

Provably robust estimation of modulo 1 samples of a smooth function with applications to phase unwrapping

Mihai Cucuringu

The Alan Turing Institute, London, UK

Department of Statistics and Mathematical Institute, University of Oxford, Oxford, UK

MIHAI.CUCURINGU@STATS.OX.AC.UK

Hemant Tyagi

INRIA Lille - Nord Europe, Lille, France

HEMANT.TYAGI@INRIA.FR

Editor: Sujay Sanghavi

Abstract

Consider an unknown smooth function $f : [0, 1]^d \rightarrow \mathbb{R}$, and assume we are given n noisy mod 1 samples of f , i.e., $y_i = (f(x_i) + \eta_i) \bmod 1$, for $x_i \in [0, 1]^d$, where η_i denotes the noise. Given the samples $(x_i, y_i)_{i=1}^n$, our goal is to recover smooth, robust estimates of the clean samples $f(x_i) \bmod 1$. We formulate a natural approach for solving this problem, which works with angular embeddings of the noisy mod 1 samples over the unit circle, inspired by the angular synchronization framework. This amounts to solving a smoothness regularized least-squares problem – a quadratically constrained quadratic program (QCQP) – where the variables are constrained to lie on the unit circle. Our proposed approach is based on solving its relaxation, which is a *trust-region sub-problem* and hence solvable efficiently. We provide theoretical guarantees demonstrating its robustness to noise for adversarial, as well as random Gaussian and Bernoulli noise models. To the best of our knowledge, these are the first such theoretical results for this problem. We demonstrate the robustness and efficiency of our proposed approach via extensive numerical simulations on synthetic data, along with a simple least-squares based solution for the unwrapping stage, that recovers the original samples of f (up to a global shift). It is shown to perform well at high levels of noise, when taking as input the denoised modulo 1 samples.

Finally, we also consider two other approaches for denoising the modulo 1 samples that leverage tools from Riemannian optimization on manifolds, including a Burer-Monteiro approach for a semidefinite programming relaxation of our formulation. For the two-dimensional version of the problem, which has applications in synthetic aperture radar interferometry (InSAR), we are able to solve instances of real-world data with a million sample points in under 10 seconds, on a personal laptop.

Keywords: quadratically constrained quadratic programming (QCQP), trust-region sub-problem, angular embedding, phase unwrapping, semidefinite programming, angular synchronization.

Contents

1	Introduction	3
2	Problem setup	6
3	Smoothness regularized least-squares in the angular domain	7

4	A trust-region based relaxation for denoising modulo 1 samples	9
4.1	Recovering the denoised mod 1 samples	11
4.2	Unwrapping stage and main algorithm	12
5	Analysis for the arbitrary bounded noise model	14
6	Analysis for random noise models	21
6.1	Analysis for the Bernoulli-Uniform random noise model	21
6.2	Analysis for Gaussian noise model	23
7	Analysis for multivariate functions and arbitrary bounded noise	24
8	Numerical experiments for the univariate case via TRS-based modulo denoising	28
8.1	Numerical experiments: Gaussian Model	29
8.2	Numerical experiments: Bernoulli Model	29
8.3	Comparison with Bhandari et al. (2017)	31
9	Modulo 1 denoising via optimization on manifolds	39
9.1	Two formulations for denoising modulo 1 samples	39
9.2	Numerical simulations	41
10	Related work	42
10.1	Function recovery from modulo samples	42
10.2	Relation to group synchronization	46
10.3	Phase unwrapping	48
11	Concluding remarks and future work	50
12	Acknowledgements	52
A	Rewriting the QCQP in real domain	59
B	Trust-region sub-problem with ℓ_2 ball/sphere constraint	59
C	Useful concentration inequalities	61
D	Proof of Proposition 13	62
E	Proof of Proposition 15	65
F	Additional numerical experiments	69
F1	Numerical experiments: Bounded Model	69
F2	Comparison with Bhandari et al. (2017)	70
F3	Additional elevation maps experiments for two-dimensional phase unwrapping	73

1. Introduction

In many domains of science and engineering, one is given access to noisy samples of a signal f , and the goal is to recover the original clean samples. The signal f is typically smooth in some sense, and one would like to have an algorithm that is not only robust to noise, but also outputs smooth estimates. More formally, we can think of the signal $f : \mathbb{R}^d \rightarrow \mathbb{R}$ for which we are typically given samples $f(\mathbf{x}_i)$; $i = 1, \dots, n$ for $\mathbf{x}_i \in \mathcal{U}$ for a compact $\mathcal{U} \subset \mathbb{R}^d$. Perhaps one of the most important applications of this problem arises in image denoising, with a rich body of literature (see for eg., Elad and Aharon (2006); Lou et al. (2010)).

Interestingly, there are applications where one does not observe the samples directly, but only the *modulo* samples of f , i.e., $f(\mathbf{x}_i) \bmod \zeta$ for some $\zeta \in \mathbb{R}^+$. For instance, when $\zeta = 1$, $f(\mathbf{x}) \bmod 1$ simply corresponds to the fractional part of $f(\mathbf{x})$. Such measurements are typically obtained due to constraints on the sampling hardware, or due to physical constraints imposed by the specific nature of the problem. To see this, we discuss two important applications below.

Self-reset analog-to-digital converters (ADCs). Traditional ADCs have voltage limits in place that cut off the signal, i.e., saturate, whenever the signal value lies outside the limits. Recently, a new generation of ADCs have adopted a different approach to this problem, wherein they simply *reset* the signal value to the other threshold value (Kester (2009); Rhee and Joo (2003); Kavusi and Abbas (2004); Sasagawa et al. (2016); Yamaguchi et al. (2016)). For example, if the voltage range is $[0, b]$, the reset operation simply corresponds to a modulo b operation on the signal value. This especially helps in working with signals whose dynamic range is much larger than what can be handled by a standard ADC. This signal acquisition mechanism was the main motivation behind the recent work of Bhandari et al. (2017) (also covered in the media¹), wherein the authors derived conditions for exact recovery of a band limited function from its samples, in the noiseless setting (Bhandari et al., 2017, Theorem 1). Let us note that band limitedness is implicitly a smoothness assumption on f , such an assumption being clearly necessary in order to be able to provably recover the original (unwrapped) samples of f .

Phase unwrapping. Phase unwrapping refers to the problem of recovering the original phase values of a signal ϕ (at different spatial locations), from their modulo 2π (in radians) versions. The case where $\phi : \mathbb{R}^2 \rightarrow \mathbb{R}$ has received considerable attention, in particular due to important applications arising for instance in synthetic aperture radar interferometry (InSAR) (Graham (1974); Zebker and Goldstein (1986)), magnetic resonance imaging (MRI) (Hedley and Rosenfeld (1992); Lauterbur (1973)), optics (Venema and Schmidt (2008)), diffraction tomography (Pratt and Worthington (1988)) and non destructive testing of components (Paoletti et al. (1994); Hung (1996)), to name a few. Generally speaking, remote sensing systems typically obtain information about the structure of an object by measuring the phase coherence between the transmitted and scattered waveforms. For instance, in radar interferometry, information about the terrain elevation is inherently present in the phase values. In MRI, information regarding the velocity of blood flow or the position of veins in tissues can be obtained from the phase values. The higher-dimensional case has received comparatively less attention. The three-dimensional version of the problem has

1. <http://news.mit.edu/2017/ultra-high-contrast-digital-sensing-cameras-0714>

applications in 3D MRI imaging (Jenkinson (2003)) and radar interferometry (Hooper and Zebker (2007); Osmanoglu et al. (2014)). Lastly, some papers have also formulated methods for the general d dimensional case (Jenkinson (2003); Fang et al. (2006)).

Our work can also be placed in the context of denoising smooth functions taking values in a nonlinear space. Rahman et al. (2005) considered multiscale representations for manifold-valued data observed on equispaced grids and taking values on manifolds such as the sphere S^2 , the special orthogonal group $SO(3)$ or the Special Euclidean Group $SE(3)$. They proposed a method that generalizes wavelet analysis from the traditional setting where functions defined on the equispaced values in a cartesian grid no longer take simple values such as numerical array, but rather arrays whose entries have highly structured values obeying nonlinear constraints, and show that such representations are successful in tasks such as denoising, compression or contrast enhancement.

As a word of caution to the reader, we note that this problem is different from the celebrated *phase retrieval*, a classical problem in optics that has attracted a surge of interest in recent years (see Candès et al. (2013); Jaganathan et al. (2015)). There, one attempts to recover an unknown signal from the magnitude (intensity) of its Fourier transform. Just like phase retrieval, the recovery of a function from mod 1 measurements is, by its very nature, an ill-posed problem, and one needs to incorporate prior structure on the signal. In our case is smoothness of f (in an analogous way to how enforcing sparsity renders the phase retrieval problem well-posed). While there have been a variety of approaches to phase retrieval, recent progress in the compressed sensing and convex optimization-based signal processing have inspired new potential research directions. The approach we pursue in this paper is inspired by developments in the *trust region sub-problem* (Adachi et al. (2017)) and *group synchronization* (Singer (2011); Cucuringu (2016)) literatures.

Overview of approach and contributions. At a high level, one would like to recover denoised samples (i.e., smooth, robust estimates) of f from its noisy mod 1 versions. A natural approach to tackle this problem is the following two-stage method. In the first stage, one recovers denoised mod 1 samples of f , while in the (unwrapping) second stage, one uses these samples to recover the original real-valued samples of f . In this paper, we mainly focus on the first stage, which is a challenging problem in itself. To the best of our knowledge, we provide the first algorithm for denoising mod1 samples of a function, which comes with robustness guarantees. In particular, we make the following contributions².

1. We formulate a general framework for denoising the mod1 samples of f which involves mapping the noisy mod 1 values (lying in $[0, 1)$) to the angular domain (i.e. in $[0, 2\pi)$) and solving a smoothness regularized least-squares problem. This amounts to a quadratically constrained quadratic program (QCQP) with non-convex constraints wherein the variables are required to lie on the unit circle. We consider solving a relaxation of this QCQP which is a *trust-region sub-problem*, and hence solvable efficiently.

2. A preliminary version of this paper containing the results of Sections 4, 5 and parts of Section 8 will appear in AISTATS 2018 (Cucuringu and Tyagi (2018)). This is a significantly expanded version containing additional theoretical results, numerical experiments and a detailed discussion.

2. We provide a detailed theoretical analysis for the above approach, proving its robustness to noise for different noise models, provided the noise level is not large. Specifically, for $d = 1$, we show this for arbitrary bounded noise (see (2.3),(5.1), Theorem 5), Bernoulli-uniform noise (see (2.4), Theorem 11) and Gaussian noise (see (2.5), Theorem 14). For the multivariate case $d \geq 1$, we show this for arbitrary bounded noise (see Theorem 16).
3. We test the above trust-region based method on synthetic data which demonstrates that it performs well for reasonably high levels of noise. To complete the picture, we also implement the second stage with a simple least-squares based method for recovering the (real-valued) samples of f , and show that it performs surprisingly well via extensive simulations.
4. Finally, we also consider two other approaches for denoising the modulo 1 samples that leverage tools from Riemannian optimization on manifolds (Absil et al. (2007)). The first one is based on a semidefinite programming (SDP) relaxation of the QCQP which we solve via the Burer-Monteiro approach. The second one involves solving the original QCQP by optimizing over the manifold associated with the constraints. We implement both approaches using the Manopt toolbox (Boumal et al. (2014)), and highlight their scalability and robustness to noise via extensive experiments on the two-dimensional version of the problem. In particular, we are able to solve instances of real-world problems containing a million samples in under 10 seconds on a personal laptop.

Outline of paper. Section 2 formulates the problem formally, and introduces notation for the $d = 1$ case. Section 3 sets up the mod 1 denoising problem as a smoothness regularized least-squares problem in the angular domain which is a QCQP. Section 4 describes its relaxation to a *trust-region sub-problem*, and some simple approaches for unwrapping, i.e., recovering the samples of f , along with our complete two-stage algorithm. Section 5 contains approximation guarantees for the arbitrary bounded noise model for the trust region based relaxation for recovering the denoised mod1 samples of f (when $d = 1$). Section 6 contains similar results for two random noise models (Gaussian and Bernoulli-uniform). Section 7 discusses the generalization of our approach to the multivariate setting when $d \geq 1$, along with approximation guarantees for the bounded noise model. Section 8 contains experiments on synthetic data with $d = 1$ for different noise models, as well a comparison with the algorithm of Bhandari et al. (2017). In Section 9, we describe two other approaches based on optimization on manifolds for denoising the modulo 1 samples. It also contains experiments involving these approaches for the $d = 2$ setting, on synthetic and real-world data. Section 10 surveys a number of related approaches and applications, with a focus on those arising in the phase unwrapping literature. Section 11, summarizes our findings, and also contains a discussion of possible future research directions. Finally, the Appendix contains supplementary material related to proofs and additional numerical experiments.

2. Problem setup

We begin with the problem setup for the univariate case $d = 1$, as much of the analysis in the ensuing sections is for this particular setting. The multivariate case where $d \geq 1$ is treated separately in Section 7.

Consider a smooth, unknown function $f : [0, 1] \rightarrow \mathbb{R}$, and a uniform grid on $[0, 1]$,

$$0 = x_1 < x_2 < \dots < x_n = 1 \text{ with } x_i = \frac{i-1}{n-1}. \quad (2.1)$$

We assume that we are given *mod 1* samples of f on the above grid. Note that for each sample we can decompose the function as

$$f(x_i) = q_i + r_i \in \mathbb{R}, \quad (2.2)$$

with $q_i \in \mathbb{Z}$ and $r_i \in [0, 1)$, we have $r_i = f(x_i) \bmod 1$. The modulus is fixed to 1 without loss of generality since $\frac{f \bmod s}{s} = \frac{f}{s} \bmod 1$. This is easily seen by writing $f = sq + r$, with $q \in \mathbb{Z}$, and observing that $\frac{f}{s} \bmod 1 = \frac{sq+r}{s} \bmod 1 = \frac{r}{s} = \frac{f \bmod s}{s}$. In particular, we assume that the mod 1 samples are *noisy*, and consider the following three noise models.

1. Arbitrary bounded noise

$$y_i = (f(x_i) + \delta_i) \bmod 1; \quad |\delta_i| \in (0, 1/2); \quad i = 1, \dots, n. \quad (2.3)$$

2. Bernoulli-Uniform noise

$$y_i = \begin{cases} f(x_i) \bmod 1 & ; \text{ w.p } 1-p \\ \sim U[0, 1] & ; \text{ w.p } p \end{cases} \quad \text{i.i.d}; \quad i = 1, \dots, n. \quad (2.4)$$

Hence for some parameter $p \in (0, 1)$ we either observe the clean sample with probability $1 - p$, or some garbage value generated uniformly at random in $[0, 1]$.

3. Gaussian noise

$$y_i = (f(x_i) + \eta_i) \bmod 1; \quad i = 1, \dots, n, \quad (2.5)$$

where $\eta_i \sim \mathcal{N}(0, \sigma^2)$ i.i.d.

We will denote $f(x_i)$ by f_i for convenience. Our aim is to recover smooth, robust estimates (up to a global shift) of the original samples $(f_i)_{i=1}^n$ from the measurements $(x_i, y_i)_{i=1}^n$. We will assume f to be Hölder continuous meaning that for constants $M > 0$, $\alpha \in (0, 1]$,

$$|f(x) - f(y)| \leq M|x - y|^\alpha; \quad \forall x, y \in [0, 1]. \quad (2.6)$$

The above assumption is quite general and reduces to Lipschitz continuity when $\alpha = 1$.

Notation. Scalars and matrices are denoted by lower case and upper cases symbols respectively, while vectors are denoted by lower bold face symbols. Sets are denoted by calligraphic symbols (eg., \mathcal{N}), with the exception of $[n] = \{1, \dots, n\}$ for $n \in \mathbb{N}$. The imaginary unit is denoted by $\iota = \sqrt{-1}$. The notation introduced throughout Sections 3 and 4 is summarized in Table 1. We will denote the ℓ_p ($1 \leq p \leq \infty$) norm of a vector $\mathbf{x} \in \mathbb{R}^n$ by $\|\mathbf{x}\|_p$ (defined as $(\sum_i |x_i|^p)^{1/p}$). In particular, $\|\mathbf{x}\|_\infty := \max_i |x_i|$. For a matrix $A \in \mathbb{R}^{m \times n}$, we will denote its spectral norm (i.e., largest singular value) by $\|A\|$ and its Frobenius norm by $\|A\|_F$ (defined as $(\sum_{i,j} A_{i,j}^2)^{1/2}$). For a square matrix, we denote its trace by $\text{tr}(\cdot)$.




Figure 1: Motivation for the angular embedding approach.

3. Smoothness regularized least-squares in the angular domain

Our algorithm essentially works in two stages.

1. **Denoising stage.** Our goal here is to denoise the mod 1 samples, which is also the main focus of this paper. In a nutshell, we map the given noisy mod 1 samples to points on the unit complex circle, and solve a smoothness regularized, constrained least-squares problem. The solution to this problem, followed by a simple post-processing step, yields the final denoised mod 1 samples of f .
2. **Unwrapping stage.** The second stage takes as input the above denoised mod 1 samples, and produces an estimate of the original real-valued samples of f (up to a global shift).

We start the denoising stage by mapping the mod 1 samples to the angular domain, with

$$h_i := \exp(2\pi\iota f_i) = \exp(2\pi\iota r_i), \quad z_i := \exp(2\pi\iota y_i), \quad (3.1)$$

denoting the respective representations of the clean mod 1 and noisy mod 1 samples on the unit circle in \mathbb{C} , where the first equality is due to the fact that $f_i = q_i + r_i$, with $q_i \in \mathbb{Z}$. The choice of representing the mod 1 samples in (3.1) is very natural for the following reason. For points x_i, x_j sufficiently close, the samples f_i, f_j will also be close (by Hölder continuity of f). While the corresponding wrapped samples $f_i \bmod 1, f_j \bmod 1$ can still be far apart, the complex numbers $\exp(\iota 2\pi f_i)$ and $\exp(\iota 2\pi f_j)$ will necessarily be close to each other³. This is illustrated in the toy example in Figure 1.

Figure 2 is the analogue of Figure 1, but for a noisy instance of the problem, making the point that the angular embedding facilitates the denoising process. For points x_i, x_j sufficiently close, the corresponding samples f_i, f_j will also be close in the real domain, by Hölder continuity of f . When measurements get perturbed by noise, the distance in the real domain between the noisy mod 1 samples can greatly increase and become close to 1 (in this example, the point B gets perturbed by noise, hits the “floor” and “resets” itself). However, in the angular embedding space, the two points still remain close to each other, as depicted in Figure 2c.

3. Indeed, $|\exp(\iota 2\pi f_i) - \exp(\iota 2\pi f_j)| = |1 - \exp(\iota 2\pi(f_j - f_i))| = 2|\sin(\pi(f_j - f_i))| \leq 2\pi|f_j - f_i|$ (since $|\sin x| \leq |x| \forall x \in \mathbb{R}$).

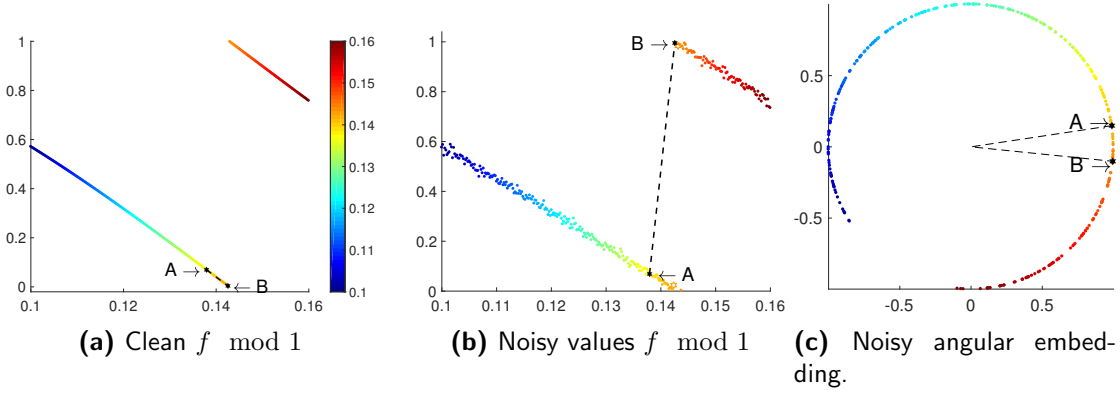


Figure 2: Motivation for the angular embedding approach. Noise perturbations may take nearby points (mod 1 samples) far away in the real domain, yet the points will remain close in the angular domain.

Consider an undirected graph $G = (V, E)$ with $V = \{1, 2, \dots, n\}$ where index i corresponds to the point x_i on our grid, and $E = \{\{i, j\} : i, j \in V, i \neq j, |i - j| \leq k\}$ denotes the set of edges for a suitable parameter $k \in \mathbb{N}$. A natural approach for recovering smooth estimates of $(h_i)_{i=1}^n$ would be to solve the following optimization problem

$$\min_{g_1, \dots, g_n \in \mathbb{C}; |g_i|=1} \sum_{i=1}^n |g_i - z_i|^2 + \lambda \sum_{\{i,j\} \in E} |g_i - g_j|^2. \quad (3.2)$$

Here, $\lambda > 0$ is a regularization parameter, which along with k , controls the smoothness of the solution. We denote by $L \in \mathbb{R}^{n \times n}$ the Laplacian matrix associated with G , defined as

$$L_{i,j} = \begin{cases} \deg(i), & i = j, \\ -1, & \{i, j\} \in E, \\ 0, & \text{otherwise.} \end{cases} \quad (3.3)$$

Denoting $\mathbf{g} = [g_1 \ g_2 \ \dots \ g_n]^T \in \mathbb{C}^n$, the second term in (3.2) can be simplified to

$$\lambda \left(\sum_{i \in V} \deg(i) |g_i|^2 - \sum_{\{i,j\} \in E} (g_i g_j^* + g_i^* g_j) \right) = \lambda \mathbf{g}^* L \mathbf{g}. \quad (3.4)$$

Next, denoting $\mathbf{z} = [z_1 \ z_2 \ \dots \ z_n]^T \in \mathbb{C}^n$, we can further simplify the first term in (3.2) as

$$\sum_{i=1}^n |g_i - z_i|^2 = \sum_{i=1}^n (|g_i|^2 + |z_i|^2 - g_i z_i^* - g_i^* z_i) \quad (3.5)$$

$$= 2n - 2\text{Re}(\mathbf{g}^* \mathbf{z}). \quad (3.6)$$

This gives us the following equivalent form of (3.2)

$$\min_{\mathbf{g} \in \mathbb{C}^n: |g_i|=1} \lambda \mathbf{g}^* L \mathbf{g} - 2\text{Re}(\mathbf{g}^* \mathbf{z}). \quad (3.7)$$

Symbol	Description
f	unknown real-valued function
r	clean $f \pmod{1}$
q	clean reminder $q = f - r$
y	noisy $f \pmod{1}$
\mathbf{h}	clean signal in angular domain
\mathbf{z}	noisy signal in angular domain
\mathbf{g}	free complex-valued variable
$\bar{\mathbf{h}}$	real-valued version of \mathbf{h}
$\bar{\mathbf{z}}$	real-valued version of \mathbf{z}
$\bar{\mathbf{g}}$	real-valued version of \mathbf{g}
L	$n \times n$ Laplacian matrix of graph G
H	$2n \times 2n$ block diagonal version of L

Table 1: Summary of frequently used symbols in the paper.

4. A trust-region based relaxation for denoising modulo 1 samples

The optimization problem in (3.7) is over a non-convex set $\mathcal{C}_n := \{\mathbf{g} \in \mathbb{C}^n : |g_i| = 1\}$. In general, the problem $\min_{\mathbf{g} \in \mathcal{C}_n} \mathbf{g}^* \mathbf{A} \mathbf{g}$, where $\mathbf{A} \in \mathbb{C}^{n \times n}$ is positive semidefinite, is known to be NP-hard (Zhang and Huang, 2006, Proposition 3.3). Note that we can get rid of the linear term $2\text{Re}(\mathbf{g}^* \mathbf{z})$ by rewriting (3.7) as

$$\min_{\mathbf{g} \in \mathcal{C}_n} [\mathbf{g}^* \ 1] \begin{bmatrix} \lambda L & -\mathbf{z} \\ -\mathbf{z}^* & 0 \end{bmatrix} \begin{bmatrix} \mathbf{g} \\ 1 \end{bmatrix}. \quad (4.1)$$

The quadratic term in (4.1) is Hermitian, and is of course structured, as L is the Laplacian of a nearest neighbor graph. However the complexity of (4.1) is still unclear. As pointed out by a reviewer of a preliminary version of this paper (Cucuringu and Tyagi (2018)), one possible approach is to discretize the angular domain, and solve (3.7) approximately via dynamic programming. Since the graph G has tree-width k , the computational cost of this approach may be⁴ exponential in k .

Our approach is to relax the constraints corresponding to \mathcal{C}_n , to one where the points lie on a sphere of radius n , which amounts to the following optimization problem

$$\min_{\mathbf{g} \in \mathbb{C}^n : \|\mathbf{g}\|_2^2 = n} \lambda \mathbf{g}^* L \mathbf{g} - 2\text{Re}(\mathbf{g}^* \mathbf{z}). \quad (4.2)$$

It is straightforward to reformulate (4.2) in terms of real variables. We do so by introducing the following notation for the real-valued versions of the variables \mathbf{h} (clean signal), \mathbf{z} (noisy signal), and \mathbf{g} (free variable)

$$\bar{\mathbf{h}} = \begin{pmatrix} \text{Re}(\mathbf{h}) \\ \text{Im}(\mathbf{h}) \end{pmatrix}, \quad \bar{\mathbf{z}} = \begin{pmatrix} \text{Re}(\mathbf{z}) \\ \text{Im}(\mathbf{z}) \end{pmatrix}, \quad \bar{\mathbf{g}} = \begin{pmatrix} \text{Re}(\mathbf{g}) \\ \text{Im}(\mathbf{g}) \end{pmatrix} \in \mathbb{R}^{2n}, \quad (4.3)$$

4. Note that naive dynamic programming will have a running time that is exponential in k . However, it is unclear whether this exponential dependence on k is unavoidable.

and the corresponding block-diagonal Laplacian

$$H = \begin{pmatrix} \lambda L & 0 \\ 0 & \lambda L \end{pmatrix} = \lambda \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \otimes L \in \mathbb{R}^{2n \times 2n}. \quad (4.4)$$

In light of this, the optimization problem (4.2) can be equivalently formulated as

$$\min_{\bar{\mathbf{g}} \in \mathbb{R}^{2n}: \|\bar{\mathbf{g}}\|_2^2 = n} \bar{\mathbf{g}}^T H \bar{\mathbf{g}} - 2\bar{\mathbf{g}}^T \bar{\mathbf{z}}, \quad (4.5)$$

which is formally shown in the appendix for completeness. Let us note that the Laplacian matrix L is positive semi-definite (p.s.d), with its smallest eigenvalue $\lambda_1(L) = 0$ with multiplicity 1 (since G is connected). Therefore, H is also p.s.d, with smallest eigenvalue $\lambda_1(H) = 0$ with multiplicity 2.

The optimization problem (4.5) is actually an instance of the so-called trust-region sub-problem (TRS) with equality constraint (which we denote by $TSR_=$ from now on), where one minimizes a general quadratic function (not necessarily convex), subject to a sphere constraint. For completeness, we also mention the closely related trust-region sub-problem with inequality constraint (denoted by TSR_{\leq}), where we have a ℓ_2 ball constraint. There exist several algorithms that efficiently solve TSR_{\leq} (Sorensen (1982); Morè and Sorensen (1983); Rojas et al. (2001); Rendl and Wolkowicz (1997); Gould et al. (1999); Adachi et al. (2017)), and also some which explicitly solve $TSR_=$ (Hager (2001); Adachi et al. (2017)). In particular, Adachi et al. (2017) recently showed that trust-region sub-problems can be solved to high accuracy via a single generalized eigenvalue problem. The computational complexity is $O(n^3)$ in the worst case, but improves when the matrices involved are sparse. In our case, the Laplacian is sparse when k is not large. In the experiments, we employ their algorithm for solving (4.5).

Remark 1 *We remark that the $TSR_=$ formulation is just one possible relaxation. One could also, for instance, consider a semidefinite programming based relaxation for the QCQP. We discuss this in Section 9, and solve it numerically via the Burer-Monteiro approach. Moreover, we also consider solving the original QCQP (3.7) using tools from optimization on manifolds (Absil et al. (2007)), as \mathcal{C}_n is a manifold (Absil et al. (2008)).*

Remark 2 *Note that (3.7) is similar to the angular synchronization problem (see for eg., Singer (2011)) as they both optimize a quadratic form subject to entries lying on the unit circle. The fundamental difference is that the matrix in the quadratic term in synchronization is formed using the given noisy pairwise angle offsets (embedded on the unit circle), and thus depends on the data. In our setup, the quadratic term is formed using the Laplacian of the smoothness regularization graph, and thus is independent of the given data (noisy mod 1 samples).*

Rather surprisingly, one can fully characterize⁵ the solutions to both $TSR_=$ and TSR_{\leq} . The following Lemma 3 characterizes the solution for (4.5); it follows directly from (Sorensen, 1982, Lemma 2.4, 2.8) (see also (Hager, 2001, Lemma 1)).

5. Discussed in detail in the appendix for completeness.

Lemma 3 $\widehat{\mathbf{g}}$ is a solution to (4.5) iff $\|\widehat{\mathbf{g}}\|_2^2 = n$ and $\exists \mu^*$ such that (a) $2H + \mu^*I \succeq 0$ and (b) $(2H + \mu^*I)\widehat{\mathbf{g}} = 2\bar{\mathbf{z}}$. Moreover, if $2H + \mu^*I \succ 0$, then the solution is unique.

Let $\{\lambda_j(H)\}_{j=1}^{2n}$, with $\lambda_1(H) \leq \lambda_2(H) \leq \dots \leq \lambda_{2n}(H)$, and $\{\mathbf{q}_j\}_{j=1}^{2n}$ denote the eigenvalues, respectively eigenvectors, of H . Note that $\lambda_1(H) = \lambda_2(H) = 0$, and $\lambda_3(H) > 0$, since G is connected. Let us denote the null space of H by $\mathcal{N}(H)$, so $\mathcal{N}(H) = \text{span}\{\mathbf{q}_1, \mathbf{q}_2\}$. Next, we analyze the solution to (4.5) with the help of Lemma 3, by considering the following two cases.

Case 1. $\bar{\mathbf{z}} \notin \mathcal{N}(H)$. The solution is given by

$$\widehat{\mathbf{g}}(\mu^*) = 2(2H + \mu^*I)^{-1}\bar{\mathbf{z}} = 2 \sum_{j=1}^{2n} \frac{\langle \bar{\mathbf{z}}, \mathbf{q}_j \rangle}{2\lambda_j(H) + \mu^*} \mathbf{q}_j, \quad (4.6)$$

for a unique $\mu^* \in (0, \infty)$ satisfying $\|\widehat{\mathbf{g}}(\mu^*)\|_2^2 = n$. Indeed, denoting $\phi(\mu) = \|\widehat{\mathbf{g}}(\mu)\|_2^2 = 4 \sum_{j=1}^{2n} \frac{\langle \bar{\mathbf{z}}, \mathbf{q}_j \rangle^2}{(2\lambda_j(H) + \mu)^2}$, we can see that $\phi(\mu)$ has a pole at $\mu = 0$ and decreases monotonically to 0 as $\mu \rightarrow \infty$. Hence, there exists a unique $\mu^* \in (0, \infty)$ such that $\|\widehat{\mathbf{g}}(\mu^*)\|_2^2 = n$. The solution $\widehat{\mathbf{g}}(\mu^*)$ will be unique by Lemma 3, since $2H + \mu^*I \succ 0$ holds.

Case 2. $\bar{\mathbf{z}} \perp \mathcal{N}(H)$. This second scenario requires additional attention. To begin with, note that

$$\phi(0) = 4 \sum_{j=1}^{2n} \frac{\langle \bar{\mathbf{z}}, \mathbf{q}_j \rangle^2}{(2\lambda_j(H))^2} = \sum_{j=3}^{2n} \frac{\langle \bar{\mathbf{z}}, \mathbf{q}_j \rangle^2}{\lambda_j(H)^2} \quad (4.7)$$

is now well defined, i.e., 0 is not a pole of $\phi(\mu)$ anymore. If $\phi(0) > n$, then as before, we can again find a unique $\mu^* \in (0, \infty)$ satisfying $\phi(\mu^*) = n$. The solution is given by $\widehat{\mathbf{g}}(\mu^*) = 2(2H + \mu^*I)^{-1}\bar{\mathbf{z}}$ and is unique since $2H + \mu^*I \succ 0$ (by Lemma 3).

In case $\phi(0) \leq n$, we set $\mu^* = 0$ and define our solution to be of the form

$$\widehat{\mathbf{g}}(\theta, \mathbf{v}) = (H)^\dagger \bar{\mathbf{z}} + \theta \mathbf{v}; \quad \mathbf{v} \in \mathcal{N}(H), \|\mathbf{v}\|_2 = 1, \quad (4.8)$$

where \dagger denotes pseudo-inverse and $\theta \in \mathbb{R}$. In particular, for any given $\mathbf{v} \in \mathcal{N}(H)$, $\|\mathbf{v}\|_2 = 1$, we obtain $\widehat{\mathbf{g}}(\theta^*, \mathbf{v}), \widehat{\mathbf{g}}(-\theta^*, \mathbf{v})$ as the solutions to (4.5), with $\pm\theta^*$ being the solutions to the equation

$$\|\widehat{\mathbf{g}}(\theta, \mathbf{v})\|_2^2 = n \Leftrightarrow \underbrace{\|(H)^\dagger \bar{\mathbf{z}}\|_2^2}_{=\phi(0) \leq n} + \theta^2 = n \quad (4.9)$$

Hence the solution is not unique if $\phi(0) < n$.

4.1. Recovering the denoised mod 1 samples

The solution to (4.5) is a vector $\widehat{\mathbf{g}} \in \mathbb{R}^{2n}$. Let $\widehat{\mathbf{g}} \in \mathbb{C}^n$ be the complex representation of $\widehat{\mathbf{g}}$ as per (4.3), so that $\widehat{\mathbf{g}} = [\text{Re}(\widehat{\mathbf{g}})^T \text{Im}(\widehat{\mathbf{g}})^T]^T$. Denoting $\widehat{g}_i \in \mathbb{C}$ to be the i^{th} component of $\widehat{\mathbf{g}}$, note that $|\widehat{g}_i|$ is not necessarily equal to one. On the other hand, recall that $h_i = \exp(i2\pi f_i \text{ mod } 1)$, $\forall i = 1, \dots, n$ for the ground truth $\mathbf{h} \in \mathbb{C}^n$. We obtain our final estimate $\widehat{f}_i \text{ mod } 1$ to $f_i \text{ mod } 1$ by projecting \widehat{g}_i onto the unit complex disk

$$\exp(i2\pi(\widehat{f}_i \text{ mod } 1)) = \frac{\widehat{g}_i}{|\widehat{g}_i|}; \quad i = 1, \dots, n. \quad (4.10)$$

In order to measure the distance between $\widehat{f}_i \bmod 1$ and $f_i \bmod 1$, we will use the so-called *wrap-around* distance on $[0, 1]$ denoted by $d_w : [0, 1]^2 \rightarrow [0, 1/2]$, where

$$d_w(t_1, t_2) := \min \{|t_1 - t_2|, 1 - |t_1 - t_2|\}, \quad (4.11)$$

for $t_1, t_2 \in [0, 1]$. We will now show that if \widehat{g}_i is sufficiently close to h_i for each $i = 1, \dots, n$, then each $d_w(\widehat{f}_i \bmod 1, f_i \bmod 1)$ will be correspondingly small. This is stated precisely in the following lemma.

Lemma 4 *For $0 < \epsilon < 1/2$, let $|\widehat{g}_i - h_i| \leq \epsilon$ hold for each $i = 1, \dots, n$. Then, for each $i = 1, \dots, n$*

$$d_w(\widehat{f}_i \bmod 1, f_i \bmod 1) \leq \frac{1}{\pi} \sin^{-1} \left(\frac{\epsilon}{1 - \epsilon} \right). \quad (4.12)$$

Proof To begin with, note that $|\widehat{g}_i - h_i| \leq \epsilon$ implies $|\widehat{g}_i| \in [1 - \epsilon, 1 + \epsilon]$. This means that $|\widehat{g}_i| > 0$ holds if $\epsilon < 1$. Consequently, we obtain

$$\left| \frac{\widehat{g}_i}{|\widehat{g}_i|} - h_i \right| = \left| \frac{\widehat{g}_i}{|\widehat{g}_i|} - \frac{h_i}{|\widehat{g}_i|} + \frac{h_i}{|\widehat{g}_i|} - h_i \right| \quad (4.13)$$

$$\leq \frac{|\widehat{g}_i - h_i|}{|\widehat{g}_i|} + |h_i| \left(\frac{||\widehat{g}_i| - 1|}{|\widehat{g}_i|} \right) \quad (4.14)$$

$$\leq \frac{2\epsilon}{|\widehat{g}_i|} \leq \frac{2\epsilon}{1 - \epsilon}. \quad (4.15)$$

We will now show that provided $0 < \epsilon < 1/2$ holds, then (4.15) implies the bound (4.12). Indeed, from the definition of h_i , and of $\widehat{g}_i/|\widehat{g}_i|$ (from (4.10)), we have

$$\left| \frac{\widehat{g}_i}{|\widehat{g}_i|} - h_i \right| = |\exp(\iota 2\pi(\widehat{f}_i \bmod 1)) - \exp(\iota 2\pi(f_i \bmod 1))| \quad (4.16)$$

$$= |1 - \exp(\iota 2\pi(f_i \bmod 1 - \widehat{f}_i \bmod 1))| \quad (4.17)$$

$$= 2|\sin[\pi \underbrace{(f_i \bmod 1 - \widehat{f}_i \bmod 1)}_{\in(-1,1)}]| \quad (4.18)$$

$$= 2\sin[\pi|(f_i \bmod 1 - \widehat{f}_i \bmod 1)|] \quad (4.19)$$

$$= 2\sin[\pi(1 - |(f_i \bmod 1 - \widehat{f}_i \bmod 1)|)]. \quad (4.20)$$

Then, (4.12) follows from (4.20), (4.15) and by noting that $0 < \epsilon/(1 - \epsilon) < 1$ for $0 < \epsilon < 1/2$. ■

4.2. Unwrapping stage and main algorithm

Having recovered the denoised mod 1 samples $\widehat{f}_i \bmod 1$ for $i = 1, \dots, n$, we now move onto the next stage of our method where the goal is to recover the samples f , for which we discuss two possible approaches.

1. Quotient tracker (QT) method. The first approach for unwrapping the mod 1 samples is perhaps the most natural one. It is based on the idea that, provided the denoised

mod 1 samples are very close estimates to the original clean mod 1 samples, then we can sequentially find the quotient terms, by checking whether $|\widehat{f}_{i+1} \bmod 1 - \widehat{f}_i \bmod 1| \geq \zeta$, for a suitable threshold parameter $\zeta \in (0, 1)$. More formally, after initializing $\widehat{q}_1 = 0$, consider the rule

$$\begin{aligned} \widehat{q}_{i+1} &= \widehat{q}_i + \text{sign}_\zeta(\widehat{f}_{i+1} \bmod 1 - \widehat{f}_i \bmod 1); \\ \text{sign}_\zeta(t) &= \begin{cases} -1; & t \geq \zeta \\ 0; & |t| < \zeta \\ 1; & t \leq -\zeta \end{cases}. \end{aligned} \quad (4.21)$$

Clearly, if $\widehat{f}_i \bmod 1 \approx f_i \bmod 1$ for each i , then for n sufficiently large, the procedure (4.21) will result in correct recovery of the quotients. However, as illustrated in Figure 3, it is also obviously sensitive to noise, and hence would not be a robust solution for high levels of noise.

2. Ordinary least-squares (OLS) based method. A robust alternative to the aforementioned approach is based on directly recovering the function via a simple least-squares problem. Recall that in the noise-free case, $f_i = q_i + r_i$, $q_i \in \mathbb{Z}$, $r_i \in [0, 1)$, and consider, for a pair of nearby points i, j , the difference

$$f_i - f_j = q_i - q_j + r_i - r_j. \quad (4.22)$$

The OLS formulation we solve stems from the observation that, if $|r_i - r_j| < \zeta$ for a small ζ , then $q_i = q_j$. This intuition can be easily gained from the top plots of Figure 4, especially 4a, which pertains to the noisy case (but in the low noise regime $\gamma = 0.15$), that plots $l_{i+1} - l_i$ versus $y_i - y_{i+1}$, where l_i denotes the noisy quotient of sample i , and y_i the noisy remainder. For small enough $|y_i - y_{i+1}|$, we observe that $|l_{i+1} - l_i| = 0$. Whenever $y_i - y_{i+1} > \zeta$, we see that $l_{i+1} - l_i = 1$, while $y_i - y_{i+1} < -\zeta$, indicates that $l_{i+1} - l_i = -1$. Throughout all our experiments we set $\zeta = 0.5$. In Figure 23 we also plot the true quotient q , which can be observed to be piecewise constant, in agreement with our above intuition.

For a graph⁶ $G = (V, E)$ with $k \in \mathbb{N}$, and for a suitable threshold parameter $\zeta \in (0, 1)$, this intuition leads us to estimate the function values f_i as the least-squares solution to an overdetermined system of linear equations without involving the quotients q_1, \dots, q_n . To this end, we consider a linear system of equations involving the differences $\widehat{f}_i - \widehat{f}_j$, $\forall \{i, j\} \in E$

$$\widehat{f}_i - \widehat{f}_j = l_i - l_j + y_i - y_j = \text{sign}_\zeta(y_i - y_j) + y_i - y_j, \quad (4.23)$$

and solve it in the least-squares sense. (4.23) is analogous to (4.21), except that we now recover $(\widehat{f}_i)_{i=1}^n$ collectively as the least-squares solution to (4.23). Denoting by T the least-squares matrix associated with the overdetermined linear system (4.23), $b_{i,j} = \text{sign}_\zeta(y_i - y_j) + y_i - y_j$, and $\widehat{f}_{i,j} = \widehat{f}_i - \widehat{f}_j$, the system of equations can be written as $T\widehat{\mathbf{f}} = \mathbf{b}$ where $\widehat{\mathbf{f}}, \mathbf{b} \in \mathbb{R}^{|E|}$. Note that the matrix T is sparse with only two non-zero entries per row, and that the all-ones vector $\mathbf{1} = (1, 1, \dots, 1)^T$ lies in the null space of T , i.e., $T\mathbf{1} = 0$. Therefore, we will find the minimum norm least-squares solution to (4.23), and recover f_i 's only up to a global shift. We remark that the above line of thought, initiated by (4.22) is very similar

6. In the sequel, we will consider the same graph G as for the denoising stage. But this is of course not necessary, one can consider a different graph here as well.




Figure 3: Recovery of the estimated quotient at each sample point, via the Quotient Tracker (QT) algorithm (4.21) and the Ordinary Least Squares (OLS) (4.23), when the input is given by the denoised modulo samples produced by our QCQP method in Algorithm 1, for which the unwrapping stage is performed via OLS (4.23). Different columns pertain to varying levels of noise under the Gaussian noise model. We keep fixed $\lambda = 0.1$, and $k = 3$. We also plot the difference between consecutive noisy mod 1 measurements, $\delta_i = y_i - y_{i+1}$, highlighting the fact that QT lacks robustness at high levels of noise.

to Step 3 of the ASAP algorithm introduced by Cucuringu et al. (2012a) (Section 3.3) in the context of the graph realization problem, that performed synchronization over \mathbb{R}^d to estimate the unknown translations. Algorithm 1 summarizes our two-stage method for recovering the samples of f (up to a global shift). Figure 4 shows additional noisy instances of the Uniform noise model. The scatter plots on the first row show that, as the noise level increases, the function (4.21) will produce more and more errors in (4.23). The remaining plots show the corresponding $f \bmod 1$ signal (clean, noisy, and denoised via Algorithm 1) for three levels of noise.

5. Analysis for the arbitrary bounded noise model

This section provides approximation guarantees for the solution $\hat{\mathbf{g}} \in \mathbb{R}^{2n}$ to (4.5) for the arbitrary bounded noise model (2.3). In particular, we consider a slightly modified version

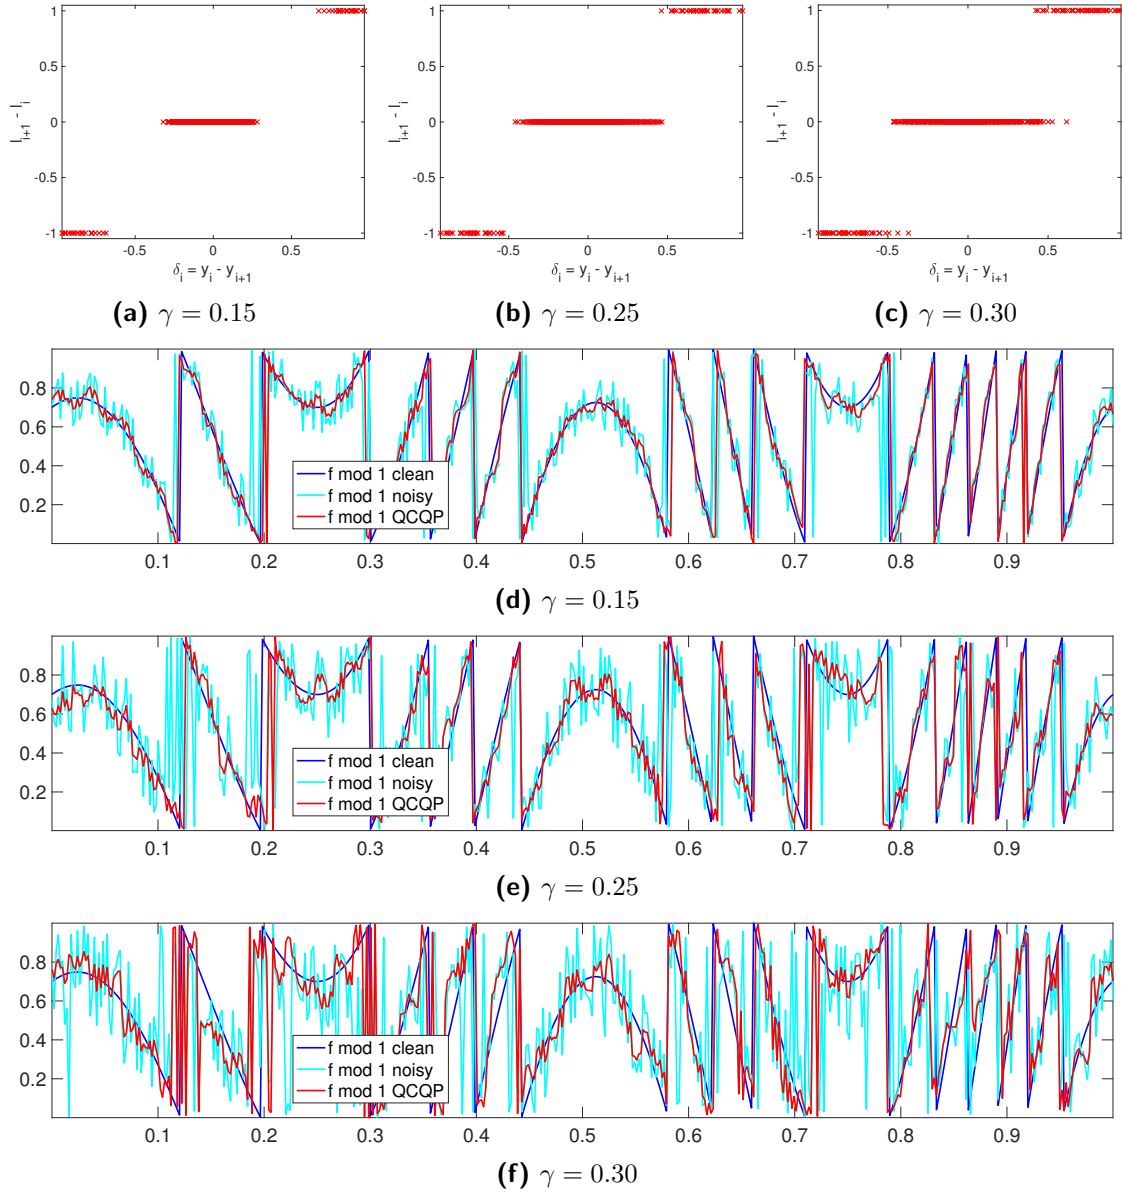


Figure 4: Noisy instances of the Uniform noise model ($n = 500$). Top row (a)-(c) shows scatter plots of change in y (the observed noisy $f \bmod 1$ values) versus change in l (the noisy quotient). Plots (d)-(f) show the clean $f \bmod 1$ values (blue), the noisy $f \bmod 1$ values (cyan) and the denoised (via **QCQP**) $f \bmod 1$ values (red), for increasing levels of noise. **QCQP** denotes Algorithm 1 without the unwrapping stage performed by **OLS** (4.23).

Algorithm 1 Algorithm for recovering the samples f_i

1: **Input:** $(y_i)_{i=1}^n$ (noisy mod 1 samples), $k, \lambda, n, G = (V, E)$.

2: **Output:** Denoised mod 1 samples $\widehat{f}_i \bmod 1; i = 1, \dots, n$.

// STAGE 1: RECOVERING DENOISED MOD 1 SAMPLES OF f .

3: Form $H \in R^{2n \times 2n}$ using λ, L as in (4.4).

4: Form $\bar{\mathbf{z}} = [\text{Re}(\mathbf{z})^T \text{Im}(\mathbf{z})^T]^T \in \mathbb{R}^{2n}$ as in (4.3).

5: Obtain $\widehat{\mathbf{g}} \in \mathbb{R}^{2n}$ as the solution to (4.5), i.e.,

$$\widehat{\mathbf{g}} = \underset{\mathbf{g} \in \mathbb{R}^{2n}: \|\mathbf{g}\|_2^2 = n}{\text{argmin}} \quad \mathbf{g}^T H \mathbf{g} - 2\mathbf{g}^T \bar{\mathbf{z}}.$$

6: Obtain $\widehat{\mathbf{g}} \in \mathbb{C}^n$ from $\widehat{\mathbf{g}}$ where $\widehat{\mathbf{g}} = [\text{Re}(\widehat{\mathbf{g}})^T \text{Im}(\widehat{\mathbf{g}})^T]^T$.

7: Recover $\widehat{f}_i \bmod 1 \in [0, 1)$ from $\frac{\widehat{g}_i}{|\widehat{g}_i|}$ for each $i = 1, \dots, n$, as in (4.10).

// STAGE 2: RECOVERING DENOISED REAL VALUED SAMPLES OF f .

8: **Input:** $(\widehat{f}_i \bmod 1)_{i=1}^n$ (denoised mod 1 samples), $G = (V, E), \zeta \in (0, 1)$.

9: **Output:** Denoised samples $\widehat{f}_i; i = 1, \dots, n$.

10: Obtain $(\widehat{f}_i)_{i=1}^n$ via the Quotient Tracker (QT) or OLS based method for suitable threshold ζ .

of this model, assuming

$$\|\bar{\mathbf{z}} - \bar{\mathbf{h}}\|_2 \leq \delta\sqrt{n} \tag{5.1}$$

holds true for some $\delta \in [0, 1]$. This is reasonable, since $\|\bar{\mathbf{z}} - \bar{\mathbf{h}}\|_2 \leq 2\sqrt{n}$ holds in general by triangle inequality. Also, note for (2.3) that $|z_i - h_i| = 2|\sin(\pi(\delta_i \bmod 1))| \leq 2\pi|\delta_i|$, and thus $\|\bar{\mathbf{z}} - \bar{\mathbf{h}}\|_2 = \|\mathbf{z} - \mathbf{h}\|_2 \leq 2\pi \max_i(|\delta_i|)\sqrt{n}$. Hence, while a small enough uniform bound on $\max_i(|\delta_i|)$ would of course imply (5.1), however, clearly (5.1) can also hold even if some of the δ_i 's are large.

Theorem 5 *Under the above notation and assumptions, consider the arbitrary bounded noise model in (2.3), with $\bar{\mathbf{z}}$ satisfying $\|\bar{\mathbf{z}} - \bar{\mathbf{h}}\|_2 \leq \delta\sqrt{n}$ for $\delta \in [0, 1]$. Let $n \geq 2$, and let $\mathcal{N}(H)$ denote the null space of H .*

1. *If $\bar{\mathbf{z}} \notin \mathcal{N}(H)$ then $\widehat{\mathbf{g}}$ is the unique solution to (4.5) satisfying*

$$\frac{1}{n} \langle \bar{\mathbf{h}}, \widehat{\mathbf{g}} \rangle \geq 1 - \frac{3\delta}{2} - \frac{\lambda\pi^2 M^2 (2k)^{2\alpha+1}}{n^{2\alpha}} + \frac{1}{(4\lambda k + 1)^2} \left(\frac{1}{2n} \bar{\mathbf{z}}^T H \bar{\mathbf{z}} \right). \tag{5.2}$$

2. *If $\bar{\mathbf{z}} \perp \mathcal{N}(H)$ and $\lambda < \frac{1}{4k}$ then $\widehat{\mathbf{g}}$ is the unique solution to (4.5) satisfying*

$$\frac{1}{n} \langle \bar{\mathbf{h}}, \widehat{\mathbf{g}} \rangle \geq 1 - \frac{3\delta}{2} - \frac{\lambda\pi^2 M^2 (2k)^{2\alpha+1}}{n^{2\alpha}} + \frac{1}{(1 + 4\lambda k - 4\lambda k \sin^2(\frac{\pi}{2n}))^2} \left(\frac{1}{2n} \bar{\mathbf{z}}^T H \bar{\mathbf{z}} \right). \tag{5.3}$$

The following useful Corollary of Theorem 5 is a direct consequence of the fact that $1/(2n)\bar{\mathbf{z}}^T H \bar{\mathbf{z}} \geq 0$ for all $\bar{\mathbf{z}} \in \mathbb{R}^{2n}$, since H is positive semi-definite.

Corollary 6 Consider the arbitrary bounded noise model defined in (2.3), with $\bar{\mathbf{z}}$ satisfying $\|\bar{\mathbf{z}} - \bar{\mathbf{h}}\|_2 \leq \delta\sqrt{n}$ for $\delta \in [0, 1]$. Let $n \geq 2$. If $\lambda < \frac{1}{4k}$ then $\hat{\mathbf{g}}$ is the unique solution to (4.5) satisfying

$$\frac{1}{n} \langle \bar{\mathbf{h}}, \hat{\mathbf{g}} \rangle \geq 1 - \frac{3\delta}{2} - \frac{\lambda\pi^2 M^2 (2k)^{2\alpha+1}}{n^{2\alpha}}. \quad (5.4)$$

Let us note that the above bounds are meaningful only when δ is small enough, specifically $\delta < 2/3$. Before presenting the proof of Theorem 5, some remarks are in order.

1. Theorem 5 give us a lower bound on the correlation between $\bar{\mathbf{h}}, \hat{\mathbf{g}} \in \mathbb{R}^{2n}$, where clearly, $\frac{1}{n} \langle \bar{\mathbf{h}}, \hat{\mathbf{g}} \rangle \in [-1, 1]$. Note that the correlation improves when the noise term δ decreases, as one would expect. The term $\frac{\lambda\pi^2 M^2 (2k)^{2\alpha+1}}{n^{2\alpha}}$ effectively arises on account of the smoothness of f , and is an upper bound on the term $\frac{1}{2n} \bar{\mathbf{h}}^T H \bar{\mathbf{h}}$ (made clear in Lemma 8). Hence, as the number of samples increases, $\frac{1}{2n} \bar{\mathbf{h}}^T H \bar{\mathbf{h}}$ goes to zero at the rate $n^{-2\alpha}$ (for fixed k, λ). Also note that the lower bound on $\frac{1}{n} \langle \bar{\mathbf{h}}, \hat{\mathbf{g}} \rangle$ readily implies the ℓ_2 norm bound $\|\hat{\mathbf{g}} - \bar{\mathbf{h}}\|_2^2 = O(\delta n + n^{1-2\alpha})$.
2. The term $\frac{1}{2n} \bar{\mathbf{z}}^T H \bar{\mathbf{z}}$ represents the smoothness of the observed noisy samples. While an increasing amount of noise would usually render $\bar{\mathbf{z}}$ to be more and more non-smooth, and thus typically increase $\frac{1}{2n} \bar{\mathbf{z}}^T H \bar{\mathbf{z}}$, note that this would be met by a corresponding increase in δ , and hence the lower bound on the correlation would not necessarily improve.
3. It is easy to verify that (5.1) implies $\langle \bar{\mathbf{z}}, \bar{\mathbf{h}} \rangle / n \geq 1 - \frac{\delta^2}{2} \geq 1 - \frac{\delta}{2}$. Thus for $\bar{\mathbf{z}}$, which is feasible for (4.5), we have a bound on correlation which is better than the bound in Corollary 6 by a $\delta + O(n^{-2\alpha})$ term. However, the solution $\hat{\mathbf{g}}$ to (4.5) is a *smooth* estimate of $\bar{\mathbf{h}}$ (and hence more interpretable), while $\bar{\mathbf{z}}$ is typically highly non-smooth.

Proof [Proof of Theorem 5] The proof of Theorem 5 relies heavily on Lemma 7 outlined below.

Lemma 7 Consider the arbitrary bounded noise model in (2.3), with $\bar{\mathbf{z}}$ satisfying $\|\bar{\mathbf{z}} - \bar{\mathbf{h}}\|_2 \leq \delta\sqrt{n}$ for $\delta \in [0, 1]$. Any solution $\hat{\mathbf{g}}$ to (4.5) satisfies

$$\frac{1}{n} \langle \bar{\mathbf{h}}, \hat{\mathbf{g}} \rangle \geq 1 - \frac{3\delta}{2} - \frac{1}{2n} \bar{\mathbf{h}}^T H \bar{\mathbf{h}} + \frac{1}{2n} \hat{\mathbf{g}}^T H \hat{\mathbf{g}}. \quad (5.5)$$

Proof To begin with, note that

$$\|\bar{\mathbf{z}} - \bar{\mathbf{h}}\|_2 \leq \delta\sqrt{n} \Leftrightarrow \langle \bar{\mathbf{z}}, \bar{\mathbf{h}} \rangle \geq n - \frac{\delta^2 n}{2}. \quad (5.6)$$

Since $\bar{\mathbf{h}} \in \mathbb{R}^{2n}$ is feasible for (4.5), we get

$$\bar{\mathbf{h}}^T H \bar{\mathbf{h}} - 2 \langle \bar{\mathbf{h}}, \bar{\mathbf{z}} \rangle \geq \hat{\mathbf{g}}^T H \hat{\mathbf{g}} - 2 \langle \hat{\mathbf{g}}, \bar{\mathbf{z}} \rangle \quad (5.7)$$

$$\Leftrightarrow \langle \hat{\mathbf{g}}, \bar{\mathbf{z}} \rangle \geq \langle \bar{\mathbf{h}}, \bar{\mathbf{z}} \rangle - \frac{1}{2} \bar{\mathbf{h}}^T H \bar{\mathbf{h}} + \frac{1}{2} \hat{\mathbf{g}}^T H \hat{\mathbf{g}} \quad (5.8)$$

$$\geq n - \frac{\delta^2 n}{2} - \frac{1}{2} \bar{\mathbf{h}}^T H \bar{\mathbf{h}} + \frac{1}{2} \hat{\mathbf{g}}^T H \hat{\mathbf{g}} \quad (\text{from (5.6)}). \quad (5.9)$$

Moreover, we can upper bound $\langle \widehat{\mathbf{g}}, \bar{\mathbf{z}} \rangle$ as follows.

$$\langle \widehat{\mathbf{g}}, \bar{\mathbf{z}} \rangle = \langle \widehat{\mathbf{g}}, \bar{\mathbf{z}} - \bar{\mathbf{h}} \rangle + \langle \widehat{\mathbf{g}}, \bar{\mathbf{h}} \rangle \quad (5.10)$$

$$\leq \| \widehat{\mathbf{g}} \|_2 \| \bar{\mathbf{z}} - \bar{\mathbf{h}} \|_2 + \langle \widehat{\mathbf{g}}, \bar{\mathbf{h}} \rangle \quad (\text{Cauchy-Schwarz}) \quad (5.11)$$

$$\leq \sqrt{n}(\sqrt{n}\delta) + \langle \widehat{\mathbf{g}}, \bar{\mathbf{h}} \rangle \quad (\text{from (5.6)}). \quad (5.12)$$

Plugging (5.12) in (5.9) and using $\delta^2 \leq \delta$ for $\delta \in [0, 1]$ completes the proof. \blacksquare

We now upper bound the term $\frac{1}{2n} \bar{\mathbf{h}}^T H \bar{\mathbf{h}}$ in (5.5) using the Hölder continuity of f . This is formally shown below in Lemma 8.

Lemma 8 *For $n \geq 2$, it holds true that*

$$\frac{1}{2n} \bar{\mathbf{h}}^T H \bar{\mathbf{h}} \leq \frac{\lambda \pi^2 M^2 (2k)^{2\alpha+1}}{n^{2\alpha}}, \quad (5.13)$$

where $\alpha \in (0, 1]$ and $M > 0$ are related to the smoothness of f and defined in (2.6), and $\lambda \geq 0$ is the regularization parameter in (3.2).

Proof Denoting $\mathbf{h} = (h_1 \dots h_n)^T \in \mathbb{C}^n$ to be the complex valued representation of $\bar{\mathbf{h}} \in \mathbb{R}^{2n}$ as per (4.3), clearly

$$\frac{1}{2n} \bar{\mathbf{h}}^T H \bar{\mathbf{h}} = \frac{1}{2n} \mathbf{h}^* (\lambda L) \mathbf{h} = \frac{\lambda}{2n} \sum_{\{i,j\} \in E} |h_i - h_j|^2 \leq \frac{\lambda}{2n} |E| \max_{\{i,j\} \in E} |h_i - h_j|^2, \quad (5.14)$$

where the first equality is shown in Appendix A. Since for each $i \in V$ we have $\deg(i) \leq 2k$, hence $|E| = (1/2) \sum_{i \in V} \deg(i) \leq kn$. Next, for any $\{i, j\} \in E$ note that by Hölder continuity of f we have

$$|f_i - f_j| \leq M |x_i - x_j|^\alpha \leq M \left(\frac{k}{n-1} \right)^\alpha \leq M \left(\frac{2k}{n} \right)^\alpha, \quad (5.15)$$

if $n \geq 2$ (since then $n-1 \geq n/2$). Finally, we can bound $|h_i - h_j|$ as follows.

$$|h_i - h_j| = |1 - \exp(i2\pi(f_j - f_i))| \quad (5.16)$$

$$= 2|\sin \pi(f_j - f_i)| \quad (5.17)$$

$$\leq 2\pi |f_j - f_i| \quad (\text{since } |\sin x| \leq |x|; \forall x \in \mathbb{R}) \quad (5.18)$$

$$\leq \frac{2\pi M (2k)^\alpha}{n^\alpha} \quad (\text{using (5.15)}). \quad (5.19)$$

Plugging (5.19) in (5.14) with the bound $|E| \leq kn$ yields the desired bound. \blacksquare

Lastly, we lower bound the term $\frac{1}{2n} \widehat{\mathbf{g}}^T H \widehat{\mathbf{g}}$ in (5.5) using knowledge of the structure of the solution $\widehat{\mathbf{g}}$. This is outlined below as Lemma 9.

Lemma 9 *Denoting $\mathcal{N}(H)$ to be the null space of H , the following holds for the solution $\widehat{\mathbf{g}}$ to (4.5).*

1. *If $\bar{\mathbf{z}} \notin \mathcal{N}(H)$ then $\widehat{\mathbf{g}}$ is unique and*

$$\frac{1}{2n} \widehat{\mathbf{g}}^T H \widehat{\mathbf{g}} \geq \frac{1}{(1 + 4\lambda k)^2} \left(\frac{1}{2n} \bar{\mathbf{z}}^T H \bar{\mathbf{z}} \right). \quad (5.20)$$

2. If $\bar{\mathbf{z}} \perp \mathcal{N}(H)$ and $\lambda < \frac{1}{4k}$, then $\widehat{\mathbf{g}}$ is unique and

$$\frac{1}{2n} \widehat{\mathbf{g}}^T H \widehat{\mathbf{g}} \geq \frac{1}{(1 + 4\lambda k - 4\lambda k \sin^2(\frac{\pi}{2n}))^2} \left(\frac{1}{2n} \bar{\mathbf{z}}^T H \bar{\mathbf{z}} \right). \quad (5.21)$$

Proof Let $\{\lambda_j(H)\}_{j=1}^{2n}$ (with $\lambda_1(H) \leq \lambda_2(H) \leq \dots$) and $\{\mathbf{q}_j\}_{j=1}^{2n}$ denote the eigenvalues and eigenvectors respectively for H . Also, let $0 = \beta_1(L) < \beta_2(L) \leq \beta_3(L) \leq \dots \leq \beta_n(L)$ denote the eigenvalues of the Laplacian L . Note that $\beta_2(L) > 0$ since G is connected. By Gershgorin's disk theorem, it is easy to see⁷ that $\beta_n(L) \leq 4k$ for the graph G . Hence,

$$0 = \lambda_1(H) = \lambda_2(H) < \lambda_3(H) \leq \dots \leq \lambda_{2n}(H) \leq 4\lambda k \quad (5.22)$$

and $\mathcal{N}(H) = \text{span}\{\mathbf{q}_1, \mathbf{q}_2\}$. We now consider the two cases separately below.

1. Consider the case where $\bar{\mathbf{z}} \not\perp \mathcal{N}(H)$. We know that $\widehat{\mathbf{g}} = 2(2H + \mu^* I)^{-1} \bar{\mathbf{z}}$ for a unique $\mu^* \in (0, \infty)$ (and so $\widehat{\mathbf{g}}$ is the unique solution to (4.5) by Lemma 3 since $2H + \mu^* I \succ 0$) satisfying

$$\|\widehat{\mathbf{g}}\|_2^2 = 4 \sum_{j=1}^{2n} \frac{\langle \bar{\mathbf{z}}, \mathbf{q}_j \rangle^2}{(2\lambda_j(H) + \mu^*)^2} = n. \quad (5.23)$$

Since $\lambda_j(H) \geq 0$ for all j , we obtain from (5.23) that

$$n \leq 4 \sum_{j=1}^{2n} \frac{\langle \bar{\mathbf{z}}, \mathbf{q}_j \rangle^2}{(\mu^*)^2} = \frac{4n}{(\mu^*)^2} \implies \mu^* \leq 2. \quad (5.24)$$

Note that equality holds in (5.24) if $\bar{\mathbf{z}} \in \mathcal{N}(H)$. We can now lower bound $\frac{1}{2n} \widehat{\mathbf{g}}^T H \widehat{\mathbf{g}}$ as follows

$$\frac{1}{2n} \widehat{\mathbf{g}}^T H \widehat{\mathbf{g}} = \frac{2}{n} \bar{\mathbf{z}}^T (2H + \mu^* I)^{-1} H (2H + \mu^* I)^{-1} \bar{\mathbf{z}} \quad (5.25)$$

$$= \frac{2}{n} \sum_{j=1}^{2n} \frac{\langle \bar{\mathbf{z}}, \mathbf{q}_j \rangle^2 \lambda_j(H)}{(2\lambda_j(H) + \mu^*)^2} \quad (5.26)$$

$$\geq \frac{2}{n(8\lambda k + 2)^2} \sum_{j=1}^{2n} \langle \bar{\mathbf{z}}, \mathbf{q}_j \rangle^2 \lambda_j(H) \quad (\text{from (5.22), (5.24)}) \quad (5.27)$$

$$= \frac{1}{(4\lambda k + 1)^2} \left(\frac{1}{2n} \bar{\mathbf{z}}^T H \bar{\mathbf{z}} \right). \quad (5.28)$$

2. Let us now consider the case where $\bar{\mathbf{z}} \perp \mathcal{N}(H)$. Denote

$$\phi(\mu) := \|\widehat{\mathbf{g}}(\mu)\|_2^2 = 4 \sum_{j=1}^{2n} \frac{\langle \bar{\mathbf{z}}, \mathbf{q}_j \rangle^2}{(2\lambda_j(H) + \mu)^2} = 4 \sum_{j=3}^{2n} \frac{\langle \bar{\mathbf{z}}, \mathbf{q}_j \rangle^2}{(2\lambda_j(H) + \mu)^2}. \quad (5.29)$$

7. Denote L_{ij} to be the (i, j) th entry of L . Then by Gershgorin's disk theorem, we know that each eigenvalue lies in $\bigcup_{i=1}^{2n} \{x : |x - L_{ii}| \leq \sum_{j \neq i} |L_{ij}|\}$. Since $L_{ii} \leq 2k$ and $\sum_{j \neq i} |L_{ij}| \leq 2k$ holds for each i , the claim follows.

Observe that ϕ does not have a pole at 0 anymore, $\phi(0)$ is well defined. In order to have a unique $\widehat{\mathbf{g}}$, it is sufficient if $\phi(0) > n$ holds. Indeed, we would then have a unique $\mu^* \in (0, \infty)$ such that $\|\widehat{\mathbf{g}}(\mu^*)\|_2^2 = n$. Consequently, $\widehat{\mathbf{g}}(\mu^*)$ will be the unique solution to (4.5) by Lemma 3 since $2H + \mu^*I \succ 0$. Now let us note that

$$\phi(0) = \sum_{j=3}^{2n} \frac{\langle \bar{\mathbf{z}}, \mathbf{q}_j \rangle^2}{\lambda_j(H)^2} \geq \frac{n}{16\lambda^2 k^2} \quad (5.30)$$

since $\lambda_j(H) \leq 4\lambda k$ for all j (recall (5.22)). Therefore clearly, the choice $\lambda < \frac{1}{4k}$ implies $\phi(0) > n$, and consequently that the solution $\widehat{\mathbf{g}}$ is unique. Assuming $\lambda < \frac{1}{4k}$ holds, we can derive an upper bound on μ^* as follows

$$n = 4 \sum_{j=3}^{2n} \frac{\langle \bar{\mathbf{z}}, \mathbf{q}_j \rangle^2}{(2\lambda_j(H) + \mu^*)^2} \leq 4 \sum_{j=3}^{2n} \frac{\langle \bar{\mathbf{z}}, \mathbf{q}_j \rangle^2}{(2\lambda_3(H) + \mu^*)^2} = \frac{4n}{(2\lambda_3(H) + \mu^*)^2} \quad (5.31)$$

$$\Rightarrow \mu^* \leq 2 - 2\lambda_3(H). \quad (5.32)$$

Hence $\mu^* \in (0, 2 - 2\lambda_3(H))$ when $\lambda < \frac{1}{4k}$. We can now lower bound $\frac{1}{2n} \widehat{\mathbf{g}}^T H \widehat{\mathbf{g}}$ in the same manner as before

$$\frac{1}{2n} \widehat{\mathbf{g}}^T H \widehat{\mathbf{g}} = \frac{2}{n} \bar{\mathbf{z}}^T (2H + \mu^*I)^{-1} H (2H + \mu^*I)^{-1} \bar{\mathbf{z}} \quad (5.33)$$

$$= \frac{2}{n} \sum_{j=1}^{2n} \frac{\langle \bar{\mathbf{z}}, \mathbf{q}_j \rangle^2 \lambda_j(H)}{(2\lambda_j(H) + \mu^*)^2} \quad (5.34)$$

$$\geq \frac{2}{n(8\lambda k + 2 - 2\lambda_3(H))^2} \sum_{j=1}^{2n} \langle \bar{\mathbf{z}}, \mathbf{q}_j \rangle^2 \lambda_j(H) \quad (\text{from (5.22), (5.32)}) \quad (5.35)$$

$$= \frac{1}{(1 + 4\lambda k - \lambda_3(H))^2} \left(\frac{1}{2n} \bar{\mathbf{z}}^T H \bar{\mathbf{z}} \right). \quad (5.36)$$

It remains to lower bound $\lambda_3(H) = \lambda\beta_2(L)$. We do this by using the following result of Fiedler (1973) (adjusted to our notation) for lower bounding the second eigenvalue of the Laplacian of a simple graph.

Theorem 10 (Fiedler (1973)) *Let G be a simple graph of order n other than a complete graph, with vertex connectivity $\kappa(G)$ and edge connectivity $\kappa'(G)$. It holds true that*

$$2\kappa'(G)(1 - \cos(\pi/n)) \leq \beta_2(L) \leq \kappa(G) \leq \kappa'(G). \quad (5.37)$$

The graph G in our setting has $\kappa(G) = k$. Indeed, there does not exist a vertex cut of size $k - 1$ or less, but there does exist a vertex cut of size k . This in turn implies that $\kappa'(G) \geq k$, and Theorem 10 yields

$$\beta_2(L) \geq 2k(1 - \cos(\pi/n)) = 4k \sin^2 \left(\frac{\pi}{2n} \right). \quad (5.38)$$

Hence $\lambda_3(H) \geq 4\lambda k \sin^2 \left(\frac{\pi}{2n} \right)$. Plugging this in to (5.32) completes the proof.

■
■

This completes the proof of Theorem 5.

6. Analysis for random noise models

We now analyze two random noise models, namely the Bernoulli-Uniform model described in (2.4), and the Gaussian noise model described in (2.5).

6.1. Analysis for the Bernoulli-Uniform random noise model

Let $u_1, \dots, u_n \sim U[0, 1]$ i.i.d, where $U[0, 1]$ denotes the uniform distribution over $[0, 1]$. Also, let $\beta_1, \beta_2, \dots, \beta_n$ be i.i.d Bernoulli random variables where $\beta_i = 1$ with probability p , and 0 with probability $1 - p$, for some $p \in [0, 1]$. Then (2.4) is equivalent to

$$y_i = \begin{cases} f_i \bmod 1 & ; & \text{if } \beta_i = 0 \\ u_i & ; & \text{if } \beta_i = 1 \end{cases} \quad ; \quad i = 1, \dots, n. \quad (6.1)$$

The following theorem is our main result for this noise model.

Theorem 11 *Consider the Bernoulli-Uniform noise model in (6.1) for some $p \in [0, 1]$. If $\lambda < \frac{1}{4k}$ then $\widehat{\mathbf{g}}$ is the unique solution to (4.5). Moreover, assuming $n \geq 2$, then for any $\varepsilon \in (0, 1/2)$ satisfying $p + \varepsilon \leq 1/2$ and absolute constants $c, c' > 0$, the following is true.*

1.

$$\frac{1}{n} \langle \bar{\mathbf{h}}, \widehat{\mathbf{g}} \rangle \geq 1 - \left(3\sqrt{\frac{p + \varepsilon}{2}} - \frac{\lambda k(1 - \varepsilon)}{6(4\lambda k + 1)^2 p} \right) - \frac{\lambda \pi^2 M^2 (2k)^{2\alpha+1}}{n^{2\alpha}} \left(1 - \frac{(1 - p)^2}{(4\lambda k + 1)^2} \right) \quad (6.2)$$

with probability at least

$$1 - e \cdot \exp\left(-\frac{c'(1 - p)^2 \varepsilon^2 n}{16}\right) - 4 \exp\left(-\frac{c\varepsilon^2 p^2 n}{512}\right) - 2e \cdot \exp\left(-\frac{c' p^2 n (1 - \varepsilon)^2}{2304}\right). \quad (6.3)$$

2.

$$\frac{1}{n} \langle \bar{\mathbf{h}}, \widehat{\mathbf{g}} \rangle \geq 1 - \left(3\sqrt{\frac{p + \varepsilon}{2}} \right) - \frac{\lambda \pi^2 M^2 (2k)^{2\alpha+1}}{n^{2\alpha}} \quad (6.4)$$

with probability at least

$$1 - e \cdot \exp\left(-\frac{c'(1 - p)^2 \varepsilon^2 n}{16}\right). \quad (6.5)$$

Both (6.2) and 6.4 give lower bounds on the correlation between $\bar{\mathbf{h}}, \widehat{\mathbf{g}}$, holding with high probability over the randomness of the noise samples provided p is small enough. However, note that the bound in (6.2) is strictly better than that in (6.4), albeit with a worse bound on success probability. The usefulness of (6.2) is seen when $p = \Theta(1)$ and for large n , in which case (6.2) will hold with a sufficiently large probability, thus giving a better result

than (6.4). In contrast, the bound in (6.4) is useful for all values of $p \in [0, 1/2)$, and moderately large values of n . This is especially important when p is close to 0 and n is moderately large. In that scenario, the bound on the success probability for (6.2) becomes trivial.

Proof [Proof of Theorem 11] Our starting point is the following (slightly restricted) version of Theorem 5, which follows in a straightforward manner from Lemma 7 and Lemma 9.

Theorem 12 *Consider the arbitrary bounded noise model defined in (2.3), with $\bar{\mathbf{z}}$ satisfying $\|\bar{\mathbf{z}} - \bar{\mathbf{h}}\|_2 \leq \delta\sqrt{n}$ for $\delta \in [0, 1]$. If $\lambda < \frac{1}{4k}$, then $\hat{\mathbf{g}}$ is the unique solution to (4.5) satisfying*

$$\frac{1}{n} \langle \bar{\mathbf{h}}, \hat{\mathbf{g}} \rangle \geq 1 - \frac{3\delta}{2} - \frac{1}{2n} \bar{\mathbf{h}}^T H \bar{\mathbf{h}} + \frac{1}{(1 + 4\lambda k)^2} \left(\frac{1}{2n} \bar{\mathbf{z}}^T H \bar{\mathbf{z}} \right). \quad (6.6)$$

Next, let us recall from (3.1) that $z_i = \exp(\iota 2\pi y_i)$, and so $(z_i)_{i=1}^n$ are independent, complex-valued random variables. Consequently, we can upper bound the noise term parameter δ , and lower bound the quadratic form $\frac{1}{2n} \bar{\mathbf{z}}^T H \bar{\mathbf{z}}$, w.h.p. This is stated precisely in the following Proposition.

Proposition 13 *Consider the Bernoulli-Uniform noise model in (6.1) for some $p \in [0, 1]$. For absolute constants $c, c' > 0$, the following is true.*

1. For any $\varepsilon \in (0, 1)$,

$$\frac{1}{2n} \bar{\mathbf{z}}^T H \bar{\mathbf{z}} \geq \frac{\lambda p k (1 - \varepsilon)}{6} + (1 - p)^2 \frac{1}{2n} \bar{\mathbf{h}}^T H \bar{\mathbf{h}} \quad (6.7)$$

holds with probability at least

$$1 - 4 \exp\left(-\frac{c\varepsilon^2 p^2 n}{512}\right) - 2e \cdot \exp\left(-\frac{c' p^2 n (1 - \varepsilon)^2}{2304}\right).$$

2. For any $\varepsilon \in (0, 1)$,

$$\|\bar{\mathbf{z}} - \bar{\mathbf{h}}\|_2 \leq \sqrt{2(p + \varepsilon)} \sqrt{n} \quad (6.8)$$

holds with probability at least $1 - e \cdot \exp\left(-\frac{c'(1-p)^2 \varepsilon^2 n}{16}\right)$.

The proof is deferred to Appendix D. Observe from (6.8) and Theorem 12 that $\delta = \sqrt{2(p + \varepsilon)} \in [0, 1]$ if $p + \varepsilon \leq 1/2$. Part 1 of Theorem 11 follows by applying the union bound to the two events in Proposition 13, and combining it with Theorem 12 and Lemma 8. To prove part 2 of Theorem 11, we simply use the fact $\frac{1}{2n} \bar{\mathbf{z}}^T H \bar{\mathbf{z}} \geq 0$ (since H is p.s.d) in Theorem 12, and then combine it with part 2 of Proposition 13 and Lemma 8. This completes the proof. \blacksquare

6.2. Analysis for Gaussian noise model

We now analyze the Gaussian noise model in (2.5). Recall that $y_i = (f(x_i) + \eta_i) \bmod 1$ where $\eta_i \sim \mathcal{N}(0, \sigma^2)$ i.i.d for $i = 1, \dots, n$. The following theorem is our main result for this noise model.

Theorem 14 *Consider the Gaussian noise model in (2.5) for some $\sigma \geq 0$. If $\lambda < \frac{1}{4k}$ then $\widehat{\mathbf{g}}$ is the unique solution to (4.5). Moreover, assuming $n \geq 2$, then for any $\varepsilon \in (0, 1/2)$ satisfying $(1 - \varepsilon)e^{-2\pi^2\sigma^2} \geq 1/2$, and absolute constants $c, c' > 0$, the following is true.*

1.

$$\begin{aligned} \frac{1}{n} \langle \bar{\mathbf{h}}, \widehat{\mathbf{g}} \rangle \geq & 1 - \left(3\sqrt{\frac{(1 - (1 - \varepsilon)e^{-2\pi^2\sigma^2})}{2}} - \frac{\lambda k}{6(4\lambda k + 1)^2} (1 - \varepsilon)(1 - e^{-4\pi^2\sigma^2})^2 \right) \\ & - \frac{\lambda\pi^2 M^2 (2k)^{2\alpha+1}}{n^{2\alpha}} \left(1 - \frac{e^{-4\pi^2\sigma^2}}{(4\lambda k + 1)^2} \right) \end{aligned} \quad (6.9)$$

with probability at least

$$\begin{aligned} & 1 - 2e \exp\left(-\frac{c'(1 - \varepsilon)^2(1 - e^{-4\pi^2\sigma^2})^4 e^{4\pi^2\sigma^2}}{2304} n\right) - 4 \exp\left(-\frac{c\varepsilon^2(1 - e^{-4\pi^2\sigma^2})^4 n}{1024}\right) \\ & - e \cdot \exp\left(-\frac{c'e^{-4\pi^2\sigma^2} n\varepsilon^2}{4}\right). \end{aligned} \quad (6.10)$$

2.

$$\frac{1}{n} \langle \bar{\mathbf{h}}, \widehat{\mathbf{g}} \rangle \geq 1 - 3\sqrt{\frac{(1 - (1 - \varepsilon)e^{-2\pi^2\sigma^2})}{2}} - \frac{\lambda\pi^2 M^2 (2k)^{2\alpha+1}}{n^{2\alpha}} \quad (6.11)$$

with probability at least

$$1 - e \cdot \exp\left(-\frac{c'e^{-4\pi^2\sigma^2} n\varepsilon^2}{4}\right). \quad (6.12)$$

The statement of Theorem 14 is very similar in flavour, to that of Theorem 11. Indeed, the correlation lower bounds hold provided the noise level (dictated by σ) is sufficiently small. Furthermore, note that the bound in (6.9) is strictly better than that in (6.11), albeit with a worse success probability. If $\sigma = \Theta(1)$, and n is sufficiently large, then (6.9) is a better bound than (6.11), both holding with respectively high probabilities. On the other hand, (6.11) is meaningful for all admissible values of σ (satisfying $(1 - \varepsilon)e^{-2\pi^2\sigma^2} \geq 1/2$), and moderately large n . In comparison, if for instance $\sigma \approx 0$ and n is moderately large, then the success probability lower bound for the event (6.9) becomes trivial.

Proof [Proof of Theorem 14] Our starting point is Theorem 12 as before. The following Proposition is an analogous version of Proposition 13 for the Gaussian noise model, where we upper bound the noise term parameter δ , and lower bound the quadratic form $\frac{1}{2n} \bar{\mathbf{z}}^T H \bar{\mathbf{z}}$, w.h.p.

Proposition 15 *Consider the Gaussian noise model in (2.5). For absolute constants $c, c' > 0$, the following is true.*

1. For any $\varepsilon \in (0, 1)$,

$$\frac{1}{2n} \bar{\mathbf{z}}^T H \bar{\mathbf{z}} \geq \frac{\lambda k}{6} (1 - \varepsilon) (1 - e^{-4\pi^2 \sigma^2})^2 + \frac{e^{-4\pi^2 \sigma^2}}{2n} \mathbf{h}^T H \mathbf{h} \quad (6.13)$$

holds with probability at least

$$1 - 2e \exp\left(\frac{-c'(1 - \varepsilon)^2 (1 - e^{-4\pi^2 \sigma^2})^4 e^{4\pi^2 \sigma^2}}{2304} n\right) - 4 \exp\left(-\frac{c\varepsilon^2 (1 - e^{-4\pi^2 \sigma^2})^4 n}{1024}\right). \quad (6.14)$$

2. For any $\varepsilon \in (0, 1)$,

$$\|\bar{\mathbf{z}} - \bar{\mathbf{h}}\|_2 \leq \sqrt{2(1 - (1 - \varepsilon)e^{-2\pi^2 \sigma^2})} \sqrt{n} \quad (6.15)$$

holds with probability at least $1 - e \cdot \exp(-\frac{c'e^{-4\pi^2 \sigma^2} n \varepsilon^2}{4})$.

The proof is deferred to Appendix E. Observe that $\delta = \sqrt{2(1 - (1 - \varepsilon)e^{-2\pi^2 \sigma^2})} \in [0, 1]$ if $(1 - \varepsilon)e^{-2\pi^2 \sigma^2} \geq 1/2$. Part 1 of Theorem 14 follows by combining Proposition 15 with Theorem 12 and Lemma 8. To prove part 2 of Theorem 14, we simply use the fact $\frac{1}{2n} \bar{\mathbf{z}}^T H \bar{\mathbf{z}} \geq 0$ (since H is p.s.d) in Theorem 12, and then combine it with part 2 of Proposition 15 and Lemma 8. This altogether completes the proof. \blacksquare

7. Analysis for multivariate functions and arbitrary bounded noise

We now consider the general setting where f is a d -variate function with $d \geq 1$. Specifically, we consider $f : [0, 1]^d \rightarrow \mathbb{R}$ where f is Hölder continuous, meaning that for constants $M > 0$ and $\alpha \in (0, 1]$,

$$|f(\mathbf{x}) - f(\mathbf{y})| \leq M \|\mathbf{x} - \mathbf{y}\|_2^\alpha; \quad \forall \mathbf{x}, \mathbf{y} \in [0, 1]^d. \quad (7.1)$$

For a positive integer m , we consider f to be sampled everywhere on the regular grid $\mathcal{G} = \left\{0, \frac{1}{m-1}, \dots, \frac{m-2}{m-1}, 1\right\}^d$, with $n = |\mathcal{G}| = m^d$ being the total number of samples. Let $V = \{1, 2, \dots, m\}^d$ be the index set for the points on the grid. We assume for concreteness the same arbitrary bounded noise model as in (2.3), i.e.,

$$\mathbf{y}_i = (f(\mathbf{x}_i) + \delta_i) \bmod 1; \quad |\delta_i| \in (0, 1/2); \quad \mathbf{i} \in V, \mathbf{x}_i \in \mathcal{G}. \quad (7.2)$$

We will work with a regularization graph $G = (V, E)$ where the set of edges now consists of vertices that are close in the ℓ_∞ metric (also known as *Chebyshev* distance). The edge set E is formally defined as

$$E = \{\{\mathbf{i}, \mathbf{j}\} : \mathbf{i}, \mathbf{j} \in V, \quad \mathbf{i} \neq \mathbf{j}, \quad \|\mathbf{i} - \mathbf{j}\|_\infty \leq k\}. \quad (7.3)$$

Note that the degree of each vertex is less than or equal to $(2k+1)^d - 1$, and is also at least $(k+1)^d - 1$ (with equality achieved by nodes at corners of the grid). As in (3.1), we denote

$$h_{\mathbf{i}} := \exp(2\pi\iota f_{\mathbf{i}}), \quad z_{\mathbf{i}} := \exp(2\pi\iota y_{\mathbf{i}}); \quad \mathbf{i} \in V, \quad (7.4)$$

to be the respective representations of the clean mod 1 and noisy mod 1 samples on the unit circle in \mathbb{C} . Let $\mathbf{h}, \mathbf{z} \in \mathbb{C}^n$ denote the respective corresponding vectors, with entries indexed by V , and the indices sorted in lexicographic increasing order. Then, we simply consider solving the relaxed (trust-region) problem in (4.2), in particular, its equivalent reformulation (4.5) with the real-valued representations as in (4.3) and (4.4). This altogether leads to the following theorem, which is a generalization of Theorem 5 to the multivariate case.

Theorem 16 *Consider the arbitrary bounded noise model defined in (7.2), with $\bar{\mathbf{z}}$ satisfying $\|\bar{\mathbf{z}} - \bar{\mathbf{h}}\|_2 \leq \delta\sqrt{n}$ for $\delta \in [0, 1]$. Let $m \geq 2$, and let $\mathcal{N}(H)$ denote the null space of H .*

1. *If $\bar{\mathbf{z}} \notin \mathcal{N}(H)$ then $\hat{\mathbf{g}}$ is the unique solution to (4.5) satisfying*

$$\begin{aligned} \frac{1}{n} \langle \bar{\mathbf{h}}, \hat{\mathbf{g}} \rangle &\geq 1 - \frac{3\delta}{2} - \frac{\lambda\pi^2 M^2 (2k)^{2\alpha} d^{2\alpha} [(2k+1)^d - 1]}{n^{2\alpha/d}} \\ &\quad + \frac{1}{(2\lambda[(2k+1)^d - 1] + 1)^2} \left(\frac{1}{2n} \bar{\mathbf{z}}^T H \bar{\mathbf{z}} \right). \end{aligned} \quad (7.5)$$

2. *If $\bar{\mathbf{z}} \perp \mathcal{N}(H)$ and $\lambda < \frac{1}{2((2k+1)^d - 1)}$ then $\hat{\mathbf{g}}$ is the unique solution to (4.5) satisfying*

$$\begin{aligned} \frac{1}{n} \langle \bar{\mathbf{h}}, \hat{\mathbf{g}} \rangle &\geq 1 - \frac{3\delta}{2} - \frac{\lambda\pi^2 M^2 (2k)^{2\alpha} d^{2\alpha} [(2k+1)^d - 1]}{n^{2\alpha/d}} \\ &\quad + \frac{1}{(1 + 2\lambda[(2k+1)^d - 1] - 4\lambda\kappa_{k,d} \sin^2(\frac{\pi}{2n}))^2} \left(\frac{1}{2n} \bar{\mathbf{z}}^T H \bar{\mathbf{z}} \right). \end{aligned} \quad (7.6)$$

Here, $\kappa_{k,d} \leq (k+1)^d - 1$ is the vertex connectivity of the graph G .

As H is positive semi-definite, we obtain the following useful Corollary of Theorem 16.

Corollary 17 *Consider the arbitrary bounded noise model defined in (2.3), with $\bar{\mathbf{z}}$ satisfying $\|\bar{\mathbf{z}} - \bar{\mathbf{h}}\|_2 \leq \delta\sqrt{n}$ for $\delta \in [0, 1]$. Let $m \geq 2$. If $\lambda < \frac{1}{2((2k+1)^d - 1)}$ then $\hat{\mathbf{g}}$ is the unique solution to (4.5) satisfying*

$$\frac{1}{n} \langle \bar{\mathbf{h}}, \hat{\mathbf{g}} \rangle \geq 1 - \frac{3\delta}{2} - \frac{\lambda\pi^2 M^2 (2k)^{2\alpha} d^{2\alpha} [(2k+1)^d - 1]}{n^{2\alpha/d}}. \quad (7.7)$$

It can be easily verified that for $d = 1$, Theorem 16 reduces to Theorem 5 since $\kappa_{k,1} = k$. We see that the bounds on the correlation involve terms depending exponentially in the dimension d . This is the curse of dimensionality typically arising in high-dimensional non-parametric estimation problems. Nevertheless, as noted in Section 1, many interesting applications in phase unwrapping arise when d is small, specifically, $d = 2, 3$. Furthermore, the lower bound on $\frac{1}{n} \langle \bar{\mathbf{h}}, \hat{\mathbf{g}} \rangle$ readily implies the ℓ_2 error bound on $\|\hat{\mathbf{g}} - \bar{\mathbf{h}}\|_2^2$ via

$$\|\hat{\mathbf{g}} - \bar{\mathbf{h}}\|_2^2 = \|\hat{\mathbf{g}}\|_2^2 + \|\bar{\mathbf{h}}\|_2^2 - 2n \frac{\langle \hat{\mathbf{g}}, \bar{\mathbf{h}} \rangle}{n} \leq 2n - 2n \left(1 - \frac{3\delta}{2} - \frac{c}{n^{2\alpha/d}} \right) = O(\delta n + n^{1 - \frac{2\alpha}{d}}). \quad (7.8)$$

Proof [Proof of Theorem 16] The proof follows the same structure as that of Theorem 5, with minor technical changes. To begin with, Lemma 7 remains unchanged, meaning that any solution $\widehat{\mathbf{g}}$ to (4.5) satisfies

$$\frac{1}{n} \langle \bar{\mathbf{h}}, \widehat{\mathbf{g}} \rangle \geq 1 - \frac{3\delta}{2} - \frac{1}{2n} \bar{\mathbf{h}}^T H \bar{\mathbf{h}} + \frac{1}{2n} \widehat{\mathbf{g}}^T H \widehat{\mathbf{g}}. \quad (7.9)$$

We then need the following analogue of Lemma 8 to upper bound the term $\frac{1}{2n} \bar{\mathbf{h}}^T H \bar{\mathbf{h}}$ using the smoothness of the ground truth $\bar{\mathbf{h}} \in \mathbb{R}^{2n}$.

Lemma 18 *For $m \geq 2$, it holds true that*

$$\frac{1}{2n} \bar{\mathbf{h}}^T H \bar{\mathbf{h}} \leq \frac{\lambda \pi^2 M^2 (2k)^{2\alpha} d^{2\alpha} [(2k+1)^d - 1]}{n^{2\alpha/d}}, \quad (7.10)$$

where $\alpha \in (0, 1]$ and $M > 0$ are related to the smoothness of f and defined in (7.1), and $\lambda \geq 0$ is the regularization parameter in (3.2).

Proof The proof is very similar to that of Lemma 8, so we only focus on the changes. To begin with, as in (5.14), we have

$$\frac{1}{2n} \bar{\mathbf{h}}^T H \bar{\mathbf{h}} \leq \frac{\lambda}{2n} |E| \max_{\{\mathbf{i}, \mathbf{j}\} \in E} |h_{\mathbf{i}} - h_{\mathbf{j}}|^2. \quad (7.11)$$

For each $\mathbf{i} \in V$ we have $\deg(\mathbf{i}) \leq (2k+1)^d - 1$, hence $|E| \leq \frac{n}{2} [(2k+1)^d - 1]$. Next, for any $\{\mathbf{i}, \mathbf{j}\} \in E$, we have

$$|f_{\mathbf{i}} - f_{\mathbf{j}}| \leq M \|\mathbf{x}_{\mathbf{i}} - \mathbf{x}_{\mathbf{j}}\|_2^\alpha \leq M \left(\frac{k}{m-1} \right)^\alpha d^{\alpha/2} \leq M \left(\frac{2k}{m} \right)^\alpha d^{\alpha/2}, \quad (7.12)$$

if $m \geq 2$ (since then $m-1 \geq m/2$). This leads to

$$|h_{\mathbf{i}} - h_{\mathbf{j}}| \leq 2\pi |f_{\mathbf{j}} - f_{\mathbf{i}}| \leq \frac{2\pi M (2k)^\alpha d^{\alpha/2}}{m^\alpha} \quad (\text{using (5.18) and (7.12)}). \quad (7.13)$$

Plugging (7.13) in (7.11) with the bound $|E| \leq \frac{n}{2} [(2k+1)^d - 1]$ yields the desired bound. ■

Finally, we lower bound the term $\frac{1}{2n} \widehat{\mathbf{g}}^T H \widehat{\mathbf{g}}$ in (7.9) via an analogue of Lemma 9 outlined below.

Lemma 19 *Denoting $\mathcal{N}(H)$ to be the null space of H , the following holds for the solution $\widehat{\mathbf{g}}$ to (4.5).*

1. *If $\bar{\mathbf{z}} \notin \mathcal{N}(H)$ then $\widehat{\mathbf{g}}$ is unique and*

$$\frac{1}{2n} \widehat{\mathbf{g}}^T H \widehat{\mathbf{g}} \geq \frac{1}{(1 + 2\lambda [(2k+1)^d - 1])^2} \left(\frac{1}{2n} \bar{\mathbf{z}}^T H \bar{\mathbf{z}} \right). \quad (7.14)$$

2. *If $\bar{\mathbf{z}} \perp \mathcal{N}(H)$ and $\lambda < \frac{1}{2[(2k+1)^d - 1]}$, then $\widehat{\mathbf{g}}$ is unique and*

$$\frac{1}{2n} \widehat{\mathbf{g}}^T H \widehat{\mathbf{g}} \geq \frac{1}{(1 + 2\lambda [(2k+1)^d - 1] - 4\lambda \kappa_{k,d} \sin^2(\frac{\pi}{2n}))^2} \left(\frac{1}{2n} \bar{\mathbf{z}}^T H \bar{\mathbf{z}} \right). \quad (7.15)$$

Here, $\kappa_{k,d} \leq (k+1)^d - 1$ is the vertex connectivity of the graph G .

Proof The outline of the proof is the same as the proof of Lemma 9, so we will only highlight the technical changes. To begin with, the largest eigenvalue of L , namely $\beta_n(L)$, now satisfies $\beta_n(L) \leq 2([2k+1]^d - 1)$. This is easily verified⁸ through Gershgorin's disk theorem. Therefore, the eigenvalues of H now satisfy

$$0 = \lambda_1(H) = \lambda_2(H) < \lambda_3(H) \leq \dots \leq \lambda_{2n}(H) \leq 2\lambda([2k+1]^d - 1) \quad (7.16)$$

Recall that $\{\mathbf{q}_j\}_{j=1}^{2n}$ denote the corresponding eigenvectors of H , with $\mathcal{N}(H) = \text{span}\{\mathbf{q}_1, \mathbf{q}_2\}$. We now consider the two cases separately below.

1. Consider the case where $\bar{\mathbf{z}} \notin \mathcal{N}(H)$. We know that $\widehat{\mathbf{g}} = 2(2H + \mu^*I)^{-1}\bar{\mathbf{z}}$ for a unique $\mu^* \in (0, \infty)$, and so $\widehat{\mathbf{g}}$ is the unique solution to (4.5) by Lemma 3 since $2H + \mu^*I \succ 0$. We again have $\mu^* \leq 2$ as in (5.24). Finally, we can lower bound $\frac{1}{2n}\widehat{\mathbf{g}}^T H \widehat{\mathbf{g}}$ as follows.

$$\frac{1}{2n}\widehat{\mathbf{g}}^T H \widehat{\mathbf{g}} = \frac{2}{n} \sum_{j=1}^{2n} \frac{\langle \bar{\mathbf{z}}, \mathbf{q}_j \rangle^2 \lambda_j(H)}{(2\lambda_j(H) + \mu^*)^2} \quad (7.17)$$

$$\geq \frac{2}{n(4\lambda([2k+1]^d - 1) + 2)^2} \sum_{j=1}^{2n} \langle \bar{\mathbf{z}}, \mathbf{q}_j \rangle^2 \lambda_j(H) \quad (\text{from (7.16) and since } \mu^* \leq 2) \quad (7.18)$$

$$= \frac{1}{(2\lambda([2k+1]^d - 1) + 1)^2} \left(\frac{1}{2n} \bar{\mathbf{z}}^T H \bar{\mathbf{z}} \right). \quad (7.19)$$

2. Let us now consider the case where $\bar{\mathbf{z}} \perp \mathcal{N}(H)$. The analogous of (5.30) to the multivariate case is given by

$$\phi(0) = \sum_{j=3}^{2n} \frac{\langle \bar{\mathbf{z}}, \mathbf{q}_j \rangle^2}{\lambda_j(H)^2} \geq \frac{n}{4\lambda^2([2k+1]^d - 1)^2}, \quad (7.20)$$

since $\lambda_j(H) \leq 2\lambda([2k+1]^d - 1)$ for all j (recall (7.16)). Therefore, the choice $\lambda < \frac{1}{2[(2k+1)^d - 1]}$ clearly implies $\phi(0) > n$, and consequently that the solution $\widehat{\mathbf{g}}$ is unique. Assuming $\lambda < \frac{1}{2[(2k+1)^d - 1]}$ holds, we arrive at $\mu^* \in (0, 2 - 2\lambda_3(H))$ as previously shown in (5.32). As a sanity check, note that $2 - 2\lambda_3(H) > 0$, in light of (7.16) stating that $\lambda_3(H) \leq 2\lambda([2k+1]^d - 1)$. Consequently, we can provide the following lower bound for $\frac{1}{2n}\widehat{\mathbf{g}}^T H \widehat{\mathbf{g}}$.

$$\frac{1}{2n}\widehat{\mathbf{g}}^T H \widehat{\mathbf{g}} = \frac{2}{n} \sum_{j=1}^{2n} \frac{\langle \bar{\mathbf{z}}, \mathbf{q}_j \rangle^2 \lambda_j(H)}{(2\lambda_j(H) + \mu^*)^2} \quad (7.21)$$

$$\geq \frac{2}{n(4\lambda([2k+1]^d - 1) + 2 - 2\lambda_3(H))^2} \sum_{j=1}^{2n} \langle \bar{\mathbf{z}}, \mathbf{q}_j \rangle^2 \lambda_j(H) \quad (7.22)$$

$$= \frac{1}{(1 + 2\lambda([2k+1]^d - 1) - \lambda_3(H))^2} \left(\frac{1}{2n} \bar{\mathbf{z}}^T H \bar{\mathbf{z}} \right), \quad (7.23)$$

8. Denote L_{ij} to be the (i, j) th entry of L . Using Gershgorin's disk theorem, each eigenvalue of L lies in $\bigcup_{i=1}^{2n} \left\{ x : |x - L_{ii}| \leq \sum_{j \neq i} |L_{ij}| \right\}$. Since $L_{ii} \leq [2k+1]^d - 1$ and $\sum_{j \neq i} |L_{ij}| \leq [2k+1]^d - 1$ holds for each i , the claim follows.

where the second inequality follows from (7.16) and the fact $\mu^* \in (0, 2 - 2\lambda_3(H))$. Finally, using Theorem 10 we readily obtain the bound $\lambda_3(H) \geq 4\lambda\kappa_{k,d} \sin^2\left(\frac{\pi}{2n}\right)$, where $\kappa_{k,d}$ denotes the vertex connectivity of G . The minimum degree of G is $(k+1)^d - 1$ and so $\kappa_{k,d} \leq (k+1)^d - 1$. Plugging the lower bound on $\lambda_3(H)$ into (7.23) completes the proof. ■

This completes the proof of Theorem 16. ■

Unwrapping the samples of f . Once we obtain the denoised $f \bmod 1$ samples, we can recover an estimate of the original samples of f (up to a global shift) by using, for instance, the OLS method outlined in Section 4.2. This is the approach we adopt in our simulations, but one could of course consider using more sophisticated unwrapping algorithms. Exploring other approaches for the unwrapping stage is an interesting direction for future work.

Extensions to random noise models. One could also consider extending the results to random noise models (Bernoulli-Uniform, Gaussian) as was shown in Sections 6.1 and 6.2 for the $d = 1$ case. We expect the proof outline to be very similar to that of Theorems 11 and 14 with minor technical changes.

8. Numerical experiments for the univariate case via TRS-based modulo denoising

This section contains numerical experiments⁹ for the univariate case ($d = 1$) for the following three noise models discussed in Section 2.

1. **Arbitrary bounded** noise model (2.3), analyzed theoretically in Section 5. In particular, we experiment with the Uniform model, with samples generated uniformly at random in $[-\gamma, \gamma]$ for bounded γ . Results for this model are shown in Appendix F1.
2. **Bernoulli-Uniform** noise model (2.4), with guarantees put forth in Section 6.1.
3. **Gaussian** noise model (2.5), analyzed theoretically in Section 6.2.

The function f we consider throughout all the experiments in this Section is given by

$$f(x) = 4x \cos^2(2\pi x) - 2 \sin^2(2\pi x). \tag{8.1}$$

For each experiment, we average the results over 20 trials, and show the RMSE error on a log scale, both for the denoised samples of $f \bmod 1$ and the final f estimates. For the latter, we compute the RMSE after we mod out the best global shift¹⁰. For each noise model, we compare the performance of three methods.

9. Code is publicly available online at <http://www.stats.ox.ac.uk/~cucuring/ModuloDenoising.htm>
 10. Any algorithm that takes as input the mod 1 samples will be able to recover f only up to a global shift. We mod out the global shift via a simple procedure, that first computes the offset at each sample point $f_i - \hat{f}_i$, and then considers the mode of this distribution, after a discretization step. More specifically, we compute a histogram of the offsets, and consider the center of the bucket that gives the mode of the distribution as our final global shift estimate with respect to the ground truth.

- **OLS** denotes the algorithm based on the least-squares formulation (4.23) used to recover samples of f , and works directly with noisy mod 1 samples. The final estimated $f \bmod 1$ values are then obtained as the corresponding mod 1 versions.
- **QCQP** denotes Algorithm 1 where the unwrapping stage is performed via **OLS** (4.23).
- **iQCQP** denotes an iterated version of **QCQP**, wherein we repeatedly denoise the noisy $f \bmod 1$ estimates (via Stage 1 of Algorithm 1) for $l \in \{3, 5, 10\}$ iterations, and finally perform the unwrapping stage via **OLS** (4.23), to recover the final sample estimates of f .

8.1. Numerical experiments: Gaussian Model

Figure 5 shows noisy instances of the Gaussian noise model, for $n = 500$ samples. The scatter plots on the top row show that, as the noise level increases, the function (4.21) will produce more and more errors in (4.23), while the remaining plots show the corresponding $f \bmod 1$ signal (clean, noisy, and denoised via **QCQP**) for three levels of noise.

Figure 6 depicts instances of the recovery process, highlighting the noise level at which each method shows a significant decrease in performance. **iQCQP** shows surprisingly good performance even at very high levels of noise, ($\sigma = 0.17$).

Figures 7, respectively 8, plot the recovery errors (averaged over 20 runs) for the denoised values of $f \bmod 1$, respectively the estimated f samples, as we scan across the input parameters $k \in \{2, 3, 5\}$ and $\lambda \in \{0.03, 0.1, 0.3, 0.5, 1\}$, at varying levels of noise $\sigma \in [0, 0.15]$. We remark that $k = 5$ and $\lambda = 1$ most often lead to the worst performance. For most parameter combinations **iQCQP** improves on **QCQP** (except for very high values of λ and k). The best recovery errors are obtained for $k = 2$ and higher values of λ .

8.2. Numerical experiments: Bernoulli Model

Figure 9 shows noisy instances of the Bernoulli-Uniform noise model, for $n = 500$ samples. As shown in the scatter plots (top row), most δ increments take value 0 (since only a small fraction of the samples are corrupted by noise). For higher noise levels, the function (4.21) will produce an increase numbers of errors in (4.23). The remaining plots show the corresponding $f \bmod 1$ signal (clean, noisy, denoised via **QCQP**) for three levels of noise.

Figure 10 shows instances of the recovery process, highlighting the noise level at which each method shows a significant decrease in performance. Note that at $p = 10\%$ noise level, the simple **OLS** approach exhibits poor performance, while **QCQP** is able to denoise most of the erroneous spikes. Interestingly, **iQCQP** performs worse than **QCQP** across all levels of noise.

Finally, Figures 11 and 12 show recovery errors (averaged over 20 runs) for the estimated values of $f \bmod 1$ and f , as we scan across the input parameters $k \in \{2, 3, 5\}$ and $\lambda \in \{0.03, 0.1, 0.3, 0.5, 1\}$, at varying levels of noise $p \in [0, 0.15]$. We remark that $k = 5$ clearly returns the worst performance, and higher values of λ yield better results for **QCQP**. Interestingly, the behavior of **iQCQP** oscillates, and its performance with respect to **QCQP** depends on k and λ . The best recovery errors are obtained for $k = 2$ and $\lambda = 1$.

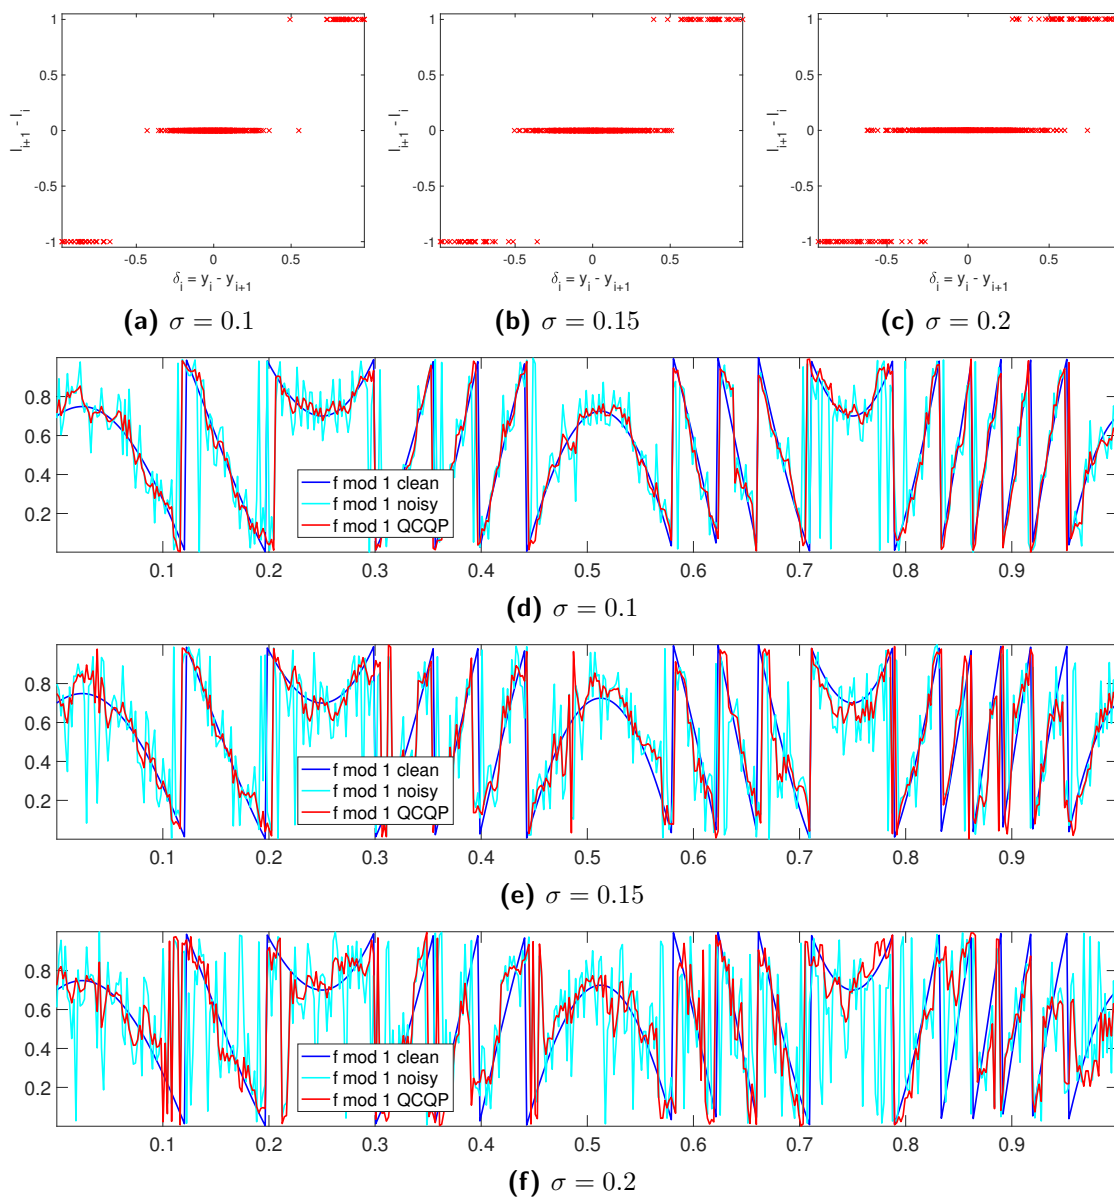


Figure 5: Noisy instances of the Gaussian noise model ($n = 500$). Top row (a)-(c) shows scatter plots of change in y (the observed noisy $f \bmod 1$ values) versus change in l (the noisy quotient). Plots (d)-(f) show the clean $f \bmod 1$ values (blue), the noisy $f \bmod 1$ values (cyan) and the denoised (via **QCQP**) $f \bmod 1$ values (red). **QCQP** denotes Algorithm 1 without the unwrapping stage performed by **OLS** (4.23).

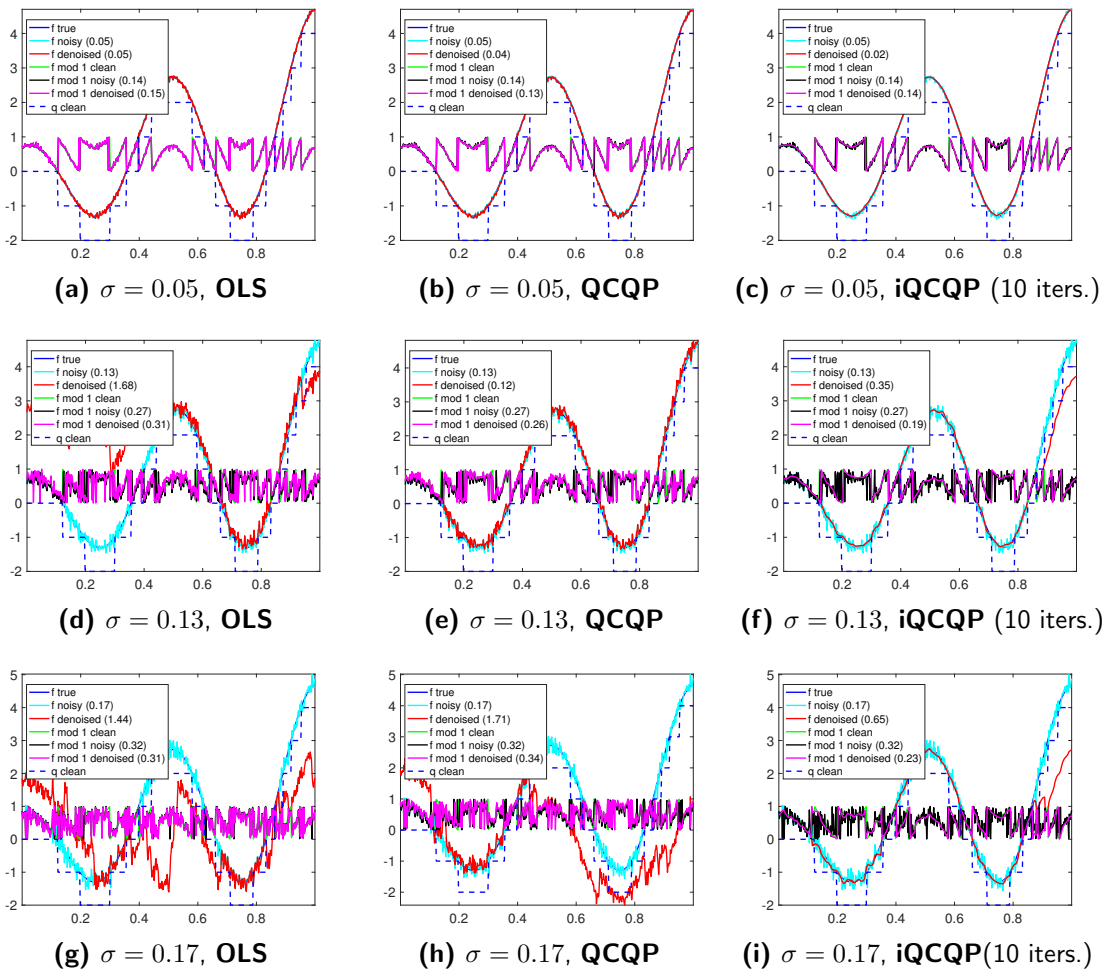


Figure 6: Denoised instances under the Gaussian noise model, for **OLS**, **QCQP** and **iQCQP**, as we increase the noise level σ . We keep fixed the parameters $n = 500$, $k = 2$, $\lambda = 0.1$. **QCQP** denotes Algorithm 1 where the unwrapping stage is performed by **OLS** (4.23).

8.3. Comparison with Bhandari et al. (2017)

This section is a comparison of **OLS**, **QCQP**, and **iQCQP** with the approach of Bhandari et al. (2017), whose algorithm we denote by **BKR** for brevity. We compare all four approaches across two different noise models, Bounded and Gaussian, on two different functions: the function (8.1) used in the experiments throughout this paper and a bandlimited function used by Bhandari et al. (2017). We defer to the appendix all the experiments, except for those reported in Figure 13, where we consider the function (8.1) under the Bounded noise model. We note that at a lower level of noise $\gamma = 0.13$, all methods perform similarly well, with relative performance in the following order: **iQCQP** (RMSE=0.25), **QCQP** (RMSE=0.29), **OLS** (RMSE=0.30), **BKR** (RMSE=0.30). However, at higher

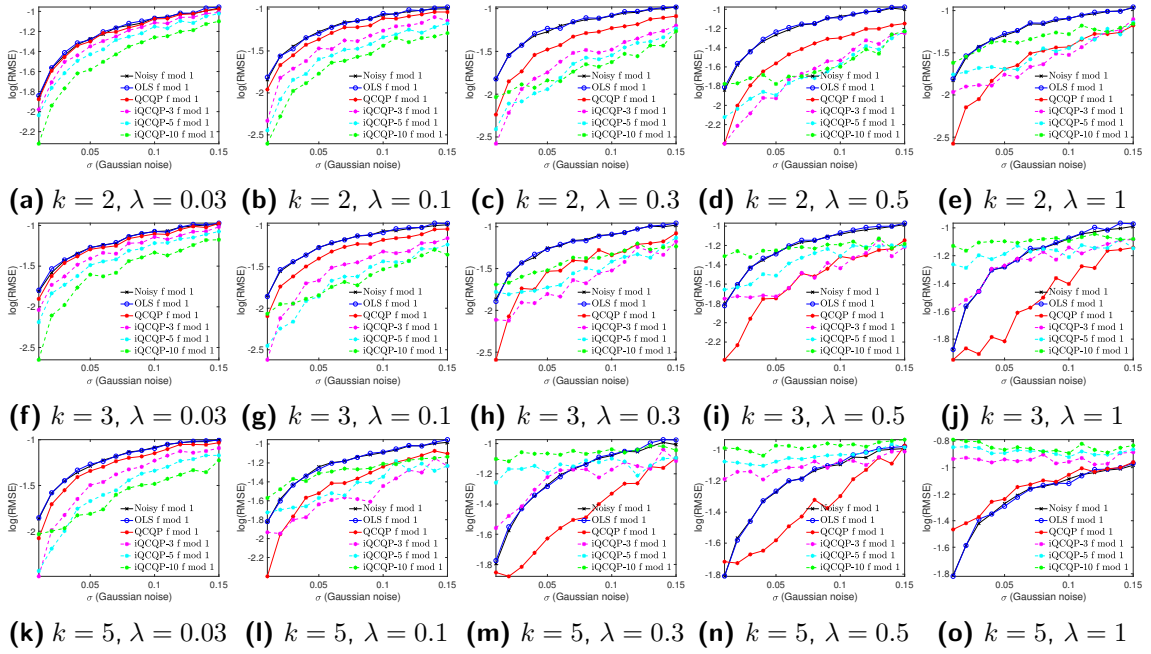


Figure 7: Recovery errors for the denoised $f \bmod 1$ samples, for $n = 500$ under the Gaussian noise model (20 trials). **QCQP** denotes Algorithm 1 without the unwrapping stage performed by **OLS** (4.23).

levels of noise, **BKR** returns meaningless results, while **QCQP**, and especially **iQCQP**, return more accurate results.

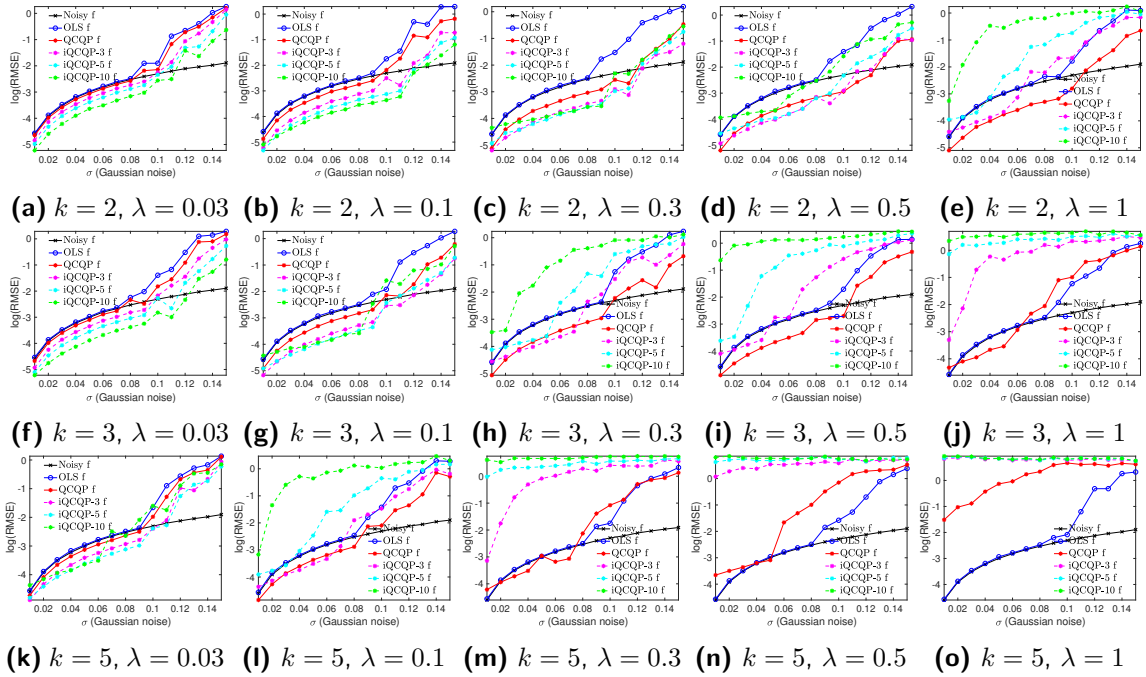


Figure 8: Recovery errors for the final estimated f samples, for $n = 500$ under the Gaussian noise model (20 trials). **QCQP** denotes Algorithm 1 where the unwrapping stage is performed via **OLS** (4.23).

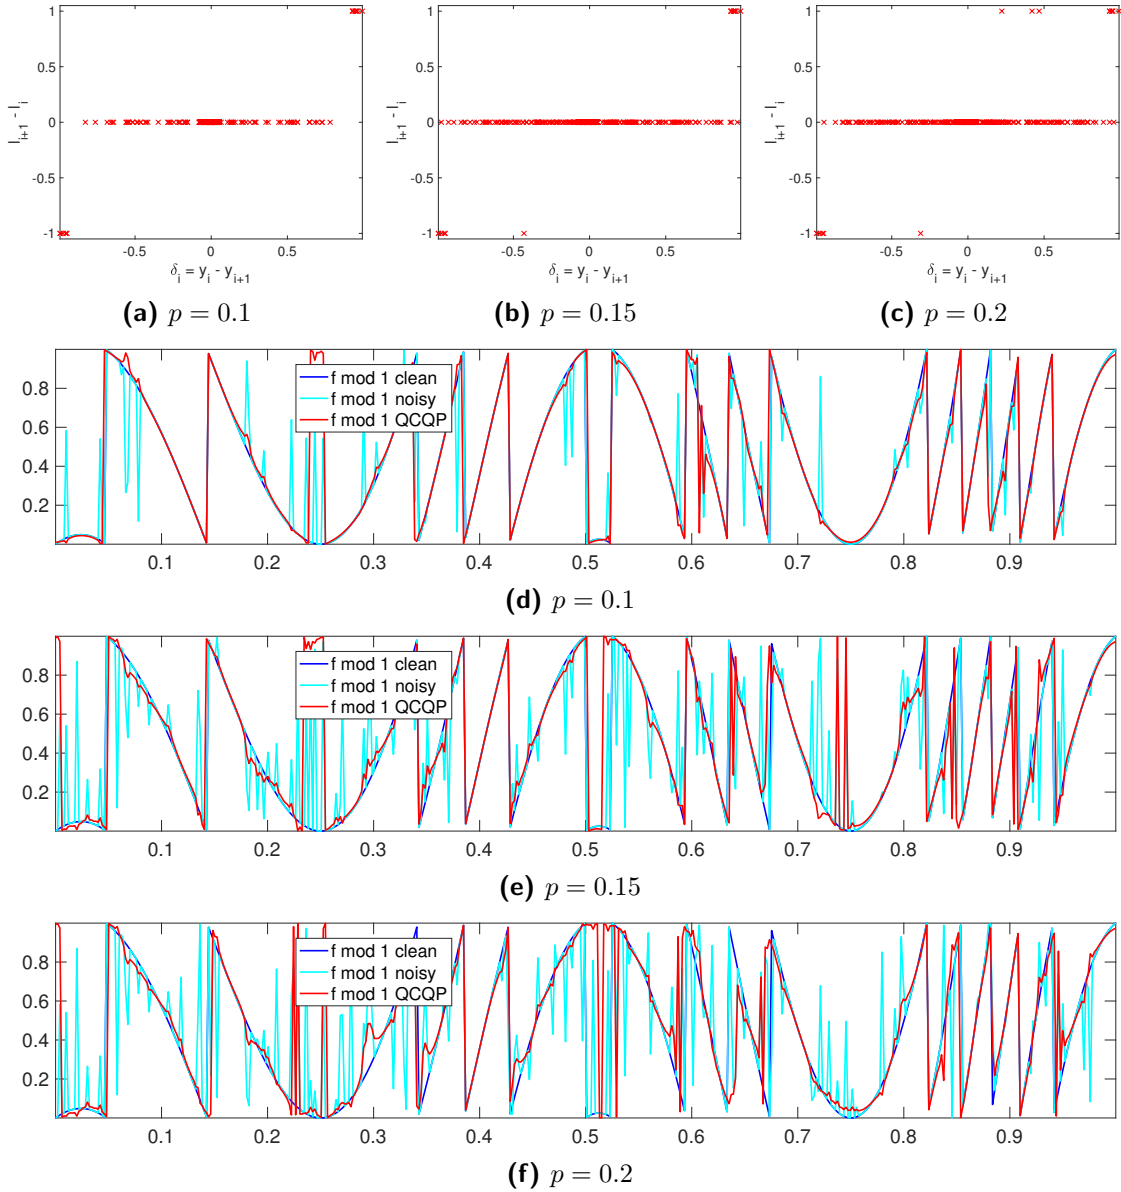


Figure 9: Noisy instances of the Bernoulli noise model ($n = 500$). Top row (a)-(c) shows scatter plots of change in y (the observed noisy $f \bmod 1$ values) versus change in l (the noisy quotient). Plots (d)-(f) show the clean $f \bmod 1$ values (blue), the noisy $f \bmod 1$ values (cyan) and the denoised (via **QCQP**) $f \bmod 1$ values (red). **QCQP** denotes Algorithm 1 without the unwrapping stage performed by **OLS** (4.23).

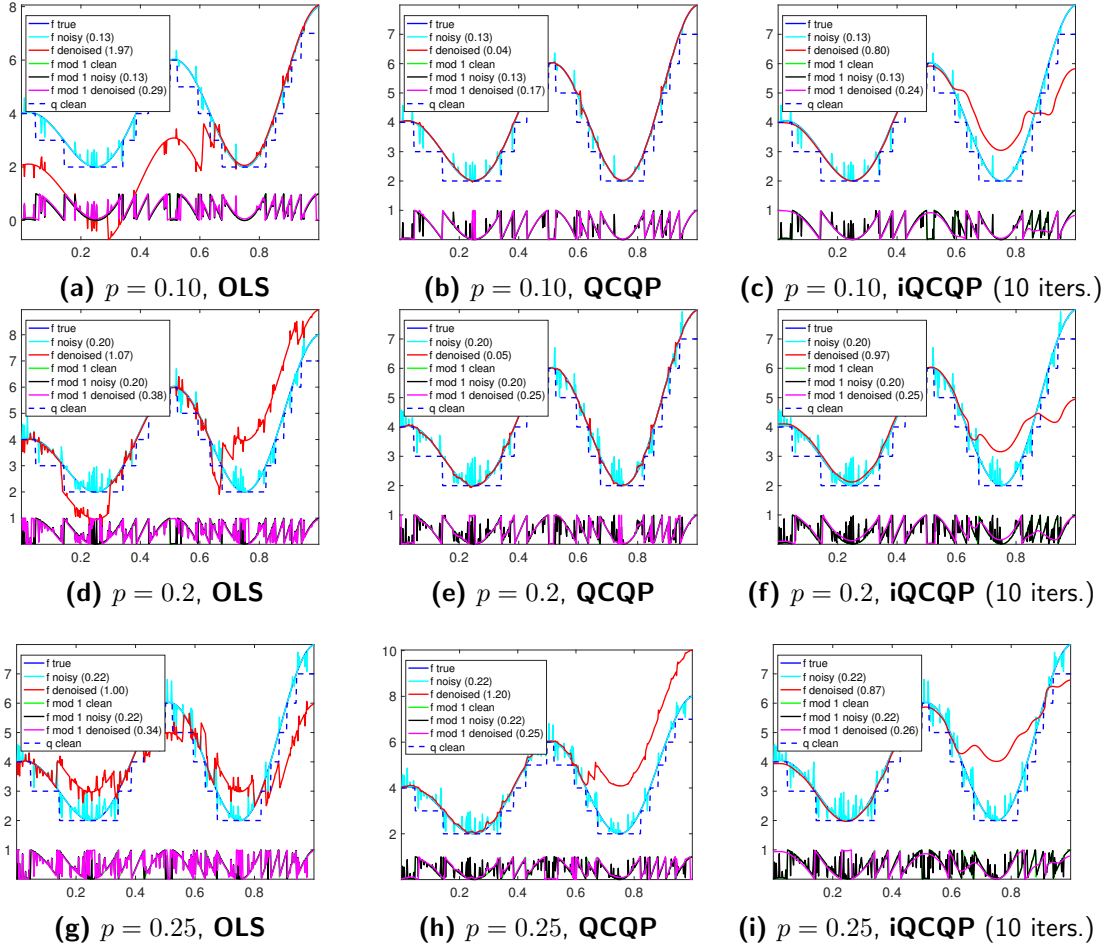


Figure 10: Denoised instances under the Bernoulli noise model, for **OLS**, **QCQP** (Algorithm 1) and **iQCQP**, as we increase the noise level σ . **QCQP** denotes Algorithm 1, for which the unwrapping stage is performed via **OLS** (4.23). We keep fixed the parameters $n = 500$, $k = 2$, $\lambda = 1$.

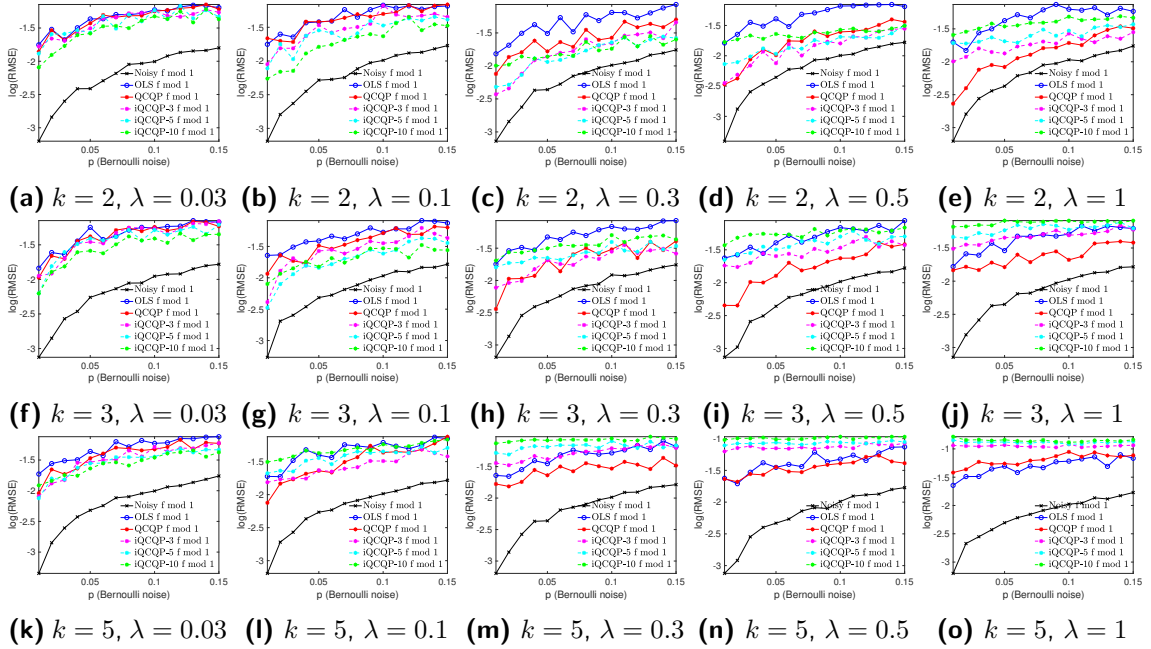


Figure 11: Recovery errors for the denoised $f \bmod 1$ samples, for $n = 500$ under the Bernoulli noise model (20 trials). **QCQP** denotes Algorithm 1 without the unwrapping stage performed by **OLS** (4.23).

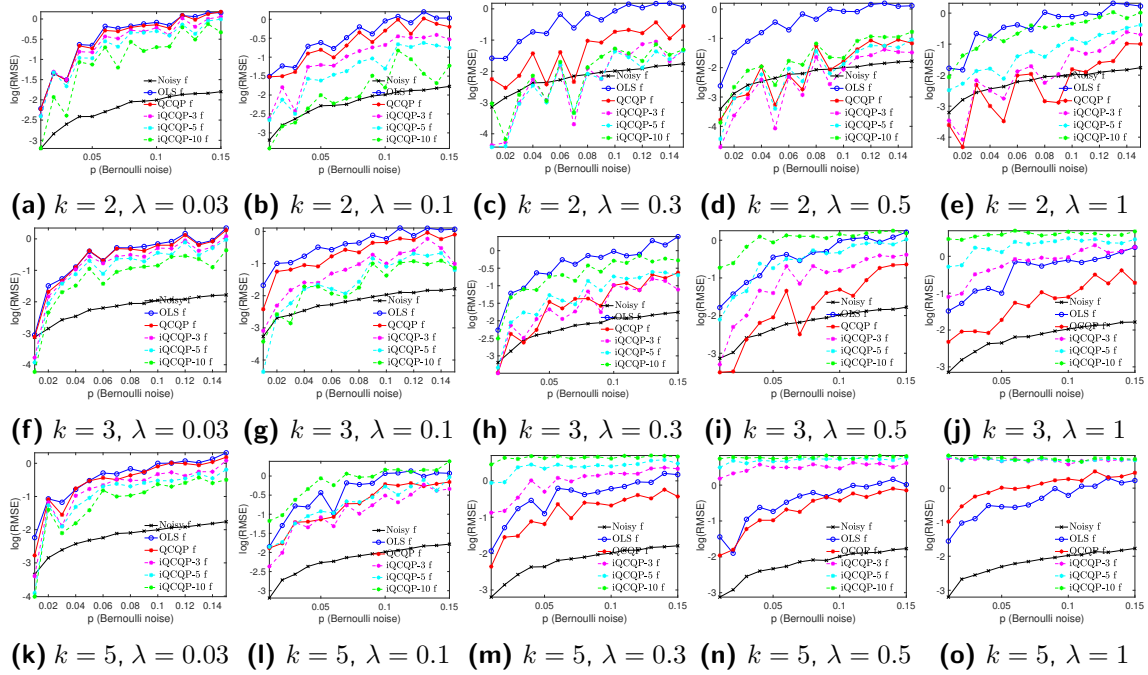


Figure 12: Recovery errors for the final estimated f samples, for $n = 500$ under the Bernoulli noise model (20 trials). Here, **QCQP** denotes Algorithm 1, for which the unwrapping stage is performed via **OLS** (4.23).

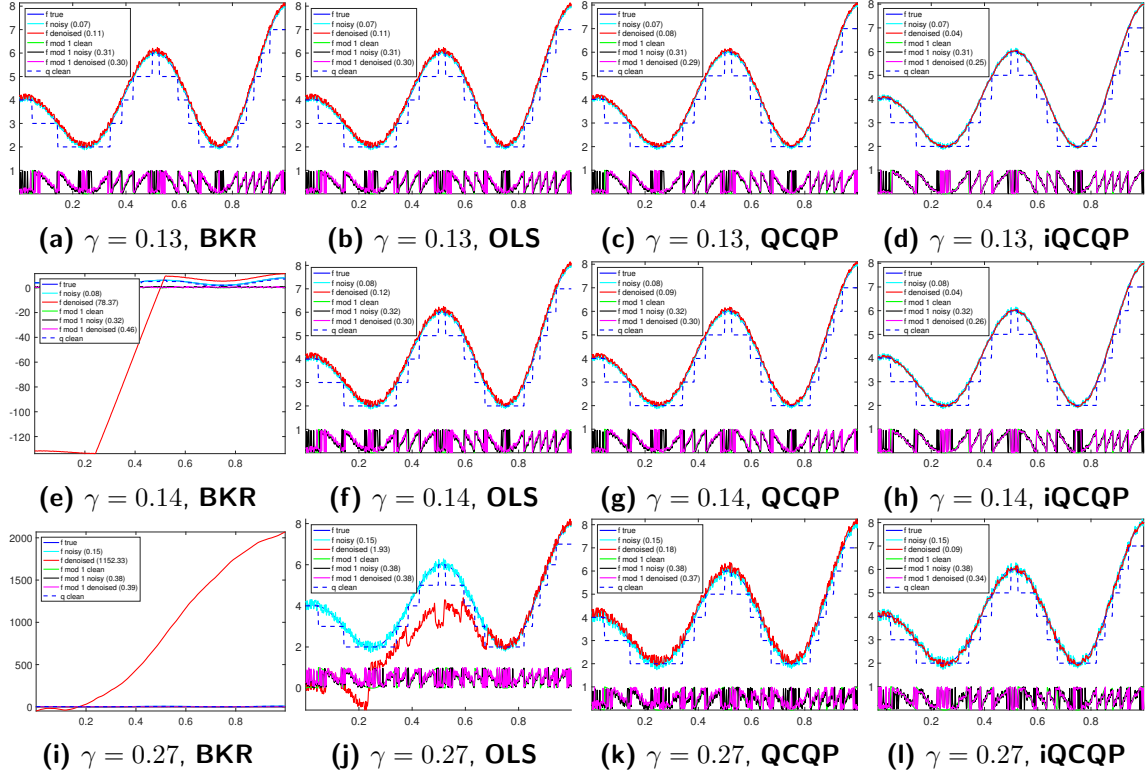


Figure 13: Denoised instances for both the $f \bmod 1$ and f values, under the Bounded-Uniform noise model for function (8.1), for **BKR**, **OLS**, **QCQP** (Algorithm 1) and **iQCQP**, as we increase the noise level γ . **QCQP** denotes Algorithm 1, for which the unwrapping stage is performed via **OLS** (4.23). We keep fixed the parameters $n = 500$, $k = 2$, $\lambda = 0.1$. The numerical values in the legend denote the RMSE.

9. Modulo 1 denoising via optimization on manifolds

Section 9.1 outlines two alternative approaches for denoising the modulo 1 samples, both of which involve optimization on a smooth manifold. Section 9.2 contains numerical simulations involving these approaches for the setting $d = 2$, on both synthetic and real data.

9.1. Two formulations for denoising modulo 1 samples

We now describe two approaches for denoising the modulo 1 samples, which are based on optimization over a smooth manifold. Both methods are implemented using Manopt (Boumal et al. (2014)), a free Matlab toolbox for optimization on manifolds. Among others, the toolbox features a Riemannian trust-region solver based on the algorithm originally developed by Absil et al. (2007), who proposed a truncated conjugate-gradient algorithm to solve trust-region subproblems. The algorithm enjoys global convergence guarantees, i.e., it will converge to stationary points for any initial guess. Under certain assumptions on the smoothness of the manifold and cost function, the method enjoys quadratic local convergence if the true Hessian is used, superlinear if the approximate Hessian is accurate enough, and linear otherwise. We refer the reader to Absil et al. (2007, 2008) for an in-depth description of the algorithm, along with its theoretical convergence guarantees, numerical results on several problems popular in the numerical linear algebra literature, and comparison with other state-of-the-art methods. Boumal et al. (2016a) recently provided the first global rates of convergence to approximate first and second-order KKT points on the manifold.

(i) Solving original QCQP via optimization over \mathcal{C}_n . Our first approach involves solving the original QCQP (3.7), i.e.,

$$\min_{\mathbf{g} \in \mathcal{C}_n} \lambda \mathbf{g}^* L \mathbf{g} - 2 \operatorname{Re}(\mathbf{g}^* \mathbf{z}). \quad (9.1)$$

The feasible set $\mathcal{C}_n := \{\mathbf{g} \in \mathbb{C}^n : |g_i| = 1\}$ is a manifold (Absil et al. (2008)). It is in fact a submanifold of the embedding space $R^2 \times \dots \times R^2$, where the complex circle is identified with the unit circle in the real plane. We employ the **complexcirclefactory** structure of Manopt that returns a manifold structure to optimize over unit-modulus complex numbers. This constitutes the input to the **trustregions** function within Manopt, along with our objective function, the gradient and the Hessian operator. In our experiments, we will denote this approach by Manopt-Phases.

(ii) SDP relaxation via Burer-Monteiro. An alternative direction we consider is to cast (9.1) as a semidefinite program (SDP), and solve it efficiently via a Burer-Monteiro approach. Denoting

$$T \stackrel{\text{def}}{=} \begin{bmatrix} \lambda L & -\mathbf{z} \\ -\mathbf{z}^* & 0 \end{bmatrix} \in \mathbb{C}^{(n+1) \times (n+1)} \quad (9.2)$$

to account for the linear term, the SDP relaxation of (9.1) is given by

$$\begin{aligned}
 & \underset{\mathbf{g} \in \mathbb{C}^n; \Upsilon \in \mathbb{C}^{n \times n}}{\text{minimize}} && \lambda \mathbf{g}^* L \mathbf{g} - 2 \text{Re}(\mathbf{g}^* \mathbf{z}) \\
 & \text{subject to} && \Upsilon_{ii} = 1, i = 1, \dots, n \\
 & && \underbrace{\begin{bmatrix} \Upsilon & \mathbf{g} \\ \mathbf{g}^* & 1 \end{bmatrix}}_W \succeq 0
 \end{aligned} \tag{9.3}$$

$$\begin{aligned}
 & \underset{W \in \mathbb{C}^{(n+1) \times (n+1)}}{\text{minimize}} && \text{tr}(TW) \\
 & \text{subject to} && W_{ii} = 1, i = 1, \dots, n + 1 \\
 & && W \succeq 0.
 \end{aligned} \tag{9.4}$$

Note that, for (9.3), if $\Upsilon = \mathbf{g}\mathbf{g}^*$, and using properties of the trace (invariance under cyclic permutations), it can be easily verified that $\text{tr}(TW) = \lambda \mathbf{g}^* L \mathbf{g} - 2 \text{Re}(\mathbf{g}^* \mathbf{z})$. Ideally, we would like to enforce the constraint $\Upsilon = \mathbf{g}\mathbf{g}^*$, which guarantees that Υ is indeed a rank-1 solution. However, since rank constraints do not lead to convex optimization problems, we relax this constraint to $\Upsilon \succeq \mathbf{g}\mathbf{g}^*$, which via Schur's lemma is equivalent to $W \succeq 0$ (see for eg., Boyd et al. (1994)).

Since semidefinite programs are computationally expensive to solve, we resort to the Burer-Monteiro approach, which consists of replacing optimization of a linear function $\langle C, X \rangle$ over the convex set $\mathcal{X} = \{X \succeq 0 : \mathcal{A}(X) = b\}$ with optimization of the quadratic function $\langle CY, Y \rangle$ over the non-convex set $\mathcal{Y} = \{Y \in \mathbb{R}^{n \times p} : \mathcal{A}(YY^T) = b\}$.

If the convex set \mathcal{X} is compact, and m denotes the number of constraints, it holds true that whenever p satisfies $\frac{p(p+1)}{2} \geq m$, the two problems share the same global optimum Barvinok (1995); Burer and Monteiro (2005). Note that over the complex domain, an analogous statement holds true as soon as $p^2 \geq m$. Building on earlier work of Burer and Monteiro (2005), Boumal et al. (2016b) show that if the set \mathcal{X} is compact and the set \mathcal{Y} is a smooth manifold, then $\frac{p(p+1)}{2} \geq m$ implies that, for almost all cost matrices, global optimality is achieved by any Y satisfying second-order necessary optimality conditions.

The problem we consider in (9.3) and (9.4) falls in the above setting with the set $\mathcal{X} := \{W \succeq 0 : \mathcal{A}(W) = 1\}$ being compact and involving $m = n + 1$ diagonal constraints, and \mathcal{Y} (defined below) being a smooth manifold. Note that in theory, choosing $p \geq \lfloor \sqrt{n+1} \rfloor + 1$ is sufficient, however in our simulations the choice $p = 3$ already returned excellent results. In the Burer-Monteiro approach, we replace Υ with YY^* in (9.3) which amounts to the following optimization problem over $Y \in \mathbb{C}^{n \times p}$ and $\mathbf{v} \in \mathbb{C}^p$ (the latter of which arises due to the linear term in the objective function)

$$\begin{aligned}
 & \underset{Y \in \mathbb{C}^{n \times p}; \mathbf{v} \in \mathbb{C}^p}{\text{minimize}} && \lambda \langle LY, Y \rangle - 2 \text{Re}(\mathbf{z}^* Y \mathbf{v}) \\
 & \text{subject to} && \text{diag}(YY^*) = 1 \\
 & && \|\mathbf{v}\|_2^2 = 1.
 \end{aligned} \tag{9.5}$$

Therefore, we optimize over the product manifold $\mathcal{Y} = \mathcal{M}_1 \times \mathcal{M}_2$, where \mathcal{M}_1 denotes the manifold of $n \times p$ complex matrices with unit 2-norm rows, and \mathcal{M}_2 the unit sphere in \mathbb{C}^p . Within Manopt, \mathcal{M}_1 is captured by the **obliquecomplexfactory**(\mathbf{n}, \mathbf{p}) structure. Its underlying geometry is a product geometry of n unit spheres in \mathbb{C}^p , with the real and imaginary parts being treated separately as $2p$ real coordinates, thus making the complex oblique manifold a Riemannian submanifold of $(\mathbb{R}^2)^{p \times n}$. Furthermore, \mathcal{M}_2 is captured within Manopt by the **spherecomplexfactory** structure. In our experiments, we will denote this approach by Manopt-Burer-Monteiro.

9.2. Numerical simulations

This section details the application of our approach to the two-dimensional phase unwrapping problem, on a variety of both synthetic and real-world examples. We empirically evaluate the performance of Algorithm 1 with three different methods for denoising the modulo 1 samples (Stage 1): Trust-region subproblem (denoted by **QCQP** for consistency of notation with previous sections), Manopt-Phases, and Manopt-Burer-Monteiro. Throughout all our experiments, we fix the neighborhood radius $k = 1$ so that each node has on average $(2k + 1)^2 - 1 = 8$ neighbors. The noise model throughout this section is fixed to be the Gaussian noise model. Note that the running times reported throughout this section are exclusively for the optimization problem itself, and exclude the time needed to build the measurement graph and its Laplacian.

Two-dimensional synthetic example. We consider a synthetic example given by the multivariate function

$$f(x, y) = 6xe^{-x^2-y^2}, \quad (9.6)$$

with $n \approx 15,000$ (i.e., a square grid of size $\sqrt{n} \times \sqrt{n}$), which we almost perfectly recover at $\sigma = 10\%$ Gaussian noise level, as depicted in Figure 14. Figure 15 is a comparison of the computational running times for the Manopt-Phases and Manopt-Burer-Monteiro approaches, as well **OLS** and **QCQP**, for varying parameters $\lambda \in \{0.01, 0.1, 1\}$, noise level $\sigma \in \{0, 0.1\}$, and $n \in \{1, 2, 4, 8, \dots, 512, 1024\} \times 1000$. For the same set of parameters, Figure 16 shows the $\log(\text{RMSE})$ errors for the recovered $f \bmod 1$ values, while Figure 17 shows similar recovery errors for the final estimate of f . For large values of n , Manopt-Phases is significantly faster than Manopt-Burer-Monteiro and **QCQP**, with the recovery errors being almost identical for all three approaches. For the noisy setup, **OLS** consistently returns much worse results compared to the other three methods, both in terms of denoising the $f \bmod 1$ values and estimating the final f samples.

In the noiseless setting, for mod 1 recovery, **OLS** consistently returns an RMSE close to machine precision as shown in the top row of Figure 16 (in other words, **OLS** preserves the mod 1 values when recovering the function f). **OLS** has a mixed performance in terms of recovering the actual function f , shown in the top row of Figure 17, where it can be seen that it either has a performance similar to that of the other three methods, or achieves an RMSE that is close to machine precision.

We remark that for the largest problem we consider, with over one million sample points, Manopt-Phases returns an almost perfect solution ($\log(\text{RMSE}) \approx -11$ for the mod 1 estimate with $\lambda = 0.01$) for the noiseless case in under 2 seconds, for a fraction of the cost of running **OLS** (Figure 15 (a)). For the noisy example ($\sigma = 0.10$, $\lambda = 1$), the same algorithm achieves an error of $\log(\text{RMSE}) \approx -4$ in under 10 seconds (Figure 15 (f)).

Two-dimensional real elevation maps. We then consider a variety of elevation maps¹¹, in particular those surrounding the Vesuvius and Etna volcanoes and Mont Blanc, at various resolution parameters, all of which we solve by Manopt-Phases. Figure 31 (shown in the

11. Data acquired via the Digital Elevation Model (DEM) of the Earth using radar interferometry. We relied on a Matlab script developed by François Beauduce (Institut de Physique du Globe de Paris, https://dds.cr.usgs.gov/srtm/version2_1), which downloads SRTM data and plots a map corresponding to the geographic area defined by latitude and longitude limits (in decimal degrees).

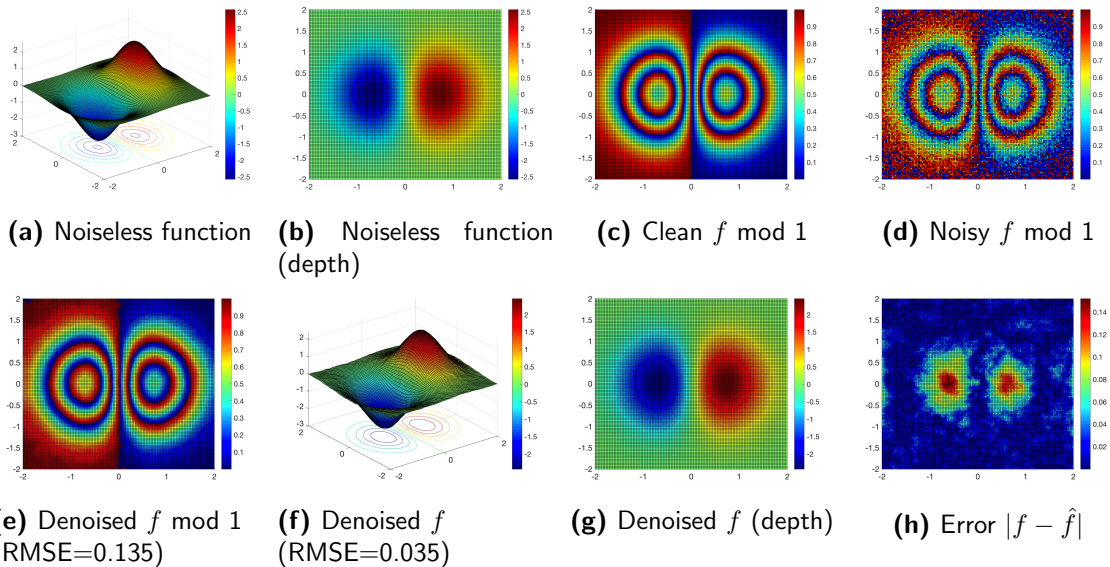


Figure 14: Synthetic example $f(x, y) = 6xe^{-x^2-y^2}$, with $n = 15000$, $k = 1$ (Chebychev distance), $\lambda = 2$, and noise level $\sigma = 0.10$ under the Gaussian model, as recovered by Manopt-Phases.

Appendix) pertains to the elevation map of Vesuvius, using $n = 3600$ sample points, and noise level $\sigma = 0.05$. Figures 18, respectively 19, also concern Mount Vesuvius, but at a higher resolution, using $n = 32400$ sample points with noise level $\sigma = 0.10$, respectively $\sigma = 0.25$. Figure 20, respectively 32 (deferred to the Appendix), pertain to noisy measurements of Mount Etna at $\sigma = 0.10$ Gaussian noise added, where the number of sample points is $n = 4100$, respectively $n = 19000$. Finally, Figure 21, respectively Figure 22, show the recovered elevation map of Mont Blanc from over 1 million sample points in under 18 seconds, for clean, respectively noisy ($\gamma = 10\%$ Bounded uniform noise), measurements.

10. Related work

This section provides a discussion of the recent closely related work of Bhandari et al. (2017), details the connection with the problem of group synchronization with anchors, and finally surveys some of the approaches in the phase unwrapping literature.

10.1. Function recovery from modulo samples

The recent work of Bhandari et al. (2017) considers the problem of recovering a bandlimited function g (the spectrum is limited to $[-\pi, \pi]$) from its centered modulo samples defined using the non-linear map $M_\lambda : \mathbb{R} \rightarrow [-\lambda, \lambda]$

$$M_\lambda(g(t)) := 2\lambda \left(\left[\frac{g(t)}{2\lambda} + \frac{1}{2} \right] - \frac{1}{2} \right), \tag{10.1}$$

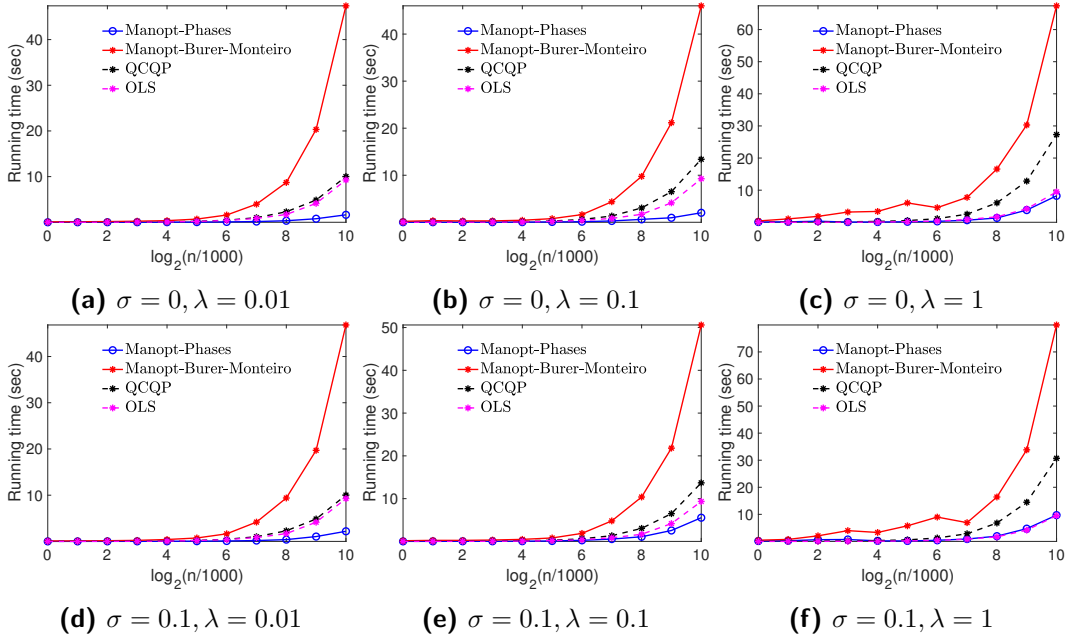


Figure 15: Comparison of running times for the synthetic example $f(x, y) = 6xe^{-x^2-y^2}$, $k = 1$ (Chebychev distance), for several noise levels σ and values of λ . **QCQP** denotes Algorithm 1, where the unwrapping stage is performed by **OLS** (4.23).

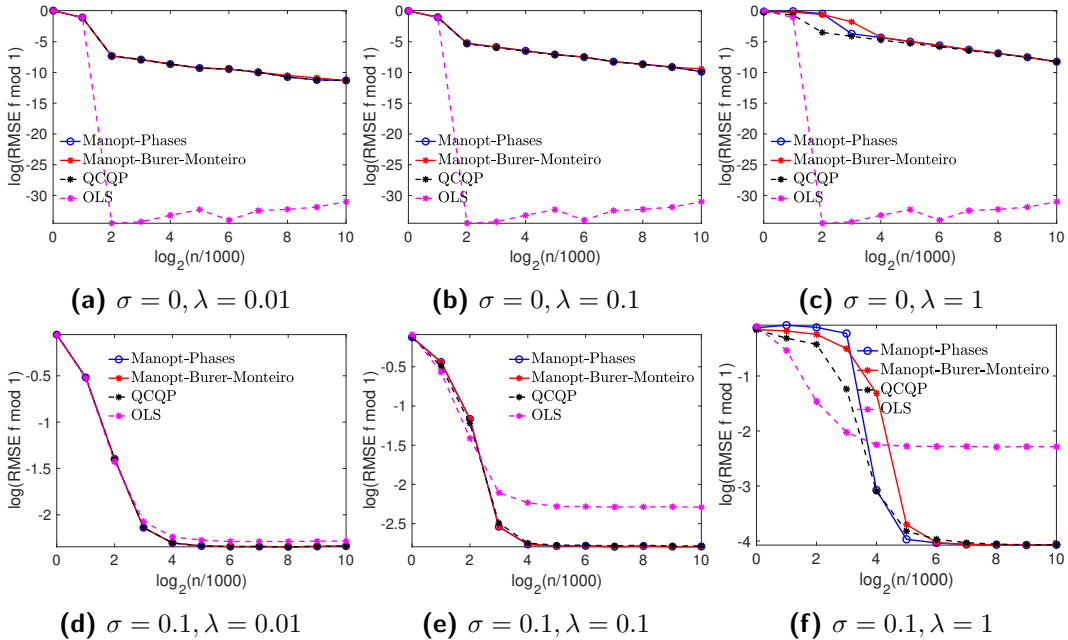


Figure 16: RMSE comparison for $f \bmod 1$ denoising, for the synthetic example $f(x, y) = 6xe^{-x^2-y^2}$, with $k = 1$ (Chebychev distance), while varying σ and λ . **QCQP** denotes Algorithm 1 where the unwrapping stage is performed via **OLS** (4.23).

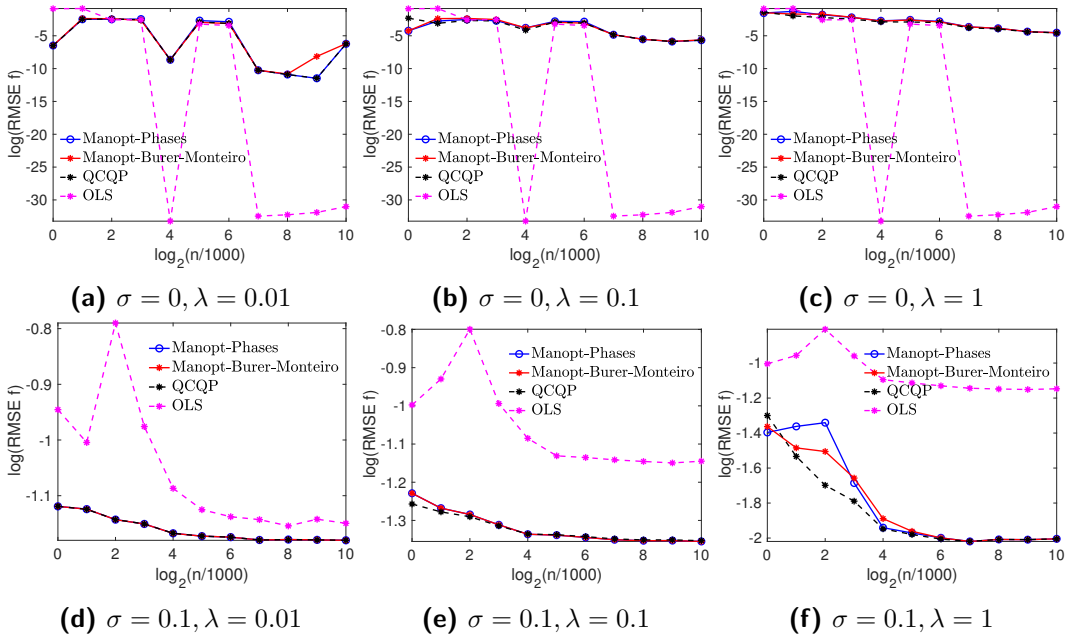


Figure 17: RMSE comparison for the final estimate f , for the synthetic example $f(x, y) = 6xe^{-x^2-y^2}$, with $k = 1$ (Chebychev distance), across several noise levels σ and values of λ . Here, **QCQP** denotes Algorithm 1, where the unwrapping stage is performed by **OLS** (4.23).

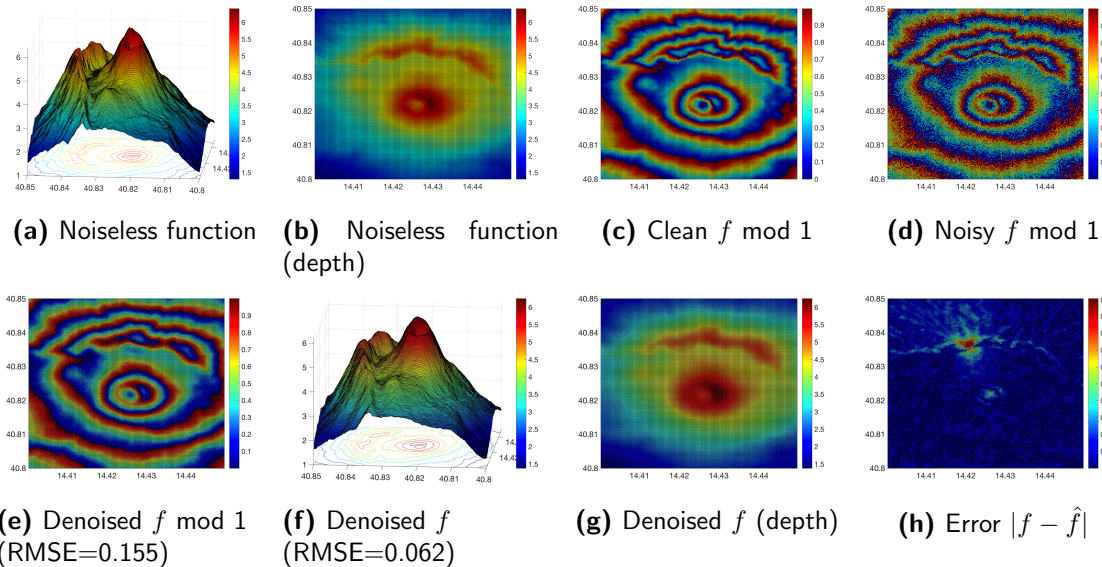


Figure 18: Elevation map of Mount Vesuvius with $n = 32400$, recovered by Manopt-Phases with $k = 1$ (Chebychev distance), $\lambda = 1$, and noise level $\sigma = 0.10$ under the Gaussian model.

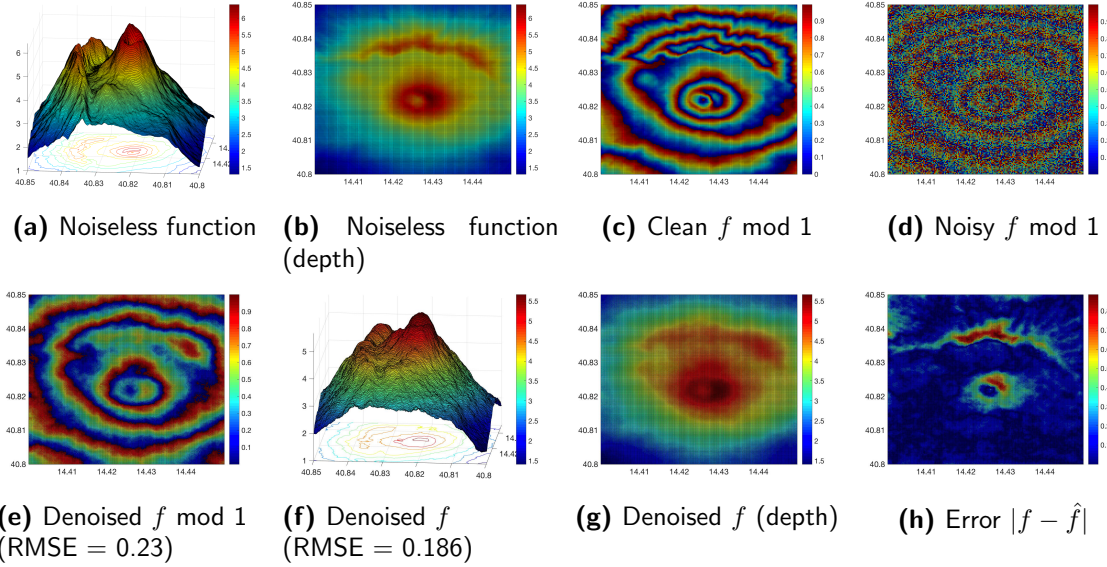


Figure 19: Elevation map of Mount Vesuvius with $n = 32400$, $k = 1$ (Chebychev distance), $\lambda = 1$, and noise level $\sigma = 0.25$ under the Gaussian model, as recovered by Manopt-Phases.

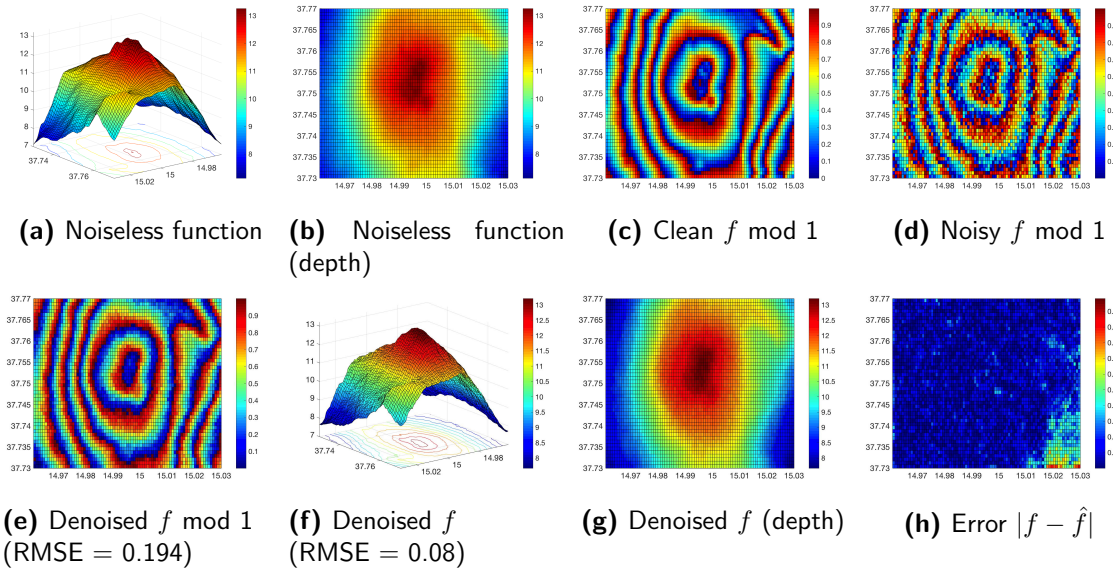


Figure 20: Elevation map of Mount Etna with $n = 4100$, recovered by Manopt-Phases with $k = 1$ (Chebychev distance), $\lambda = 0.3$, and noise level $\sigma = 0.10$ under the Gaussian model.

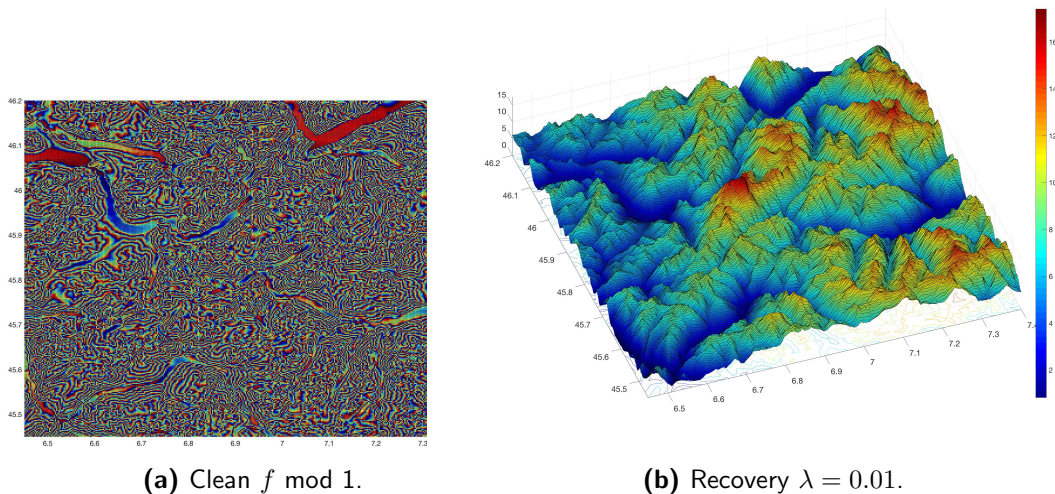


Figure 21: Recovery of the elevation map of Mont Blanc from over 1 million clean sample points, as recovered by Manopt-Phases. Despite the challenging topology, we are able to almost perfectly recover (unwrap) the shape of the mountain, in just under 18 seconds. RMSE for the mod 1 recovery is 0.0457, while the RMSE for the final f estimates was 0.69.

where $[x]$ denotes the fractional part of x , and $t \in \mathbb{R}$. By considering a regular sampling on \mathbb{R} with step length T , it is shown (Bhandari et al., 2017, Theorem 1) that if $T \leq 1/(2\pi e)$, then their Algorithm recovers the samples of g exactly (assuming no noise) from N^{th} order finite differences of the modulo samples (with N suitably large). Consequently, as g is bandlimited, one can recover the function g itself via a low pass filter. The analysis for unwrapping the samples mainly relies on the additional assumption that $g \in C^\infty(\mathbb{R}) \cap L^\infty(\mathbb{R})$. This assumption is required in order to be able to control the N^{th} order finite differences of the modulo samples for appropriately large N . Moreover, Theorem 1 only holds in the case of *noiseless* modulo samples, while the conditions under which the method stably unwraps the samples are unclear. As shown in the experiments, our finding is that their algorithm is stable only for very low levels of noise. Finally, note that for $\lambda = 0.5$, (10.1) corresponds to a centered modulo 1 function, with range in $[-0.5, 0.5)$. It is easy to verify that

$$M_{0.5}(g(t)) = \left(\left[g(t) + \frac{1}{2} \right] - \frac{1}{2} \right) = \begin{cases} g(t) \bmod 1 & ; \text{ if } g(t) \bmod 1 < 1/2 \\ g(t) \bmod 1 - 1 & ; \text{ if } g(t) \bmod 1 \geq 1/2 \end{cases} \quad (10.2)$$

Therefore, (10.2) implies that there is a one-to-one correspondence between $M_{0.5}(g(t))$ and $g(t) \bmod 1$, the latter of which we adopt throughout our paper.

10.2. Relation to group synchronization

The problem we consider is closely related to the well-known *group synchronization* problem introduced by Singer (2011). The synchronization of clocks in a distributed network

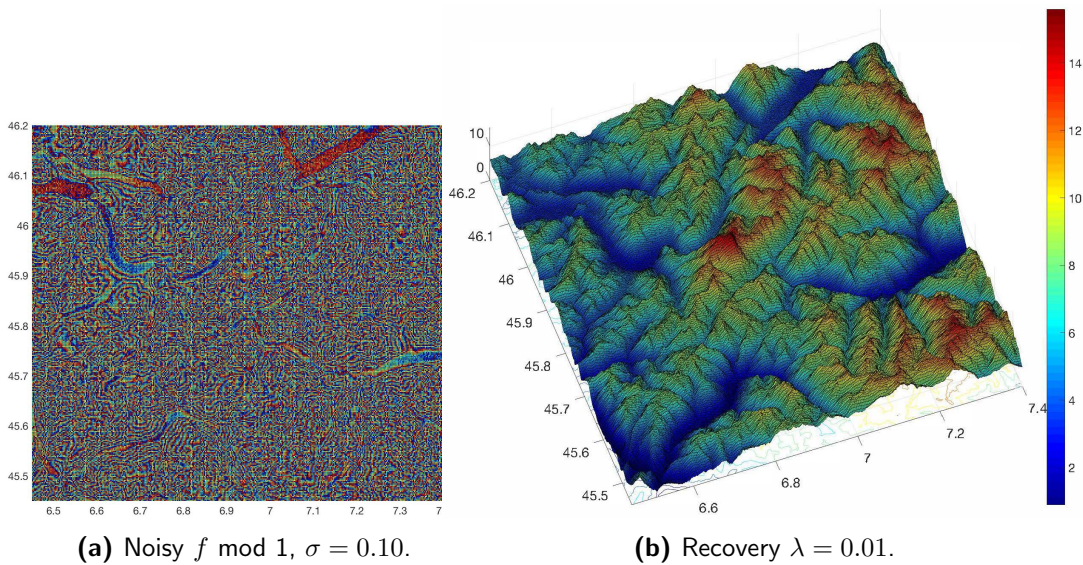


Figure 22: Recovery of the elevation map of Mont Blanc from over 1 million sample points under the Gaussian noise model with $\sigma = 10\%$, as recovered by Manopt-Phases. Despite the noise level and the topology, our approach recovers the details of the topographical relief, in under 20 seconds. RMSE for the mod 1 recovery is 0.26, while the RMSE for the final f estimates equals 1.06.

from noisy pairwise measurements of their time offsets is one such example of synchronization, where the underlying group is the real line \mathbb{R} . The eigenvector and semidefinite programming methods for solving an instance of the synchronization problem were initially introduced by Singer (2011) in the context of angular synchronization (over the group $\text{SO}(2)$ of planar rotations) where one is asked to estimate n unknown angles $\theta_1, \dots, \theta_n \in [0, 2\pi)$ given m noisy measurements δ_{ij} of their offsets $\theta_i - \theta_j \pmod{2\pi}$. The difficulty of the problem is amplified on one hand by the amount of noise in the pairwise offset measurements, and on the other hand by the fact that $m \ll \binom{n}{2}$, i.e., only a very small subset of all possible pairwise offsets are measured. Given the noisy measurements, Singer (2011) introduced an approach that first embeds the noisy angle differences on the unit complex circle via $Z_{ij} = \exp(i2\pi\delta_{ij})$, leading to the Hermitian $n \times n$ matrix \mathbf{Z} . Thereafter, one recovers the angle estimates by maximizing a quadratic form $\mathbf{x}^* \mathbf{Z} \mathbf{x}$, subject to \mathbf{x} lying on a sphere. This is similar to the formulation in (4.2), but without the linear term. However, there is a *fundamental difference* between the two problems. The matrix in the quadratic term in (4.2) is formed using the Laplacian of the smoothness regularization graph, and thus is independent of the data (given noisy mod 1 samples). On the other hand, as described above, the matrix \mathbf{Z} in synchronization is formed using the given noisy pairwise angle offsets embedded on the unit circle, and hence depends on the data. Additionally, the angular representations of the modulo samples bear a smoothness property arising due to the function f , while the angle offsets in synchronization are typically arbitrary.

We point out that a similar objective function encompassing both a quadratic and a linear term arises in the setting of synchronization with *anchors*, which are nodes whose corresponding group element is known a priori. The goal is to combine information given by the anchors with the matrix of noisy pairwise group ratios, in order to estimate the corresponding group elements of the non-anchor nodes. Such a variation of the synchronization problem has been explored in prior work by a subset of the authors (see Cucuringu et al. (2012b); Cucuringu (2015)), who introduced several methods for incorporating anchor information in the synchronization problem over \mathbb{Z}_2 , in the context of the molecule problem from structural biology. The first approach relies on casting the problem as a quadratically constrained quadratic program (QCQP) similar to the approach we pursue for denoising the modulo 1 samples, while a second one relies on an SDP formulation.

10.3. Phase unwrapping

We now discuss methods from the phase unwrapping literature to put our methodology in context. Phase unwrapping is a field in its own with a vast body of work, therefore this section is by no means a comprehensive overview of the literature.

Phase unwrapping is a classical and challenging problem in signal processing with a long line of work wherein one recovers the original phase values (in radians) from their wrapped (modulo 2π) versions. Formally, for an unknown sequence of phase values $(\phi_i)_{i=1}^n$ (in radians), one is given the wrapped values $\psi_i = w(\phi_i)$, where $w : \mathbb{R} \rightarrow [-\pi, \pi)$ is the wrap function that outputs centered modulo 2π values. We consider $\phi_i = \phi(\mathbf{x}_i)$ for some unknown function $\phi : \mathbb{R}^d \rightarrow \mathbb{R}$, with $\mathbf{x}_i \in \mathbb{R}^d$. The goal is to recover the original ϕ_i 's from the ψ_i 's. The problem is, of course, ill posed in general so one typically needs to make some additional smoothness assumptions on ϕ , although some methods are developed with an eye towards dealing with *abrupt* phase changes. For $d = 1$, Itoh (1982) showed in 1982 that if the condition

$$|\phi_i - \phi_{i-1}| \leq \pi \tag{10.3}$$

holds for each i , then this implies $\phi_i - \phi_{i-1} = w(\psi_i - \psi_{i-1})$ for all $i = 1, \dots, n$. Thereafter, one can easily recover the ϕ_i 's in a sequential manner by integration, up to a global shift by an integer multiple of 2π . We remark that this approach is essentially the Quotient Tracker Algorithm (QTA), discussed in Section 4.2. A similar argument extends to the case $d = 2$ (cf., (Ying, 2006, Lemma 1.2)). Of course, (10.3) will typically not hold in the presence of noise. Therefore, several robust alternatives have been proposed to date, and we briefly review some of them for the $d = 2$ case (which is also one of the most studied cases due to the many applications). The methods can be broadly classified into the following categories. We refer the reader to (Ying, 2006, Section 2.1) for a more detailed overview of these methods.

- *Least squares methods.* This class of methods seeks to find the phase function $\phi : \mathbb{R}^2 \rightarrow \mathbb{R}$ that minimizes a cost function of the form

$$\| \partial_x \phi - w(\partial_x \psi) \|_p + \| \partial_y \phi - w(\partial_y \psi) \|_p . \tag{10.4}$$

Here, $\| \cdot \|_p$ is the L_p norm and ∂_x, ∂_y denote partial derivatives w.r.t x, y respectively. Specifically, one solves a discrete version of (10.4) with discrete partial derivatives. For

$p = 2$, the solution actually has an analytical form given by the discrete form of Poisson’s equation with Neumann boundary conditions. This is known to be efficiently implementable, using for instance methods based on fast Fourier transform (FFT) (Pritt and Shipman (1994)), discrete cosine transform (DCT) (Kerr et al. (1996)), multigrid techniques (Pritt (1996)) etc. More broadly, the least-squares formulation of phase unwrapping (both the weighted and unweighted versions) lead to discretized partial differentiable equations, which are solvable efficiently by standard methods from the numerical PDE literature (see for eg., Takajo and Takahashi (1988a,b); Hunt (1979); Ghiglia and Romero (1989, 1994); Buzbee et al. (1970)). We remark that smoothness constraints have been incorporated by introducing a Tikhonov regularization term, which increases robustness to noise and missing data (Marroquin and Rivera (1995); Press et al. (2007)). Huang et al. (2012) proposed a TV (Total variation) minimization based method wherein the unwrapped phase gradients are estimated from the wrapped counterparts via a TV norm regularized least-squares method. Rivera and Marroquin (2004) proposed solving a “half quadratic” regularized-least squares problem, and proposed convex as well as non-convex formulations of the problem. Finally, we note that Ghiglia and Romero (1996) proposed an algorithm to minimize a more general ℓ_p norm based objective function for 2D phase unwrapping; they showed that this framework is equivalent to weighted least squares phase unwrapping where the weights are data dependent.

- *Branch cut methods.* A common problem that arises in 2D phase unwrapping is the so-called *path dependent problem*, wherein the integration result depends on the path chosen between two points. In particular, it is known (Kreysig (1966)) that if the unwrapping is path dependent then it implies that there exists a closed path C such that $\sum_C [w(\Delta_x \psi_{m,n}) + w(\Delta_y \psi_{m,n})]$ is not zero; here Δ_x, Δ_y denote discrete partial derivatives. This non-zero value is referred to as the *residue*, and arises for instance due to noise or (near) discontinuities in the original phase signal (see for eg., (Ying, 2006, Section 1.1.2)). Branch cut methods rely on the assumption that paths from positive residues to negative residues (referred to as branch cuts) cross the regions corresponding to the discontinuities. To this end, many algorithms have been proposed to find the optimal branch cuts in the phase image (Huntley (1989); Bone (1991); Prati et al. (1990); Chavez et al. (2002)).
- *Network flow methods.* These methods are similar to branch-cut methods and rely on the same assumptions. The difference is that they explicitly try to quantify the discontinuities at each sample and then find the optimal unwrapped phases that minimize the overall discontinuities. Broadly speaking, the idea is to construct a network with the nodes defined by the residues, and with a directed arc from the positive residues to the negative ones. With a suitable cost per unit flow assigned to each arc, the algorithms then seek to find a flow that minimizes the cumulative cost over all the arcs. (Towers et al. (1991); Ching et al. (1992); Takeda and Abe (1996)).

More recently, Bioucas-Dias and Valadao (2007) proposed a method that involves the minimization of an energy functional based on a generalization of the classical ℓ_p norm, using graph cut algorithms. Bioucas-Dias et al. (2008); Valadão and Bioucas-Dias (2009) proposed

more noise tolerant versions of PUMA, which together with PUMA constitute the current state of art in phase unwrapping. Bioucas-Dias and Valadao (2007) proposed a two-stage algorithm similar in spirit to ours. Specifically, the first stage involves obtaining the denoised modulo 2π samples via a local adaptive denoising scheme; the second stage then takes as input these denoised samples and unwraps them using PUMA. Assuming the original phase to be piecewise smooth that is well approximated by a polynomial in the neighborhood of the estimation point, their denoising scheme is based on local polynomial approximations in sliding windows of varying sizes. Valadão and Bioucas-Dias (2009) adopted a Bayesian framework assuming a first order Markov random field prior. The proposed algorithm first unwraps the phase image using PUMA, and then in the second stage, it denoises the unwrapped phase to form the final estimate.

Gonzlez and Jacques (2014) considered placing a sparsity prior on the phase signal in the wavelet domain, to model discontinuities. They proposed an unwrapping algorithm based on solving a ℓ_1 minimization problem. Kamilov et al. (2015) proposed a method based on the minimization of an energy functional that includes a weighted ℓ_1 norm based data fidelity term, along with a regularizer term based on the higher order total variation (TV) norm.

Extensions to higher dimensions have also attracted significant interest from the community, in light of numerous applications. In SAR interferometry, the data typically consists of a sequence of time-dependent consecutive interferograms. Existing approaches utilize the time dimension to perform 1-D phase unwrapping along the time axis component (as opposed to performing standard 2-D phrase unwrapping for each time stamped interferogram), an approach which increases robustness of the recovery process for scenarios of steep gradients in the phase maps (Huntley and Saldner (1993)). This approach has been subsequently extended to the case of 3-D phrase unwrapping in a 3-D space-time domain by Costantini et al. (2002). They formulate the phase unwrapping problem as a linear integer minimization, and report increased robustness when compared to standard 2-D phase unwrapping on simulated and real SAR images. More recently, Osmanoglu et al. (2014) introduced a novel 3-D phase unwrapping approach for generating digital elevation models, by combining multiple SAR acquisitions using an extended Kalman filter.

Very recently, Dardikman et al. (2018) introduced a new 4-D phase unwrapping approach for time-lapse quantitative phase microscopy, which allows for the reconstruction of optically thick objects that are optically thin in a certain temporal point and angular view. The authors leverage both the angular dimension and the temporal dimension, in addition to the usual $x - y$ dimensions, to enhance the reconstruction process. They report improved numerical results over other state-of-the-art methods. Finally, we remark that Jenkinson (2003) introduced a fast algorithm for the N -dimensional phase unwrapping problem, built around a cost function that leads to an integer programming problem, solved via a *greedy* optimization approach, and tested on 3D MRI medical data for venogram studies.

11. Concluding remarks and future work

We considered the problem of denoising noisy modulo 1 samples of a smooth function, arising in the context of the popular phase unwrapping problem. We proposed a method centered around solving a trust-region subproblem with a sphere constraint, and provided extensive

empirical evidence as well as theoretical analysis, that altogether highlight its robustness to noise. We mostly focused on the one-dimensional case, but also provided extensions to the two-dimensional setting, whose applications include the phase unwrapping problem. In addition, we have also explored two formulations that leverage tools from Riemannian optimization on manifolds, including a relaxation based on semidefinite programming, solvable fast via a Burer-Monteiro approach.

There are several possible research directions worth exploring in future work, that we detail below.

- An interesting direction would be to better understand the unwrapping stage of our approach, either via the simple OLS (potentially with an additional smoothness regularization term), or by proposing a new method based on group synchronization (Singer (2011)). For the latter direction, finding the right global shift of each sample point (i.e., the quotient q_i) can be cast a synchronization problem over the real line \mathbb{R} , given pairwise (potentially noisy) offset measurements $q_i - q_j$ between nearby points. The non-compactness of \mathbb{R} renders the eigenvector or semidefinite programming relaxations no longer applicable directly, thus motivating an approach that first compactifies the real line by wrapping it over the unit circle, as explored recently by Cucuringu (2016) in the context of ranking from noise incomplete pairwise information.
- Another potentially interesting direction to explore would be a patch-based divide-and-conquer method that first decomposes the graph into many (either overlapping or non-overlapping) subgraphs (potentially with an eye towards the *jumps* in the mod 1 measurements), solves the problem locally and then integrates the local solutions into a globally consistent framework, in the spirit of existing methods from the group synchronization literature (see for eg., Cucuringu et al. (2012a,b)).
- As discussed briefly at the start of Section 4, an interesting approach for solving (3.7) would be to first discretize the angular domain, and then solve (3.7) approximately via dynamic programming. In general, the graph G has tree-width $\Omega(k^d)$, rendering the computational cost of naive dynamic programming to be exponential in k^d .
- Another potential direction would be to explore a variation of (3.2), with the last term replaced by a TV norm

$$\min_{g_1, \dots, g_n \in \mathbb{C}; |g_i|=1} \sum_{i=1}^n |g_i - z_i|^2 + \lambda \sum_{\{i,j\} \in E} |g_i - g_j|. \quad (11.1)$$

At first sight, the TV norm may not seem beneficial in light of the smoothness assumptions in our setting; however, in certain imaging applications (Fang et al. (2006)), it is desirable to preserve any sharp discontinuities that may arise in the boundaries between classes within the image. Rudin et al. (1992) introduced **ROF**, a total variation (TV)-based approach as a regularizing functional for image denoising. The intuition is that the TV term in the minimization discourages the solution from having an oscillatory behavior, while allowing it to have discontinuities. Inspired by the above seminal variational framework, and building on their prior work, such as Bergmann

et al. (2014), Bergmann and Weinmann (2016) introduced variational models for various nonlinear data spaces. In particular, they developed proximal point algorithms for the solution of the second order TV-based problems for denoising, inpainting, and a combination of both, for combined *cyclic* (i.e., modulo measurements) and linear space valued images, along with a convergence analysis and a number of applications to real-world problems.

- Another natural direction to explore is the development and theoretical analysis of a “single-stage method” that directly outputs denoised estimates of the original real-valued samples.
- Finally, an interesting question would be to consider the regression problem, where one attempts to learn a smooth, robust estimate to the underlying mod 1 function itself (not just the samples) and/or the original f , from noisy mod 1 samples of f .

12. Acknowledgements

We thank Joel Tropp for bringing this problem to our attention. We are grateful to Afonso Bandeira and Joel Tropp for insightful discussions on the topic and pointing out the relevant literature; Nicolas Boumal for helping us in using the Manopt optimization toolbox and providing useful comments; Yuji Nakatsukasa for many helpful discussions on the trust-region subproblem, and for kindly providing us the code for the algorithm in his paper Adachi et al. (2017); Ayush Bhandari for sharing the code from his paper Bhandari et al. (2017). This work was done while HT was affiliated to the Alan Turing Institute, London, and the School of Mathematics, University of Edinburgh, UK. Both MC and HT were supported by EPSRC grant EP/N510129/1 at the Alan Turing Institute.

Bibliography

- P.-A Absil, C.G Baker, and K.A Gallivan. Trust-region methods on riemannian manifolds. *Foundations of Computational Mathematics*, 7(3):303–330, 2007.
- P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, Princeton, NJ, 2008. ISBN 978-0-691-13298-3.
- S. Adachi, S. Iwata, Y. Nakatsukasa, and A. Takeda. Solving the trust-region subproblem by a generalized eigenvalue problem. *SIAM Journal on Optimization*, 27(1):269–291, 2017.
- A. I. Barvinok. Problems of distance geometry and convex properties of quadratic maps. *Discrete & Computational Geometry*, 13(2):189–202, Mar 1995.
- Ronny Bergmann and Andreas Weinmann. A second order TV-type approach for inpainting and denoising higher dimensional combined cyclic and vector space data. *Journal of Mathematical Imaging and Vision*, 55(3):401–427, 2016.

- Ronny Bergmann, Friederike Laus, Gabriele Steidl, and Andreas Weinmann. Second order differences of cyclic data and applications in variational denoising. *SIAM Journal on Imaging Sciences*, 7(4):29162953, 2014.
- A. Bhandari, F. Kraemer, and R. Raskar. On unlimited sampling. ArXiv e-prints, arXiv:1707.06340v1, 2017.
- J. Bioucas-Dias, V. Katkovnik, J. Astola, and K. Egiazarian. Absolute phase estimation: adaptive local denoising and global unwrapping. *Appl. Opt.*, 47(29):5358–5369, 2008.
- J.M. Bioucas-Dias and G. Valadao. Phase unwrapping via graph cuts. *IEEE Transactions on Image Processing*, 16(3):698–709, 2007.
- D.J. Bone. Fourier fringe analysis: the two-dimensional phase unwrapping problem. *Appl. Opt.*, 30(25):3627–3632, 1991.
- N. Boumal, B. Mishra, P.-A. Absil, and R. Sepulchre. Manopt, a Matlab toolbox for optimization on manifolds. *Journal of Machine Learning Research*, 15:1455–1459, 2014. URL <http://www.manopt.org>.
- N. Boumal, P.-A. Absil, and C. Cartis. Global rates of convergence for nonconvex optimization on manifolds. *arXiv preprint arXiv:1605.08101*, 2016a.
- N. Boumal, V. Voroninski, and A.S. Bandeira. The non-convex Burer-Monteiro approach works on smooth semidefinite programs. In *Advances in Neural Information Processing Systems 29*, pages 2757–2765. 2016b.
- S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan. *Linear matrix inequalities in system and control theory*. SIAM, 1994.
- S. Burer and R.D.C. Monteiro. Local minima and convergence in low-rank semidefinite programming. *Mathematical Programming*, 103(3):427–444, 2005.
- B.L Buzbee, G.H Golub, and C.W Nielson. On direct methods for solving poissons equations. *SIAM Journal on Numerical analysis*, 7(4):627–656, 1970.
- E.J. Candès, Y.C. Eldar, T. Strohmer, and V. Voroninski. Phase retrieval via matrix completion. *SIAM Journal on Imaging Sciences*, 6(1):199–225, 2013.
- S. Chavez, Q.S Xiang, and L. An. Understanding phase maps in mri: a new cutline phase unwrapping method. *IEEE Transactions on Medical Imaging*, 21(8):966–977, 2002.
- N.H Ching, R. Rosenfeld, and M. Braun. Two-dimensional phase unwrapping using a minimum spanning tree algorithm. *IEEE Transactions on Image Processing*, 1(3):355–365, 1992.
- M. Costantini, F. Malvarosa, F. Minati, L. Pietranera, and G. Milillo. A three-dimensional phase unwrapping algorithm for processing of multitemporal sar interferometric measurements. In *IEEE International Geoscience and Remote Sensing Symposium*, volume 3, pages 1741–1743 vol.3, 2002.

- M. Cucuringu. Synchronization over Z_2 and community detection in multiplex networks with constraints. *Journal of Complex Networks*, 3:469–506, 2015.
- M. Cucuringu. Sync-Rank: Robust Ranking, Constrained Ranking and Rank Aggregation via Eigenvector and Semidefinite Programming Synchronization. *IEEE Transactions on Network Science and Engineering*, 3(1):58–79, 2016.
- M. Cucuringu and H. Tyagi. On denoising modulo 1 samples of a function. *arXiv:1710.10210*, (to appear in *AISTATS 2018*), 2018.
- M. Cucuringu, Y. Lipman, and A. Singer. Sensor network localization by eigenvector synchronization over the Euclidean group. *ACM Trans. Sen. Netw.*, 8(3):19:1–19:42, 2012a.
- M. Cucuringu, A. Singer, and D. Cowburn. Eigenvector synchronization, graph rigidity and the molecule problem. *Information and Inference*, 1(1):21–67, 2012b.
- G. Dardikman, G. Singh, and N.T. Shaked. Four dimensional phase unwrapping of dynamic objects in digital holography. *Opt. Express*, 26(4):3772–3778, 2018.
- M. Elad and M. Aharon. Image denoising via sparse and redundant representations over learned dictionaries. *IEEE Transactions on Image processing*, 15(12):3736–3745, 2006.
- Q. Fang, P. M. Meaney, and K. D. Paulsen. The multidimensional phase unwrapping integral and applications to microwave tomographical image reconstruction. *IEEE Transactions on Image Processing*, 15(11):3311–3324, 2006.
- M. Fiedler. Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal*, 23(2):298–305, 1973.
- D.C. Ghiglia and L.A. Romero. Direct phase estimation from phase differences using fast elliptic partial differential equation solvers. *Opt. Lett.*, 14(20):1107–1109, 1989.
- D.C. Ghiglia and L.A. Romero. Robust two-dimensional weighted and unweighted phase unwrapping that uses fast transforms and iterative methods. *J. Opt. Soc. Am. A*, 11(1):107–117, 1994.
- D.C. Ghiglia and L.A. Romero. Minimum lp-norm two-dimensional phase unwrapping. *J. Opt. Soc. Am. A*, 13(10):1999–2013, 1996.
- A. Gonzalez and L. Jacques. Robust phase unwrapping by convex optimization. In *2014 IEEE International Conference on Image Processing (ICIP)*, pages 1713–1717, 2014.
- N.I.M Gould, S. Lucidi, M. Roma, and P.L. Toint. Solving the trust-region subproblem using the lanczos method. *SIAM Journal on Optimization*, 9(2):504–525, 1999.
- L. C. Graham. Synthetic interferometer radar for topographic mapping. *Proceedings of the IEEE*, 62(6):763–768, 1974.
- W.W. Hager. Minimizing a quadratic over a sphere. *SIAM J. on Optimization*, 12(1):188–208, 2001.

- D.L Hanson and E.T Wright. Hanson-wright inequality and sub-gaussian concentration. *Ann. Math. Statist.*, 42:1079–1083, 1971.
- M. Hedley and D. Rosenfeld. A new two-dimensional phase unwrapping algorithm for mri images. *Magnetic Resonance in Medicine*, 24(1):177–181, 1992.
- A. Hooper and H.A. Zebker. Phase unwrapping in three dimensions with application to insar time series. *J. Opt. Soc. Am. A*, 24(9):2737–2747, 2007.
- H.Y.H Huang, L. Tian, Z. Zhang, Y. Liu, Z. Chen, and G. Barbastathis. Path-independent phase unwrapping using phase gradient and total-variation (tv) denoising. *Optics express*, 20(13):14075–14089, 2012.
- Y.Y. Hung. Shearography for non-destructive evaluation of composite structures. *Optics and Lasers in Engineering*, 24(2):161 – 182, 1996.
- B. R. Hunt. Matrix formulation of the reconstruction of phase values from phase differences. *J. Opt. Soc. Am.*, 69(3):393–399, Mar 1979.
- J.M Huntley. Noise-immune phase unwrapping algorithm. *Appl. Opt.*, 28(16):3268–3270, 1989.
- J.M Huntley and H. Saldner. Temporal phase-unwrapping algorithm for automated interferogram analysis. *Appl. Opt.*, 32(17):3047–3052, 1993.
- K. Itoh. Analysis of the phase unwrapping algorithm. *Appl. Opt.*, 21(14):2470–2470, 1982.
- K. Jaganathan, Y.C. Eldar, and B. Hassibi. Phase retrieval: An overview of recent developments. *CoRR*, abs/1510.07713, 2015.
- M. Jenkinson. Fast, automated, n-dimensional phase-unwrapping algorithm. *Magnetic Resonance in Medicine*, 49(1):193–197, 2003.
- U.S. Kamilov, I.N. Papadopoulos, M.H. Shoreh, D. Psaltis, and M. Unser. Isotropic inverse-problem approach for two-dimensional phase unwrapping. *J. Opt. Soc. Am. A*, 32(6):1092–1100, 2015.
- S. Kavusi and E. Abbas. Quantitative study of high-dynamic-range image sensor architectures. *Proc.SPIE*, 5301:5301 – 5301 – 12, 2004.
- D. Kerr, G.H. Kaufmann, and G.E. Galizzi. Unwrapping of interferometric phase-fringe maps by the discrete cosine transform. *Appl. Opt.*, 35(5):810–816, 1996.
- W. Kester. Mt-025 tutorial adc architectures vi: Folding adcs. Analog Devices, Tech. report, 2009.
- E. Kreysig. *Advanced Engineering Mathematics*. Wiley, New York, 1966.
- P. Lauterbur. Image formation by induced local interactions: examples employing nuclear magnetic resonance. *Nature*, 242:190–191, 1973.

- Y. Lou, X. Zhang, S. Osher, and A. Bertozzi. Image recovery via nonlocal operators. *Journal of Scientific Computing*, 42(2):185–197, 2010.
- J.L. Marroquin and M. Rivera. Quadratic regularization functionals for phase unwrapping. *J. Opt. Soc. Am. A*, 12(11):2393–2400, 1995.
- J.J. Morè and D.C. Sorensen. Computing a trust region step. *SIAM Journal on Scientific and Statistical Computing*, 4(3):553–572, 1983.
- B. Osmanoglu, T. H. Dixon, and S. Wdowinski. Three-Dimensional Phase Unwrapping for Satellite Radar Interferometry, I: DEM Generation. *IEEE Transactions on Geoscience and Remote Sensing*, 52(2):1059–1075, 2014.
- D. Paoletti, G.S. Spagnolo, P. Zanetta, M. Facchini, and D. Albrecht. Manipulation of speckle fringes for non-destructive testing of defects in composites. *Optics and Laser Technology*, 26(2):99 – 104, 1994.
- C. Prati, M. Giani, and N. Leuratti. Sar interferometry: A 2-d phase unwrapping technique based on phase and absolute values informations. In *10th Annual International Symposium on Geoscience and Remote Sensing*, pages 2043–2046, 1990.
- R.G. Pratt and M.H. Worthington. The application of diffraction tomography to crosshole seismic data. *GEOPHYSICS*, 53(10):1284–1294, 1988.
- W.H. Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery. *Numerical Recipes 3rd Edition: The Art of Scientific Computing*. Cambridge University Press, New York, NY, USA, 3 edition, 2007.
- M. D Pritt and J.S. Shipman. Least-squares two-dimensional phase unwrapping using fft’s. *IEEE Transactions on Geoscience and Remote Sensing*, 32(3):706–708, 1994.
- M.D. Pritt. Phase Unwrapping by Means of Multigrid Techniques for Interferometric SAR. *Geoscience and Remote Sensing, IEEE Transactions on*, 34(3):728–738, 1996.
- Inam Ur Rahman, Iddo Drori, Victoria C. Stodden, David L. Donoho, and Peter Schrder. Multiscale representations for manifold-valued data. *Multiscale Modeling & Simulation*, 4(4):1201–1232, 2005. doi: 10.1137/050622729.
- F. Rendl and H. Wolkowicz. A semidefinite framework for trust region subproblems with applications to large scale minimization. *Mathematical Programming*, 77(1):273–299, 1997.
- J. Rhee and Y. Joo. Wide dynamic range cmos image sensor with pixel level adc. *Electronics Letters*, 39(4):360–361, 2003.
- Mariano Rivera and Jose L. Marroquin. Half-quadratic cost functions for phase unwrapping. *Opt. Lett.*, 29(5):504–506, 2004.
- M. Rojas, A.S. Santos, and D.C. Sorensen. A new matrix-free algorithm for the large-scale trust-region subproblem. *SIAM Journal on Optimization*, 11(3):611–646, 2001.

- M. Rudelson and R. Vershynin. Hanson-wright inequality and sub-gaussian concentration. *Electron. Commun. Probab.*, 18:9 pp., 2013.
- Leonid I. Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Phys. D*, 60(1-4):259–268, November 1992. ISSN 0167-2789. doi: 10.1016/0167-2789(92)90242-F. URL [http://dx.doi.org/10.1016/0167-2789\(92\)90242-F](http://dx.doi.org/10.1016/0167-2789(92)90242-F).
- K. Sasagawa, T. Yamaguchi, M. Haruta, Y. Sunaga, H. Takehara, H. Takehara, T. Noda, T. Tokuda, and J. Ohta. An implantable cmos image sensor with self-reset pixels for functional brain imaging. *IEEE Transactions on Electron Devices*, 63(1):215–222, 2016.
- A. Singer. Angular synchronization by eigenvectors and semidefinite programming. *Appl. Comput. Harmon. Anal.*, 30(1):20–36, 2011.
- D.C. Sorensen. Newtons method with a model trust region modification. *SIAM Journal on Numerical Analysis*, 19(2):409–426, 1982.
- H. Takajo and T. Takahashi. Least-squares phase estimation from the phase difference. *J. Opt. Soc. Am. A*, 5(3):416–425, 1988a.
- H. Takajo and T. Takahashi. Noniterative method for obtaining the exact solution for the normal equation in least-squares phase estimation from the phase difference. *J. Opt. Soc. Am. A*, 5(11):1818–1827, Nov 1988b.
- M. Takeda and T. Abe. Phase unwrapping by a maximum cross-amplitude spanning tree algorithm: a comparative study. *Optical Engineering*, 35:35 – 35 – 7, 1996.
- D. P. Towers, T. R. Judge, and P. J. Bryanston-Cross. Automatic interferogram analysis techniques applied to quasi-heterodyne holography and ESPI. *Optics and Lasers in Engineering*, 14:239–281, 1991.
- G. Valadão and J. Bioucas-Dias. Cape: combinatorial absolute phase estimation. *J. Opt. Soc. Am. A*, 26(9):2093–2106, 2009.
- T.M. Venema and J.D. Schmidt. Optical phase unwrapping in the presence of branch points. *Opt. Express*, 16(10):6985–6998, May 2008.
- R. Vershynin. *Introduction to the non-asymptotic analysis of random matrices*, pages 210–268. Cambridge University Press, 2012.
- T. Yamaguchi, H. Takehara, Y. Sunaga, M. Haruta, M. Motoyama, Y. Ohta, T. Noda, K. Sasagawa, T. Tokuda, and J. Ohta. Implantable self-reset cmos image sensor and its application to hemodynamic response detection in living mouse brain. *Japanese Journal of Applied Physics*, 55(4S):04EM02, 2016.
- L. Ying. *Phase Unwrapping*. John Wiley and Sons, Inc., 2006.
- H.A. Zebker and R.M. Goldstein. Topographic mapping from interferometric synthetic aperture radar observations. *Journal of Geophysical Research: Solid Earth*, 91(B5):4993–4999, 1986.

- S. Zhang and Y. Huang. Complex quadratic optimization and semidefinite programming.
SIAM Journal on Optimization, 16(3):871–890, 2006.

Supplementary Material: Provably robust estimation of modulo 1 samples of a smooth function with applications to phase unwrapping

A. Rewriting the QCQP in real domain

We first show that $\lambda \mathbf{g}^* L \mathbf{g} = \bar{\mathbf{g}}^T H \bar{\mathbf{g}}$. Indeed, it holds true that

$$\bar{\mathbf{g}}^T H \bar{\mathbf{g}} = (\operatorname{Re}(\mathbf{g})^T \operatorname{Im}(\mathbf{g})^T) \begin{pmatrix} \lambda L & 0 \\ 0 & \lambda L \end{pmatrix} \begin{pmatrix} \operatorname{Re}(\mathbf{g}) \\ \operatorname{Im}(\mathbf{g}) \end{pmatrix} \quad (\text{A.1})$$

$$= (\operatorname{Re}(\mathbf{g})^T \operatorname{Im}(\mathbf{g})^T) \begin{pmatrix} \lambda L \operatorname{Re}(\mathbf{g}) \\ \lambda L \operatorname{Im}(\mathbf{g}) \end{pmatrix} \quad (\text{A.2})$$

$$= \operatorname{Re}(\mathbf{g})^T (\lambda L) \operatorname{Re}(\mathbf{g}) + \operatorname{Im}(\mathbf{g})^T (\lambda L) \operatorname{Im}(\mathbf{g}) \quad (\text{A.3})$$

$$= \lambda (\operatorname{Re}(\mathbf{g}) - \iota \operatorname{Im}(\mathbf{g}))^T L (\operatorname{Re}(\mathbf{g}) + \iota \operatorname{Im}(\mathbf{g})) \quad (\text{A.4})$$

$$= \lambda \mathbf{g}^* L \mathbf{g}. \quad (\text{A.5})$$

Next, we can verify that

$$\operatorname{Re}(\mathbf{g}^* \mathbf{z}) = \operatorname{Re}((\operatorname{Re}(\mathbf{g}) - \iota \operatorname{Im}(\mathbf{g}))^T (\operatorname{Re}(\mathbf{z}) + \iota \operatorname{Im}(\mathbf{z}))) \quad (\text{A.6})$$

$$= \operatorname{Re}(\mathbf{g})^T \operatorname{Re}(\mathbf{z}) + \operatorname{Im}(\mathbf{g})^T \operatorname{Im}(\mathbf{z}) \quad (\text{A.7})$$

$$= \bar{\mathbf{g}}^T \bar{\mathbf{z}}. \quad (\text{A.8})$$

Lastly, it is trivially seen that $\|\bar{\mathbf{g}}\|_2^2 = \|\mathbf{g}\|_2^2 = n$. Hence (4.2) and (4.5) are equivalent.

B. Trust-region sub-problem with ℓ_2 ball/sphere constraint

Consider the following two optimization problems

$$\left. \begin{array}{l} \min_{\mathbf{x}} \quad \mathbf{b}^T \mathbf{x} + \frac{1}{2} \mathbf{x}^T \mathbf{P} \mathbf{x} \\ \text{s.t.} \quad \|\mathbf{x}\|_2 \leq r \end{array} \right\} (\text{P1}) \quad \left| \quad \left. \begin{array}{l} \min_{\mathbf{x}} \quad \mathbf{b}^T \mathbf{x} + \frac{1}{2} \mathbf{x}^T \mathbf{P} \mathbf{x} \\ \text{s.t.} \quad \|\mathbf{x}\|_2 = r \end{array} \right\} (\text{P2}) \quad (\text{B.1})$$

with $\mathbf{P} \in \mathbb{R}^{n \times n}$ being a symmetric matrix. (P1) is known as the trust-region sub-problem in the optimization literature and has been studied extensively with a rich body of work. (P2) is closely related to (P1), and has a non-convex equality constraint. There exist algorithms that efficiently find the global solution of (P1) and (P2), to arbitrary accuracy. In this section, we provide a discussion on the characterization of the solution of these two problems.

To begin with, it is useful to note for (P1) that

- If the solution lies in the interior of the feasible domain, then it implies $\mathbf{P} \succeq 0$. This follows from the second order necessary condition for a local minimizer.
- In the other direction, if $\mathbf{P} \not\succeq 0$ then the solution will always lie on the boundary.

Surprisingly, we can characterize the solution of (P1), as shown in the following lemma.

Lemma 20 (Sorensen (1982)) \mathbf{x}^* is a solution to (P1) iff $\|\mathbf{x}^*\|_2 \leq r$ and $\exists \mu^* \geq 0$ such that (a) $\mu^* (\|\mathbf{x}^*\|_2 - r) = 0$, (b) $(\mathbf{P} + \mu^* I) \mathbf{x}^* = -\mathbf{b}$ and (c) $\mathbf{P} + \mu^* I \succeq 0$. Moreover, if $\mathbf{P} + \mu^* I \succ 0$, then the solution is unique.

Note that if the solution lies in the interior, and if \mathbf{P} is p.s.d and singular, then there will also be a pair of solutions on the boundary of the ball. This is easily verified. The solution to (P2) is characterized by the following lemma.

Lemma 21 (Hager (2001); Sorensen (1982)) \mathbf{x}^* is a solution to (P2) iff $\|\mathbf{x}^*\|_2 = r$ and $\exists \mu^*$ such that (a) $\mathbf{P} + \mu^*I \succeq 0$ and (b) $(\mathbf{P} + \mu^*I)\mathbf{x}^* = -\mathbf{b}$. Moreover, if $\mathbf{P} + \mu^*I \succ 0$, then the solution is unique.

The solution to (P1) and (P2) is closely linked to solving the nonlinear equation $\|(\mathbf{P} + \mu I)^{-1}\mathbf{b}\|_2 = r$. Let

$$\mathbf{x}(\mu) = -(\mathbf{P} + \mu I)^{-1}\mathbf{b} = -\sum_{j=1}^n \frac{\langle \mathbf{b}, \mathbf{q}_j \rangle}{\mu + \mu_j} \mathbf{q}_j, \quad (\text{B.2})$$

where $\mu_1 \leq \mu_2 \leq \dots \leq \mu_n$ and $\{\mathbf{q}_j\}$ are the eigenvalues and eigenvectors of \mathbf{P} respectively. Let us define $\phi(\mu)$ as

$$\phi(\mu) := \|\mathbf{x}(\mu)\|_2^2 = \sum_{j=1}^n \frac{\langle \mathbf{b}, \mathbf{q}_j \rangle^2}{(\mu + \mu_j)^2}. \quad (\text{B.3})$$

Denoting $\mathcal{S} = \{\mathbf{q} \in \mathbb{R}^n : \mathbf{P}\mathbf{q} = \mu_1\mathbf{q}\}$, there are two cases to consider.

1. $\langle \mathbf{b}, \mathbf{q} \rangle \neq 0$ for some $\mathbf{q} \in \mathcal{S}$

This is the easy case. $\phi(\mu)$ has a pole at $-\mu_1$ and is monotonically decreasing in $(-\mu_1, \infty)$ with $\lim_{\mu \rightarrow \infty} \phi(\mu) = 0$ and $\lim_{\mu \rightarrow -\mu_1} \phi(\mu) = \infty$. Hence there is a unique $\mu^* \in (-\mu_1, \infty)$ such that $\phi(\mu) = r^2$, and $\mathbf{x}(\mu^*) = -\sum_{j=1}^n \frac{\langle \mathbf{b}, \mathbf{q}_j \rangle}{\mu^* + \mu_j} \mathbf{q}_j$ will be the unique solution to (P2). Some remarks are in order.

- If \mathbf{P} was p.s.d and singular, then there is no solution to $\mathbf{P}\mathbf{x} = -\mathbf{b}$, since $\mathbf{b} \notin \text{colspan}(\mathbf{P})$. Also, since $\mu_1 = 0$, we would have $\mu^* \in (0, \infty)$. Hence the corresponding solution $\mathbf{x}(\mu^*)$ would be the same for (P1), (P2) and would lie on the boundary. Moreover, the solution will be unique due to Lemma 21 since $\mathbf{P} + \mu^*I \succ 0$.
- If \mathbf{P} was p.d and $\phi(0) < r^2$, then this would mean that the global solution to the unconstrained problem is a feasible point for (P1). In other words, $\mathbf{x}(0) = -\mathbf{P}^{-1}\mathbf{b}$ would be the unique solution to (P1) with $\mu^* = 0$. Moreover, $\mathbf{x}(\mu^*)$ with $\mu^* \in (-\mu_1, \infty)$ satisfying $\phi(\mu^*) = r^2$, would be the unique solution to (P2) (with $\mu^* < 0$); the uniqueness follows from Lemma 21 since $\mathbf{P} + \mu^*I \succ 0$.

2. $\langle \mathbf{b}, \mathbf{q} \rangle = 0$ for all $\mathbf{q} \in \mathcal{S}$

This is referred to as the “hard” case in the literature - $\phi(\mu)$ does not have a pole at $-\mu_1$, so $\phi(-\mu_1)$ is well defined. There are two possibilities.

- (a) If $\phi(-\mu_1) \geq r^2$ then the solution is straightforward – simply find the unique $\mu^* \in [-\mu_1, \infty)$ such that $\phi(\mu^*) = r^2$. This is possible since $\phi(\mu)$ is monotonically decreasing in $[-\mu_1, \infty)$. Hence,

$$\mathbf{x}(\mu^*) = -\sum_{j: \mu_j \neq \mu_1} \frac{\langle \mathbf{b}, \mathbf{q}_j \rangle}{\mu^* + \mu_j} \mathbf{q}_j$$

is the unique solution to (P2). If \mathbf{P} was p.s.d and singular, then $\mu_1 = 0$, and so $\mu^* \geq 0$. Hence, $\mathbf{x}(\mu^*)$ would be the solution to both (P1) and (P2) and would lie on the boundary.

- (b) If $\phi(-\mu_1) < r^2$, then slightly more work is needed. For $\theta \in \mathbb{R}$ and any $\mathbf{z} \in \mathcal{S}$ with $\|\mathbf{z}\|_2 = 1$, define

$$\mathbf{x}(\theta) := - \underbrace{\sum_{j:\mu_j \neq \mu_1} \frac{\langle \mathbf{b}, \mathbf{q}_j \rangle}{\mu_j - \mu_1} \mathbf{q}_j}_{\mathbf{x}(-\mu_1)} + \theta \mathbf{z}.$$

Then, it holds true that

$$\|\mathbf{x}(\theta)\|_2^2 = \sum_{j:\mu_j \neq \mu_1} \frac{\langle \mathbf{b}, \mathbf{q}_j \rangle^2}{(\mu_j - \mu_1)^2} + \theta^2 \quad (\text{B.4})$$

$$= \phi(-\mu_1) + \theta^2. \quad (\text{B.5})$$

Solving $\|\mathbf{x}(\theta)\|_2^2 = r^2$ for θ , we see that for any solution θ^* , we will also have $-\theta^*$ as a solution. Hence, $\mathbf{x}(\theta^*), \mathbf{x}(-\theta^*)$ will be solutions to (P2) with $\mu^* = -\mu_1$. If \mathbf{P} was p.s.d and singular, then $\mu^* = 0$, and $\mathbf{x}(\pm\theta^*) = \mathbf{x}(0) \pm \theta^* \mathbf{z}$ would be solutions to both (P1) and (P2). Note that $\mathbf{x}(-\mu_1)$ is a solution to (P1). In fact, any point in the interior of the form $\mathbf{x}(-\mu_1) + \theta \mathbf{z}$ is a solution to (P1).

C. Useful concentration inequalities

We present in this section some useful concentration inequalities, that will be employed as a tool to prove our other results. Recall that for a random variable X , its sub-Gaussian norm $\|X\|_{\psi_2}$ is defined as

$$\|X\|_{\psi_2} := \sup_{p \geq 1} \frac{(\mathbb{E}|X|^p)^{1/p}}{\sqrt{p}}. \quad (\text{C.1})$$

Moreover, X is a sub-Gaussian random variable if $\|X\|_{\psi_2}$ is finite. For instance, consider a bounded random variable X with $|X| \leq M$. Then, X is a sub-Gaussian random variable with $\|X\|_{\psi_2} \leq M$ (Vershynin, 2012, Example 5.8).

We begin with the well-known Hanson-Wright inequality (Hanson and Wright (1971)) for concentration of random quadratic forms. The following version is taken from Rudelson and Vershynin (2013).

Theorem 22 (Rudelson and Vershynin (2013)) *Let $(X_1, X_2, \dots, X_n) \in \mathbb{R}^n$ be a random vector with independent components X_i satisfying $\mathbb{E}[X_i] = 0$ and $\|X_i\|_{\psi_2} \leq K$. Let \mathbf{A} be a $n \times n$ matrix. Then for every $t \geq 0$*

$$\mathbb{P}(|X^T \mathbf{A} X - \mathbb{E}[X^T \mathbf{A} X]| \geq t) \leq 2 \exp\left(-c \min\left(\frac{t^2}{K^4 \|\mathbf{A}\|_F^2}, \frac{t}{K^2 \|\mathbf{A}\|}\right)\right) \quad (\text{C.2})$$

for an absolute constant $c > 0$.

Next, we recall the following Hoeffding type inequality for sums of independent sub-Gaussian random variables.

Proposition 23 (*Vershynin, 2012, Proposition 5.10*) *Let X_1, \dots, X_n be independent centered sub-Gaussian random variables and let $K = \max_i \|X_i\|_{\psi_2}$. Then for every $\mathbf{a} \in \mathbb{R}^n$, and every $t \geq 0$, we have*

$$\mathbb{P}\left(\left|\sum_{i=1}^n a_i X_i\right| \geq t\right) \leq e \cdot \exp\left(-\frac{c't^2}{K^2 \|\mathbf{a}\|_2^2}\right), \quad (\text{C.3})$$

where $c' > 0$ is an absolute constant.

D. Proof of Proposition 13

Let us first recall the definition of $\bar{\mathbf{z}}$ from (4.3). For clarity of notation, we will denote $\bar{\mathbf{z}}_R = \text{Re}(\mathbf{z}) \in \mathbb{R}^n$ and $\bar{\mathbf{z}}_I = \text{Im}(\mathbf{z}) \in \mathbb{R}^n$, and thus $\bar{\mathbf{z}} = [\bar{\mathbf{z}}_R^T \ \bar{\mathbf{z}}_I^T]^T \in \mathbb{R}^{2n}$. Clearly, $(z_i)_{i=1}^n$ are independent, complex-valued random variables. Note that

$$(\bar{\mathbf{z}}_R)_i = \cos(2\pi y_i) = \begin{cases} \cos(2\pi(f_i \bmod 1)) & ; \text{ if } \beta_i = 0 \\ \cos(2\pi u_i) & ; \text{ if } \beta_i = 1 \end{cases} \quad ; \quad i = 1, \dots, n. \quad (\text{D.1})$$

and

$$(\bar{\mathbf{z}}_I)_i = \sin(2\pi y_i) = \begin{cases} \sin(2\pi(f_i \bmod 1)) & ; \text{ if } \beta_i = 0 \\ \sin(2\pi u_i) & ; \text{ if } \beta_i = 1 \end{cases} \quad ; \quad i = 1, \dots, n. \quad (\text{D.2})$$

Since $(\beta_i)_{i=1}^n$ and $(u_i)_{i=1}^n$ are i.i.d random variables, hence the components of $\bar{\mathbf{z}}_R$ are independent real-valued random variables. The same is true for the components of $\bar{\mathbf{z}}_I$.

1. Lower bounding $\frac{1}{2^n} \bar{\mathbf{z}}^T H \bar{\mathbf{z}}$

To begin with, note that $\bar{\mathbf{z}}^T H \bar{\mathbf{z}} = \bar{\mathbf{z}}_R^T (\lambda L) \bar{\mathbf{z}}_R + \bar{\mathbf{z}}_I^T (\lambda L) \bar{\mathbf{z}}_I$. Denote $\bar{\boldsymbol{\mu}}_R = \mathbb{E}[\bar{\mathbf{z}}_R] \in \mathbb{R}^n$, and $\bar{\boldsymbol{\mu}}_I = \mathbb{E}[\bar{\mathbf{z}}_I] \in \mathbb{R}^n$. We see that

$$(\bar{\boldsymbol{\mu}}_R)_i = \mathbb{E}[(\bar{\mathbf{z}}_R)_i] \quad (\text{D.3})$$

$$= (1-p) \cos(2\pi(f_i \bmod 1)) + p \mathbb{E}[\cos(2\pi u_i)] \quad (\text{D.4})$$

$$= (1-p) \cos(2\pi(f_i \bmod 1)), \quad (\text{D.5})$$

since $\mathbb{E}[\cos(2\pi u_i)] = 0$ for $i = 1, \dots, n$. Similarly, we have that $(\bar{\boldsymbol{\mu}}_I)_i = (1-p) \sin(2\pi(f_i \bmod 1))$. Hence, we may write

$$\bar{\boldsymbol{\mu}}_R = (1-p) \text{Re}(\mathbf{h}) \quad \text{and} \quad \bar{\boldsymbol{\mu}}_I = (1-p) \text{Im}(\mathbf{h}). \quad (\text{D.6})$$

We now focus on lower bounding the term $\bar{\mathbf{z}}_R^T L \bar{\mathbf{z}}_R$. We observe that

$$\bar{\mathbf{z}}_R^T L \bar{\mathbf{z}}_R = (\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R + \bar{\boldsymbol{\mu}}_R)^T L (\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R + \bar{\boldsymbol{\mu}}_R) \quad (\text{D.7})$$

$$= (\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R)^T L (\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R) + 2(\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R)^T L \bar{\boldsymbol{\mu}}_R + \bar{\boldsymbol{\mu}}_R^T L \bar{\boldsymbol{\mu}}_R. \quad (\text{D.8})$$

The first two terms in (D.8) are random, and we proceed to lower bound them w.h.p starting with the first term. Let us note that

$$\begin{aligned} \mathbb{E}[(\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R)^T L(\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R)] &= \sum_{i=1}^n \mathbb{E}[(\bar{\mathbf{z}}_R)_i - (\bar{\boldsymbol{\mu}}_R)_i]^2 L_{ii} \\ &+ \sum_{i \neq j} \underbrace{\mathbb{E}[(\bar{\mathbf{z}}_R)_i - (\bar{\boldsymbol{\mu}}_R)_i][(\bar{\mathbf{z}}_R)_j - (\bar{\boldsymbol{\mu}}_R)_j]}_{=0} L_{ij} \end{aligned} \quad (\text{D.9})$$

$$= \sum_{i=1}^n (\mathbb{E}[(\bar{\mathbf{z}}_R)_i]^2 - (\bar{\boldsymbol{\mu}}_R)_i^2) \deg(i), \quad (\text{D.10})$$

since the cross terms in (D.9) are zero. We now obtain

$$\begin{aligned} \mathbb{E}[(\bar{\mathbf{z}}_R)_i]^2 &= \mathbb{E} \left[\frac{1 + \cos(4\pi y_i)}{2} \right] \\ &= \frac{1}{2} + \frac{1}{2} [(1-p) \cos(4\pi(f_i \bmod 1)) + p \mathbb{E}[\cos(4\pi u_i)]] \\ &= \frac{1}{2} + \frac{1}{2} (1-p) \cos(4\pi(f_i \bmod 1)). \end{aligned} \quad (\text{D.11})$$

Concerning the second term in (D.10), we note that

$$(\bar{\boldsymbol{\mu}}_R)_i^2 = (1-p)^2 \cos^2(2\pi(f_i \bmod 1)) = \frac{(1-p)^2}{2} + \frac{(1-p)^2}{2} \cos(4\pi(f_i \bmod 1)), \quad (\text{D.12})$$

for $i = 1, \dots, n$. Plugging (D.11), (D.12) in (D.10), and observing that $k \leq \deg(i) \leq 2k$, one can easily verify that

$$\frac{pnk}{2} \leq \mathbb{E}[(\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R)^T L(\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R)] \leq 3pnk. \quad (\text{D.13})$$

Next, for each $i = 1, \dots, n$, the random variables $(\bar{\mathbf{z}}_R)_i - (\bar{\boldsymbol{\mu}}_R)_i$ are zero mean, and also uniformly bounded since

$$|(\bar{\mathbf{z}}_R)_i - (\bar{\boldsymbol{\mu}}_R)_i| = |\cos(2\pi y_i) - (1-p) \cos(2\pi(f_i \bmod 1))| \leq 2. \quad (\text{D.14})$$

Hence $\|(\bar{\mathbf{z}}_R)_i - (\bar{\boldsymbol{\mu}}_R)_i\|_{\psi_2} \leq 2$ for each i . Therefore, applying Hanson-Wright inequality to $(\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R)^T L(\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R)$ yields

$$\begin{aligned} \mathbb{P}(|(\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R)^T L(\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R) - \mathbb{E}[(\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R)^T L(\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R)]| \geq t) \\ \leq 2 \exp \left(-c \min \left(\frac{t^2}{16 \|L\|_F^2}, \frac{t}{4 \|L\|} \right) \right). \end{aligned} \quad (\text{D.15})$$

Since $\deg(i) \leq 2k$ for each $i = 1, \dots, n$, therefore Gershgorins disk theorem yields the estimate $\|L\| \leq 4k$. Moreover, $\|L\|_F^2 \leq \sum_i ((2k)^2 + 2k) \leq 8k^2 n$. Plugging $t = \frac{\varepsilon pnk}{2}$ for $\varepsilon \in (0, 1)$, we observe that

$$\frac{t^2}{16 \|L\|_F^2} \geq \frac{(\varepsilon^2 p^2 n^2 k^2)/4}{128 k^2 n} = \frac{\varepsilon^2 p^2 n}{512}, \quad (\text{D.16})$$

$$\text{and } \frac{t}{4 \|L\|} \geq \frac{(\varepsilon pnk)/2}{16k} = \frac{\varepsilon pn}{32}. \quad (\text{D.17})$$

Thus, $\min \left\{ \frac{t^2}{16\|L\|_F^2}, \frac{t}{4\|L\|} \right\} \geq \frac{\varepsilon^2 p^2 n}{512}$. Plugging this estimate in (D.15) for the aforementioned choice of t , and using the bounds in (D.13), we have with probability at least $1 - 2 \exp\left(-\frac{c\varepsilon^2 p^2 n}{512}\right)$ that

$$\frac{pnk}{2}(1 - \varepsilon) \leq (\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R)^T L (\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R) \leq pnk \left(3 + \frac{\varepsilon}{2}\right). \quad (\text{D.18})$$

We now turn our attention to the second term in (D.8) namely $2(\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R)^T L \bar{\boldsymbol{\mu}}_R$. Recall that $(\bar{\mathbf{z}}_R)_i - (\bar{\boldsymbol{\mu}}_R)_i$ are independent, zero-mean sub-Gaussian random variables with $\|(\bar{\mathbf{z}}_R)_i - (\bar{\boldsymbol{\mu}}_R)_i\|_{\psi_2} \leq 2$, for each i . Hence, invoking Proposition 23 yields

$$\mathbb{P}(|\langle \bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R, 2L\bar{\boldsymbol{\mu}}_R \rangle| \geq t) \leq e \cdot \exp\left(-\frac{c't^2}{4\|2L\bar{\boldsymbol{\mu}}_R\|_2^2}\right) \quad (\text{D.19})$$

$$\leq e \cdot \exp\left(-\frac{c't^2}{16\|L\|^2\|\bar{\boldsymbol{\mu}}_R\|_2^2}\right) \quad (\text{D.20})$$

$$\leq e \cdot \exp\left(-\frac{c't^2}{256k^2n}\right), \quad (\text{D.21})$$

where in (D.21), we used the bounds $\|L\| \leq 4k$ (as shown earlier) and $\|\bar{\boldsymbol{\mu}}_R\|_2^2 \leq n$ (using (D.6)). Plugging $t = \frac{pnk(1-\varepsilon)}{3}$ in (D.21), we have that the following holds with probability at least $1 - e \cdot \exp\left(-\frac{c'p^2n(1-\varepsilon)^2}{2304}\right)$

$$|\langle \bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R, 2L\bar{\boldsymbol{\mu}}_R \rangle| \leq \frac{pnk(1-\varepsilon)}{3}. \quad (\text{D.22})$$

Combining (D.18), (D.22) and applying the union bound, we have with probability at least $1 - e \cdot \exp\left(-\frac{c'p^2n(1-\varepsilon)^2}{2304}\right) - 2 \exp\left(-\frac{c\varepsilon^2 p^2 n}{512}\right)$ that the following bound holds

$$\bar{\mathbf{z}}_R^T L \bar{\mathbf{z}}_R \geq \frac{pnk}{6}(1 - \varepsilon) + \bar{\boldsymbol{\mu}}_R^T L \bar{\boldsymbol{\mu}}_R. \quad (\text{D.23})$$

By proceeding as above, one can verify that with probability at least $1 - e \cdot \exp\left(-\frac{c'p^2n(1-\varepsilon)^2}{2304}\right) - 2 \exp\left(-\frac{c\varepsilon^2 p^2 n}{512}\right)$, the following bound holds

$$\bar{\mathbf{z}}_I^T L \bar{\mathbf{z}}_I \geq \frac{pnk}{6}(1 - \varepsilon) + \bar{\boldsymbol{\mu}}_I^T L \bar{\boldsymbol{\mu}}_I. \quad (\text{D.24})$$

Combining (D.23), (D.24) and applying the union bound, we have with probability at least $1 - 2e \cdot \exp\left(-\frac{c'p^2n(1-\varepsilon)^2}{2304}\right) - 4 \exp\left(-\frac{c\varepsilon^2 p^2 n}{512}\right)$ that

$$\frac{1}{2n} \bar{\mathbf{z}}^T H \bar{\mathbf{z}} = \frac{1}{2n} \lambda (\bar{\mathbf{z}}_R^T L \bar{\mathbf{z}}_R + \bar{\mathbf{z}}_I^T L \bar{\mathbf{z}}_I) \quad (\text{D.25})$$

$$\geq \frac{\lambda pnk}{6n} (1 - \varepsilon) + \frac{1}{2n} (\bar{\boldsymbol{\mu}}_R^T (\lambda L) \bar{\boldsymbol{\mu}}_R + \bar{\boldsymbol{\mu}}_I^T (\lambda L) \bar{\boldsymbol{\mu}}_I) \quad (\text{D.26})$$

$$= \frac{\lambda pk}{6} (1 - \varepsilon) + (1 - p)^2 \frac{1}{2n} \bar{\mathbf{h}} H \bar{\mathbf{h}} \quad (\text{D.27})$$

holds true. To go from (D.26) to (D.27), we used (D.6) along with the definition of H (see also Appendix A). This completes the derivation of the lower bound on $\frac{1}{2n} \bar{\mathbf{z}}^T H \bar{\mathbf{z}}$.

2. **Upper bounding $\|\bar{\mathbf{z}} - \bar{\mathbf{h}}\|_2$** By recalling the definition of $\bar{\mathbf{z}}, \bar{\mathbf{h}} \in \mathbb{R}^{2n}$ from (4.3), we note that

$$\|\bar{\mathbf{z}} - \bar{\mathbf{h}}\|_2^2 = \left\| \begin{pmatrix} \operatorname{Re}(\mathbf{z}) \\ \operatorname{Im}(\mathbf{z}) \end{pmatrix} - \begin{pmatrix} \operatorname{Re}(\mathbf{h}) \\ \operatorname{Im}(\mathbf{h}) \end{pmatrix} \right\|_2^2 \quad (\text{D.28})$$

$$= \|\mathbf{z} - \mathbf{h}\|_2^2 \quad (\text{D.29})$$

$$= 2n - (\mathbf{z}^* \mathbf{h} + \mathbf{h}^* \mathbf{z}) \quad (\text{Since } |z_i|, |h_i| = 1, \text{ for each } i) \quad (\text{D.30})$$

$$= 2n - \sum_{i=1}^n (\exp(\iota 2\pi(f_i \bmod 1 - y_i)) + \exp(-\iota 2\pi(f_i \bmod 1 - y_i))) \quad (\text{D.31})$$

$$= 2n - \sum_{i=1}^n \underbrace{(2 \cos(2\pi(f_i \bmod 1 - y_i)))}_{M_i}. \quad (\text{D.32})$$

Clearly, $(M_i)_{i=1}^n$ are independent sub-Gaussian random variables with $|M_i| \leq 2$; hence $\|M_i\|_{\psi_2} \leq 2$ for each i . Moreover,

$$\mathbb{E}[M_i] = 2(1-p) + p\mathbb{E}[(2 \cos(2\pi(f_i \bmod 1 - u_i)))] \quad (\text{D.33})$$

$$= 2(1-p) + p \int_0^1 2 \cos(2\pi(f_i \bmod 1 - u_i)) du_i \quad (\text{D.34})$$

$$= 2(1-p) + p \left[\frac{2 \sin(2\pi u_i - 2\pi(f_i \bmod 1))}{2\pi} \right]_0^1 \quad (\text{D.35})$$

$$= 2(1-p). \quad (\text{D.36})$$

Therefore, $(M_i - \mathbb{E}[M_i])_{i=1}^n$ are centered, independent sub-Gaussian random variables with $\|M_i - \mathbb{E}[M_i]\|_{\psi_2} \leq 4$. Thus, by applying Proposition 23 to $\sum_{i=1}^n (M_i - \mathbb{E}[M_i])$, we obtain

$$\mathbb{P}(|\sum_i M_i - 2n(1-p)| \geq t) \leq e \cdot \exp\left(-\frac{c't^2}{16n}\right). \quad (\text{D.37})$$

Plugging $t = 2(1-p)n\varepsilon$ in (D.37) for $\varepsilon \in (0, 1)$, we have with probability at least $1 - e \cdot \exp\left(-\frac{c'(1-p)^2\varepsilon^2}{4n}\right)$ that the following bound holds

$$2(1-p)n(1-\varepsilon) \leq \sum_{i=1}^n M_i \leq 2(1-p)n(1+\varepsilon). \quad (\text{D.38})$$

Conditioning on the event in (D.38), we finally obtain from (D.32) the bound

$$\|\bar{\mathbf{z}} - \bar{\mathbf{h}}\|_2^2 \leq 2n - 2(1-p)n(1-\varepsilon) \leq 2n(p+\varepsilon). \quad (\text{D.39})$$

E. Proof of Proposition 15

We consider the noise model $y_i = (f_i + \eta_i) \bmod 1$, where $\eta_i \sim \mathcal{N}(0, \sigma^2)$ are i.i.d. Recall the definition of $\bar{\mathbf{z}}$ from (4.3), where upon denoting $\bar{\mathbf{z}}_R = \operatorname{Re}(\mathbf{z}) \in \mathbb{R}^n$, $\bar{\mathbf{z}}_I = \operatorname{Im}(\mathbf{z}) \in \mathbb{R}^n$, we have $\bar{\mathbf{z}} = [\bar{\mathbf{z}}_R^T \ \bar{\mathbf{z}}_I^T]^T \in \mathbb{R}^{2n}$. Clearly $(\bar{\mathbf{z}}_R)_i = \cos(2\pi((f_i + \eta_i) \bmod 1))$, $i = 1, \dots, n$ are independent random variables, the same being true for $(\bar{\mathbf{z}}_I)_i = \sin(2\pi((f_i + \eta_i) \bmod 1))$.

1. **Lower bounding $\frac{1}{2n}\bar{\mathbf{z}}^T H \bar{\mathbf{z}}$.** To begin with, note that $\bar{\mathbf{z}}^T H \bar{\mathbf{z}} = \bar{\mathbf{z}}_R^T(\lambda L)\bar{\mathbf{z}}_R + \bar{\mathbf{z}}_I^T(\lambda L)\bar{\mathbf{z}}_I$. Denote $\bar{\boldsymbol{\mu}}_R = \mathbb{E}[\bar{\mathbf{z}}_R] \in \mathbb{R}^n$, and $\bar{\boldsymbol{\mu}}_I = \mathbb{E}[\bar{\mathbf{z}}_I] \in \mathbb{R}^n$. We then have

$$(\bar{\boldsymbol{\mu}}_R)_i = \mathbb{E}[\cos(2\pi((f_i + \eta_i) \bmod 1))] \quad (\text{E.1})$$

$$= \mathbb{E}\left[\frac{e^{\iota 2\pi((f_i + \eta_i) \bmod 1)} + e^{-\iota 2\pi((f_i + \eta_i) \bmod 1)}}{2}\right] \quad (\text{E.2})$$

$$= \frac{e^{\iota 2\pi(f_i \bmod 1)}\mathbb{E}[e^{\iota 2\pi\eta_i}] + e^{-\iota 2\pi(f_i \bmod 1)}\mathbb{E}[e^{-\iota 2\pi\eta_i}]}{2}. \quad (\text{E.3})$$

Using the fact¹² $\mathbb{E}[e^{\iota 2\pi\eta_i}], \mathbb{E}[e^{-\iota 2\pi\eta_i}] = e^{-2\pi^2\sigma^2}$ in (E.3), we get $(\bar{\boldsymbol{\mu}}_R)_i = e^{-2\pi^2\sigma^2} \cos(2\pi(f_i \bmod 1))$. In an analogous manner, one can show that $(\bar{\boldsymbol{\mu}}_I)_i = e^{-2\pi^2\sigma^2} \sin(2\pi(f_i \bmod 1))$. To summarize, we have shown that

$$\bar{\boldsymbol{\mu}}_R = e^{-2\pi^2\sigma^2} \text{Re}(\mathbf{h}), \quad \bar{\boldsymbol{\mu}}_I = e^{-2\pi^2\sigma^2} \text{Im}(\mathbf{h}). \quad (\text{E.4})$$

Now recall that $\bar{\mathbf{z}}^T H \bar{\mathbf{z}} = \bar{\mathbf{z}}_R^T(\lambda L)\bar{\mathbf{z}}_R + \bar{\mathbf{z}}_I^T(\lambda L)\bar{\mathbf{z}}_I$. We will lower bound $\bar{\mathbf{z}}_R^T(\lambda L)\bar{\mathbf{z}}_R$, respectively $\bar{\mathbf{z}}_I^T(\lambda L)\bar{\mathbf{z}}_I$ individually, w.h.p., via similar arguments as for the Bernoulli noise, via a Hanson-Wright inequality, respectively, a Hoeffding type inequality arguments.

Consider $\bar{\mathbf{z}}_R^T L \bar{\mathbf{z}}_R$, and recall its decomposition as in (D.8). We will lower bound $(\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R)^T L (\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R)$ and $2(\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R)^T L \bar{\boldsymbol{\mu}}_R$ w.h.p.

To begin with, recall from (D.10) that

$$\mathbb{E}[(\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R)^T L (\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R)] = \sum_{i=1}^n (\mathbb{E}[(\bar{\mathbf{z}}_R)_i]^2 - (\bar{\boldsymbol{\mu}}_R)_i^2) \deg(i). \quad (\text{E.5})$$

Next, using the following observation

$$\mathbb{E}[(\bar{\mathbf{z}}_R)_i]^2 = \mathbb{E}[\cos^2(2\pi((f_i + \eta_i) \bmod 1))] = \mathbb{E}[\cos^2(2\pi(f_i + \eta_i))] = \mathbb{E}\left[\frac{1 + \cos(4\pi(f_i + \eta_i))}{2}\right] \quad (\text{E.6})$$

together with the fact that $\mathbb{E}[\cos(4\pi(f_i + \eta_i))] = e^{-8\pi^2\sigma^2} \cos(4\pi(f_i \bmod 1))$, we remark that, for each $i = 1, \dots, n$, it holds true that

$$\mathbb{E}[(\bar{\mathbf{z}}_R)_i]^2 = \frac{1 + e^{-8\pi^2\sigma^2} \cos(4\pi(f_i \bmod 1))}{2}. \quad (\text{E.7})$$

Note that the last equality follows along similar steps as those in the derivation of $(\bar{\boldsymbol{\mu}}_R)_i$ earlier. A similar calculation yields that

$$(\bar{\boldsymbol{\mu}}_R)_i^2 = e^{-4\pi^2\sigma^2} \cos^2(2\pi(f_i \bmod 1)) = e^{-4\pi^2\sigma^2} \left[\frac{1 + \cos(4\pi(f_i \bmod 1))}{2}\right]. \quad (\text{E.8})$$

12. Indeed, for $X \sim \mathcal{N}(\mu, \sigma^2)$, we have $\mathbb{E}[e^{\iota t X}] = e^{\iota t \mu - \frac{\sigma^2 t^2}{2}}$ for any $t \in \mathbb{R}$.

Combining (E.7) with (E.8) yields

$$\mathbb{E}[(\bar{\mathbf{z}}_R)_i]^2 - (\bar{\boldsymbol{\mu}}_R)_i^2 = \frac{(1 - e^{-4\pi^2\sigma^2}) + (e^{-8\pi^2\sigma^2} - e^{-4\pi^2\sigma^2}) \cos(4\pi(f_i \bmod 1))}{2} \quad (\text{E.9})$$

$$= \frac{(1 - e^{-4\pi^2\sigma^2})(1 - e^{-4\pi^2\sigma^2} \cos(4\pi(f_i \bmod 1)))}{2}. \quad (\text{E.10})$$

Since $k \leq \deg(i) \leq 2k$, we readily conclude that (E.5) amounts to

$$\frac{(1 - e^{-4\pi^2\sigma^2})^2}{2} kn \leq \mathbb{E}[(\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R)^T L (\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R)] \leq (1 - e^{-8\pi^2\sigma^2}) kn. \quad (\text{E.11})$$

Note that for each $i = 1, \dots, n$, the random variables $(\bar{\mathbf{z}}_R)_i - (\bar{\boldsymbol{\mu}}_R)_i$ are zero-mean, and are also uniformly bounded as

$$|(\bar{\mathbf{z}}_R)_i - (\bar{\boldsymbol{\mu}}_R)_i| = |\cos(2\pi((f_i + \eta_i) \bmod 1)) - e^{-2\pi^2\sigma^2} \cos(2\pi(f_i \bmod 1))| \leq 2. \quad (\text{E.12})$$

Hence $\|(\bar{\mathbf{z}}_R)_i - (\bar{\boldsymbol{\mu}}_R)_i\|_{\psi_2} \leq 2$ for each i , therefore allowing us to apply the Hanson-Wright inequality to $(\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R)^T L (\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R)$ yields

$$\begin{aligned} \mathbb{P}(|(\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R)^T L (\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R) - \mathbb{E}[(\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R)^T L (\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R)]| \geq t) \\ \leq 2 \exp\left(-c \min\left(\frac{t^2}{16 \|L\|_F^2}, \frac{t}{4 \|L\|}\right)\right). \end{aligned} \quad (\text{E.13})$$

We saw earlier that $\|L\| \leq 4k$ (see the proof of Lemma 9). Since $L_{ii} \leq 2k$, and $L_{i,j} \leq 1$, for $i, j \in E$, it holds true that $\|L\|_F^2 \leq 8k^2n$. Plugging in $t = \varepsilon kn \frac{(1 - e^{-4\pi^2\sigma^2})^2}{2}$ for $\varepsilon \in (0, 1)$ we observe that

$$\frac{t^2}{16 \|L\|_F^2} \geq \frac{(\varepsilon^2(1 - e^{-4\pi^2\sigma^2})^4 n^2 k^2)/4}{128k^2n} = \frac{\varepsilon^2(1 - e^{-4\pi^2\sigma^2})^4 n}{1024}, \quad (\text{E.14})$$

$$\text{and } \frac{t}{4 \|L\|} \geq \frac{(\varepsilon(1 - e^{-4\pi^2\sigma^2})^2 nk)/2}{16k} = \frac{\varepsilon(1 - e^{-4\pi^2\sigma^2})^2 n}{32}. \quad (\text{E.15})$$

Thus, $\min\left\{\frac{t^2}{16\|L\|_F^2}, \frac{t}{4\|L\|}\right\} \geq \frac{\varepsilon^2(1 - e^{-4\pi^2\sigma^2})^4 n}{1024}$. Plugging this estimate in (E.13) for the aforementioned choice of t , and using the bounds in (E.11), we have with probability at least $1 - 2 \exp\left(-\frac{c\varepsilon^2(1 - e^{-4\pi^2\sigma^2})^4 n}{1024}\right)$ that

$$(1 - \varepsilon) \frac{(1 - e^{-4\pi^2\sigma^2})^2}{2} kn \leq (\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R)^T L (\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R) \leq kn \left[(1 - e^{-8\pi^2\sigma^2}) + \frac{\varepsilon}{2} (1 - e^{-4\pi^2\sigma^2})^2 \right]. \quad (\text{E.16})$$

Next, we turn our attention to lower-bounding the second term, namely $2(\bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R)^T L \bar{\boldsymbol{\mu}}_R$. Recall that $(\bar{\mathbf{z}}_R)_i - (\bar{\boldsymbol{\mu}}_R)_i$ are independent, zero-mean sub-Gaussian random variables with $\|(\bar{\mathbf{z}}_R)_i - (\bar{\boldsymbol{\mu}}_R)_i\|_{\psi_2} \leq 2$, for each i . Hence, invoking Proposition 23 along with the facts $\|L\| \leq 4k$ and $\|\bar{\boldsymbol{\mu}}_R\|_2^2 \leq e^{-4\pi^2\sigma^2} n$, we obtain

$$\mathbb{P}(|\langle \bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R, 2L\bar{\boldsymbol{\mu}}_R \rangle| \geq t) \leq e \cdot \exp\left(-\frac{ct^2}{256k^2 \|\bar{\boldsymbol{\mu}}_R\|_2^2}\right) \leq e \cdot \exp\left(-\frac{ct^2 e^{4\pi^2\sigma^2}}{256k^2 n}\right). \quad (\text{E.17})$$

Plugging in $t = \frac{(1-\varepsilon)(1-e^{-4\pi^2\sigma^2})^2}{3}kn$, we have with probability at least

$$1 - e \exp\left(\frac{-c'(1-\varepsilon)^2(1-e^{-4\pi^2\sigma^2})^4 e^{4\pi^2\sigma^2}}{2304}n\right)$$

that

$$|\langle \bar{\mathbf{z}}_R - \bar{\boldsymbol{\mu}}_R, 2L\bar{\boldsymbol{\mu}}_R \rangle| \leq \frac{(1-\varepsilon)(1-e^{-4\pi^2\sigma^2})^2}{3}kn. \quad (\text{E.18})$$

Recalling the decomposition of $\bar{\mathbf{z}}_R^T L \bar{\mathbf{z}}_R$ as in (D.8), combining the lower bound from (E.16) with the lower bound implied by (E.18), and applying the union bound, we have that, with probability at least

$$1 - e \exp\left(\frac{-c'(1-\varepsilon)^2(1-e^{-4\pi^2\sigma^2})^4 e^{4\pi^2\sigma^2}}{2304}n\right) - 2 \exp\left(-\frac{c\varepsilon^2(1-e^{-4\pi^2\sigma^2})^4 n}{1024}\right) \quad (\text{E.19})$$

the following bound holds

$$\bar{\mathbf{z}}_R^T L \bar{\mathbf{z}}_R \geq \frac{(1-\varepsilon)(1-e^{-4\pi^2\sigma^2})^2}{6}kn + \bar{\boldsymbol{\mu}}_R^T L \bar{\boldsymbol{\mu}}_R. \quad (\text{E.20})$$

By proceeding as above, we obtain the same lower bound on $\bar{\mathbf{z}}_I^T L \bar{\mathbf{z}}_I$ with the same lower bound on the success probability. Hence, by applying the union bound to (E.20) and its imaginary counterpart, we finally have with probability at least

$$1 - 2e \exp\left(\frac{-c'(1-\varepsilon)^2(1-e^{-4\pi^2\sigma^2})^4 e^{4\pi^2\sigma^2}}{2304}n\right) - 4 \exp\left(-\frac{c\varepsilon^2(1-e^{-4\pi^2\sigma^2})^4 n}{1024}\right) \quad (\text{E.21})$$

that the following bound holds

$$\frac{1}{2n} \bar{\mathbf{z}}^T H \bar{\mathbf{z}} = \frac{1}{2n} \lambda (\bar{\mathbf{z}}_R^T L \bar{\mathbf{z}}_R + \bar{\mathbf{z}}_I^T L \bar{\mathbf{z}}_I) \quad (\text{E.22})$$

$$\geq \frac{\lambda}{2n} \left[\frac{(1-\varepsilon)(1-e^{-4\pi^2\sigma^2})^2}{3}kn \right] + \frac{1}{2n} (\bar{\boldsymbol{\mu}}_R^T (\lambda L) \bar{\boldsymbol{\mu}}_R + \bar{\boldsymbol{\mu}}_I^T (\lambda L) \bar{\boldsymbol{\mu}}_I) \quad (\text{E.23})$$

$$= \frac{\lambda k}{6} (1-\varepsilon)(1-e^{-4\pi^2\sigma^2})^2 + \frac{e^{-4\pi^2\sigma^2}}{2n} \mathbf{h}^T H \mathbf{h}. \quad (\text{E.24})$$

In the last step, we used (E.4) along with the definition of H (see also Appendix A). This completes the derivation of the lower bound on $\frac{1}{2n} \bar{\mathbf{z}}^T H \bar{\mathbf{z}}$.

2. **Upper bounding $\|\bar{\mathbf{z}} - \bar{\mathbf{h}}\|_2$.** By recalling the definition of $\bar{\mathbf{z}}, \bar{\mathbf{h}} \in \mathbb{R}^{2n}$ from (4.3), we observe that

$$\|\bar{\mathbf{z}} - \bar{\mathbf{h}}\|_2^2 = \|\mathbf{z} - \mathbf{h}\|_2^2 \quad (\text{E.25})$$

$$= 2n - (\mathbf{z}^* \mathbf{h} + \mathbf{h}^* \mathbf{z}) \quad (\text{E.26})$$

$$= 2n - \sum_{i=1}^n [e^{-\iota 2\pi((f_i + \eta_i) \bmod 1 - f_i \bmod 1)} + e^{\iota 2\pi((f_i + \eta_i) \bmod 1 - f_i \bmod 1)}] \quad (\text{E.27})$$

$$= 2n - \sum_{i=1}^n [e^{\iota 2\pi(\eta_i \bmod 1)} + e^{-\iota 2\pi(\eta_i \bmod 1)}] = 2n - \sum_{i=1}^n \underbrace{[e^{\iota 2\pi\eta_i} + e^{-\iota 2\pi\eta_i}]}_{M_i}. \quad (\text{E.28})$$

Clearly $(M_i)_{i=1}^n$ are i.i.d real-valued sub-Gaussian random variables with $|M_i| \leq 2$; hence $\|M_i\|_{\psi_2} \leq 2$ for each i . Moreover, we readily obtain $\mathbb{E}[M_i] = 2e^{-2\pi^2\sigma^2}$, and consequently, $(M_i - \mathbb{E}[M_i])_{i=1}^n$ are centered, i.i.d sub-Gaussian random variables with $\|M_i - \mathbb{E}[M_i]\|_{\psi_2} \leq 4$. Therefore, by applying Proposition 23 to $\sum_{i=1}^n (M_i - \mathbb{E}[M_i])$, we obtain

$$\mathbb{P}\left(\left|\sum_i M_i - 2ne^{-2\pi^2\sigma^2}\right| \geq t\right) \leq e \cdot \exp\left(-\frac{c't^2}{16n}\right). \quad (\text{E.29})$$

Plugging in $t = 2n\epsilon e^{-2\pi^2\sigma^2}$ for $\epsilon \in (0, 1)$, we have with probability at least $1 - e \cdot \exp(-\frac{c'e^{-4\pi^2\sigma^2}n\epsilon^2}{4})$ that the following bound holds

$$2n\epsilon e^{-2\pi^2\sigma^2}(1 - \epsilon) \leq \sum_{i=1}^n M_i \leq 2n\epsilon e^{-2\pi^2\sigma^2}(1 + \epsilon). \quad (\text{E.30})$$

Conditioning on the event in (E.30), we finally obtain from (E.28) the bound

$$\|\bar{\mathbf{z}} - \bar{\mathbf{h}}\|_2^2 \leq 2n - 2n(1 - \epsilon)e^{-2\pi^2\sigma^2} = 2n(1 - (1 - \epsilon)e^{-2\pi^2\sigma^2}). \quad (\text{E.31})$$

F. Additional numerical experiments

F1. Numerical experiments: Bounded Model

Figure 23 shows several denoising instances as we increase the noise level in the Uniform noise model ($\gamma \in \{0.15, 0.27, 0.30\}$). Note that **OLS** starts failing at $\gamma = 0.27$, while **QCQP** still estimates the samples of f well. Interestingly, **iQCQP** performs quite well, even for $\gamma = 0.30$ (where **QCQP** starts failing) and produces highly smooth, and accurate estimates. It would be interesting to investigate the properties of **iQCQP** in future work.

Figures 24, 25 plot RMSE (on a log scale) for denoised $f \bmod 1$ and f samples versus the noise level, for the Uniform noise model. They illustrate the importance of the choice of the regularization parameters λ, k . If λ is too small (eg., $\lambda = 0.03$), then **QCQP** has negligible improvement in performance, and sometimes also has worse RMSE than the raw noisy samples. However, for a larger λ ($\lambda \in \{0.3, 0.5\}$), **QCQP** has a strictly smaller error than **OLS** and the raw noisy samples. Interestingly, **iQCQP** typically performs very well, even for $\lambda = 0.03$.

Figure 26 plots the RMSE (on a log scale) for both the denoised $f \bmod 1$ samples, and samples of f , versus n (for Uniform noise model). Observe that for large enough n , **QCQP** shows strictly smaller RMSE than both the initial input noisy data, and **OLS**. Furthermore, we remark that **iQCQP** typically has superior performance to **QCQP** except for small values of n .

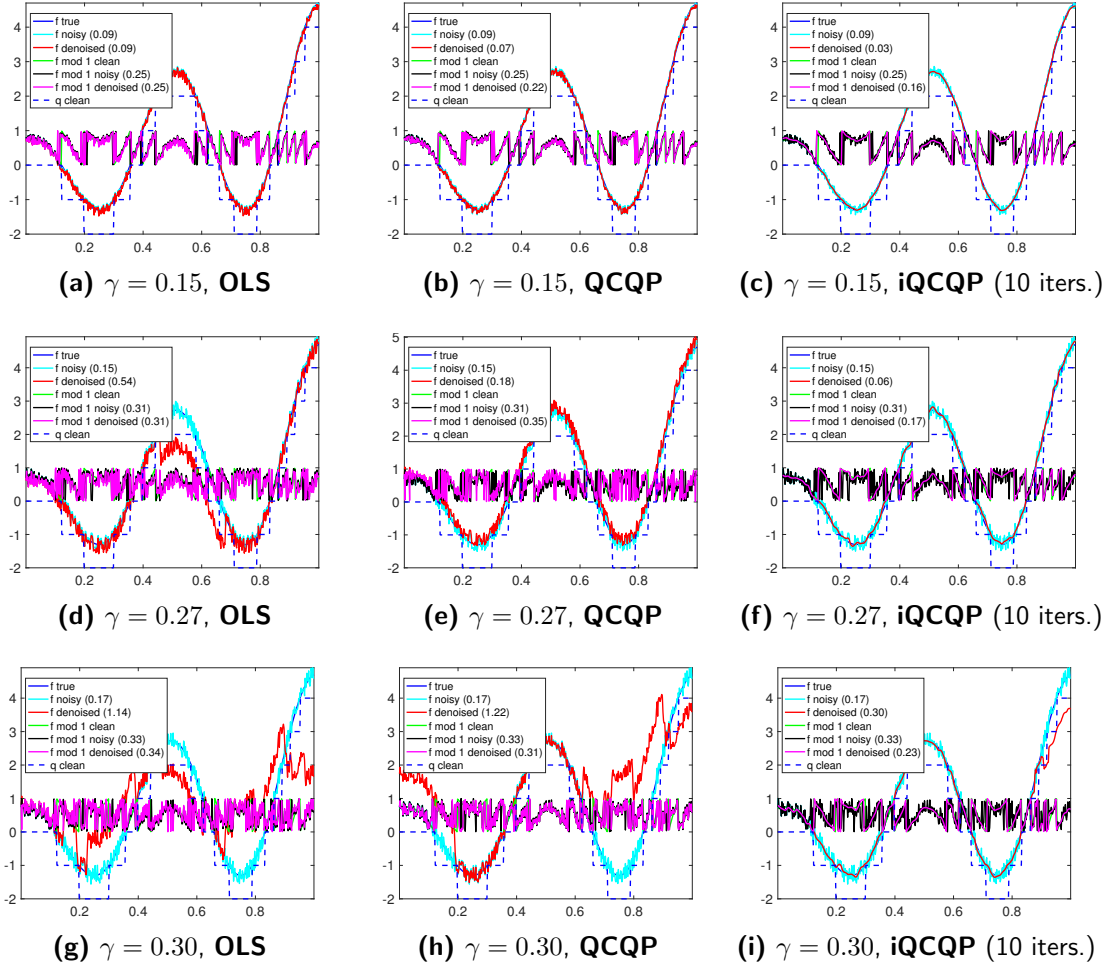


Figure 23: Denoised instances under the Uniform noise model, for **OLS**, **QCQP** and **iQCQP**, as we increase the noise level γ . We keep fixed the parameters $n = 500$, $k = 2$, $\lambda = 0.1$. The numerical values in the legend denote the RMSE. **QCQP** denotes Algorithm 1, for which the unwrapping stage is performed via **OLS** (4.23).

F2. Comparison with Bhandari et al. (2017)

This section is a comparison of **OLS**, **QCQP**, and **iQCQP** with the approach introduced by Bhandari et al. (2017), whose algorithm we denote by **BKR** for brevity. We compare all four approaches across two different noise models, Bounded and Gaussian, on two different

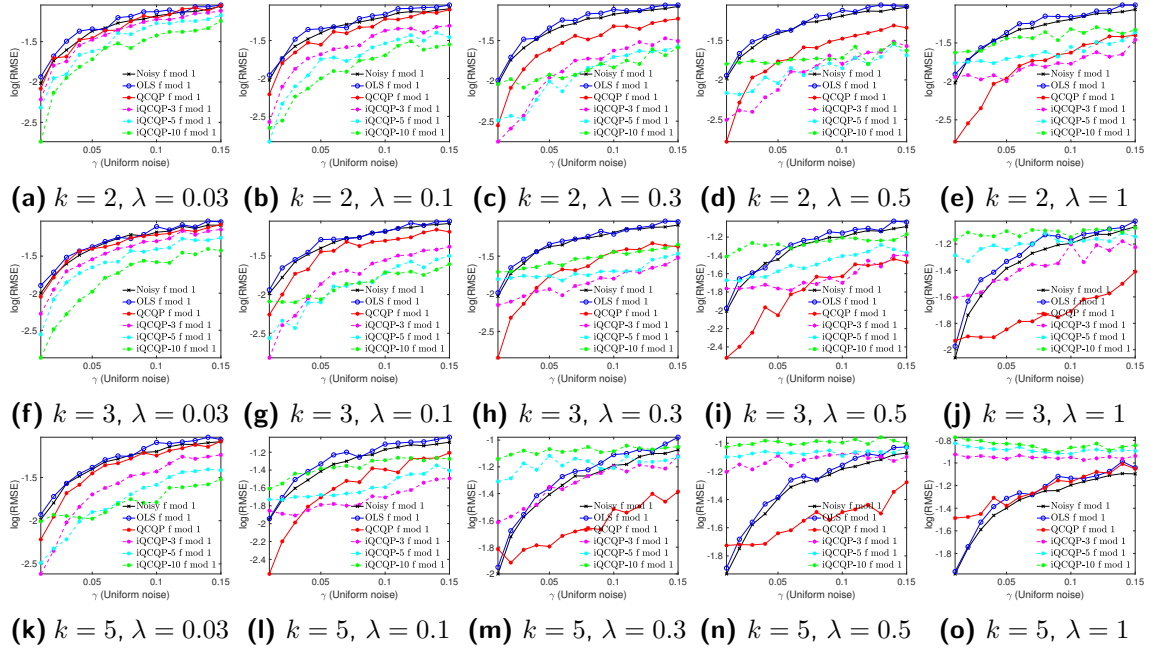


Figure 24: Recovery errors for the denoised $f \bmod 1$ samples, for the Bounded noise model (20 trials). **QCQP** denotes Algorithm 1 without the unwrapping stage performed by **OLS** (4.23).

functions: the function (8.1) used thus far in the experiments throughout this paper, and a bandlimited function used by Bhandari et al. (2017).

Figure 27 pertains to function (8.1) but for the Gaussian noise model. At a lower level of noise $\sigma = 5\%$, all methods are able to recover the function fairly well, with **iQCQP** and **QCQP** yielding the most accurate reconstructions in terms of the RMSE of both the denoised $f \bmod 1$ samples, and the final estimates for f . At higher levels of noise $\sigma \in \{6\%, 13\%\}$, **BKR** breaks down, while the other methods (in order of accuracy) **iQCQP**, **QCQP**, **OLS** return meaningful estimates.

Figures 28, respectively 29, compare the results of the four methods on the same bandlimited function considered by the authors of Bhandari et al. (2017), under the Bounded, respectively Gaussian, noise models with $n = 494$ sample points. The bandlimited function is generated by multiplying the Fourier spectrum of the sinc-function with weights drawn from the standard distribution, rescaled such that the largest magnitude entry is equal to 1. To make the problem more challenging, we also scale the function by a factor of 3, such that the function wraps itself more often when taking modulo 1 values. Finally, to make the resulting figures more visually appealing, we shift the function values by $+3$ in order to minimize the overlap with the modulo 1 values (clean, noisy and denoised samples). For the Bounded noise model, Figure 28 illustrates the fact that at a lower noise level $\gamma = 15\%$ all methods recover a good approximation of the function, the winners being **iQCQP** and **QCQP** both in terms of denoising the $f \bmod 1$ samples and estimating f . At $\gamma = 17\%$ **BKR** breaks down, while the remaining methods return good approximations. Finally, at

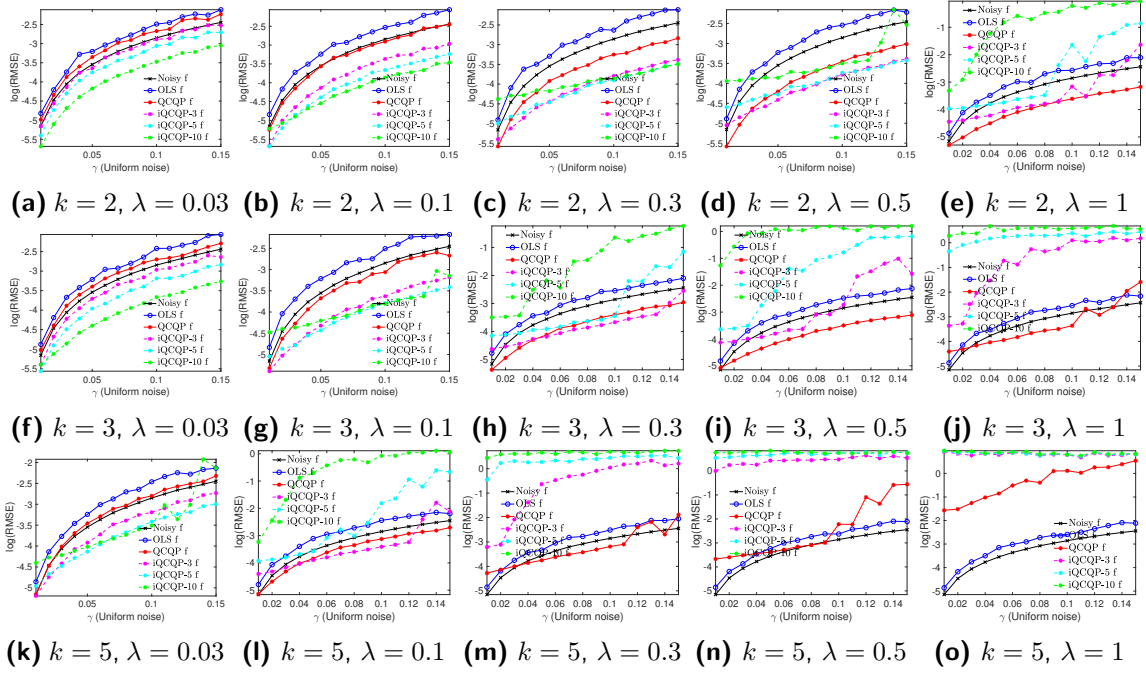


Figure 25: Recovery errors for the final estimated f samples, for the Bounded noise model (20 trials). **QCQP** denotes Algorithm 1, for which the unwrapping stage is performed via **OLS** (4.23).

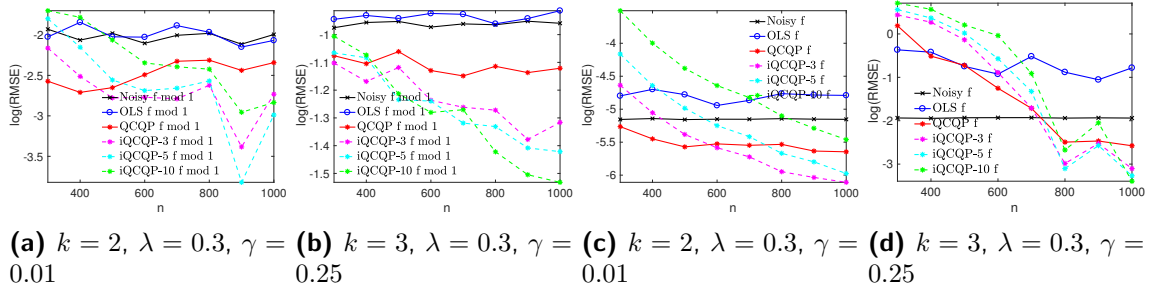


Figure 26: Recovery errors for **OLS**, **QCQP**, and **iQCQP** as a function of n (number of samples), for both the $f \bmod 1$ samples (leftmost two plots) and the final f estimates (rightmost two plots) under the Uniform noise model, for different values of k , λ and γ . Results are averaged over 20 runs. **QCQP** denotes Algorithm 1, for which the unwrapping stage is performed via **OLS** (4.23).

$\gamma = 20\%$ the latter three methods have difficulties in the unwrapping stage. Figure 29 shows analogous results and conclusions under the more challenging Gaussian noise model, with $\sigma \in \{5\%, 8\%, 12\%\}$.

Finally, Figure 30, shows analogous results as in Figures 28 under the Bounded noise model, but with a sparser sampling pattern $n = 194$, similar to the setup used by the

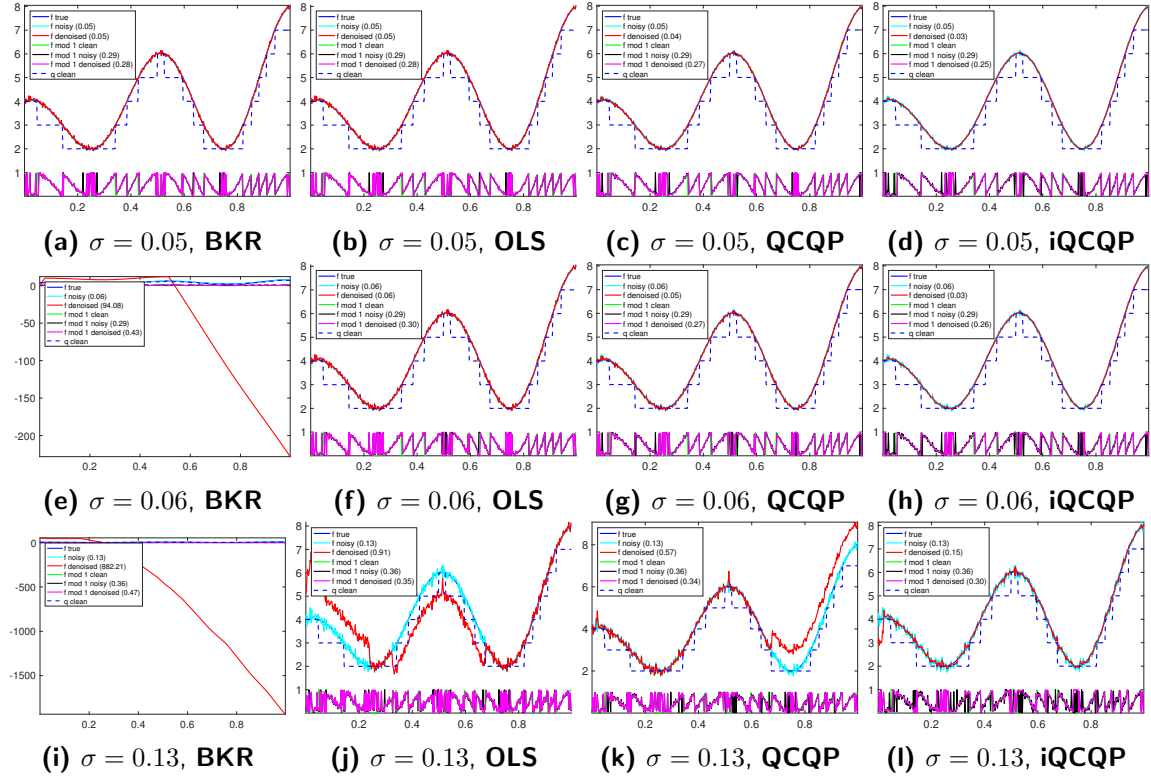


Figure 27: Denoised instances for both the $f \bmod 1$ and f values, under the Gaussian noise model for function (8.1), for **BKR**, **OLS**, **QCQP** and **iQCQP**, as we increase the noise level σ . **QCQP** denotes Algorithm 1, for which the unwrapping stage is performed via **OLS** (4.23). We keep fixed the parameters $n = 500$, $k = 2$, $\lambda = 0.1$. The numerical values in the legend denote the RMSE.

authors of Bhandari et al. (2017). At a lower noise level, $\gamma = 10\%$, all methods recover a good approximation of the function, the performance ranking being **iQCQP** (RMSE=0.23), **QCQP** (RMSE=0.23), **OLS** (RMSE=0.31), **BKR** (RMSE=0.31) for the denoised $f \bmod 1$ samples, and **iQCQP** (RMSE=0.06), **QCQP** (RMSE=0.11), **OLS** (RMSE=0.11), **BKR** (RMSE=0.11) for the final f estimates. At higher levels of noise, **iQCQP**, **QCQP**, **OLS** return meaningful estimates, though all methods experience difficulties in the unwrapping stage, in the interval where the function f has the highest slope.

F3. Additional elevation maps experiments for two-dimensional phase unwrapping

Figure 31 pertains to the elevation map of Vesuvius, using $n = 3600$ samples, under the Gaussian noise model ($\sigma = 0.05$). Figure 32, pertains to noisy measurements of Mount Etna under the Gaussian noise model ($\sigma = 0.10$), using $n = 19000$ samples.

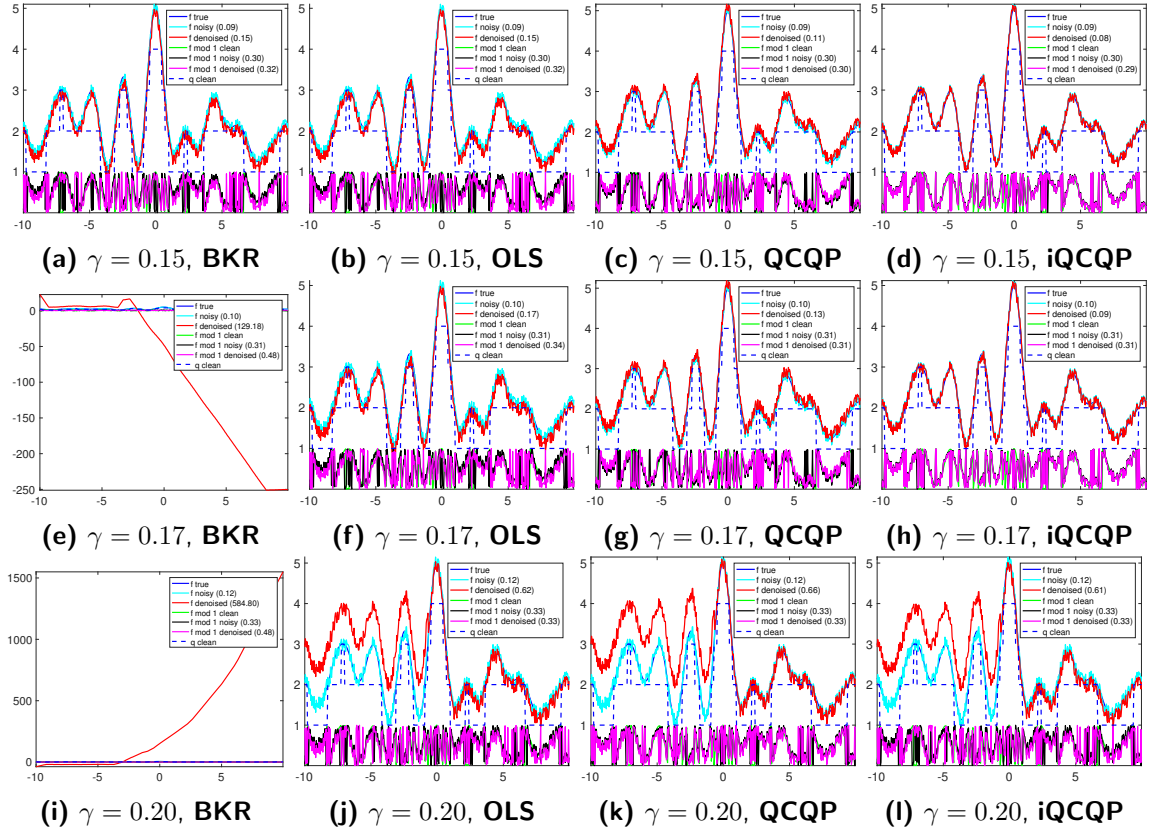


Figure 28: Denoised instances for the bandlimited function considered in Bhandari et al. (2017), under the Bounded-Uniform noise model, for **BKR**, **OLS**, **QCQP** and **iQCQP**, as we increase the noise level γ . We keep fixed the parameters $n = 484$, $k = 2$, $\lambda = 0.01$. The numerical values in the legend denote the RMSE. **QCQP** denotes Algorithm 1, for which the unwrapping stage is performed via **OLS** (4.23).

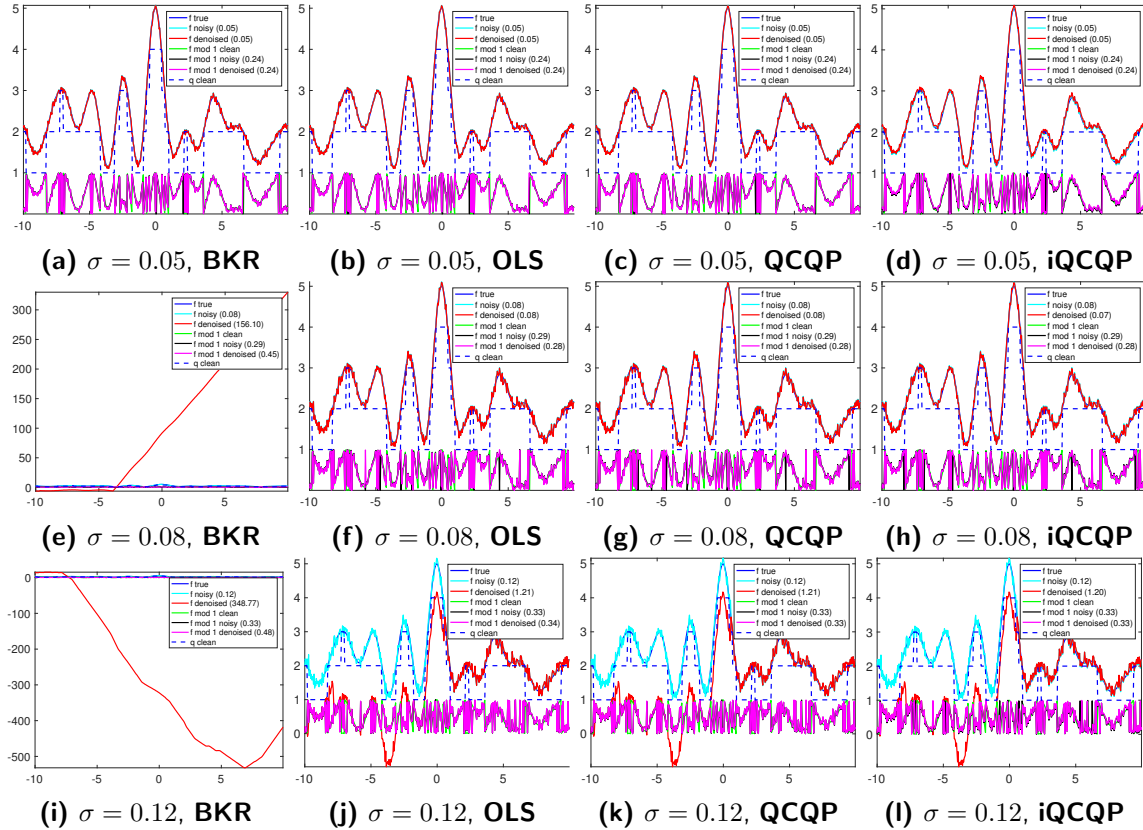


Figure 29: Denoised instances for the bandlimited function considered in Bhandari et al. (2017), under the Gaussian noise model, for **BKR**, **OLS**, **QCQP** and **iQCQP**, as we increase the noise level σ . **QCQP** denotes Algorithm 1, for which the unwrapping stage is performed via **OLS** (4.23). We keep fixed the parameters $n = 484$, $k = 2$, $\lambda = 0.01$. The numerical values in the legend denote the RMSE.

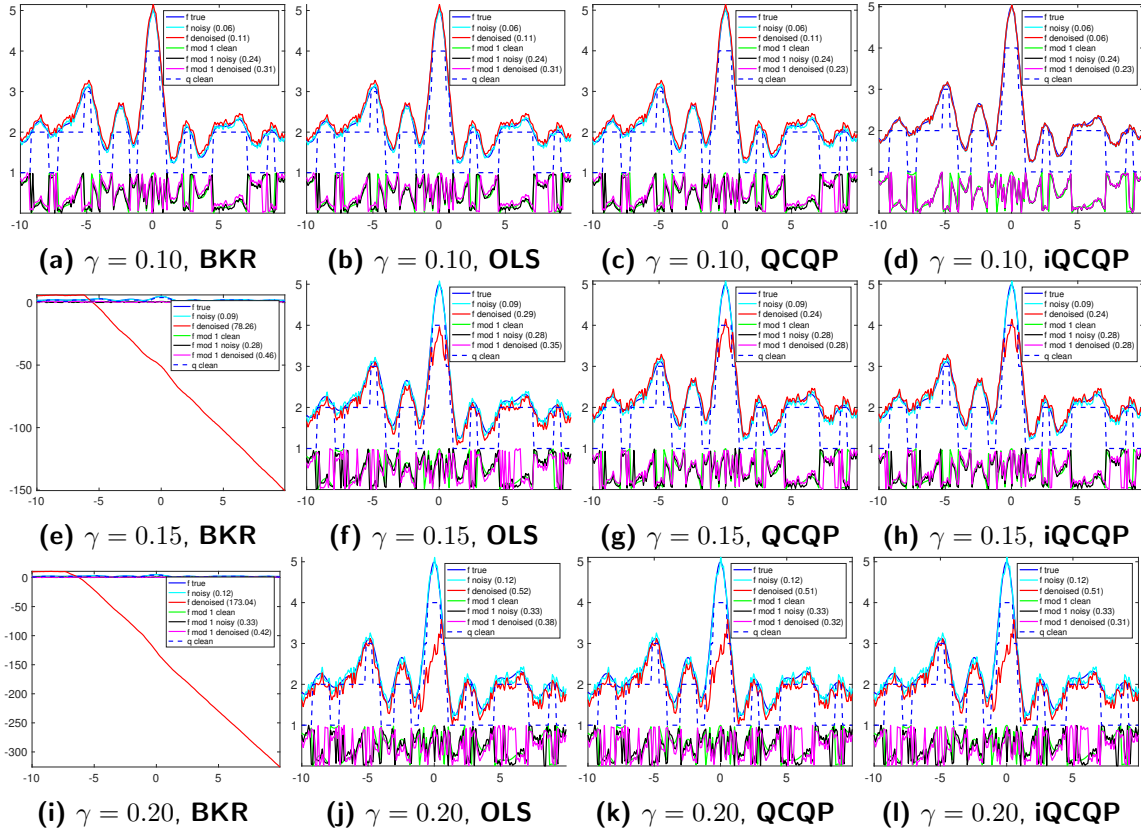


Figure 30: Denoised instances for the bandlimited function considered in Bhandari et al. (2017), under the Bounded-Uniform noise model, for **BKR**, **OLS**, **QCQP** and **iQCQP**, as we increase the noise level γ . **QCQP** denotes Algorithm 1, for which the unwrapping stage is performed via **OLS** (4.23). The function is sampled less frequently compared to the earlier Figure 28, which in this experiment leads to only $n = 194$ sample points. We keep fixed $k = 2$ and $\lambda = 0.01$. The numerical values in the legend denote the RMSE.

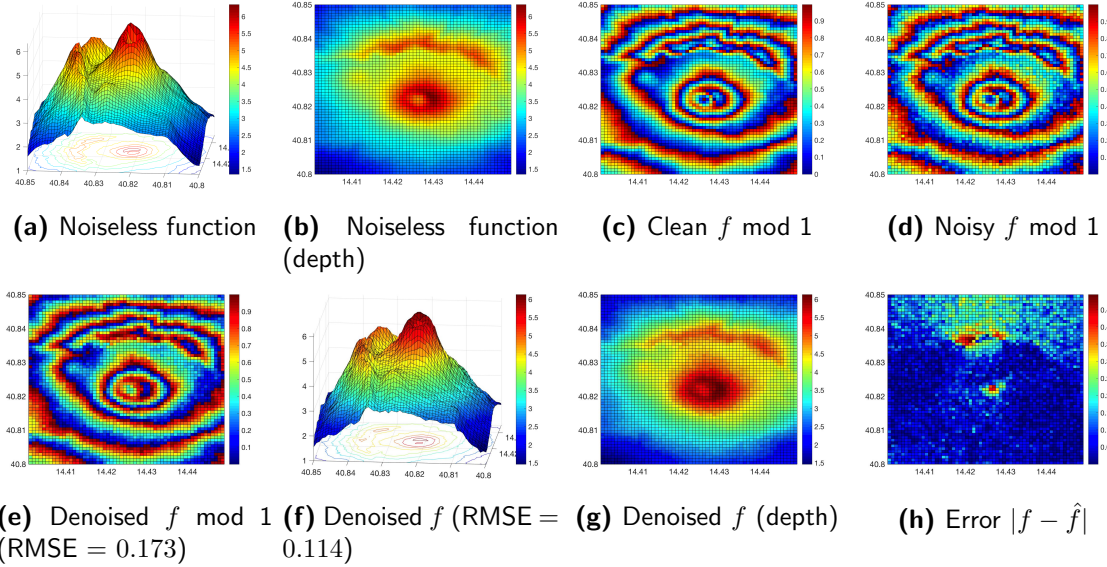


Figure 31: Elevation map of Mount Vesuvius with $n = 3600$, $k = 1$ (Chebychev distance), $\lambda = 0.1$, and noise level $\sigma = 0.05$ under the Gaussian model, as recovered by Manopt-Phases.

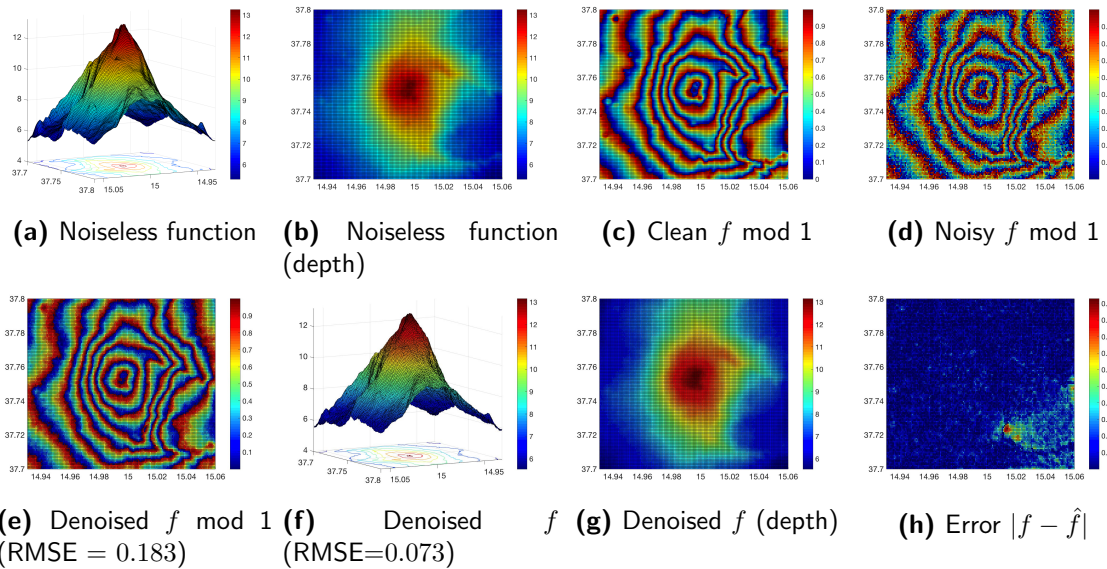


Figure 32: Elevation map of Mount Etna with $n = 19000$, $k = 1$ (Chebychev distance), $\lambda = 0.3$, and noise level $\sigma = 0.10$ under the Gaussian model, as recovered by Manopt-Phases.