Proof. Let us consider a point configuration $\Xi \in \mathbb{R}^M$ and a perturbation $\epsilon \in \mathbb{R}^M$. Given a vector of measurements $\hat{u}(\Xi) \in \mathbb{C}^M$, we let $\nabla \hat{u}(\Xi) = \begin{pmatrix} \hat{u}'(\xi_1) \\ \vdots \\ \hat{u}'(\xi_M) \end{pmatrix}$ denote the vector of derivatives at the sampling locations. Elementary calculus leads to the following identities for every ϵ direction of variation:

$$\begin{aligned} (\text{Jac}_A(\Xi)\epsilon)^* &= \text{Jac}_{A^*}(\Xi)\epsilon \\ \nabla \hat{u}(\Xi) \odot \epsilon &= \text{Jac}_{A^*}(\Xi)\epsilon \cdot u. \end{aligned}$$

Then, we apply standard calculus of variations:

$$\begin{aligned} J_1(\Xi + \epsilon) &= J_1(\Xi) + \text{Re}\langle \text{Jac}_A(\Xi)\epsilon \cdot \hat{u}(\Xi), r(\Xi) \rangle \\ &\quad + \text{Re}\langle A(\Xi)\text{Jac}_{A^*}(\Xi)\epsilon \cdot u, r(\Xi) \rangle + o(\|\epsilon\|_2^2) \\ &= J_1(\Xi) + \text{Re}\langle \hat{u}(\Xi), (\text{Jac}_A(\Xi)\epsilon)^* r(\Xi) \rangle \\ &\quad + \text{Re}\langle \nabla \hat{u}(\Xi) \odot \epsilon, \hat{r}(\Xi) \rangle + o(\|\epsilon\|_2^2) \\ &= J_1(\Xi) + \text{Re}\langle \hat{u}(\Xi), \nabla \hat{r}(\Xi) \odot \epsilon \rangle \\ &\quad + \text{Re}\langle \epsilon, \overline{\nabla \hat{u}(\Xi)} \odot \hat{r}(\Xi) \rangle + o(\|\epsilon\|_2^2) \\ &= J_1(\Xi) + \text{Re}\langle \overline{\nabla \hat{r}(\Xi)} \odot \hat{u}(\Xi), \epsilon \rangle \\ &\quad + \text{Re}\langle \epsilon, \overline{\nabla \hat{u}(\Xi)} \odot \hat{r}(\Xi) \rangle + o(\|\epsilon\|_2^2). \end{aligned}$$

Hence, by identification

$$\begin{aligned} \nabla J_1(\Xi) &= \text{Re}\left(\overline{\nabla \hat{r}(\Xi)} \odot \hat{u}(\Xi) + \nabla \hat{u}(\Xi) \odot \overline{\hat{r}(\Xi)}\right) \\ &= \text{Re}\left(\nabla\left(\hat{u}(\Xi) \odot \overline{\hat{r}(\Xi)}\right)\right). \end{aligned}$$

□

7.4 Proof of Theorem 2

In order to simplify the notation, let $L(\Xi) \stackrel{\text{def}}{=} A(\Xi)^*A(\Xi)$. By Proposition 7, we have

$$\left| \frac{\partial J_1(\Xi)}{\partial \xi_m} \right| \leq |\hat{u}'(\xi_m)| \cdot |\hat{r}(\xi_m)| + |\hat{u}(\xi_m)| \cdot |\hat{r}'(\xi_m)|.$$

By definition, we have $\hat{r}(\Xi) = (L(\Xi) - \text{Id})\hat{u}(\Xi)$, hence

$$|\hat{r}(\xi_m)| \leq \|\hat{r}(\Xi)\|_2 \leq \frac{\|\hat{u}(\Xi)\|_2}{\text{md}(\Xi)}, \quad (38)$$

where we used Proposition 8 to obtain the last inequality. Now, we also wish to control $|\hat{r}'(\xi_m)|$. To

this end, first notice that

$$\begin{aligned} \hat{r}'(\xi_m) &= \sum_{m'=1}^M \left(\frac{\partial L(\Xi)_{m,m'}}{\partial \xi_m} \hat{u}(\xi_{m'}) \right. \\ &\quad \left. + L(\Xi)_{m,m'} \hat{u}'(\xi_{m'}) \mathbb{1}_{m=m'} \right) - \hat{u}'(\xi_m) \\ &= \sum_{m'=1}^M \frac{\partial L(\Xi)_{m,m'}}{\partial \xi_m} \hat{u}(\xi_{m'}). \end{aligned}$$

We start with an analytical expression of the matrix $L(\Xi)$.

Proposition 9 (The expression of A^*A). *Let $L(\Xi) \stackrel{\text{def}}{=} A(\Xi)^*A(\Xi)$. We have*

$$[L(\Xi)]_{m,m'} = \begin{cases} 1 & \text{if } m = m', \\ \frac{1}{N} \exp\left(\frac{\iota\pi(\xi_m - \xi_{m'})}{N}\right) & \\ \quad \times \frac{\sin(\pi(\xi_m - \xi_{m'}))}{\sin\left(\frac{\pi(\xi_m - \xi_{m'})}{N}\right)} & \text{otherwise.} \end{cases} \quad (39)$$

Proof. We have:

$$\begin{aligned} [L(\Xi)]_{m,m'} &= \frac{1}{N} \sum_n e^{2\iota\frac{\pi}{N}\langle \xi_{m'} - \xi_m, n \rangle} \\ &= \frac{1}{N} e^{-\iota\pi(\xi_{m'} - \xi_m)} \frac{1 - e^{2\iota\pi(\xi_{m'} - \xi_m)}}{1 - e^{2\iota\frac{\pi}{N}(\xi_{m'} - \xi_m)}} \\ &= \frac{1}{N} e^{-\iota\pi(\xi_{m'} - \xi_m)} \frac{e^{\iota\pi(\xi_{m'} - \xi_m)}}{e^{\iota\frac{\pi}{N}(\xi_{m'} - \xi_m)}} \\ &\quad \times \frac{e^{-\iota\pi(\xi_{m'} - \xi_m)} - e^{\iota\pi(\xi_{m'} - \xi_m)}}{e^{-\iota\frac{\pi}{N}(\xi_{m'} - \xi_m)} - e^{\iota\frac{\pi}{N}(\xi_{m'} - \xi_m)}} \\ &= \frac{1}{N} e^{-\iota\frac{\pi}{N}(\xi_{m'} - \xi_m)} \frac{\sin(\pi(\xi_{m'} - \xi_m))}{\sin\left(\frac{\pi}{N}(\xi_{m'} - \xi_m)\right)}. \end{aligned}$$

□

Now, we will use the following lemma.

Lemma 1. *The following bound holds:*

$$\left| \frac{\partial L(\Xi)_{m,m'}}{\partial \xi_m} \right| \leq \frac{\pi}{N} + \frac{4}{\text{dist}(\xi_{m'}, \xi_m)} \leq \frac{\pi}{N} + \frac{4}{\text{md}(\Xi)}.$$

Proof. Letting $\delta = \xi_m - \xi_{m'}$, we have

$$\begin{aligned} \frac{\partial L(\Xi)_{m,m'}}{\partial \xi_m} &= \frac{\pi}{N^2} \times \frac{\iota e^{\iota\frac{\pi}{N}\delta} \sin(\pi\delta)}{\sin\left(\frac{\pi}{N}\delta\right)} \\ &\quad + \frac{\pi}{N} \times \frac{e^{-\iota\frac{\pi}{N}\delta}}{\sin\left(\frac{\pi}{N}\delta\right)} \left(\cos(\pi\delta) - \frac{\sin(\pi\delta)}{N} \times \frac{\cos\left(\frac{\pi}{N}\delta\right)}{\sin\left(\frac{\pi}{N}\delta\right)} \right). \end{aligned}$$

Without loss of generality we consider the case $0 \leq \delta \leq N/2$. Using $\left| \frac{\sin(\pi\delta)}{N \sin\left(\frac{\pi}{N}\delta\right)} \right| \leq 1$ let us remark that

$$\begin{aligned} \left| \frac{\partial L(\Xi)_{m,m'}}{\partial \xi_m} \right| &\leq \frac{\pi}{N} \\ &\quad + \frac{\pi}{N} \left| \frac{1}{\sin\left(\frac{\pi}{N}\delta\right)} \left(\frac{\sin(\pi\delta) \cos\left(\frac{\pi}{N}\delta\right)}{N \sin\left(\frac{\pi}{N}\delta\right)} - \cos(\pi\delta) \right) \right|. \end{aligned}$$

Using the inequality $\left| \frac{\sin(\pi\delta)}{N \sin(\frac{\pi}{N}\delta)} \right| \leq 1$ again, we obtain

$$\left| \frac{\sin(\pi\delta) \cos(\frac{\pi}{N}\delta)}{N \sin(\frac{\pi}{N}\delta)} - \cos(\pi\delta) \right| \leq \left| \cos(\frac{\pi}{N}\delta) \right| + 1 \leq 2.$$

Finally, using the inequality $\sin(x) \geq x/2$ for $x \in (0, \pi/2)$, we get $\left| \frac{\partial L(\Xi)_{m',m}}{\partial \xi_{m'}} \right| \leq \frac{\pi}{N} + \frac{4}{\delta}$. \square

Lemma 1 and a Cauchy-Schwarz inequality provides the following bound:

$$|\hat{r}'(\xi_m)| \leq \left(\frac{\pi}{N} + \frac{4}{\text{md}(\Xi)} \right) \|\hat{u}(\Xi)\|_1.$$

Combining everything finally yields:

$$\left| \frac{\partial J_1(\Xi)}{\partial \xi_m} \right| \leq |\hat{u}'(\xi_m)| \cdot \frac{\|\hat{u}(\Xi)\|_2}{\text{md}(\Xi)} + |\hat{u}(\xi_m)| \cdot \|\hat{u}(\Xi)\|_1 \cdot \left(\frac{\pi}{N} + \frac{4}{\text{md}(\Xi)} \right).$$

Under the decay assumptions of Theorem 2, we obtain

$$\left| \frac{\partial J_1(\Xi)}{\partial \xi_m} \right| \lesssim \frac{\|\hat{u}(\Xi)\|_1}{\text{md}(\Xi) |\xi_m|^\alpha}.$$

References

- [1] Hemant Kumar Aggarwal and Mathews Jacob. J-modl: Joint model-based deep learning for optimized sampling and reconstruction. *IEEE journal of selected topics in signal processing*, 14(6):1151–1162, 2020.
- [2] Céline Aubel and Helmut Bölcskei. Vandermonde matrices with nodes in the unit disk and the large sieve. *Applied and Computational Harmonic Analysis*, 47(1):53–86, 2019.
- [3] Cagla Deniz Bahadir, Adrian V Dalca, and Mert R Sabuncu. Learning-based optimization of the under-sampling pattern in mri. In *International Conference on Information Processing in Medical Imaging*, pages 780–792. Springer, 2019.
- [4] Dmitry Batenkov, Laurent Demanet, Gil Goldman, and Yosef Yomdin. Conditioning of partial nonuniform fourier matrices with clustered nodes. *SIAM Journal on Matrix Analysis and Applications*, 41(1):199–220, 2020.
- [5] Enrico Bombieri. On the large sieve. In *Goldbach Conjecture*, pages 227–252. World Scientific, 1984.
- [6] Léon Bottou. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010*, pages 177–186. Springer, 2010.
- [7] Donald Goldfarb. A family of variable-metric methods derived by variational means. *Mathematics of computation*, 24(109):23–26, 1970.
- [8] Alban Gossard, Frédéric de Gournay, and Pierre Weiss. Off-the-grid data-driven optimization of sampling schemes in mri. *arXiv preprint arXiv:2010.01817*, 2020.
- [9] Baran Gözcü, Rabeeh Karimi Mahabadi, Yen-Huan Li, Efe Ilcak, Tolga Cukur, Jonathan Scarlett, and Volkan Cevher. Learning-based compressive mri. *IEEE transactions on medical imaging*, 37(6):1394–1406, 2018.
- [10] Kyong Hwan Jin, Michael Unser, and Kwang Moo Yi. Self-supervised deep active accelerated mri. *arXiv preprint arXiv:1901.04547*, 2019.
- [11] Alexander Loktyushin, Kai Herz, Nam Dang, Felix Glang, Anagha Deshmane, Simon Weinmüller, Arnd Doerfler, Bernhard Schölkopf, Klaus Scheffler, and Moritz Zaiss. Mrzero-automated discovery of mri sequences using supervised learning. *Magnetic Resonance in Medicine*, 86(2):709–724, 2021.
- [12] Ankur Moitra. Super-resolution, extremal functions and the condition number of vandermonde matrices. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 821–830, 2015.
- [13] Alan Oppenheim, Don Johnson, and Kenneth Steiglitz. Computation of spectra with unequal resolution using the fast fourier transform. *Proceedings of the IEEE*, 59(2):299–301, 1971.
- [14] Wei Peng, Li Feng, Guoying Zhao, and Fang Liu. Learning optimal k-space acquisition and reconstruction using physics-informed neural networks. *CVPR*, 2022.
- [15] Ferdia Sherry, Martin Benning, Juan Carlos De los Reyes, Martin J Graves, Georg Maierhofer, Guy Williams, Carola-Bibiane Schönlieb, and Matthias J Ehrhardt. Learning the sampling pattern for mri. *IEEE Transactions on Medical Imaging*, 39(12):4310–4321, 2020.
- [16] Guanhua Wang, Tianrui Luo, Jon-Fredrik Nielsen, Douglas C Noll, and Jeffrey A Fessler. B-spline parameterized joint optimization of reconstruction and k-space trajectories (bjork) for accelerated 2d mri. *IEEE Transactions on Medical Imaging*, 2021.
- [17] Tomer Weiss, Ortal Senouf, Sanketh Vedula, Oleg Michailovich, Michael Zibulevsky, and Alex Bronstein. Pilot: Physics-informed learned optimal trajectories for accelerated mri. *arXiv preprint arXiv:1909.05773*, 2019.
- [18] Jure Zbontar, Florian Knoll, Anuroop Sriram, Tullie Murrell, Zhengnan Huang, Matthew J Muckley, Aaron Defazio, Ruben Stern, Patricia Johnson, Mary Bruno, et al. fastmri: An open dataset and benchmarks for accelerated mri. *arXiv preprint arXiv:1811.08839*, 2018.
- [19] Marcelo VW Zibetti, Gabor T Herman, and Ravinder R Regatte. Fast data-driven learning of parallel mri sampling patterns for large scale problems. *Scientific Reports*, 11(1):1–19, 2021.