

On the Global Geometry of Sphere-Constrained Sparse Blind Deconvolution

Yuqian Zhang¹, Yenson Lau², Han-Wen Kuo²,
Sky Cheung³, Abhay Pasupathy³, John Wright^{2,4}

¹Department of Computer Science, Cornell University

²Department of Electrical Engineering and Data Science Institute, Columbia University

³Department of Physics, Columbia University

⁴Department of Applied Physics and Applied Mathematics, Columbia University

January 7, 2019

Abstract

Blind deconvolution is the problem of recovering a convolutional kernel \mathbf{a}_0 and an activation signal \mathbf{x}_0 from their convolution $\mathbf{y} = \mathbf{a}_0 \circledast \mathbf{x}_0$. This problem is ill-posed without further constraints or priors. This paper studies the situation where the nonzero entries in the activation signal are sparsely and randomly populated. We normalize the convolution kernel to have unit Frobenius norm and cast the sparse blind deconvolution problem as a nonconvex optimization problem over the sphere. With this spherical constraint, every spurious local minimum turns out to be close to some signed shift truncation of the ground truth, under certain hypotheses. This benign property motivates an effective two stage algorithm that recovers the ground truth from the partial information offered by a suboptimal local minimum. This geometry-inspired algorithm recovers the ground truth for certain microscopy problems, also exhibits promising performance in the more challenging image deblurring problem. Our insights into the global geometry and the two stage algorithm extend to the convolutional dictionary learning problem, where a superposition of multiple convolution signals is observed.

1 Introduction

Blind deconvolution aims to recover two unknown signals: a kernel \mathbf{a}_0 and some underlying signal \mathbf{x}_0 from their convolution $\mathbf{y} = \mathbf{a}_0 \circledast \mathbf{x}_0$. Blind deconvolution is ill-posed in general: there are infinitely many pairs of signals rendering the same convolution. To render the problem well-posed, one may exploit prior knowledge about the structure of \mathbf{a}_0 and \mathbf{x}_0 . For example, the underlying signal \mathbf{x}_0 is *sparse* in many engineering and scientific applications:

Microscopy data analysis: In the crystal lattice of nanoscale materials, there exist randomly and sparsely distributed “defects”, whose locations and signatures encode crucial information about the electronic structure of the material. Accurate recovery of such information can facilitate investigation of the detailed structure of materials [CLC⁺17].

Neural spike sorting: Neurons communicate by firing brief voltage spikes, whose characteristics reflect important features of the neuron. These spikes occur randomly and sparsely in time. Neurophysiologists are interested in assigning stereotyped spikes to putative cells, as well as in knowing their respective spike times [ETS11, Lew98].

Image deblurring: Motion blur can be modeled as the convolution of a latent sharp image and a kernel capturing the motion of the camera, usually assumed to be invariant across the image [FSH⁺06]. The inverse process of recovering the original sharp image from a blurry image has been widely studied [CL09,

[KH96]. Many well-performing approaches leverage the observation that sharp natural images typically have (approximately) sparse gradients [CW98, LWDF11, PF14].

All of these applications lead to instances of the sparse blind deconvolution problem. The dominant algorithmic approach to sparse blind deconvolution involves *nonconvex optimization*¹. Nonconvex formulations for deconvolution can be derived via several probabilistic formalisms (ML/MAP, VB, ect.), or simply from heuristics. For example, in image deblurring, the kernel \mathbf{a} can be modeled as residing on a simplex [LWDF11, GTZ⁺16, KTF11, LCM14]. This is natural from a modeling perspective², but problematic for optimization: natural formulations of deconvolution over the simplex admit trivial global minimizers (corresponding to spiky convolutional kernels $\mathbf{a} = \delta$) [PF14, BVG13], which provide no information about the ground truth. Practical remedies for this problem include exploiting additional data priors [GTZ⁺16, LCM14, XWHH16] or careful initialization via edge restoration or multi-scale refinement [KTF11, XJ10], to avoid the trivial spiky global minima.

In contrast, motivated by a careful comparison of MAP and VB approaches, [WZ13, ZWZ13] propose to instead constrain \mathbf{a} to have unit Frobenius norm – i.e., to reside on a high-dimensional sphere.³ This choice is arguably more appropriate for certain scientific applications – such as microscopy – in which the kernel \mathbf{a} can have negative entries. For image deblurring, \mathbf{a} can be assumed to be nonnegative, and the sphere constraint seems less natural from a modeling perspective.

In this paper, we study the geometry of sphere-constrained sparse blind deconvolution. Our goal is to understand when simple algorithms based on nonconvex optimization can exactly recover the convolutional kernel \mathbf{a} and the sparse signal \mathbf{x} . This goal is motivated by the applications described above – in particular, microscopy data analysis – in which there is a strong, physical sparsity prior and a clear, physical notion of the ground truth. We develop our theory and algorithms under the assumption that \mathbf{a} is a short kernel, and that \mathbf{x} is sparsely and randomly supported. We demonstrate through a theoretical analysis of certain (idealized) cases and many numerical experiments that when these assumptions are satisfied, the proposed algorithm correctly recovers \mathbf{a} , and hence \mathbf{x} . These results stem from a striking geometric property of sphere-constrained sparse blind deconvolution: although the problem is still nonconvex, every local minimizer $\bar{\mathbf{a}}$ is very close to a signed shift-truncation of the ground truth kernel \mathbf{a}_0 . This observation provides a geometric explanation of how the sphere constraint can facilitate sparse blind deconvolution.

The remainder of this paper is organized as follows. Section 2 discusses the intrinsic symmetries associated with the convolutional operator and their implication on the geometry of sphere-constrained sparse blind deconvolution. Section 3 introduces the optimization-based two stage algorithm and some related technical details. Section 4 discusses two other important extensions in image deblurring and convolutional dictionary learning. Section 5 gives experimental corroboration of our theory, and shows promising results on microscopy data analysis, image deblurring, and convolutional dictionary learning. Section 6 discusses directions for future work.

For simplicity, we assume that the convolutional signals are one dimensional in both our problem formulation and technical proofs; all of our results extend naturally to higher-dimensional signals. Throughout this paper, vectors $\mathbf{v} \in \mathbb{R}^k$ are indexed as $\mathbf{v} = [v_0, v_1, \dots, v_{k-1}]$, and $[\cdot]_m$ denotes the modulo- m operation. We use $\|\cdot\|$ to denote the operator norm, and $\|\cdot\|_p$ to denote the entry wise ℓ^p norm. A projection onto the Frobenius sphere is denoted with $\mathcal{P}_{\mathbb{S}}[\cdot] = \frac{\cdot}{\|\cdot\|_2}$, and a projection onto subset I is denoted with $(\cdot)_I$.

¹In signal processing, a number of elegant *convex* relaxations of the problem have been developed [ARR12, Chi16, LLSW18]. However, these approaches typically require stronger prior information (subspace constraint rather than sparsity) or exhibit suboptimal scalings.

²Since entries of \mathbf{a} roughly represent the fraction of the camera exposure time at a given location.

³[WZ13] contains a wealth of additional ideas about the role of sparsity-promoting priors in obtaining good local minima, and on the probabilistic underpinnings of deconvolution. Our experiments support the viewpoint that the key insight in [WZ13] is the role of the spherical constraint in avoiding bad minimizers.

2 Symmetry and Global Geometry

Without loss of generality, we assume the observation data \mathbf{y} is generated via a circular convolution \otimes of the ground truth $\mathbf{a}_0 \in \mathbb{R}^k$ and $\mathbf{x}_0 \in \mathbb{R}^m$:

$$\mathbf{y}(\mathbf{a}_0, \mathbf{x}_0) = \mathbf{a}_0 \otimes \mathbf{x}_0 = \widetilde{\mathbf{a}}_0 \otimes \mathbf{x}_0 \in \mathbb{R}^m. \quad (2.1)$$

Here, $\widetilde{\mathbf{a}}_0 \in \mathbb{R}^m$ denote the zero padded m -length version of \mathbf{a}_0 , which can be expressed as $\widetilde{\mathbf{a}}_0 = \iota \mathbf{a}_0$ with $\iota : \mathbb{R}^k \rightarrow \mathbb{R}^m$ be a zero padding operator. Its adjoint $\iota^* : \mathbb{R}^m \rightarrow \mathbb{R}^k$ acts as a projection to lower dimensional space by keeping the first k components. Equivalently, we can write

$$\mathbf{y}(\mathbf{a}_0, \mathbf{x}_0) = \mathbf{C}_{\widetilde{\mathbf{a}}_0} \mathbf{x}_0 = \mathbf{C}_{\mathbf{x}_0} \widetilde{\mathbf{a}}_0. \quad (2.2)$$

Here, $\mathbf{C}_v \in \mathbb{R}^{m \times m}$ is the circulant matrix generated from vector \mathbf{v} , whose j -th column is a cyclic shift $s_{j-1}[\mathbf{v}]$ of the vector \mathbf{v} :

$$s_\tau[\mathbf{v}](i) = \mathbf{v}([i - \tau]_m), \quad \forall i \in [0, \dots, m - 1]. \quad (2.3)$$

2.1 Symmetries and Symmetry Breaking

The SBD problem exhibits a *scaled-shift symmetry*, which derives from the symmetries of the convolution operator. Namely, given a pair $(\mathbf{a}_0, \mathbf{x}_0)$ satisfying $\mathbf{y} = \mathbf{a}_0 \otimes \mathbf{x}_0$, for any nonzero scalar α and integer τ

$$\mathbf{y} = (\alpha s_\tau[\widetilde{\mathbf{a}}_0]) \otimes (\alpha^{-1} s_{-\tau}[\mathbf{x}_0]). \quad (2.4)$$

Note that a scaled shift $\alpha^{-1} s_{-\tau}[\mathbf{x}_0]$ of a sparse signal \mathbf{x}_0 remains sparse, and that a scaled shift $\alpha s_\tau[\widetilde{\mathbf{a}}_0]$ of a length- k kernel \mathbf{a}_0 still has k -nonzero entries. So, these symmetries are intrinsic to the SBD problem, as formulated here. We can only hope to recover $(\mathbf{a}_0, \mathbf{x}_0)$ up to this symmetry.

The presence of nontrivial symmetries is a hallmark of bilinear problems arising in practice – see, e.g., [SQW17a, SQW17b] for examples from dictionary learning and generalized phase retrieval. Symmetries render straightforward approaches to convexify the problem ineffective.⁴ They also raise challenges for nonconvex optimization: equivalent symmetric solutions correspond to multiple disconnected global optima. This creates a very complicated objective landscape, which could potentially also contain spurious local optimizers. Certain highly symmetric nonconvex problems arising in signal processing do not exhibit spurious minimizers [SQW17a, SQW17b], however, proving this can be challenging.

2.1.1 Symmetry breaking.

We employ a weak symmetry breaking mechanism by constraining $\mathbf{a} \in \mathbb{S}^{k-1}$ ⁵: we reduce the scale ambiguity to a sign ambiguity, by constraining \mathbf{a} to have unit Frobenius norm; we mitigate the shift ambiguity by constraining \mathbf{a} to be supported on the first k entries.

In general, $s_\tau[\widetilde{\mathbf{a}}_0]$ are not be supported on the first k entries, hence constraining \mathbf{a} to be supported on the first k entries removes the shift symmetry. However, effects of such shift symmetry still persist. Since the restriction $\iota^* s_\tau[\widetilde{\mathbf{a}}_0]$ to the first k entries can be convolved with the sparse signal $s_{-\tau}[\mathbf{x}_0]$ to produce a near approximation to \mathbf{y} :

$$(\iota^* s_\tau[\widetilde{\mathbf{a}}_0]) \otimes s_{-\tau}[\mathbf{x}_0] \approx \mathbf{y}, \quad (2.5)$$

especially when the shift $|\tau|$ is small. We will see that (i) these symmetric solutions $\iota^* s_\tau[\widetilde{\mathbf{a}}_0]$ persist as local minima of a natural optimization formulation of the SBD problem, but that (ii) under conditions, these are the *only* local minima.

⁴Given any set of points where the convex objective function achieves equal values, the function value will be no larger at any convex combination of them.

⁵This is motivated in part by [SQW17a], which demonstrates that a certain formulation of the dictionary learning problem over the sphere has no spurious local minimizers, even for relatively dense target representations. The “simplex constrained” analogue of that work, which optimizes over hyperplanes, requires the target solution to be much sparser [SWW12].

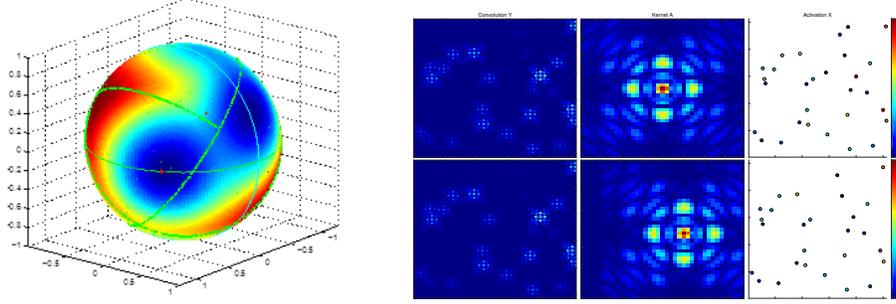


Figure 1: Geometry on the ℓ^2 ball for fixed \mathbf{a}_0 and generic \mathbf{x}_0 . Left: the objective $\varphi(\mathbf{a})$ in a low dimensional setting $\mathbf{a} \in \mathbb{S}^2$ – dark blue represents small values while dark red represents large values. All local minima are close to signed shift truncations of the ground truth \mathbf{a}_0 , with \mathbf{a}_0 itself achieving global minimum. The green lines indicate regions where \mathbf{a} are ill-posed as convolutional kernels. Right: a shift truncation \mathbf{a} achieves a local minimum of $\varphi(\mathbf{a})$ in a high dimensional setting. Shown here is the ground truth $\mathbf{y} = \mathbf{a}_0 \circledast \mathbf{x}_0$, \mathbf{a}_0 , and \mathbf{x}_0 (top right) versus their respective recovered quantities $\mathbf{a} \circledast \mathbf{x}$, \mathbf{a} , and \mathbf{x} (bottom right).

2.2 Global Geometry on the Sphere

We study the following objective function, which can be viewed as balancing sparsity of \mathbf{x} with fidelity to the observation \mathbf{y} :⁶

$$\min_{\mathbf{a} \in \mathbb{S}^{k-1}, \mathbf{x}} \psi(\mathbf{a}, \mathbf{x}) \doteq \frac{1}{2} \|\mathbf{y} - \mathbf{a} \circledast \mathbf{x}\|_2^2 + \lambda r(\mathbf{x}). \quad (2.6)$$

When \mathbf{x}_0 is long and random, it is more convenient to study this function through its “marginalization”

$$\varphi(\mathbf{a}) \doteq \min_{\mathbf{x}} \psi(\mathbf{a}, \mathbf{x}), \quad (2.7)$$

which is defined over the sphere \mathbb{S}^{k-1} .

In Figure 1, we plot the function value of $\varphi(\mathbf{a})$ on the sphere $\mathbf{a} \in \mathbb{S}^2$: red and blue imply larger and smaller objective value respectively and there are several local minima. For this highly nonconvex function, the ground truth \mathbf{a}_0 achieves the global minimum, while other local minima $\bar{\mathbf{a}}$ are very close to certain signed shift truncations of the ground truth. Figure 1 (right) exhibits an example of a local minimum in a higher-dimensional problem.

2.2.1 Analysis under Restricted Settings

Demonstrating that this observation holds in general is challenging: for most reasonable choices of the regularizer r , there is no closed form expression for the objective $\varphi(\mathbf{a})$. We develop an analysis under several simplifying assumptions. Throughout, we let $r(\mathbf{x})$ be the ℓ^1 norm, although similar conclusions hold for other sparsifying regularizers. With this choice, we can simplify the objective φ by dividing the sphere \mathbb{S}^{k-1} via the sign-support pattern of the minimizing $\mathbf{x}^*(\mathbf{a})$:

$$\mathbf{x}^*(\mathbf{a}) = \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{y} - \mathbf{a} \circledast \mathbf{x}\|_2^2 + \lambda \|\mathbf{x}\|_1. \quad (2.8)$$

Let I and σ denote the support and sign of \mathbf{x}^*

$$I = \text{supp}(\mathbf{x}^*), \quad \sigma = \text{sign}(\mathbf{x}^*), \quad (2.9)$$

then the whole sphere can be divided via the sign support pattern

$$\mathbb{S}^{k-1} = \bigcup_{\sigma} \mathcal{R}_{\sigma}, \quad \mathcal{R}_{\sigma} = \{\mathbf{a} \mid \text{sign}(\mathbf{x}^*(\mathbf{a})) = \sigma\}. \quad (2.10)$$

⁶Similar formulation can be found in lot of sparse representation problems [MBP14].

On each \mathcal{R}_σ where the sign support pattern σ remains the same, the stationarity condition for minimizer $\mathbf{x}^*(\mathbf{a})$ implies

$$\mathbf{x}_I^*(\mathbf{a}) = (\mathbf{C}_a^* \mathbf{C}_a)_I^{-1} (\mathbf{C}_a^* \mathbf{y} - \lambda \sigma)_I. \quad (2.11)$$

Plugging above expression back to the original objective function $\varphi(\mathbf{a})$ yields

$$\varphi_\sigma(\mathbf{a}) = -\frac{1}{2} (\mathbf{C}_a^* \mathbf{y} - \lambda \sigma)_I^* (\mathbf{C}_a^* \mathbf{C}_a)_I^{-1} (\mathbf{C}_a^* \mathbf{y} - \lambda \sigma)_I + \frac{1}{2} \|\mathbf{y}\|_2^2. \quad (2.12)$$

Although the objective function $\varphi(\mathbf{a})$ can be substantially simplified by removing the \mathbf{x} variable in this way, it still maintains a complicated dependence on \mathbf{a} . To obtain some preliminary insights, we make two simplifications for easier calculation while preserving important characteristics of the geometry of φ :

Simplification I: $\mathbf{x}_0 = \delta$. We maximally simplify the underlying sparse signal as a single spike δ and the observation will be $\mathbf{y} = \mathbf{a}_0 \otimes \mathbf{x}_0 = \widetilde{\mathbf{a}}_0$. This case itself is trivial, but its function geometry is a basic but important case to be understood. This simple case also yields intuitions that carry over to the less trivial situation in which \mathbf{x}_0 is a long random vector.

Simplification II: $\mathbf{C}_a^* \mathbf{C}_a \rightarrow \mathbf{I}$. For a random $\mathbf{a} \in \mathbb{S}^{k-1}$, its expectation satisfies $\mathbb{E}[\mathbf{C}_a^* \mathbf{C}_a] = \mathbf{I}$. Here, we simply use the identity matrix to replace any $\mathbf{C}_a^* \mathbf{C}_a$ and therefore reduce the complexity of Equation 2.12.

With these two simplifications, the original objective problem can be replaced with the following:

$$\text{minimize } \widehat{\varphi}(\mathbf{a}) \quad \text{subject to } \mathbf{a} \in \mathbb{S}^{k-1}, \quad (2.13)$$

where

$$\widehat{\varphi}(\mathbf{a}) \doteq \min_{\mathbf{x}} \frac{1}{2} \|\widetilde{\mathbf{a}}_0\|_2^2 + \frac{1}{2} \|\mathbf{x}\|_2^2 - \langle \widetilde{\mathbf{a}} \otimes \mathbf{x}, \widetilde{\mathbf{a}}_0 \rangle + \lambda \|\mathbf{x}\|_1.$$

In this case, the minimizing $\mathbf{x}^*(\mathbf{a})$ has a simple closed form solution:

$$\mathbf{x}^*(\mathbf{a}) = \text{SOFT}_\lambda [\mathbf{C}_a^* \widetilde{\mathbf{a}}_0] = \text{SOFT}_\lambda [\check{\mathbf{C}}_{\mathbf{a}_0}^* \boldsymbol{\iota} \mathbf{a}], \quad (2.14)$$

here, $\text{SOFT}_\lambda [u] = \text{sign}(u) \max\{|u| - \lambda, 0\}$ is the entry-wise soft-thresholding operator and $\check{\mathbf{C}}_{\mathbf{a}_0} \in \mathbb{R}^{m \times m}$ is the reversed circulant matrix for \mathbf{a}_0 defined via

$$\check{\mathbf{C}}_{\mathbf{a}_0} = [s_0 [\widetilde{\mathbf{a}}_0] \mid s_{-1} [\widetilde{\mathbf{a}}_0] \mid \dots \mid s_{-(m-1)} [\widetilde{\mathbf{a}}_0]]. \quad (2.15)$$

On a constant sign support pattern σ , $\widehat{\varphi}(\mathbf{a})$ can be written into a simpler quadratic form:

$$\widehat{\varphi}_\sigma(\mathbf{a}) = -\frac{1}{2} \|(\mathbf{C}_a^* \widetilde{\mathbf{a}}_0 - \lambda \sigma)_I\|_2^2 + \frac{1}{2} \|\widetilde{\mathbf{a}}_0\|_2^2. \quad (2.16)$$

For this surrogate $\widehat{\varphi}(\mathbf{a})$, we can show that if λ is sufficiently large compared to the magnitude of \mathbf{x}_0 , every strict local minimizer is a signed shift truncation of the ground truth:

Theorem 2.1. Define the set of possible supports of minimizer \mathbf{x}^* with

$$\mathcal{I} = \{\text{supp}(\text{SOFT}_\lambda [\check{\mathbf{C}}_{\mathbf{a}_0}^* \boldsymbol{\iota} \mathbf{a}]) \mid \mathbf{a} \in \mathbb{S}^{k-1}\}. \quad (2.17)$$

For each nonempty support $I = \{i_1 < i_2 < \dots < i_{|I|}\}$, let

$$\mathbf{W}_I = \left[\frac{\boldsymbol{\iota}^* s_{-i_1} [\widetilde{\mathbf{a}}_0]}{\|\boldsymbol{\iota}^* s_{-i_1} [\widetilde{\mathbf{a}}_0]\|_F} \mid \frac{\boldsymbol{\iota}^* s_{-i_2} [\widetilde{\mathbf{a}}_0]}{\|\boldsymbol{\iota}^* s_{-i_2} [\widetilde{\mathbf{a}}_0]\|_F} \mid \dots \mid \frac{\boldsymbol{\iota}^* s_{-i_{|I|}} [\widetilde{\mathbf{a}}_0]}{\|\boldsymbol{\iota}^* s_{-i_{|I|}} [\widetilde{\mathbf{a}}_0]\|_F} \right]$$

Suppose that $\lambda < 1$ and that for every nonempty $I \in \mathcal{I}$,

$$\|\mathbf{W}_I^* \mathbf{W}_I - \mathbf{I}\| < \frac{\lambda^2}{6}, \quad (2.18)$$

then every local minimum $\bar{\mathbf{a}}$ of $\widehat{\varphi}$ over \mathbb{S}^{k-1} satisfies either $\bar{\mathbf{a}} \in \mathcal{R}_0$ (in which case $\bar{\mathbf{a}}$ is also a global maximum), or

$$\bar{\mathbf{a}} = \pm \frac{\boldsymbol{\iota}^* s_\tau [\widetilde{\mathbf{a}}_0]}{\|\boldsymbol{\iota}^* s_\tau [\widetilde{\mathbf{a}}_0]\|_F} = \pm \mathcal{P}_\mathbb{S} [\boldsymbol{\iota}^* s_{-\tau} [\widetilde{\mathbf{a}}_0]] \quad (2.19)$$

with $\mathbf{x}^*(\bar{\mathbf{a}}) = \pm \text{SOFT}_\lambda [\|\boldsymbol{\iota}^* s_\tau [\widetilde{\mathbf{a}}_0]\|_F] s_{-\tau} [\mathbf{x}_0]$ for some shift τ .

Proof Please refer to the supplement. ■

This theorem says that the only local minima in this idealized case are signed shift truncations of the ground truth \mathbf{a}_0 , with certain choice of λ . Moreover, on those local minima $\bar{\mathbf{a}}$, the minimizing sparse $\mathbf{x}^*(\bar{\mathbf{a}})$ correspond to the soft thresholded, oppositely shifted ground truth \mathbf{x}_0 . The quantity $\|\mathbf{W}_I^* \mathbf{W}_I - \mathbf{I}\|$ measures the orthogonality of different shifts of \mathbf{a}_0 . In particular, if \mathbf{a}_0 is benign enough in the sense that any two different shifts of \mathbf{a}_0 are uncorrelated, or $\|\mathbf{W}_I^* \mathbf{W}_I - \mathbf{I}\| \rightarrow 0$ for any I , then any nonzero λ guarantees the desired geometry. On the other hand, for a fixed \mathbf{a}_0 , both $|\mathcal{I}|$ and $|I|$ becomes smaller as λ increases, therefore the constraint $\|\mathbf{W}_I^* \mathbf{W}_I - \mathbf{I}\| < \lambda^2/6$ is more likely to be satisfied.

A similar result holds when \mathbf{x}_0 is separated enough that copies of the kernel do not overlap. It also holds if \mathbf{x}_0 is a long, sufficiently sparse random vector. For example, if the entries of \mathbf{x}_0 satisfy a Bernoulli-Gaussian distribution $x_0(i) = \Omega(i)v(i)$, with $\Omega(i) \sim \text{Ber}(\theta)$ and $v(i) \sim \mathcal{N}(0, 1)$, and the probability θ diminishes sufficiently quickly with k . We conjecture that this phenomenon holds much more broadly. In particular, determining how slowly θ can diminish with k is an open problem.

2.3 Global Geometry on the Simplex

In the application to image deblurring, the blur kernel is always positive and sums to 1, which naturally leads to a simplex-constrained optimization problem. However, the optimization landscape changes drastically when the convolutional kernel \mathbf{a} is constrained to live on the ℓ^1 norm ball. The objective value of the same objective function over the ℓ^1 ball is shown in Figure 2.

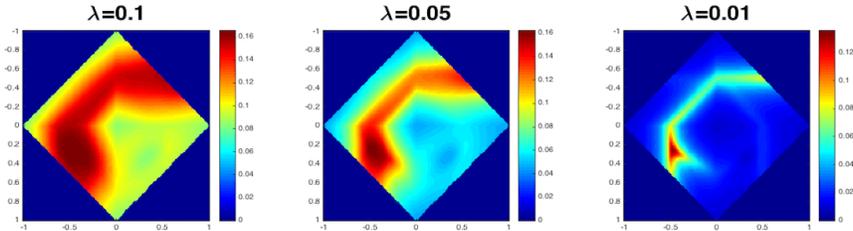


Figure 2: Geometry on the ℓ^1 ball: The trivial spike convolutional kernel is the global minimizer, while the ground truth $[1/3, 1/3, 1/3]$ becomes a local minimizer.

There is a significant difference induced by these two constraints: the trivial spike kernel ($\mathbf{a} = \delta$) becomes the global minima and other meaningful solutions become local minima with the ℓ^1 norm constraint [BVG13], while the spherical constraint always renders local minima close to some signed shift truncation of the ground truth. This important empirical knowledge of the structure of the local minima enables us to infer the ground truth from any local minimizer.

3 A Two-Stage Algorithm

Inspired by the geometric property that every local minimum of the simplified problem $\hat{\varphi}$ is a signed shift-truncation of the ground truth \mathbf{a}_0 , we present a two stage algorithm for reliable recovery of the ground truth \mathbf{a}_0 in this section. In the first stage, the algorithm recovers some signed shift truncation of the ground truth, and the following stage infers the ground truth from this partial recovery.

3.1 Stage I: Find the Signed Shift Truncation

Theorem 2.1 suggests that λ needs to be relatively large to guarantee that all local minimizers are signed shift truncations of the ground truth. However, this is not sufficient to guarantee the success of an optimization algorithm due to the non-differentiability of the ℓ^1 regularizer at $\mathbf{x}^* = \mathbf{0}$. Because of this non-differentiability, when λ is too large, there is a nonzero measure set of \mathbf{a} where $\|\mathbf{C}_a^* \mathbf{y}\|_\infty \leq \lambda$ and therefore $\mathbf{x}^*(\mathbf{a}) = \mathbf{0}$. These

\mathbf{a} are not correlated with any signed shift truncation of \mathbf{a}_0 and are the global maxima of φ . The objective function is constant over this region, so there is no way to escape using only local information.

One way to cope with this flat global maxima region is to replace the nondifferentiable ℓ^1 sparsity penalty with a differentiable one. A natural choice is the huber- μ function, which can be seen as an ℓ^1 penalty but with a rounded bottom for $|x_i| \leq \mu$:

$$h_\mu(\mathbf{x}) = \sum_{|x_i| \leq \mu} \left(\frac{x_i^2}{2\mu} + \frac{\mu}{2} \right) + \sum_{|x_i| > \mu} |x_i| \quad (3.1)$$

As we choose $\mu \ll \lambda$, the μ -huber function closely approximates the ℓ^1 norm, which still maintaining the effect of ‘‘smoothing’’ the flat region. The flat region for the ℓ^1 penalty objective occurs when $\mathbf{x}^*(\mathbf{a}) = \mathbf{0}$, correspondingly we define a region as $\mathcal{R}_{h,0}$ with small $\mathbf{x}^*(\mathbf{a})$ such that

$$\mathcal{R}_{h,0} := \{\mathbf{a} \in \mathbb{S}^{k-1} : \|\mathbf{x}^*(\mathbf{a})\|_\infty \leq \mu\}. \quad (3.2)$$

Within the $\mathcal{R}_{h,0}$ region, the original objective function can be rewritten into a simpler form:

$$\varphi_{h_\mu}(\mathbf{a}) = \frac{1}{2} \|\mathbf{y} - \mathbf{a} \otimes \mathbf{x}\|_2^2 + \frac{\lambda}{2\mu} \|\mathbf{x}\|_2^2 + \frac{\mu n}{2}, \quad (3.3)$$

thus the optimality condition for \mathbf{x}^* implies

$$\mathbf{x}^*(\mathbf{a}) = \left(\mathbf{C}_a^* \mathbf{C}_a + \frac{\lambda}{\mu} \mathbf{I} \right)^{-1} \mathbf{C}_a^* \mathbf{y} \approx \frac{\mu}{\lambda} \mathbf{C}_a^* \mathbf{y}. \quad (3.4)$$

Plugging $\mathbf{x}^*(\mathbf{a})$ back to (2.6) and ignoring the higher order term $O(\frac{\mu^2}{\lambda^2})$ yields

$$\varphi_{h_\mu}(\mathbf{a}) \approx -\frac{\mu}{2\lambda} \|\mathbf{y} \otimes \mathbf{a}\|_2^2 + \frac{1}{2} \|\mathbf{y}\|_2^2 + \frac{\mu n}{2}. \quad (3.5)$$

In this case, minimization of the objective function $\varphi_{h_\mu}(\mathbf{a})$ within the region $\mathcal{R}_{h,0}$ is equivalent to finding the maximum eigenvalue of the matrix $\boldsymbol{\nu}^* \mathbf{C}_y^* \mathbf{C}_y \boldsymbol{\nu}$ with the corresponding leading eigenvectors $e_1(\boldsymbol{\nu}^* \mathbf{C}_y^* \mathbf{C}_y \boldsymbol{\nu})$ achieving the local minima. However, these points can be excluded from $\mathcal{R}_{h,0}$ by setting $\lambda < \min_{\boldsymbol{\nu} \in e_1(\boldsymbol{\nu}^* \mathbf{C}_y^* \mathbf{C}_y \boldsymbol{\nu})} \|\mathbf{C}_y^* \boldsymbol{\nu}\|_\infty$.⁷

With above modifications, the original flat local maxima region $\mathcal{R}_{h,0}$ becomes concave and always have a direction of negative curvature for the algorithm to escape $\mathcal{R}_{h,0}$. Hence, the first stage of the algorithm can find a signed shift-truncation of the ground truth $\bar{\mathbf{a}} = \pm \frac{\boldsymbol{\nu}^* s_\tau[\bar{\mathbf{a}}_0]}{\|\boldsymbol{\nu}^* s_\tau[\bar{\mathbf{a}}_0]\|_F}$ as desired.

3.2 Stage II: Infer the Ground Truth

The second stage of the algorithm aims to recover the ground truth from its signed shift truncation $\bar{\mathbf{a}}$. To recover the truncated part, we first put $\bar{\mathbf{a}}$ in a higher dimensional sphere by zero padding (Figure 3), and then recover the ground truth \mathbf{a}_0 on this higher dimensional sphere. Intuitively, as $\bar{\mathbf{a}}$ still captures a considerable portion of the ground truth \mathbf{a}_0 (the zero padded $\bar{\mathbf{a}}$ is close to the shifted \mathbf{a}_0 in a higher dimensional space), the zero padded $\bar{\mathbf{a}}$ serves as a good initialization. This intuition is made rigorous in the following lemma:

Lemma 3.1. *Let $\lambda_{rel} = \lambda / \|\mathbf{x}_0\|_\infty$, suppose the ground truth \mathbf{a}_0 satisfies*

$$|\langle \mathbf{a}_0, \boldsymbol{\nu}_{s_\tau \neq 0}[\bar{\mathbf{a}}_0] \rangle| < \lambda_{rel}^2 - (2 + 1/\lambda_{rel}^2) \sqrt{1 - \lambda_{rel}^2} \quad (3.6)$$

for any nonzero shift τ , and \mathbf{x}_0 is separated enough such that any two nonzero components are at least $2k$ entries away from each other. If initialized at some $\mathbf{a} \in \mathbb{S}^{k-1}$ that $|\langle \mathbf{a}, \mathbf{a}_0 \rangle| > \lambda_{rel}$, a small-stepping projected gradient method minimizing $\varphi(\mathbf{a})$ recovers the signed ground truth $\pm \mathbf{a}_0$.

⁷A computationally easier upper bound would be $\sqrt{\frac{\lambda_1(\boldsymbol{\nu}^* \mathbf{C}_y^* \mathbf{C}_y \boldsymbol{\nu})}{k}}$. Note that in some scenario, there exists a local minima appearing in either region $\mathcal{R}_{h,0}$ or $\mathcal{R}_{h,0}^c$ regardless of how we set λ . Such extreme case happens when the ground truth convolutional kernel is only supported on a small consecutive portion of its full size, hence a tight estimate of the kernel size is preferred.

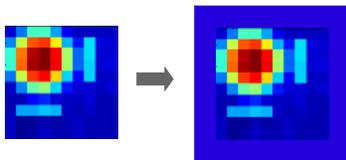


Figure 3: Zero Padding a Signed Shift Truncation The original signed shift truncation (left) and the corresponding zero padded one (right).

Proof Please refer to the supplement. ■

This lemma says when the initial point \mathbf{a} is close enough to the ground truth, the gradient always points to \mathbf{a}_0 as long as $|\langle \mathbf{a}_0, \iota_{S_{\tau \neq 0}}[\widehat{\mathbf{a}}_0] \rangle|$ is sufficiently small. Theorem 2.1 suggests that the first stage of the algorithm finds one local minimum $\widehat{\mathbf{a}}$ that $|\langle \widehat{\mathbf{a}}, \mathbf{a}_0 \rangle| \geq \lambda_{rel}$. Hence, the second stage of the algorithm, which minimizes the same objective function but on a higher dimensional sphere, recovers the ground truth up to sign shift ambiguity as desired.

To ensure accurate recovery, it is important to take the effect of λ on the function geometry into consideration. A larger λ encourages a sparser \mathbf{x} and induces a simpler and smoother function landscape, which effectively eliminates undesirable local minima that are not close to any signed shift truncations, as shown in Figure 4. On the other hand, a smaller λ emphasizes more on the accurate recovery of the signal, therefore the global minima of (2.6) will be closer to the ground truth when λ decreases.

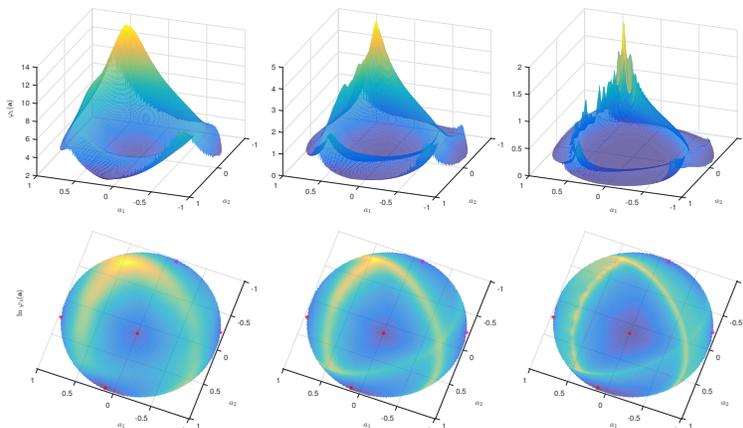


Figure 4: Function Geometry with Varying λ : The objective $\varphi(\mathbf{a})$ over the hemisphere for $\lambda = 10^{-1}, 10^{-3}, 10^{-6}$. Here $\mathbf{a}_0 = \mathcal{P}_{\mathbb{S}^2}[[1, 8, 2]]$ and $\mathbf{x}_0 \sim \text{Ber}(.1) \odot \mathcal{N}(0, 1)$. The ground truth kernel \mathbf{a}_0 and its shift-truncations $\mathcal{P}_{\mathbb{S}^2}[[8, 2, 0]]$, $\mathcal{P}_{\mathbb{S}^2}[[0, 1, 8]]$ are shown in red, and sign-flips $\mathcal{P}_{\mathbb{S}^2}[-[8, 2, 0]]$, $\mathcal{P}_{\mathbb{S}^2}[-[0, 1, 8]]$ are shown in magenta. Notice that each signed shift truncation shown on the hemisphere is close to a corresponding local minima, while as the objective landscape becomes less regularized as λ shrinks.

This geometric effect induced by λ suggests a continuation method in the second stage of the algorithm. We start with a relatively big λ for smoother function geometry, which encourages the algorithm to converge to one meaningful local minimum close to some signed shift truncation of the ground truth. Then run the same algorithm with decreasing sequence of λ to produce a finer approximation of the ground truth. The overall algorithm is described in Algorithm 1.

We need to note that solving $\mathbf{a} = \arg \min \varphi_\lambda(\mathbf{a})$ in Algorithm 1 involves iteration between (i) finding the marginalization over $\mathbf{x}^*(\mathbf{a})$ step, and (ii) updating \mathbf{a} based on the gradient/Hessian of $\varphi_\lambda(\mathbf{a})$. This could be very computationally consuming, a more efficient variant would be to optimize over the cross space of \mathbf{a} and \mathbf{x} together. The corresponding algorithm can be easily adapted to fit into the same general framework. The only things we want to emphasize in the proposed algorithm are the dimension lifting of the sphere and the continuation of λ .

Algorithm 1 Nonconvex Sparse Blind Deconvolution

Ensure: Observation data \mathbf{y} , regularization parameter λ_0 and λ_{\min} , continuation parameter $\beta > 1$

- 1: Solve $\mathbf{a}^{(0)} = \arg \min \varphi_{\lambda_0}(\mathbf{a})$ on \mathbb{S}^{k-1} with random initialization
 - 2:
 - 3: Set $\lambda_1 = \lambda_0$, zero pad $\mathbf{a}^{(0)}$ to $\mathbf{a}^{(1)}$ and $\mathbf{a}^{(1)} \in \mathbb{S}^{k'-1}$ ($k' > k$).
 - 4: **while** $\lambda_k > \lambda_{\min}$ **do**
 - 5: Solve $\mathbf{a}^{(k+1)} = \arg \min \varphi_{\lambda_k}(\mathbf{a})$ on $\mathbb{S}^{k'-1}$ with initialization $\mathbf{a}^{(k)}$.
 - 6:
 - 7: $\lambda_{k+1} = \lambda_k / \beta$
 - 8:
 - 9: **end while**
-

4 Further Extensions

In this section, we extend our algorithm to handle two other deconvolution problems of practical interests: image deblurring and convolutional dictionary learning. The proposed two stage algorithm can be modified and applied to these more complicated applications.

4.1 Image Deblurring

Image deblurring aims to recover a sharp natural image from its blurred observation due to unknown photographic processes such as camera shake or defocus. Although the natural images are not necessarily sparse, it is widely acknowledged that their gradients are approximately sparse. Let $\mathbf{y} = \mathbf{a}_0 \circledast \mathbf{x}_0$ denote the observed blurry image, which is the convolution of the original sharp image \mathbf{x}_0 and the blurring kernel \mathbf{a}_0 . Because of the linearity of the convolution operator, the gradient of the blurred image equals the convolution of the kernel and gradient of the original sharp image, which is usually sparse as desired

$$\nabla_x \mathbf{y} = \mathbf{a}_0 \circledast \nabla_x \mathbf{x}_0, \quad \nabla_y \mathbf{y} = \mathbf{a}_0 \circledast \nabla_y \mathbf{x}_0. \quad (4.1)$$

Here, ∇_x and ∇_y denote derivatives in the x and y directions. In this application, $\nabla_x \mathbf{x}_0$ and $\nabla_y \mathbf{x}_0$ are the underlying sparse signals, and the blind image deblurring problem can be cast as solving:

$$\min_{\mathbf{a} \in \mathbb{S}_+^{k-1}, \mathbf{x}_1, \mathbf{x}_2} \left\{ \begin{aligned} & \frac{1}{2} \|\nabla_x \mathbf{y} - \mathbf{a} \circledast \mathbf{x}_1\|_2^2 + \lambda r(\mathbf{x}_1) \\ & + \frac{1}{2} \|\nabla_y \mathbf{y} - \mathbf{a} \circledast \mathbf{x}_2\|_2^2 + \lambda r(\mathbf{x}_2) \end{aligned} \right\}. \quad (4.2)$$

Here, \mathbb{S}_+^{k-1} denotes the intersection of the unit sphere and the positive orthant. In this application, the non-negativity of the blurring kernel removes the sign ambiguity. We observe in experiments that local minimizers are all near some shift truncation of the ground truth kernel. The same two stage algorithm can therefore be applied to infer the ground truth.

4.2 Convolutional Dictionary Learning

Convolutional dictionary learning (CDL) is an important problem in machine learning for images, speeches, as well as scientific problems like microscopy data analysis and neural spike sorting. The observation signal \mathbf{y} is the superposition of convolutions of N pairs of kernels \mathbf{a}_{0n} and corresponding coefficients \mathbf{x}_{0n} :

$$\mathbf{y} = \sum_{n=1}^N \mathbf{a}_{0n} \circledast \mathbf{x}_{0n}. \quad (4.3)$$

Blind deconvolution can be seen as a special case of CDL with $N = 1$. If the coefficients \mathbf{x}_{0n} are sparse, a natural way to extend our knowledge of SBD would be to assume all N convolutional kernels having unit

Frobenius norm and cast it as minimizing following objective function over the product of N spheres:

$$\min_{\mathbf{a}_n \in \mathbb{S}^{k-1}} \min_{\mathbf{x}_n} \frac{1}{2} \|\mathbf{y} - \sum_{n=1}^N \mathbf{a}_n \otimes \mathbf{x}_n\|_2^2 + \lambda \sum_{n=1}^N r(\mathbf{x}_n). \quad (4.4)$$

We anticipate that all the local minima are near signed shift truncations of the ground truth, provided the target kernels \mathbf{a}_{0n} are sufficiently diverse. The modified two stage algorithm still manages to capture the partial information offered by local minima and hence recovers the ground truth. Experimental results are provided in Section 5.4 to corroborate this claim.

5 Experiments

In this section, we investigate the performance of our algorithm on both synthetic and real data. We first report a systematic investigation, performed in [CLC⁺17], of the performance of our algorithm on synthetic data, which are designed to mimic properties of the microscopy data analysis problem. In Sections 5.2-5.4, we present experiment results showing how our method performs on real data from microscopy and image deblurring.

5.1 Evaluation on Synthetic Data

Noise-free data: we generate the noise-free observation signal of size $m = 256 \times 256$ through circular convolution between a kernel of size k and a random underlying activation signal with a Bernoulli distribution with sparsity θ , i.e. $x_i \stackrel{\text{i.i.d.}}{\sim} \text{Ber}(\theta)$, or $x \sim \text{Ber}(\theta)$. We plot the kernel recovery error for varying kernel size k and sparsity level θ in the left of Figure 5 [CLC⁺17]. Each point on the diagram is the average of 20 independent measurements. The algorithm performs excellently in the blue regions, but begins to fail in the red regions, where either the kernel size is large or the underlying activation signal is dense. The region where typical STM measurements are performed are bounded below by the white dashed line, where the proposed algorithm achieves satisfying performance.

Noisy data: we generate convolutional signals by convolving fixed kernel of dimension k , $k/m = 0.14$ with the random activation map $x \sim \text{Ber}(\theta)$ of dimension m , and applying additive Gaussian noise. We test the performance of our algorithm for varying sparsity θ and noise power. The result is shown in Figure 5 (right): the algorithm achieves noise-robust recovery when the sparsity constraint is satisfied.

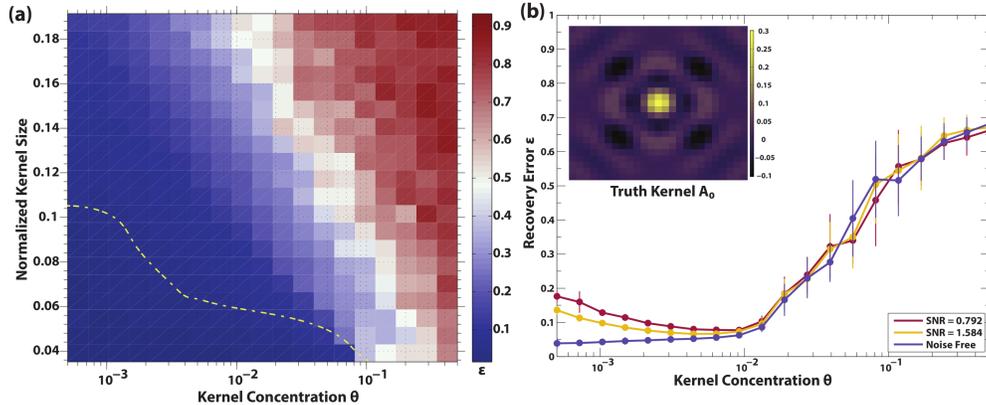


Figure 5: Recovery accuracy [CLC⁺17]. Left: phase transition diagram from noise-free simulated results. Right: performance of algorithm 1 in the presence of additive noise in the measurement; the error increases for small θ due to a lack of samples, whereas extremely large θ leads to algorithmic failure.

5.2 Microscopy Data Analysis

We apply our algorithm on experimental microscopy data obtained from a NaFeCoAs sample. Our results shown in Figure 6 indicate that the proposed algorithm manages to recover the missing details of the ripples in the Fourier domain of the defect, which encode the physical scattering processes of electrons at work.

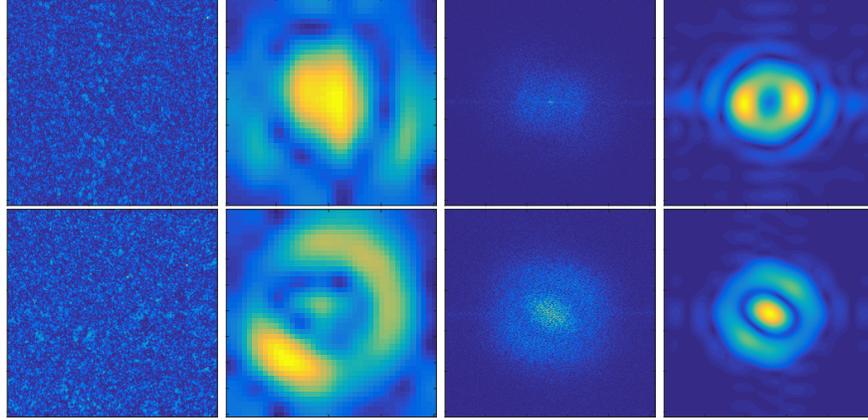


Figure 6: STM Data Analysis. From left to right: the microscopy images, extracted convolutional kernels (defect patterns), and their respective Fourier magnitude images.

5.3 Image Deblurring

We test our algorithm on the image deblurring dataset from [LWDF11], solving (4.3) to recover the convolutional kernel. To clearly separate the inaccuracy of the algorithm and the universal blurring kernel model, all the experiments are done on three kinds of blurred images: (i) synthetic blurred images generated by the convolution of sharp images and blurry kernels; (ii) noisy blurred images generated by adding Gaussian noise to the clean synthetic blurred images (SNR=100); and (iii) real blurry images taken with camera shakes [LWDF11].

We compare with algorithms by Zhang et al.[ZWZ13], Krishnan et al.[KTF11], Sun et al.[SCWH13], and Liu et al.[LCM14].⁸

Because of the shift ambiguity, we evaluate the accuracy of the recovered blurring kernel considering all possible shifts. The kernel recovery error is defined as $\min_{\tau} \|\mathbf{t}^* s_{\tau} [\tilde{\mathbf{a}}] / \|\mathbf{a}\|_1 - \mathbf{a}_0 / \|\mathbf{a}_0\|_1\|_F$, and the cumulative distribution is shown in Figure 7.

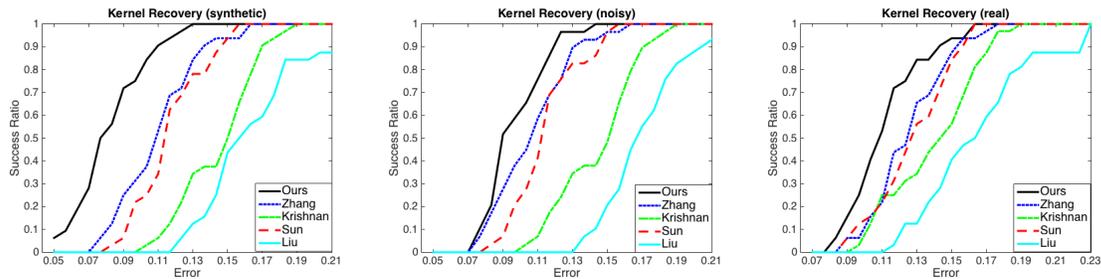


Figure 7: Blur Kernel Recovery Error: Cumulative distributions of recovered blur kernel error from synthetic (left), noised(middle) and real (right) blurry images.

We use the same non-blind deblurring algorithm from [KF09], with the same parameter. We consider the blurred image using the ground truth kernel to be the bench mark, and evaluate the quality of the deblurred

⁸We use the default parameters for these algorithms. It's possible that better performance could be obtained by tuning the parameters more carefully. In our algorithm, we fix the λ 's to be 0.1, 0.01, 0.001, 0.001 for all the instances.

image by calculating the Frobenius norm of its difference to such benchmark. Results are shown in Figure 8.

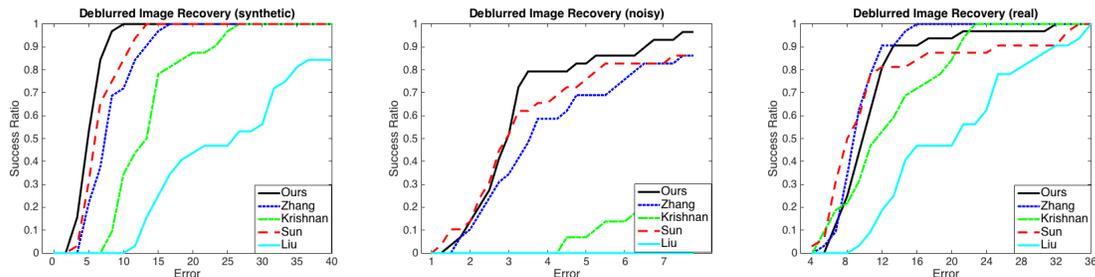


Figure 8: Non-blind Restoration Error: Cumulative distributions of deblurred image error from synthetic (left), noised(middle) and real (right) blurry images.

Our algorithm achieves better convolutional kernel recovery for all three types of images, but its improvement on deblurred image is less obvious, especially for real images. This could be due to (i) the convolutional kernel in this dataset is not strictly uniform across the image, and (ii) the non-blind deconvolution algorithm exploits the heavy-tailed distribution of a natural image’s gradient and becomes less sensitive to the accuracy of the recovered convolutional kernel.

5.4 Convolutional Dictionary Learning

We show results of recovering multiple convolutional kernels on both synthetic data (Figure 9) and real STM data (Figure 10). In the synthetic data, the three convolutional kernels are of size 16×16 and their corresponding activation signals are generated through a Bernoulli model of sparsity 0.005. Results of both stages of the algorithm are shown in Figure 9: the first stage returns kernels close to some shift truncations of the ground truth, and the second stage recovers the ground truth on a higher dimensional space.

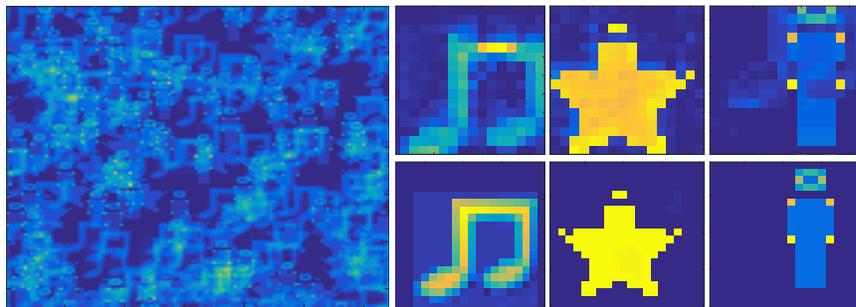


Figure 9: Multi Kernel Blind Deconvolution on Synthetic Data: Input image (left) and the recovered convolutional kernels of Stage I and Stage II of the algorithm (right).

We repeat this experiment with microscopy data obtained from a NaFeAs sample. The algorithm manages to differentiate the two convolutional kernels (defect patterns), as shown in Figure 10. For this material, the kernel orientations depend on the history of the material (stress, temperature, etc.), and using convolutional dictionary learning can be used to automatically detect these features.

6 Generalizations: Matching Loss and Constraints

In the Lasso-like objective function studied in this paper, the approximation error of $y - a \otimes x$ are measured in the squared Frobenius norm (or entry-wise ℓ^2 norm), which is usually adopted to penalize gaussian additive noise in the observation. To obtain reliable solutions, a spherical constraint on the short kernel turned out

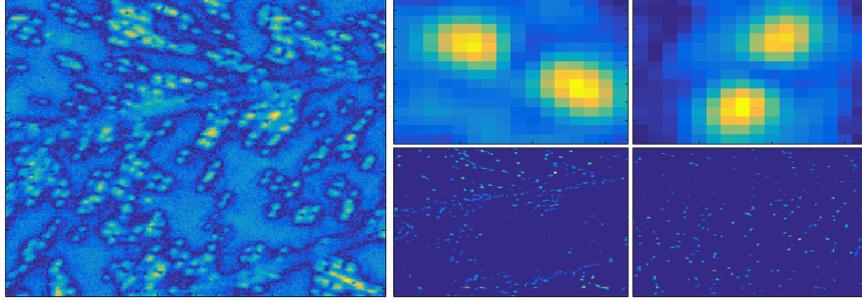


Figure 10: Multi Kernel Blind Deconvolution on Real STM Image: Input image (left) and recovered convolutional kernels and their corresponding activation signals (right).

to be crucial. Moreover, several theoretical papers [ZKW18, KZLW18] studying variations of the Lasso-like objective function also adopt the unit Frobenius constraint, and show that the claimed geometry continues to hold under much more general conditions. However, noise could be more complicated and does not always satisfy the gaussian model in real applications. If given prior knowledge of the noise, other penalties for the loss term $\mathbf{y} - \mathbf{a} \circledast \mathbf{x}$ could be preferable. For example, $\|\mathbf{y} - \mathbf{a} \circledast \mathbf{x}\|_p^p$ with $p < 2$ performs better to penalize heavy tailed noise, and $\|\mathbf{y} - \mathbf{a} \circledast \mathbf{x}\|_p^p$ with $p > 2$ performs better control on the magnitude of the noise.

In this section, we briefly discuss the following more general formulation for the short-and-sparse deconvolution problem

$$\begin{aligned} \min \quad & \psi_p(\mathbf{a}, \mathbf{x}) \doteq \frac{1}{p} \|\mathbf{y} - \mathbf{a} \circledast \mathbf{x}\|_p^p + \lambda r \|\cdot\|_1 \\ \text{s. t.} \quad & \|\mathbf{a}\|_q = 1. \end{aligned} \quad (6.1)$$

Similarly, we write $\varphi_p(\mathbf{a}) \doteq \min_{\mathbf{x}} \psi_p(\mathbf{a}, \mathbf{x})$. In this case, we use entrywise ℓ^p norm to measure the approximation error of $\mathbf{a} \circledast \mathbf{x}$, and assume the optimization constraint for \mathbf{a} to be a unit ℓ^q ($2 < q < \infty$) sphere. Comprehensive study of the function landscape of above general formulation is even more challenging. However, we can demonstrate that once $p = q$, there exist local solutions that share the same geometric property (a local optimum is close to some scaled shifted truncation) through local analysis. As in the objective optimization problem (2.6), the inherent shift ambiguity leads to these structured local minimizers. Because of the randomness in \mathbf{x}_0 , the derived results hold for more general \mathbf{x}_0 .

Lemma 6.1. *Suppose $\mathbf{y} = \mathbf{a}_0 \circledast \mathbf{x}_0$ with $\mathbf{x}_0 = \mathbf{e}_0$, and $p = q \geq 2$. Then for any shift τ , positive scalar ε and λ such that at every $\mathbf{a} \in \mathbb{S}_q \cap \mathbb{B} \left(\frac{\iota^* s_{-\tau}[\mathbf{a}_0]}{\|\iota^* s_{-\tau}[\mathbf{a}_0]\|_q}, \varepsilon \right)$ the solution $\mathbf{x}^*(\mathbf{a}) \doteq \arg \min_{\mathbf{x}} \psi_p(\mathbf{a}, \mathbf{x})$ is (i) unique and (ii) is supported on the τ -th entry, the point $\bar{\mathbf{a}} \doteq \frac{\iota^* s_{-\tau}[\mathbf{a}_0]}{\|\iota^* s_{-\tau}[\mathbf{a}_0]\|_q}$ is a strict local minimizer of the cost $\varphi_p(\mathbf{a})$ over the manifold $\|\mathbf{a}\|_p = 1$.*

Proof Please refer to the supplement. ■

This lemma says that *as long as the constraint matches the loss*, i.e., we choose $q = p \geq 2$, a scaled shift truncation $\bar{\mathbf{a}} = \iota^* s_{-\tau}[\mathbf{a}_0] / \|\iota^* s_{-\tau}[\mathbf{a}_0]\|_q$ achieves the local minimum. In contrast, *when the constraint does not match the loss*, i.e., $q \neq p$, this property is not satisfied: the Riemannian gradient vanishes at a stationary point \mathbf{a} , which satisfies

$$\text{sign}(\iota^* s_{-\tau}[\bar{\mathbf{a}}_0] - \alpha \mathbf{a}) = \text{sign}(\mathbf{a}), \quad (6.2)$$

and

$$(\iota^* s_{-\tau}[\bar{\mathbf{a}}_0] - \alpha \mathbf{a})^{\circ(p-1)} = \alpha' \mathbf{a}^{\circ(q-1)}, \quad (6.3)$$

with α' denoting another non-zero scalar of arbitrary value. This point is not a shift truncation of the ground truth.⁹

⁹When $p \neq q$, with λ decreasing, (6.2) can be rewritten as $\alpha \mathbf{a} = \iota^* s_{-\tau}[\bar{\mathbf{a}}_0] + \alpha' \mathbf{a}^{\circ \frac{q-1}{p-1}}$. With the stationary condition for $\mathbf{x}^*(\mathbf{a})$, we have $\alpha \mathbf{a} \rightarrow \iota^* s_{-\tau}[\bar{\mathbf{a}}_0]$ as $\lambda \rightarrow 0$. Algorithmically, this helps to explain why practical algorithms solving (6.1) with $p \neq q$ can sometimes recover the kernel, where a careful continuation in λ is always necessary.

Hence, when the constraint matches the loss, there is good *local* geometry. In this paper, we have empirically demonstrated that when $p = q = 2$, there is good *global* geometry, under appropriate conditions. We can potentially leverage this property to find global solutions to the general ℓ^p ($p \geq 2$) problem: one first minimizes ℓ^2 to obtain a point near the ground truth and then locally minimizes φ_p over an ℓ^p ball to obtain an estimate that uses the statistical characteristics of the noise.

Although our analysis only pertains to $q \geq 2$, we believe that there is good local geometry even when $q < 2$. In this situation, the constraint set is no longer twice differentiable and the shift truncations are always on the non-smooth region of the constraint set. As the p decrease, the normal cone at the nonsmooth points gets wider. For example, if $q \geq 2$ the normal cone is the normal direction. Therefore, as q decreases, the normal cone is larger and hence more like to contain the gradient direction, then this shift truncation is likely to be a stationary point.

7 Discussions

This work studies the global geometry of a nonconvex optimization problem for SBD when the kernel is assumed to have unit Frobenius norm. In this setting, we find that all the local minima are benign, in the sense that they are close to some signed shift truncation of the ground truth. With this insight, we propose a two stage algorithm that recovers the ground truth by exploiting the information hidden in local minima.

This problem reveals the challenges faced when analyzing the SBD problem via a geometrical approach. For problems enjoying stronger symmetry properties [SQW17a, SQW17b], similar approaches yield a global understanding of the function geometry and recovery guarantees. We expect that the weak symmetries in SBD make a major contribution to the difficulties encountered for this problem.

There are lots of additional further directions could be of great interests for both theory and application: Our empirical results show that the our characterization of local minima carries through to the convolutional dictionary learning problem, which can also be efficiently solved by slight adaptation of the proposed algorithm. However, the theory part is open, it would be interesting to know how many kinds of kernels, or what kinds of kernels are recoverable, probably by some measures of incoherence, which is a common assumption in dictionary learning problem.

Two other imperfections we encounter in scientific measurement are resolution limit and measurement error, which inspire us to consider (i) if it's possible to integrate blind deconvolution and super-resolution process together; (ii) if we can come up with a robust blind deconvolution algorithm to automatically rule out noisy entries.

Acknowledgement

The authors gratefully acknowledge support from NSF 1343282, NSF CCF 1527809, and NSF IIS 1546411.

References

- [AMS07] P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, Princeton, NJ, USA, 2007.
- [AMS09] Pierre-Antoine Absil, Robert Mahoney, and Rodolphe Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, 2009.
- [ARR12] Ali Ahmed, Benjamin Recht, and Justin Romberg. Blind deconvolution using convex programming. *arXiv preprint:1211.5608*, 2012.
- [BVG13] Alexis Benichoux, Emmanuel Vincent, and Remi Gribonval. A fundamental pitfall in blind deconvolution with sparse and shift-invariant priors. *38th International Conference on Acoustics, Speech, and Signal Processing*, May 2013.
- [Chi16] Yuejie Chi. Guaranteed blind sparse spikes deconvolution via lifting and convex optimization. *IEEE Journal of Selected Topics in Signal Processing*, 10(4):782–794, June 2016.

- [CL09] Sunghyun Cho and Seungyong Lee. Fast motion deblurring. *ACM Transactions on Graphics (SIGGRAPH ASIA 2009)*, 28(5):article no. 145, 2009.
- [CLC⁺17] Sky Cheung, Yenson Lau, Zhengyu Chen, Ju Sun, Yuqian Zhang, John Wright, and Abhay Pasupathy. Beyond the fourier transform: A nonconvex optimization approach to microscopy analysis. *Submitted*, 2017.
- [CW98] Tony F. Chan and Chiu-Kwong Wong. Total variation blind deconvolution. *IEEE Transactions on Image Processing*, 7(3):370–375, Mar 1998.
- [ETS11] Chaitanya Ekanadham, Daniel Tranchina, and Eero P. Simoncelli. A blind sparse deconvolution method for neural spike identification. In *Advances in Neural Information Processing Systems 24*, pages 1440–1448. 2011.
- [FSH⁺06] Rob Fergus, Barun Singh, Aaron Hertzmann, Sam T. Roweis, and William T. Freeman. Removing camera shake from a single photograph. *ACM Transactions on Graphics*, 25(3):787–794, June 2006.
- [GTZ⁺16] Dong Gong, Mingkui Tan, Yanning Zhang, Anton van den Hengel, and Qinfeng Shi. Blind image deconvolution by automatic gradient activation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [KF09] Dilip Krishnan and Rob Fergus. Fast image deconvolution using hyper-laplacian priors. In *Advances in Neural Information Processing Systems 22*, pages 1033–1041. 2009.
- [KH96] Deepa Kundur and Dimitrios Hatzinakos. Blind image deconvolution. *Signal Processing Magazine, IEEE*, 13(3):43–64, May 1996.
- [KTF11] Dilip Krishnan, Terence Tay, and Rob Fergus. Blind deconvolution using a normalized sparsity measure. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 233–240, June 2011.
- [KZLW18] Han-Wen Kuo, Yuqian Zhang, Yenson Lau, and John Wright. Geometry and symmetry in short-and-sparse deconvolution. In *Preparation*, 2018.
- [LCM14] Guangcan Liu, Shiyu Chang, and Yi Ma. Blind image deblurring using spectral properties of convolution operators. *IEEE Transactions on Image Processing*, 23(12):5047–5056, Dec 2014.
- [Lew98] Michael S. Lewicki. A review of methods for spike sorting: the detection and classification of neural action potentials. *Network: Computation in Neural Systems*, 9(4):53–78, 1998.
- [LLSW18] Xiaodong Li, Shuyang Ling, Thomas Strohmer, and Ke Wei. Rapid, robust, and reliable blind deconvolution via nonconvex optimization. *Applied and computational harmonic analysis*, June 2018.
- [LWDF11] Anat Levin, Yair Weiss, Fredo Durand, and William T. Freeman. Understanding blind deconvolution algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(12):2354–2367, Dec 2011.
- [MBP14] Julien Mairal, Francis Bach, and Jean Ponce. Sparse modeling for image and vision processing. *Foundations and Trends in Computer Graphics and Vision*, 8(2-3), December 2014.
- [PF14] Daniele Perrone and Paolo Favaro. Total variation blind deconvolution: The devil is in the details. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [SCWH13] Libin Sun, Sunghyun Cho, Jue Wang, and James Hays. Edge-based blur kernel estimation using patch priors. In *Proc. IEEE International Conference on Computational Photography*, 2013.
- [SQW17a] Ju Sun, Qing Qu, and John Wright. Complete dictionary recovery over the sphere. *IEEE Transactions on Information Theory*, 2017.
- [SQW17b] Ju Sun, Qing Qu, and John Wright. A geometric analysis of phase retrieval. *Foundations of Computational Mathematics*, 2017.
- [SWW12] Daniel Spielman, Huan Wang, and John Wright. Exact recovery of sparsely-used dictionaries. *COLT*, 2012.
- [WZ13] David Wipf and Haichao Zhang. Revisiting bayesian blind deconvolution. *arXiv preprint:1305.2362*, 2013.
- [XJ10] Li Xu and Jiaya Jia. Two-phase kernel estimation for robust motion deblurring. In *European Conference on Computer Vision (ECCV)*, pages 157–170, 2010.
- [XWHH16] Lei Xiao, Jue Wang, Wolfgang Heidrich, and Michael Hirsch. Learning high-order filters for efficient blind deconvolution of document photographs. In *14th European Conference on Computer Vision (ECCV)*, pages 734–749, 2016.
- [ZKW18] Yuqian Zhang, Han-Wen Kuo, and John Wright. Structured local minimum is sparse blind deconvolution. *NIPS*, 2018.
- [ZWZ13] Haichao Zhang, David Wipf, and Yanning Zhang. Multi-image blind deblurring using a coupled adaptive sparse prior. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, January 2013.

A Notations

This supplementary material contains a complete proof of Theorem 2.1 and Lemma 3.1 of the main paper. Section 1 discusses some technical notations and geometric observations, Section 2 and 3 presents a complete proof for the theorem and lemma respectively, and Section 4 contains two auxiliary lemmas used in the main proof.

Before commencing our main proof, we introduce some notations, recap the simplified objective function $\widehat{\varphi}$ and a few basic observations on its geometry. In this note, we use C_a or $C_{\bar{a}}$ exchangeably to denote the $m \times m$ circulant matrix generated by a k -length short convolutional kernel a without ambiguity.

Reversed Circulant Matrix Let \check{C}_{a_0} be the reversed circulant matrix for a_0 such that

$$\check{C}_{a_0} = [s_0[\widehat{a_0}] \mid s_{-1}[\widehat{a_0}] \mid s_{-2}[\widehat{a_0}] \mid \dots \mid s_{-(m-1)}[\widehat{a_0}]] \in \mathbb{R}^{m \times m}, \quad (\text{A.1})$$

Here, $s_\tau[\widehat{a_0}]$ denotes a cyclic shift as defined in Equation (3) of the main paper. Since the i -th entry of $C_a^* \widehat{a_0}$ satisfies

$$[C_a^* \widehat{a_0}]_i = \langle s_i[\tilde{a}], \widehat{a_0} \rangle = \langle \tilde{a}, s_{-i}[\widehat{a_0}] \rangle = \langle a, \iota^* s_{-i}[\widehat{a_0}] \rangle, \quad (\text{A.2})$$

therefore following equation always holds

$$C_a^* \widehat{a_0} = \check{C}_{a_0}^* \tilde{a}. \quad (\text{A.3})$$

Projection onto I Let e_0, \dots, e_{m-1} denote the standard basis vectors. For index set $I = \{i_1, \dots, i_{|I|}\} \neq \emptyset$, let

$$V_I \doteq [e_{i_1} \mid e_{i_2} \mid \dots \mid e_{i_{|I|}}] \in \mathbb{R}^{m \times |I|}, \quad (\text{A.4})$$

and then

$$P_I = V_I V_I^* \quad (\text{A.5})$$

denote the projection onto the coordinates indexed by I while setting other entries zero.

Partial Signed Support Let u, v be two vectors of the same dimension. If $\text{supp}(u) \subseteq \text{supp}(v)$ and $u(i)v(i) \geq 0$ for all $i \in \text{supp}(v)$, then u attains the partial signed support of v , denoted as $u \preceq v$.

Piecewise Quadratic Function Let $x^*(a)$ denote the minimizer for simplified objective function

$$\widehat{\varphi}(a) = \min_x \frac{1}{2} \|\widehat{a_0}\|_2^2 + \frac{1}{2} \|x\|_2^2 - \langle a \circledast x, \widehat{a_0} \rangle + \lambda \|x\|_1, \quad (\text{A.6})$$

with sign σ and support I defined as

$$\sigma \doteq \text{sign } x^*(a) \in \{-1, 0, 1\}^m, \quad I \doteq \text{supp}(\sigma) \subseteq \{0, 1, \dots, m-1\}^{10}. \quad (\text{A.7})$$

By stationary condition for $x^*(a)$, we obtain

$$x^*(a) = \text{SOFT}_\lambda [C_a^* \widehat{a_0}] = \text{SOFT}_\lambda [\check{C}_{a_0}^* \iota a], \quad (\text{A.8})$$

where $\text{SOFT}_\lambda [u] = \text{sign}(u) \max\{|u| - \lambda, 0\}$ is the entry-wise soft-thresholding operator.

For each sign pattern $\sigma = \text{supp}(x^*) \in \{-1, 0, 1\}^m$, there exists corresponding region on the sphere such that

$$R_\sigma = \{a \mid \text{sign } \text{SOFT}_\lambda [\check{C}_{a_0}^* \iota a] = \sigma\}. \quad (\text{A.9})$$

On the relative interior of each R_σ , the function $\widehat{\varphi}$ has a simple expression:

$$\widehat{\varphi}(a) = \widehat{\varphi}_\sigma(a) \doteq -\frac{1}{2} a^* \iota^* \check{C}_{a_0} P_I \check{C}_{a_0}^* \iota a + \lambda \sigma^* P_I \check{C}_{a_0}^* \iota a + \frac{1}{2} - \frac{\lambda^2 |I|}{2}. \quad (\text{A.10})$$

Therefore, the objective function is piecewise quadratic and can be rewritten as

$$\widehat{\varphi}_\sigma = \frac{1}{2} a^* M_\sigma a + b_\sigma^* a + \text{Const}_\sigma, \quad (\text{A.11})$$

with

$$M_\sigma = -\iota^* \check{C}_{a_0} P_I \check{C}_{a_0}^* \iota, \quad b_\sigma = \lambda \iota^* \check{C}_{a_0}^* P_I \sigma. \quad (\text{A.12})$$

With above notations clarified, we are now ready to present a proof for the main theorem.

¹⁰As in the main paper, we assume that m dimensional vectors are indexed by the integers $0, 1, \dots, m-1$.

B Proof of Theorem 2.1

Proof On the relative interior of each R_σ , the simplified objective function

$$\widehat{\varphi}_\sigma = \frac{1}{2} \mathbf{a}^* \mathbf{M}_\sigma \mathbf{a} + \mathbf{b}_\sigma^* \mathbf{a} + c_\sigma \quad (\text{B.1})$$

has Euclidean derivative and Hessian

$$\nabla \widehat{\varphi}_\sigma(\mathbf{a}) = \mathbf{M}_\sigma \mathbf{a} + \mathbf{b}_\sigma, \quad (\text{B.2})$$

$$\nabla^2 \widehat{\varphi}_\sigma(\mathbf{a}) = \mathbf{M}_\sigma. \quad (\text{B.3})$$

As we assume \mathbf{a} to have unit Frobenius norm, or to live on a sphere, the more natural Riemannian gradient and Hessian are defined as

$$\text{grad} [\widehat{\varphi}_\sigma] (\mathbf{a}) = \mathbf{P}_{\mathbf{a}^\perp} \nabla \widehat{\varphi}_\sigma(\mathbf{a}) \quad (\text{B.4})$$

$$= \mathbf{M}_\sigma \mathbf{a} + \mathbf{b}_\sigma - \mathbf{a}(\mathbf{a}^* \mathbf{M}_\sigma \mathbf{a} + \mathbf{b}_\sigma^* \mathbf{a}), \quad (\text{B.5})$$

$$\text{Hess} [\widehat{\varphi}_\sigma] (\mathbf{a}) = \mathbf{P}_{\mathbf{a}^\perp} \left(\nabla^2 \widehat{\varphi}_\sigma(\mathbf{a}) - \langle \nabla \widehat{\varphi}_\sigma(\mathbf{a}), \mathbf{a} \rangle \mathbf{I} \right) \mathbf{P}_{\mathbf{a}^\perp} \quad (\text{B.6})$$

$$= \mathbf{P}_{\mathbf{a}^\perp} \left(\mathbf{M}_\sigma - (\mathbf{a}^* \mathbf{M}_\sigma \mathbf{a} + \mathbf{b}_\sigma^* \mathbf{a}) \mathbf{I} \right) \mathbf{P}_{\mathbf{a}^\perp}. \quad (\text{B.7})$$

Here, $\mathbf{P}_{\mathbf{a}^\perp} = \mathbf{I} - \mathbf{a} \mathbf{a}^*$ denotes projection onto the tangent space over the sphere at \mathbf{a} . As in the Euclidean space, a stationary point on the sphere needs to satisfy $\text{grad} [\widehat{\varphi}_\sigma] (\mathbf{a}) = \mathbf{0}$. At a stationary point $\bar{\mathbf{a}}$, if $\text{Hess} [\widehat{\varphi}_\sigma] (\bar{\mathbf{a}})$ is positive semidefinite, the function is convex and $\bar{\mathbf{a}}$ is a local minimum; if $\text{Hess} [\widehat{\varphi}_\sigma] (\bar{\mathbf{a}})$ has a negative eigenvalue, then there exists a direction along which the objective value decreases and hence $\bar{\mathbf{a}}$ is a saddle point [AMS07].

Let $I = \{i_1 < i_2 < \dots < i_{|I|}\}$ and define

$$\eta_i = \|\boldsymbol{\iota}^* s_{-i} [\widetilde{\mathbf{a}}_0]\|_2 \quad \forall i \in I, \quad (\text{B.8})$$

$$\boldsymbol{\eta} = (\eta_{i_1}, \eta_{i_2}, \dots, \eta_{i_{|I|}}) \in \mathbb{R}^{|I|}, \quad (\text{B.9})$$

and

$$\mathbf{U} = \left[\begin{array}{c|c|c} \frac{\sigma_{i_1} \boldsymbol{\iota}^* s_{-i_1} [\widetilde{\mathbf{a}}_0]}{\eta_{i_1}} & \frac{\sigma_{i_2} \boldsymbol{\iota}^* s_{-i_2} [\widetilde{\mathbf{a}}_0]}{\eta_{i_2}} & \dots \\ \hline & & \frac{\sigma_{i_{|I|}} \boldsymbol{\iota}^* s_{-i_{|I|}} [\widetilde{\mathbf{a}}_0]}{\eta_{i_{|I|}}} \end{array} \right] \in \mathbb{R}^{k \times |I|}. \quad (\text{B.10})$$

Here, columns of \mathbf{U} have unit ℓ^2 norm. Then we have

$$\mathbf{M}_\sigma = -\mathbf{U} \text{diag}(\boldsymbol{\eta})^2 \mathbf{U}^*, \quad \mathbf{b}_\sigma = \lambda \mathbf{U} \boldsymbol{\eta}. \quad (\text{B.11})$$

As I is defined via soft thresholding, we have $|\check{\mathbf{C}}_{\mathbf{a}_0}^* \boldsymbol{\iota} \mathbf{a}|_i > \lambda$ holds for every $i \in I$. Hence,

$$\mathbf{a}^* \mathbf{M}_\sigma \mathbf{a} + \mathbf{b}_\sigma^* \mathbf{a} = -\mathbf{a}^* \boldsymbol{\iota}^* \check{\mathbf{C}}_{\mathbf{a}_0}^* \mathbf{P}_I \check{\mathbf{C}}_{\mathbf{a}_0}^* \boldsymbol{\iota} \mathbf{a} + \lambda \boldsymbol{\sigma}^* \mathbf{P}_I \check{\mathbf{C}}_{\mathbf{a}_0}^* \boldsymbol{\iota} \mathbf{a} \quad (\text{B.12})$$

$$= -\|\mathbf{P}_I \check{\mathbf{C}}_{\mathbf{a}_0}^* \boldsymbol{\iota} \mathbf{a}\|_2^2 + \lambda \|\mathbf{P}_I \check{\mathbf{C}}_{\mathbf{a}_0}^* \boldsymbol{\iota} \mathbf{a}\|_1 \quad (\text{B.13})$$

$$< 0 \quad (\text{B.14})$$

holds at any $\mathbf{a} \in \text{cl}(R_\sigma) \setminus R_0$.

Stationary point and implications

Consider any stationary point $\bar{\mathbf{a}} \in \text{cl}(R_\sigma) \setminus R_0$ of $\widehat{\varphi}$. By continuity of the gradient of $\widehat{\varphi}$ (proved in Lemma B.1), $\bar{\mathbf{a}}$ is also a stationary point of $\widehat{\varphi}_\sigma$. By definition, $\text{grad} [\widehat{\varphi}_\sigma] (\bar{\mathbf{a}}) = \mathbf{0}$, which implies that

$$(\bar{\mathbf{a}}^* \mathbf{M}_\sigma \bar{\mathbf{a}} + \mathbf{b}_\sigma^* \bar{\mathbf{a}}) \bar{\mathbf{a}} = \mathbf{M}_\sigma \bar{\mathbf{a}} + \mathbf{b}_\sigma. \quad (\text{B.15})$$

Note that since $\bar{\mathbf{a}}^* \mathbf{M}_\sigma \bar{\mathbf{a}} + \mathbf{b}_\sigma^* \bar{\mathbf{a}} \neq 0$ and $\mathbf{b}_\sigma \in \text{range}(\mathbf{M}_\sigma)$, this implies that $\bar{\mathbf{a}} \in \text{range}(\mathbf{M}_\sigma)$.

Let $\gamma = -(\mathbf{a}^* \mathbf{M}_\sigma \mathbf{a} + \mathbf{b}_\sigma^* \mathbf{a}) > 0$, then the condition for a stationary point $\bar{\mathbf{a}}$ becomes

$$\gamma \bar{\mathbf{a}} = \mathbf{U} \text{diag}(\boldsymbol{\eta})^2 \mathbf{U}^* \bar{\mathbf{a}} - \lambda \mathbf{U} \boldsymbol{\eta}. \quad (\text{B.16})$$

Let $\boldsymbol{\alpha} = \mathbf{U}^* \bar{\mathbf{a}}$, and note that for each j , $\alpha_j > 0$ and $\alpha_j \eta_j > \lambda$. In terms of \mathbf{U} , the stationarity condition becomes

$$\gamma \boldsymbol{\alpha} = \mathbf{U}^* \mathbf{U} \text{diag}(\boldsymbol{\eta})^2 \boldsymbol{\alpha} - \lambda \mathbf{U}^* \mathbf{U} \boldsymbol{\eta}. \quad (\text{B.17})$$

Since the diagonal elements of $\mathbf{U}^* \mathbf{U}$ are all ones, and hence can be written as

$$\mathbf{U}^* \mathbf{U} = \mathbf{I} + \boldsymbol{\Delta}. \quad (\text{B.18})$$

We have

$$\gamma \boldsymbol{\alpha} = \text{diag}(\boldsymbol{\eta})^2 \boldsymbol{\alpha} - \lambda \boldsymbol{\eta} + \boldsymbol{\Delta} \left\{ \text{diag}(\boldsymbol{\eta})^2 \boldsymbol{\alpha} - \lambda \boldsymbol{\eta} \right\}. \quad (\text{B.19})$$

As $\boldsymbol{\alpha} \succ \mathbf{0}$ and $\text{diag}(\boldsymbol{\eta}) \boldsymbol{\alpha} \succ \lambda \cdot \mathbf{1}$ ¹¹, together with an auxiliary Lemma E.2 proved in Section E, we have

$$\left\| \text{diag}(\boldsymbol{\eta})^2 \boldsymbol{\alpha} - \lambda \boldsymbol{\eta} \right\|_2 \leq \left\| \boldsymbol{\alpha} \right\|_2 \quad (\text{B.20})$$

$$= \left\| \mathbf{U}^* \bar{\mathbf{a}} \right\|_2 \quad (\text{B.21})$$

$$\leq \left\| \mathbf{U} \right\|_{\ell^2 \rightarrow \ell^2} \quad (\text{B.22})$$

$$\leq \sqrt{3/2}, \quad (\text{B.23})$$

whence

$$\left\| \boldsymbol{\Delta} \left\{ \text{diag}(\boldsymbol{\eta})^2 \boldsymbol{\alpha} - \lambda \boldsymbol{\eta} \right\} \right\|_\infty \leq \sqrt{3/2} \times \left\| \boldsymbol{\Delta} \right\|_{\ell^2 \rightarrow \ell^\infty}. \quad (\text{B.24})$$

Suppose that $\boldsymbol{\Delta}$ is small enough that the right hand side of (B.24) is bounded by $\lambda^2/2$, i.e.,

$$\left\| \boldsymbol{\Delta} \right\|_{\ell^2 \rightarrow \ell^\infty} \leq \frac{\lambda^2}{\sqrt{6}}. \quad (\text{B.25})$$

Plugging back into the stationary condition $\text{diag}(\boldsymbol{\eta})^2 \boldsymbol{\alpha} - \gamma \boldsymbol{\alpha} = \lambda \boldsymbol{\eta} - \boldsymbol{\Delta} \left\{ \text{diag}(\boldsymbol{\eta})^2 \boldsymbol{\alpha} - \lambda \boldsymbol{\eta} \right\}$ gives

$$(\text{diag}(\boldsymbol{\eta})^2 - \gamma) \boldsymbol{\alpha} \succ \lambda \boldsymbol{\eta} - \lambda^2/2 \succ \mathbf{0}. \quad (\text{B.26})$$

Since $\alpha_i < 1$ and $\eta_i > \lambda$ for all i , which implies that

$$\gamma < \eta_{\min}^2 - \lambda \eta_{\min} + \lambda^2/2, \quad (\text{B.27})$$

where η_{\min} is the smallest of the η_i .

Negative curvature in Hessian Recall the Riemannian Hessian on the sphere is defined as

$$\text{Hess}[\widehat{\varphi}_\sigma](\mathbf{a}) = \mathbf{P}_{\mathbf{a}^\perp} \left(\mathbf{M}_\sigma - (\mathbf{a}^* \mathbf{M}_\sigma \mathbf{a} + \mathbf{b}_\sigma^* \mathbf{a}) \mathbf{I} \right) \mathbf{P}_{\mathbf{a}^\perp} \quad (\text{B.28})$$

$$= \mathbf{P}_{\mathbf{a}^\perp} \left(-\mathbf{U} \text{diag}(\boldsymbol{\eta})^2 \mathbf{U}^* + \gamma \mathbf{I} \right) \mathbf{P}_{\mathbf{a}^\perp}. \quad (\text{B.29})$$

Below argument shows that this Riemannian Hessian has negative eigenvalues. Let $\tilde{\mathbf{U}}$ be an orthonormal matrix generated via

$$\tilde{\mathbf{U}} \doteq \mathbf{U} (\mathbf{U}^* \mathbf{U})^{-1/2} \quad (\text{B.30})$$

Whenever $\left\| \boldsymbol{\Delta} \right\|_{\ell^2 \rightarrow \ell^2} < 1/2$ holds, Lemma E.2 guarantees $\left\| \mathbf{U} - \tilde{\mathbf{U}} \right\|_{\ell^2 \rightarrow \ell^2} < 3 \left\| \boldsymbol{\Delta} \right\|_{\ell^2 \rightarrow \ell^2}$. Under this condition, we can lower bound the smallest nonzero eigenvalue of $\mathbf{U} \text{diag}(\boldsymbol{\eta})^2 \mathbf{U}^*$, as

$$\lambda_{\min}(\mathbf{U} \text{diag}(\boldsymbol{\eta})^2 \mathbf{U}^*) = \sigma_{\min}(\mathbf{U} \text{diag}(\boldsymbol{\eta})^2) \quad (\text{B.31})$$

¹¹Here, \succ denotes element-wise inequality between vectors.

$$\geq \left(\max \left\{ \sigma_{\min} \left(\tilde{U} \text{diag}(\boldsymbol{\eta}) \right) - \left\| \mathbf{U} - \tilde{U} \right\| \left\| \text{diag}(\boldsymbol{\eta}) \right\|, 0 \right\} \right)^2 \quad (\text{B.32})$$

$$= \left(\max \left\{ \eta_{\min} - \left\| \mathbf{U} - \tilde{U} \right\| \eta_{\max}, 0 \right\} \right)^2 \quad (\text{B.33})$$

$$\geq \eta_{\min}^2 - 3\eta_{\max}\eta_{\min} \left\| \boldsymbol{\Delta} \right\|_{\ell^2 \rightarrow \ell^2}. \quad (\text{B.34})$$

Since $\lambda < \eta_{\min} \leq \eta_{\max} \leq 1$, additionally if $\left\| \boldsymbol{\Delta} \right\|_{\ell^2 \rightarrow \ell^2} \leq \frac{\lambda}{6}$, we have

$$3\eta_{\max}\eta_{\min} \left\| \boldsymbol{\Delta} \right\|_{\ell^2 \rightarrow \ell^2} \leq \lambda\eta_{\min} - \lambda^2/2, \quad (\text{B.35})$$

Together with (B.25) and (B.27), we can obtain

$$\lambda_{\min}(\mathbf{U} \text{diag}(\boldsymbol{\eta})^2 \mathbf{U}^*) > \gamma, \quad (\text{B.36})$$

or

$$\lambda_{\max}(\mathbf{M}_{\boldsymbol{\sigma}}) < -\gamma. \quad (\text{B.37})$$

Thus, whenever the following conditions are satisfied

$$\lambda < 1, \quad \left\| \boldsymbol{\Delta} \right\|_{\ell^2 \rightarrow \ell^\infty} \leq \frac{\lambda^2}{\sqrt{6}}, \quad \left\| \boldsymbol{\Delta} \right\|_{\ell^2 \rightarrow \ell^2} \leq \frac{\lambda}{6}, \quad (\text{B.38})$$

we have $\lambda_{\max}(\mathbf{M}_{\boldsymbol{\sigma}}) < -\gamma$ as desired.

Above calculations imply that for every $\boldsymbol{\xi} \in \text{range}(\mathbf{M}_{\boldsymbol{\sigma}}) \subseteq \mathbb{R}^k$,

$$\boldsymbol{\xi}^* (\mathbf{M}_{\boldsymbol{\sigma}} + \gamma \mathbf{I}) \boldsymbol{\xi} < 0. \quad (\text{B.39})$$

Since

$$\text{Hess} [\hat{\varphi}_{\boldsymbol{\sigma}}] (\bar{\mathbf{a}}) = \mathbf{P}_{\bar{\mathbf{a}}^\perp} \left(\mathbf{M}_{\boldsymbol{\sigma}} + \gamma \mathbf{I} \right) \mathbf{P}_{\bar{\mathbf{a}}^\perp}, \quad (\text{B.40})$$

for $\boldsymbol{\xi} \in \bar{\mathbf{a}}^\perp \cap \text{range}(\mathbf{M}_{\boldsymbol{\sigma}})$,

$$\boldsymbol{\xi}^* \text{Hess} [\hat{\varphi}_{\boldsymbol{\sigma}}] (\bar{\mathbf{a}}) \boldsymbol{\xi} < 0. \quad (\text{B.41})$$

Hence, on the relative interior $\text{relint}(R_{\boldsymbol{\sigma}})$, $\hat{\varphi} \equiv \hat{\varphi}_{\boldsymbol{\sigma}}$ obtains, and so this implies that for $\|\boldsymbol{\sigma}\|_0 > 1$, there are no local minima in $\text{relint}(R_{\boldsymbol{\sigma}})$.

Relative boundaries We first note that if $\|\boldsymbol{\sigma}\|_0 = 1$ and $I = \{i\}$, either $R_{\boldsymbol{\sigma}}$ is empty when $\|\boldsymbol{\iota}^* s_{-i} [\bar{\mathbf{a}}_0]\|_2 \leq \lambda$, or it contains an open ball around $\text{range}(\mathbf{M}_{\boldsymbol{\sigma}}) \cap \mathbb{S}^{k-1} = \pm \frac{\boldsymbol{\iota}^* s_{-i} [\bar{\mathbf{a}}_0]}{\|\boldsymbol{\iota}^* s_{-i} [\bar{\mathbf{a}}_0]\|_2}$. Hence, if $\bar{\mathbf{a}} \in \text{relbdy}(R_{\boldsymbol{\sigma}})$ is a stationary point and $\boldsymbol{\sigma} \neq \mathbf{0}$, we necessarily have $\|\boldsymbol{\sigma}\|_0 \geq 2$.

Since $\bar{\mathbf{a}}$ is on the boundary of $R_{\boldsymbol{\sigma}}$, it is also in $\text{relbdy}(\text{cl}(R_{\boldsymbol{\sigma}'}))$ for some $\boldsymbol{\sigma}' \neq \boldsymbol{\sigma}$. Let

$$\Xi = \{ \boldsymbol{\sigma}' \mid \bar{\mathbf{a}} \in \text{relbdy}(\text{cl}(R_{\boldsymbol{\sigma}'})) \}. \quad (\text{B.42})$$

Suppose that for every $\boldsymbol{\sigma}' \in \Xi$, $\boldsymbol{\sigma} \preceq \boldsymbol{\sigma}'$. Hence, $\text{range}(\mathbf{M}_{\boldsymbol{\sigma}}) \subseteq \text{range}(\mathbf{M}_{\boldsymbol{\sigma}'})$ for every $\boldsymbol{\sigma}' \in \Xi$ and

$$\boldsymbol{\xi}^* \text{Hess} [\hat{\varphi}_{\boldsymbol{\sigma}'}] (\bar{\mathbf{a}}) \boldsymbol{\xi} < 0, \quad \forall \boldsymbol{\xi} \in \text{range}(\mathbf{M}_{\boldsymbol{\sigma}}), \boldsymbol{\sigma}' \in \Xi. \quad (\text{B.43})$$

By continuity of the gradients, $\bar{\mathbf{a}}$ is a stationary point for every $\hat{\varphi}_{\boldsymbol{\sigma}'}$, such that $\boldsymbol{\sigma}' \in \Xi$. If we choose an arbitrary nonzero $\boldsymbol{\xi} \in \text{range}(\mathbf{M}_{\boldsymbol{\sigma}})$, we have that for every $\boldsymbol{\sigma}' \in \Xi$,

$$\hat{\varphi}_{\boldsymbol{\sigma}'}(\mathcal{P}_{\mathbb{S}^{k-1}}[\bar{\mathbf{a}} + t\boldsymbol{\xi}]) < \hat{\varphi}(\bar{\mathbf{a}}) - \Omega(t^2). \quad (\text{B.44})$$

There exists a neighborhood N of $\bar{\mathbf{a}}$ for which, at every $\mathbf{a} \in N \cap R_{\boldsymbol{\sigma}'}$, $\hat{\varphi}(\mathbf{a}) = \hat{\varphi}_{\boldsymbol{\sigma}'}(\mathbf{a}) \leq \hat{\varphi}(\bar{\mathbf{a}})$ for some $\boldsymbol{\sigma}' \in \Xi$. Hence, $\bar{\mathbf{a}}$ is not a local minimum of $\hat{\varphi}$.

Local minima If $\|\boldsymbol{\sigma}\|_0 = 1$ and $I = \{i\}$, then the simplified objective function is

$$\hat{\varphi}_{\boldsymbol{\sigma}} = -\frac{1}{2} \langle \sigma_i \boldsymbol{\iota}^* s_{-i} [\bar{\mathbf{a}}_0], \mathbf{a} \rangle^2 + \lambda \langle \sigma_i \boldsymbol{\iota}^* s_{-i} [\bar{\mathbf{a}}_0], \mathbf{a} \rangle + c_{\boldsymbol{\sigma}}. \quad (\text{B.45})$$

The minimizer appears at the boundary for $\langle \sigma_i \iota^* s_{-i} [\widehat{\mathbf{a}}_0], \mathbf{a} \rangle$, namely $\|\iota^* s_{-i} [\widehat{\mathbf{a}}_0]\|_2$ ¹² obtained by

$$\bar{\mathbf{a}} = \sigma_i \frac{\iota^* s_\tau [\widehat{\mathbf{a}}_0]}{\|\iota^* s_\tau [\widehat{\mathbf{a}}_0]\|_2}. \quad (\text{B.46})$$

It can be easily verified that

$$\begin{aligned} \text{grad} [\widehat{\varphi}_\sigma] (\bar{\mathbf{a}}) &= \left(-\|\iota^* s_\tau [\widehat{\mathbf{a}}_0]\|_2^2 + \lambda \|\iota^* s_\tau [\widehat{\mathbf{a}}_0]\|_2 \right) \times (\mathbf{I} - \bar{\mathbf{a}} \bar{\mathbf{a}}^*) \bar{\mathbf{a}} \\ &= \mathbf{0} \end{aligned} \quad (\text{B.47})$$

$$= \mathbf{0} \quad (\text{B.48})$$

$$\text{Hess} [\widehat{\varphi}_\sigma] (\bar{\mathbf{a}}) = \mathbf{P}_{\mathbf{a}_\perp} \left(-\|\iota^* s_\tau [\widehat{\mathbf{a}}_0]\|_2^2 \bar{\mathbf{a}} \bar{\mathbf{a}}^* + (1 - \lambda \|\iota^* s_\tau [\widehat{\mathbf{a}}_0]\|_2) \mathbf{I} \right) \mathbf{P}_{\mathbf{a}_\perp} \quad (\text{B.49})$$

$$= (1 - \lambda \|\iota^* s_\tau [\widehat{\mathbf{a}}_0]\|_2) \mathbf{P}_{\mathbf{a}_\perp} \mathbf{P}_{\mathbf{a}_\perp} \quad (\text{B.50})$$

$$\succeq 0 \quad (\text{B.51})$$

Global maxima If $\|\sigma\|_0 = 0$, then the objective remains constant and achieves the global maximum. \blacksquare

Lemma B.1 (Continuity of the Gradient of $\widehat{\varphi}$). $\nabla \widehat{\varphi}$ is a continuous function of \mathbf{a} .

Proof Recall that for a given σ , the gradient

$$\nabla \widehat{\varphi}_\sigma (\mathbf{a}) = -\iota^* \check{\mathbf{C}}_{\mathbf{a}_0} \mathbf{P}_I \check{\mathbf{C}}_{\mathbf{a}_0}^* \iota \mathbf{a} + \lambda \iota^* \check{\mathbf{C}}_{\mathbf{a}_0} \mathbf{P}_I \sigma \quad (\text{B.52})$$

$$= -\iota^* \check{\mathbf{C}}_{\mathbf{a}_0} \mathbf{P}_I (\check{\mathbf{C}}_{\mathbf{a}_0}^* \iota \mathbf{a} - \lambda \sigma) \quad (\text{B.53})$$

This is a continuous function within the relative interior of R_σ . Next, we show this function is continuous at the relative boundary of R_σ . Let $\mathbf{a}' \in \text{relbdy}(R_\sigma)$, and $\sigma' = \text{sign}(\mathbf{a}')$, $I = \text{supp}(\sigma')$ are the corresponding sign and support. Without loss of generality, we assume $\sigma' \trianglelefteq \sigma$, denote $\mathbf{a} = \mathbf{a}' + \varepsilon \delta$ ($\|\delta\|_2 = 1$) and $I_\delta = I \setminus I'$, then

$$\nabla \widehat{\varphi}_\sigma (\mathbf{a}) - \nabla \widehat{\varphi}_{\sigma'} (\mathbf{a}') = -\iota^* \check{\mathbf{C}}_{\mathbf{a}_0} \mathbf{P}_I (\check{\mathbf{C}}_{\mathbf{a}_0}^* \iota \mathbf{a} - \lambda \sigma) + \iota^* \check{\mathbf{C}}_{\mathbf{a}_0} \mathbf{P}_{I'} (\check{\mathbf{C}}_{\mathbf{a}_0}^* \iota \mathbf{a}' - \lambda \sigma') \quad (\text{B.54})$$

$$= -\iota^* \check{\mathbf{C}}_{\mathbf{a}_0} (\mathbf{P}_{I'} + \mathbf{P}_{I_\delta}) (\check{\mathbf{C}}_{\mathbf{a}_0}^* \iota \mathbf{a} - \lambda \sigma) + \iota^* \check{\mathbf{C}}_{\mathbf{a}_0} \mathbf{P}_{I'} (\check{\mathbf{C}}_{\mathbf{a}_0}^* \iota \mathbf{a}' - \lambda \sigma') \quad (\text{B.55})$$

$$= -\iota^* \check{\mathbf{C}}_{\mathbf{a}_0} \mathbf{P}_{I_\delta} (\check{\mathbf{C}}_{\mathbf{a}_0}^* \iota \mathbf{a} - \lambda \sigma) - \varepsilon \iota^* \check{\mathbf{C}}_{\mathbf{a}_0} \mathbf{P}_{I'} \check{\mathbf{C}}_{\mathbf{a}_0}^* \iota \delta \quad (\text{B.56})$$

Since $\|\mathbf{P}_{I_\delta} (\check{\mathbf{C}}_{\mathbf{a}_0}^* \iota \mathbf{a} - \lambda \sigma)\|_\infty = \varepsilon \|\mathbf{P}_{I_\delta} \check{\mathbf{C}}_{\mathbf{a}_0}^* \iota \delta\|_\infty$, we have $\|\nabla \widehat{\varphi}_\sigma (\mathbf{a}) - \nabla \widehat{\varphi}_{\sigma'} (\mathbf{a}')\|_\infty \leq \mathcal{O}(\varepsilon)$. \blacksquare

C Proof of Lemma 3.1

Lemma C.1. Let $\lambda_{rel} = \lambda / \|\mathbf{x}_0\|_\infty$, suppose the ground truth \mathbf{a}_0 satisfies

$$|\langle \mathbf{a}_0, \iota s_{\tau \neq 0} [\widehat{\mathbf{a}}_0] \rangle| < \lambda_{rel}^2 - (2 + 1/\lambda_{rel}^2) \sqrt{1 - \lambda_{rel}^2} \quad (\text{C.1})$$

for any nonzero shift τ , and \mathbf{x}_0 is separated enough such that any two nonzero components are at least $2k$ entries away from each other. If initialized at some $\mathbf{a} \in \mathbb{S}^{k-1}$ that $|\langle \mathbf{a}, \mathbf{a}_0 \rangle| > \lambda / \|\mathbf{x}_0\|_\infty$, a small-stepping projected gradient method minimizing $\varphi(\mathbf{a})$ recovers the signed ground truth $\pm \mathbf{a}_0$.

Proof Without loss of generality, we are going to assume $\|\mathbf{x}_0\|_\infty = 1$ for simplicity. Given that

$$|\langle \mathbf{a}_0, \iota s_{\tau \neq 0} [\widehat{\mathbf{a}}_0] \rangle| < \lambda^2 - \sqrt{1 - \lambda^2} (2 + 1/\lambda^2) \quad (\text{C.2})$$

and $\mathbf{a} = \langle \mathbf{a}, \mathbf{a}_0 \rangle \mathbf{a}_0 + \delta$ with $\|\delta\|_2 \leq \sqrt{1 - \lambda^2}$, therefore

$$|\langle \mathbf{a}, \iota s_\tau [\mathbf{a}] \rangle| = |\langle \langle \mathbf{a}, \mathbf{a}_0 \rangle \mathbf{a}_0 + \delta, \iota s_\tau [\langle \mathbf{a}, \mathbf{a}_0 \rangle \mathbf{a}_0 + \delta] \rangle|$$

¹²The other boundary point is $\langle \sigma_i \iota^* s_{-i} [\widehat{\mathbf{a}}_0], \mathbf{a} \rangle = \lambda$, which achieves a smaller objective value.

$$\begin{aligned}
&\leq \langle \mathbf{a}, \mathbf{a}_0 \rangle^2 |\langle \mathbf{a}_0, \iota_{s_\tau} [\mathbf{a}_0] \rangle| + 2 \langle \mathbf{a}, \mathbf{a}_0 \rangle \|\boldsymbol{\delta}\|_2 + \|\boldsymbol{\delta}\|_2^2 \\
&< 1 - \sqrt{1 - \lambda^2} / \lambda^2
\end{aligned}$$

Moreover, as \mathbf{x}_0 is sufficiently separated, we have

$$\begin{aligned}
|\langle \mathbf{a}, \iota_{s_\tau} [\mathbf{a}] \rangle \|\mathbf{x}^*\|_\infty - \langle \mathbf{a}, \iota_{s_\tau} [\mathbf{a}_0] \rangle \|\mathbf{x}_0\|_\infty| &\leq |\langle \mathbf{a}, \iota_{s_\tau} [\mathbf{a}] \rangle| \|\mathbf{x}_0 - \mathbf{x}^*\|_\infty + |\langle \mathbf{a}, \iota_{s_\tau} [\mathbf{a}_0 - \mathbf{a}] \rangle| \|\mathbf{x}_0\|_\infty \\
&< \lambda |\langle \mathbf{a}, \iota_{s_\tau} [\mathbf{a}] \rangle| + \|\mathbf{a}_0 - \mathbf{a}\|_2 \|\mathbf{x}_0\|_\infty \\
&< \lambda.
\end{aligned}$$

Hence, there exists a unique nonzero minimizer satisfying $\text{supp}(\mathbf{x}^*) \subset \text{supp}(\mathbf{x}_0)$, and the optimality condition for \mathbf{x}^* implies

$$\mathbf{x}^* = \text{SOFT}_\lambda [\langle \mathbf{a}, \mathbf{a}_0 \rangle \mathbf{x}_0], \quad (\text{C.3})$$

In this case, we can calculate the Euclidean gradient

$$\nabla \varphi(\mathbf{a}) = \iota_{\mathbf{C}_{\mathbf{x}^*}} (\mathbf{a} \otimes \mathbf{x}^* - \mathbf{a}_0 \otimes \mathbf{x}_0) \quad (\text{C.4})$$

$$= \|\mathbf{x}^*\|_2^2 \mathbf{a} - \langle \mathbf{x}^*, \mathbf{x}_0 \rangle \mathbf{a}_0, \quad (\text{C.5})$$

and the Riemannian gradient

$$\text{grad}[\varphi](\mathbf{a}) = (\mathbf{I} - \mathbf{a}\mathbf{a}^*) \nabla \varphi(\mathbf{a}) \quad (\text{C.6})$$

$$= -\langle \mathbf{x}^*, \mathbf{x}_0 \rangle (\mathbf{I} - \mathbf{a}\mathbf{a}^*) \mathbf{a}_0. \quad (\text{C.7})$$

It is easy to check that at any point along the geodesic curve between \mathbf{a}_0 and \mathbf{a} , support recovery of \mathbf{x}^* is achieved. A small-stepping gradient descent algorithm moves towards the signed ground truth $\pm \mathbf{a}_0$, as desired. \blacksquare

D Proof of Lemma 6.1

Lemma D.1. Suppose $\mathbf{y} = \mathbf{a}_0 \otimes \mathbf{x}_0$ with $\mathbf{x}_0 = \mathbf{e}_0$, and $p = q \geq 2$. Then for any shift τ , positive scalar ε and λ such that at every $\mathbf{a} \in \mathbb{S}_q \cap \mathbb{B}\left(\frac{\iota^*_{s_{-\tau}}[\mathbf{a}_0]}{\|\iota^*_{s_{-\tau}}[\mathbf{a}_0]\|_q}, \varepsilon\right)$ the solution $\mathbf{x}^*(\mathbf{a}) \doteq \arg \min_{\mathbf{x}} \psi_p(\mathbf{a}, \mathbf{x})$ is (i) unique and (ii) supported on the τ -th entry, the point $\bar{\mathbf{a}} \doteq \frac{\iota^*_{s_{-\tau}}[\mathbf{a}_0]}{\|\iota^*_{s_{-\tau}}[\mathbf{a}_0]\|_q}$ is a strict local minimizer of the cost $\varphi_p(\mathbf{a})$ over the manifold $\|\mathbf{a}\|_q = 1$.

Proof Write $\mathbf{x}^*(\mathbf{a}) = \alpha \mathbf{e}_\tau$ at point $\mathbf{a} \in \mathbb{S}_p \cap \mathbb{B}\left(\frac{\iota^*_{s_{-\tau}}[\mathbf{a}_0]}{\|\iota^*_{s_{-\tau}}[\mathbf{a}_0]\|_q}, \varepsilon\right)$, and suppose without loss of generality that $\alpha > 0$. Then using the KKT conditions, we have

$$\lambda = \left\langle s_\tau[\tilde{\mathbf{a}}], \text{diag sign}(\tilde{\mathbf{a}}_0 - \alpha s_\tau[\tilde{\mathbf{a}}]) |\tilde{\mathbf{a}}_0 - \alpha s_\tau[\tilde{\mathbf{a}}]|^{o(p-1)} \right\rangle \quad (\text{D.1})$$

$$= \left\langle \mathbf{a}, \text{diag sign}(\iota^*_{s_{-\tau}}[\tilde{\mathbf{a}}_0] - \alpha \mathbf{a}) |\iota^*_{s_{-\tau}}[\tilde{\mathbf{a}}_0] - \alpha \mathbf{a}|^{o(p-1)} \right\rangle. \quad (\text{D.2})$$

Magnitude of \mathbf{x}^* At point $\bar{\mathbf{a}} = \frac{\iota^*_{s_{-\tau}}[\mathbf{a}_0]}{\|\iota^*_{s_{-\tau}}[\mathbf{a}_0]\|_q}$, there exists closed form solution

$$\mathbf{x}^*(\bar{\mathbf{a}}) = \bar{\alpha} \mathbf{e}_\tau = \left(\|\iota^*_{s_{-\tau}}[\mathbf{a}_0]\|_q - \lambda^{\frac{1}{p-1}} \right) \mathbf{e}_\tau. \quad (\text{D.3})$$

As the shifted support recovery of $\mathbf{x}^*(\mathbf{a})$ holds at point $\mathbf{a} \in \mathbb{B}(\bar{\mathbf{a}}, \varepsilon)$, with implicit function theorem, we can obtain that α is a differentiable function of \mathbf{a} , and

$$\nabla_{\alpha} \alpha(\bar{\mathbf{a}}) = - \left(\|\iota^*_{s_{-\tau}}[\mathbf{a}_0]\|_q - 2\lambda^{\frac{1}{p-1}} \right) \text{sign } \bar{\alpha} \circ |\bar{\mathbf{a}}|^{o(p-1)}. \quad (\text{D.4})$$

Euclidean Gradient and Hessian The Euclidean gradient of φ_p at $\bar{\mathbf{a}}$ is

$$\nabla \varphi_p(\bar{\mathbf{a}}) = \nabla_{\mathbf{a}} \psi_p \Big|_{(\bar{\mathbf{a}}, \mathbf{x}^*(\bar{\mathbf{a}}))} \quad (\text{D.5})$$

$$= \boldsymbol{\iota}^* \mathbf{C}_{\mathbf{x}^*(\bar{\mathbf{a}})}^* \text{diag sign}(\bar{\mathbf{a}} \otimes \mathbf{x}^*(\bar{\mathbf{a}}) - \mathbf{y}) |\bar{\mathbf{a}} \otimes \mathbf{x}^*(\bar{\mathbf{a}}) - \mathbf{y}|^{\circ(p-1)} \quad (\text{D.6})$$

$$= \bar{\alpha} \boldsymbol{\iota}^* \text{diag sign}(s_{-\tau} [\bar{\alpha} s_{\tau} [\bar{\mathbf{a}}] - \widetilde{\mathbf{a}}_0]) |s_{-\tau} [\bar{\alpha} s_{\tau} [\bar{\mathbf{a}}] - \widetilde{\mathbf{a}}_0]|^{\circ(p-1)} \quad (\text{D.7})$$

$$= \bar{\alpha} \text{diag sign}(\bar{\alpha} \bar{\mathbf{a}} - \boldsymbol{\iota}^* s_{-\tau} [\widetilde{\mathbf{a}}_0]) |\bar{\alpha} \bar{\mathbf{a}} - \boldsymbol{\iota}^* s_{-\tau} [\widetilde{\mathbf{a}}_0]|^{\circ(p-1)}. \quad (\text{D.8})$$

And the Euclidean Hessian of φ_p at $\bar{\mathbf{a}}$ is

$$\begin{aligned} \nabla^2 \varphi_p(\bar{\mathbf{a}}) &= (p-1) \bar{\alpha} \text{diag} |\bar{\alpha} \bar{\mathbf{a}} - \boldsymbol{\iota}^* s_{-\tau} [\widetilde{\mathbf{a}}_0]|^{\circ(p-2)} \left(\bar{\alpha} \mathbf{I} + \bar{\mathbf{a}} (\nabla_{\mathbf{a}} \alpha)^T \right) \\ &\quad + \text{sign}(\bar{\alpha} \bar{\mathbf{a}} - \boldsymbol{\iota}^* s_{-\tau} [\widetilde{\mathbf{a}}_0]) \circ |\bar{\alpha} \bar{\mathbf{a}} - \boldsymbol{\iota}^* s_{-\tau} [\widetilde{\mathbf{a}}_0]|^{\circ(p-1)} (\nabla_{\mathbf{a}} \alpha)^T. \end{aligned} \quad (\text{D.9})$$

Riemannian Gradient and Hessian We call the optimization domain manifold \mathcal{M}_q where $\|\mathbf{a}\|_q = 1$ with $2 \leq q < \infty$, the Riemannian gradient at \mathbf{a} can be written as

$$\text{grad}[\varphi_p](\mathbf{a}) = \mathbb{P}_{\mathbf{a}^\perp} \nabla \varphi_p, \quad (\text{D.10})$$

with

$$\mathbb{P}_{\mathbf{a}^\perp}(\mathbf{v}) = \mathbf{v} - \left\langle \mathbf{v}, \frac{\text{sign } \mathbf{a} \circ |\mathbf{a}|^{\circ(q-1)}}{\| |\mathbf{a}|^{\circ(q-1)} \|_2} \right\rangle \frac{\text{sign } \mathbf{a} \circ |\mathbf{a}|^{\circ(q-1)}}{\| |\mathbf{a}|^{\circ(q-1)} \|_2}. \quad (\text{D.11})$$

Therefore, if $p = q$, $\bar{\mathbf{a}}$ is a stationary point with vanishing Riemannian gradient

$$\text{grad}[\varphi_p](\bar{\mathbf{a}}) = \mathbf{0}. \quad (\text{D.12})$$

The Riemannian Hessian $\text{Hess}[\varphi_p](\mathbf{a}) : T_{\mathbf{a}} \mathcal{M}_q \rightarrow T_{\mathbf{a}} \mathcal{M}_q$ is defined via the formula

$$\text{Hess}[\varphi_p](\mathbf{a}) \boldsymbol{\delta} = \tilde{\nabla}_{\boldsymbol{\delta}} \text{grad}[\varphi_p](\mathbf{a}) \quad (\boldsymbol{\delta} \in T_{\mathbf{a}} \mathcal{M}_q) \quad (\text{D.13})$$

$$= \mathbf{P}_{T_{\mathbf{a}} \mathcal{M}} \nabla_{\boldsymbol{\delta}} \text{grad}[\varphi_p](\mathbf{a}), \quad (\text{D.14})$$

where in the second line we have used Proposition 5.3.2 of the [AMS09]. Here, $\tilde{\nabla}$ is the Riemannian connection on \mathcal{M}_q , ∇ is the standard Euclidean connection (directional derivative) on \mathbb{R}^k and $\text{grad}[\varphi_p]$ is any smooth extension of the Riemannian gradient to \mathbb{R}^k . Plugging in and using the definition of Euclidean gradient ∇ , we have

$$\text{Hess}[\varphi_p](\mathbf{a}) \boldsymbol{\delta} = \mathbb{P}_{\mathbf{a}^\perp} \nabla_{\boldsymbol{\delta}} [\mathbb{P}_{\mathbf{a}^\perp} \nabla \varphi_p(\mathbf{a})] \quad (\text{D.15})$$

$$= \mathbb{P}_{\mathbf{a}^\perp} \frac{d}{dt} [\mathbb{P}_{(\mathbf{a}+t\boldsymbol{\delta})^\perp} \nabla \varphi_p(\mathbf{a} + t\boldsymbol{\delta})] \Big|_{t=0} \quad (\text{D.16})$$

$$= \mathbb{P}_{\mathbf{a}^\perp} \left(\nabla^2 \varphi_p(\mathbf{a}) \boldsymbol{\delta} + \frac{d}{dt} \mathbb{P}_{(\mathbf{a}+t\boldsymbol{\delta})^\perp} \Big|_{t=0} \nabla \varphi_p(\mathbf{a}) \right) \quad (\text{D.17})$$

For simplicity, we write

$$\boldsymbol{\xi}(\mathbf{a}) = \frac{\text{sign}(\mathbf{a}) \circ |\mathbf{a}|^{\circ(q-1)}}{\| |\mathbf{a}|^{\circ(q-1)} \|_2}, \quad (\text{D.18})$$

and

$$\mathbb{P}_{\mathbf{a}^\perp} = \mathbf{I} - \boldsymbol{\xi}(\mathbf{a}) \boldsymbol{\xi}(\mathbf{a})^*. \quad (\text{D.19})$$

Notice that the vector-valued function $\boldsymbol{\xi}(\mathbf{a})$ is differentiable away from $\mathbf{a} = \mathbf{0}$, and that

$$\frac{d}{dt} \mathbb{P}_{(\mathbf{a}+t\boldsymbol{\delta})^\perp} \Big|_{t=0} = -\frac{d}{dt} \boldsymbol{\xi}(\mathbf{a}+t\boldsymbol{\delta}) \Big|_{t=0} \boldsymbol{\xi}(\mathbf{a})^* - \boldsymbol{\xi}(\mathbf{a}) \left[\frac{d}{dt} \boldsymbol{\xi}(\mathbf{a}+t\boldsymbol{\delta}) \Big|_{t=0} \right]^*. \quad (\text{D.20})$$

Since that $\mathbb{P}_{\mathbf{a}^\perp} \boldsymbol{\xi}(\mathbf{a}) = \mathbf{0}$, we have

$$\text{Hess}[\varphi_p](\mathbf{a})\boldsymbol{\delta} = \mathbb{P}_{\mathbf{a}^\perp} \left(\nabla^2 \varphi_p(\mathbf{a})\boldsymbol{\delta} - \left[\frac{d}{dt} \Big|_{t=0} \boldsymbol{\xi}(\mathbf{a}+t\boldsymbol{\delta}) \right] \boldsymbol{\xi}(\mathbf{a})^* \nabla \varphi_p(\mathbf{a}) \right). \quad (\text{D.21})$$

Calculus gives

$$\frac{d}{dt} \Big|_{t=0} \boldsymbol{\xi}(\mathbf{a}+t\boldsymbol{\delta}) = (q-1) \frac{|\mathbf{a}|^{\circ(q-2)} \circ \boldsymbol{\delta}}{\| |\mathbf{a}|^{\circ(q-1)} \|_2} + \text{sign}(\mathbf{a}) \circ |\mathbf{a}|^{\circ(q-1)} \left\langle \frac{1}{|\mathbf{a}|^{\circ(q-1)}}, \boldsymbol{\delta} \right\rangle, \quad (\text{D.22})$$

hence

$$\mathbb{P}_{\mathbf{a}^\perp} \frac{d}{dt} \Big|_{t=0} \boldsymbol{\xi}(\mathbf{a}+t\boldsymbol{\delta}) = \mathbb{P}_{\mathbf{a}^\perp} \left[(q-1) \frac{|\mathbf{a}|^{\circ(q-2)} \circ \boldsymbol{\delta}}{\| |\mathbf{a}|^{\circ(q-1)} \|_2} \right]. \quad (\text{D.23})$$

At last, we can obtain the expression

$$\text{Hess}[\varphi_p](\mathbf{a})\boldsymbol{\delta} = \mathbb{P}_{\mathbf{a}^\perp} \left(\nabla^2 \varphi_p(\mathbf{a}) - (q-1) \frac{\langle \text{sign}(\mathbf{a}) \circ |\mathbf{a}|^{\circ(q-1)}, \nabla \varphi_p(\mathbf{a}) \rangle}{\| |\mathbf{a}|^{\circ(q-1)} \|_2^2} \text{diag}(|\mathbf{a}|^{\circ(q-2)}) \right) \boldsymbol{\delta}, \quad (\text{D.24})$$

for $\boldsymbol{\delta} \in T_{\mathbf{a}}\mathcal{M}_q$. Noting that for any such $\boldsymbol{\delta}$, $\mathbb{P}_{\mathbf{a}^\perp} \boldsymbol{\delta} = \boldsymbol{\delta}$, we can identify the Riemannian Hessian with the matrix

$$\text{Hess}[\varphi_p](\mathbf{a})\boldsymbol{\delta} = \mathbb{P}_{\mathbf{a}^\perp} \left(\nabla^2 \varphi_p(\mathbf{a}) - (q-1) \frac{\langle \text{sign}(\mathbf{a}) \circ |\mathbf{a}|^{\circ(q-1)}, \nabla \varphi_p(\mathbf{a}) \rangle}{\| |\mathbf{a}|^{\circ(q-1)} \|_2^2} \text{diag}(|\mathbf{a}|^{\circ(q-2)}) \right) \mathbb{P}_{\mathbf{a}^\perp}. \quad (\text{D.25})$$

Local Convexity Plugging in $\bar{\mathbf{a}} = \frac{\boldsymbol{\iota}^* s_{-\tau}[\mathbf{a}_0]}{\| \boldsymbol{\iota}^* s_{-\tau}[\mathbf{a}_0] \|_q}$ with $p = q$, we have

$$\text{Hess}[\varphi_p] \left(\frac{\boldsymbol{\iota}^* s_{-\tau}[\mathbf{a}_0]}{\| \boldsymbol{\iota}^* s_{-\tau}[\mathbf{a}_0] \|_q} \right) = \mathbb{P}_{\bar{\mathbf{a}}^\perp} \left[\nabla^2 \varphi_p + (q-1) \bar{\alpha} \lambda \text{diag}(|\bar{\mathbf{a}}|^{\circ(q-2)}) \right] \mathbb{P}_{\bar{\mathbf{a}}^\perp} \quad (\text{D.26})$$

$$= \mathbb{P}_{\bar{\mathbf{a}}^\perp} \left[(q-1) \bar{\alpha} \left(\lambda^{\frac{p-2}{p-1}} + \lambda \right) \text{diag} |\bar{\mathbf{a}}|^{\circ(q-2)} \right] \mathbb{P}_{\bar{\mathbf{a}}^\perp}. \quad (\text{D.27})$$

Hence, the function is strictly convex at point $\bar{\mathbf{a}} = \frac{\boldsymbol{\iota}^* s_{-\tau}[\mathbf{a}_0]}{\| \boldsymbol{\iota}^* s_{-\tau}[\mathbf{a}_0] \|_q}$ along direction $\boldsymbol{\delta} \in T_{\bar{\mathbf{a}}}\mathcal{M}_q$ as long as whose support are contained in the support of \mathbf{a} . We further consider a perturbation $\boldsymbol{\delta}_c$ such that $\text{supp}(\boldsymbol{\delta}_c) \cap \text{supp}(\bar{\mathbf{a}}) = \emptyset$, and $\| \boldsymbol{\delta}_c \|_2 \leq \varepsilon$, then the objective value at $\frac{\bar{\mathbf{a}} + \boldsymbol{\delta}_c}{\| \bar{\mathbf{a}} + \boldsymbol{\delta}_c \|_q}$ satisfies

$$\varphi_p \left(\frac{\bar{\mathbf{a}} + \boldsymbol{\delta}_c}{\| \bar{\mathbf{a}} + \boldsymbol{\delta}_c \|_q} \right) = \frac{1}{p} \left\| \mathbf{y} - \frac{\bar{\mathbf{a}} + \boldsymbol{\delta}_c}{\| \bar{\mathbf{a}} + \boldsymbol{\delta}_c \|_q} \circledast \mathbf{x}^*(\bar{\mathbf{a}} + \boldsymbol{\delta}_c) \right\|_p^p + \lambda \| \mathbf{x}^*(\bar{\mathbf{a}} + \boldsymbol{\delta}_c) \|_1 \quad (\text{D.28})$$

$$\geq \frac{1}{p} \left\| \mathbf{y} - \frac{\bar{\mathbf{a}}}{\| \bar{\mathbf{a}} + \boldsymbol{\delta}_c \|_q} \circledast \mathbf{x}^*(\bar{\mathbf{a}} + \boldsymbol{\delta}_c) \right\|_p^p + \lambda \| \mathbf{x}^*(\bar{\mathbf{a}} + \boldsymbol{\delta}_c) \|_1 \quad (\text{D.29})$$

$$\geq \varphi_p \left(\frac{\bar{\mathbf{a}}}{\| \bar{\mathbf{a}} + \boldsymbol{\delta}_c \|_q} \right) \quad (\text{D.30})$$

$$\geq \varphi_p(\bar{\mathbf{a}}). \quad (\text{D.31})$$

This completes the proof. ■

E Auxiliary Lemmas

Lemma E.1. (Lemma B.2 of [SQW17a]) Suppose that $\mathbf{A} \succ \mathbf{0}$ is a positive definite matrix. For any symmetric matrix Δ with $\|\Delta\|_{\ell^2 \rightarrow \ell^2} \leq \sigma_{\min}(\mathbf{A})/2$,

$$\left\| (\mathbf{A} + \Delta)^{-1/2} - \mathbf{A}^{-1/2} \right\|_{\ell^2 \rightarrow \ell^2} \leq \frac{2 \|\mathbf{A}\|_{\ell^2 \rightarrow \ell^2}^{1/2} \|\Delta\|_{\ell^2 \rightarrow \ell^2}}{\sigma_{\min}(\mathbf{A})^2}. \quad (\text{E.1})$$

Lemma E.2. Let \mathbf{U} be a matrix such that $\mathbf{U}^* \mathbf{U} = \mathbf{I} + \Delta$, with $\|\Delta\|_{\ell^2 \rightarrow \ell^2} < 1/2$. Then $\mathbf{U}^* \mathbf{U}$ is invertible,

$$\|\mathbf{U}\|_{\ell^2 \rightarrow \ell^2} < \sqrt{3/2}, \quad (\text{E.2})$$

and

$$\left\| \mathbf{U} - \mathbf{U}(\mathbf{U}^* \mathbf{U})^{-1/2} \right\|_{\ell^2 \rightarrow \ell^2} < 3 \|\Delta\|_{\ell^2 \rightarrow \ell^2}. \quad (\text{E.3})$$

Proof Upper bound for the first quantity can be derived

$$\|\mathbf{U}\|_{\ell^2 \rightarrow \ell^2} = \sqrt{\|\mathbf{U}^* \mathbf{U}\|_{\ell^2 \rightarrow \ell^2}} \quad (\text{E.4})$$

$$\leq \sqrt{\|\mathbf{I}\|_{\ell^2 \rightarrow \ell^2} + \|\Delta\|_{\ell^2 \rightarrow \ell^2}} \quad (\text{E.5})$$

$$< \sqrt{3/2}. \quad (\text{E.6})$$

Applying Lemma E.1 for the second term

$$\left\| \mathbf{U} - \mathbf{U}(\mathbf{U}^* \mathbf{U})^{-1/2} \right\|_{\ell^2 \rightarrow \ell^2} \leq \|\mathbf{U}\|_{\ell^2 \rightarrow \ell^2} \left\| \mathbf{I} - (\mathbf{U}^* \mathbf{U})^{-1/2} \right\|_{\ell^2 \rightarrow \ell^2} \quad (\text{E.7})$$

$$= \|\mathbf{U}\|_{\ell^2 \rightarrow \ell^2} \left\| \mathbf{I}^{-1/2} - (\mathbf{I} + \Delta)^{-1/2} \right\|_{\ell^2 \rightarrow \ell^2} \quad (\text{E.8})$$

$$\leq \sqrt{3/2} \times 2 \|\Delta\|_{\ell^2 \rightarrow \ell^2}, \quad (\text{E.9})$$

Hence, we can obtain the claim by using $2\sqrt{3/2} < 3$ to simplify the constant. ■