

Signal Recovery on Incoherent Manifolds

Chinmay Hegde and Richard G. Baraniuk *

Department of Electrical and Computer Engineering
Rice University

Abstract

Suppose that we observe noisy linear measurements of an unknown signal that can be modeled as the sum of two component signals, each of which arises from a nonlinear sub-manifold of a high-dimensional ambient space. We introduce *SPIN*, a first-order projected gradient method to recover the signal components. Despite the nonconvex nature of the recovery problem and the possibility of underdetermined measurements, *SPIN provably* recovers the signal components, provided that the signal manifolds are *incoherent* and that the measurement operator satisfies a certain *restricted isometry property*. *SPIN* significantly extends the scope of current recovery models and algorithms for low-dimensional linear inverse problems and matches (or exceeds) the current state of the art in terms of performance.

1 Introduction

Estimation of an unknown signal from linear observations is a core problem in signal processing, statistics, and information theory. Particular energy has been invested in problem instances where the available information is *limited and noisy* and where the signals of interest possess a *low-dimensional* geometric structure. Indeed, focused efforts on certain instances of the linear inverse problem framework have spawned entire research subfields, encompassing both theoretical and algorithmic advances. Examples include signal separation and morphological component analysis [1, 2]; sparse approximation and compressive sensing [3–5]; affine rank minimization [6]; and robust principal component analysis [7, 8].

In this work, we study a very general version of the linear inverse problem. Suppose that the signal of interest \mathbf{x}^* can be written as the sum of two constituent signals $\mathbf{a}^* \in \mathcal{A}$ and $\mathbf{b}^* \in \mathcal{B}$, where \mathcal{A}, \mathcal{B} are nonlinear sub-manifolds of the signal space \mathbb{R}^N . We are given access to noisy linear measurements of \mathbf{x}^* :

$$\mathbf{z} = \Phi(\mathbf{a}^* + \mathbf{b}^*) + \mathbf{e}, \quad (1)$$

where $\Phi \in \mathbb{R}^{M \times N}$ is the measurement matrix. Our objective is to recover the pair of signals $(\mathbf{a}^*, \mathbf{b}^*)$, and thus also \mathbf{x}^* , from \mathbf{z} . At the outset, numerous obstacles arise while trying to solve (1), some of which appear to be insurmountable:

*Email: {chinmay, richb}@rice.edu. Web: dsp.rice.edu/cs. This work was supported by the grants NSF CCF-0431150, CCF-0926127, and CCF-1117939; DARPA/ONR N66001-11-C-4092 and N66001-11-1-4090; ONR N00014-08-1-1112, N00014-10-1-0989, and N00014-11-1-0714; AFOSR FA9550-09-1-0432; ARO MURI W911NF-07-1-0185 and W911NF-09-1-0383; and the TI Leadership University Program. Thanks to Christoph Studer for valuable comments on an early draft of the manuscript.

1. (Identifiability I) Consider the simplest case, where the measurements are noiseless and the measurement operator is the identity, i.e., we observe $\mathbf{x} \in \mathbb{R}^N$ such that

$$\mathbf{x} = \mathbf{a}^* + \mathbf{b}^*, \quad (2)$$

where $\mathbf{a}^* \in \mathcal{A}, \mathbf{b}^* \in \mathcal{B}$. This expression for \mathbf{x} contains $2N$ unknowns but only N observations and hence is fundamentally ill-posed. Unless we make additional assumptions on the geometric structure of the component manifolds \mathcal{A} and \mathcal{B} , a unique decomposition of \mathbf{x} into its constituent signals $(\mathbf{a}^*, \mathbf{b}^*)$ may not exist.

2. (Identifiability II) To complicate matters, in more general situations the linear operator Φ in (1) might have fewer rows than columns, so that $M < N$. Thus, Φ possesses a nontrivial nullspace. Indeed, we are particularly interested in cases where $M \ll N$, in which case the nullspace of Φ is extremely large relative to the ambient space. This further obscures the issue of identifiability of the ordered pair $(\mathbf{a}^*, \mathbf{b}^*)$, given the available observations \mathbf{z} .
3. (Nonconvexity) Even if the above two identifiability issues were resolved, the manifolds \mathcal{A}, \mathcal{B} might be extremely nonconvex, or even non-differentiable. Thus, classical numerical methods, such as Newton's method or steepest descent, cannot be successfully applied; neither can the litany of convex optimization methods that have been specially designed for linear inverse problems with certain types of signal priors [1, 6].

In this paper, we propose a simple method to recover the component signals $(\mathbf{a}^*, \mathbf{b}^*)$ from \mathbf{z} in (1). We dub our method *Successive Projections onto INcoherent manifolds* (SPIN) (see Algorithm 1). Despite the highly nonconvex nature of the problem and the possibility of underdetermined measurements, SPIN *provably* recovers the signal components $(\mathbf{a}^*, \mathbf{b}^*)$. For this to hold true, we will require that (i) the signal manifolds \mathcal{A}, \mathcal{B} are *incoherent* in the sense that the secants of \mathcal{A} are almost orthogonal to the secants of \mathcal{B} ; and (ii) the measurement operator Φ satisfies a certain *restricted isometry property* (RIP) on the secants of the direct sum manifold $\mathcal{C} = \mathcal{A} \oplus \mathcal{B}$. We will formally define these conditions in Section 2. We prove the following theoretical statement below in Section 3.

Theorem 1 (Signal recovery) *Let \mathcal{A}, \mathcal{B} be incoherent manifolds in \mathbb{R}^N . Let Φ be a measurement matrix that satisfies the RIP on the direct sum $\mathcal{C} = \mathcal{A} \oplus \mathcal{B}$. Suppose we observe linear measurements $\mathbf{z} = \Phi(\mathbf{a}^* + \mathbf{b}^*)$, where $\mathbf{a}^* \in \mathcal{A}$ and $\mathbf{b}^* \in \mathcal{B}$. Then, given any precision parameter $\nu > 0$, there exists an iterative algorithm that outputs a sequence of iterates $(\mathbf{a}_k, \mathbf{b}_k) \in \mathcal{A} \times \mathcal{B}, k = 1, 2, \dots$ such that $\max\{\|\mathbf{a}_k - \mathbf{a}^*\|, \|\mathbf{b}_k - \mathbf{b}^*\|\} \leq C\nu$ for all $k > T$. Here C and T are global positive constants.*

Our proposed algorithm (SPIN) is iterative in nature. Each iteration consists of three steps: computation of the gradient of the error function $\psi(\mathbf{a}, \mathbf{b}) = \frac{1}{2} \|\mathbf{z} - \Phi(\mathbf{a} + \mathbf{b})\|^2$, forming signal proxies for \mathbf{a} and \mathbf{b} , and orthogonally projecting the proxies onto the manifolds \mathcal{A} and \mathcal{B} . The projection operators onto the component manifolds play a crucial role in algorithm stability and performance; some manifolds admit stable, efficient projection operators while others do not. We discuss this in detail in Section 3. Additionally, we demonstrate that our proposed algorithm is stable to measurement noise (the quantity \mathbf{e} in (1)), as well as numerical inaccuracies (such as finite precision arithmetic). Interestingly, SPIN exhibits a convergence rate comparable to several state-of-the-art algorithms [9, 10], despite the nonlinear and nonconvex nature of the reconstruction problem.

The core essence of our proposed approach has been extensively studied in a number of different contexts. Methods such as Projected Landweber iterations [11], iterative hard thresholding (IHT) [9], and singular value projection (SVP) [10] are all instances of the same basic framework. SPIN subsumes and generalizes these methods. In particular, it is an iterative projected gradient method with the same basic approach as two recent signal recovery algorithms — Gradient Descent with Sparsification (GraDeS) [12], and Manifold Iterative Pursuit (MIP) [13]. We generalize these approaches to situations where the signal of interest is a linear mixture of signals arising from a pair of nonlinear manifolds. Due to this particular structure of our setting, SPIN consists of *two* projection steps (instead of one), and the analysis is more involved (see Section 4). We also explore the interplay between the geometric structure of the component manifolds, the linear measurement operator, and the stability of the recovery algorithm.

SPIN exhibits a strong geometric convergence rate comparable to many state-of-the-art first-order methods. We duly note that for the case of certain special manifolds, sophisticated higher-order methods with stronger stability guarantees have been proposed (e.g., approximate message passing (AMP) [14] for sparse signal recovery and augmented Lagrangian multiplier (ALM) methods for low-rank matrix recovery [15]). However, an appealing feature of SPIN is its conceptual simplicity plus its ability to generalize to mixtures of arbitrary nonlinear manifolds, provided these manifolds satisfy certain geometric assumptions.

2 Geometric Assumptions

The analysis and proof of accuracy of SPIN (Algorithm 1) requires three core ingredients: (i) a geometric notion of *manifold incoherence* that crystallizes the approximate orthogonality between secants of submanifolds of \mathbb{R}^N ; (ii) a *restricted isometry* condition satisfied by the measurement operator Φ over the secants of a submanifold; and (iii) the availability of *projection operators* that compute the orthogonal projection of any point $x \in \mathbb{R}^N$ onto a submanifold of \mathbb{R}^N .

2.1 Manifold incoherence

In linear inverse problems such as sparse signal approximation and compressive sensing, the assumption of incoherence between linear subspaces, bases, or dictionary elements is common. We introduce a nonlinear generalization of this concept.

Definition 1 *Given a manifold $\mathcal{A} \subset \mathbb{R}^N$, a normalized secant, or simply, a secant, $\mathbf{u} \in \mathbb{R}^N$ of \mathcal{A} is a unit vector such that*

$$\mathbf{u} = \frac{\mathbf{a} - \mathbf{a}'}{\|\mathbf{a} - \mathbf{a}'\|}, \quad \mathbf{a}, \mathbf{a}' \in \mathcal{A}, \mathbf{a} \neq \mathbf{a}'.$$

The secant manifold $\mathcal{S}(\mathcal{A})$ is the family of all unit vectors \mathbf{u} generated by pairs \mathbf{a}, \mathbf{a}' belonging to \mathcal{A} .

Definition 2 *Suppose \mathcal{A}, \mathcal{B} are submanifolds of \mathbb{R}^N . Let*

$$\sup_{\mathbf{u} \in \mathcal{S}(\mathcal{A}), \mathbf{u}' \in \mathcal{S}(\mathcal{B})} |\langle \mathbf{u}, \mathbf{u}' \rangle| = \epsilon, \tag{3}$$

where $\mathcal{S}(\mathcal{A}), \mathcal{S}(\mathcal{B})$ are the secant manifolds of \mathcal{A}, \mathcal{B} respectively. Then, \mathcal{A} and \mathcal{B} are called ϵ -incoherent manifolds.

By definition, the quantity ϵ is always positive; further, $\epsilon \leq 1$, due to the Cauchy-Schwartz inequality. We prove that any signal \mathbf{x} belonging to the direct sum $\mathcal{A} \oplus \mathcal{B}$ can be *uniquely* decomposed into its constituent signals when the upper bound on ϵ holds with strict inequality.

Lemma 1 (Uniqueness) *Suppose that \mathcal{A}, \mathcal{B} are ϵ -incoherent with $0 < \epsilon < 1$. Consider $\mathbf{x} = \mathbf{a} + \mathbf{b} = \mathbf{a}' + \mathbf{b}'$, where $\mathbf{a}, \mathbf{a}' \in \mathcal{A}$ and $\mathbf{b}, \mathbf{b}' \in \mathcal{B}$. Then, $\mathbf{a} = \mathbf{a}'$, $\mathbf{b} = \mathbf{b}'$.*

Proof. It is clear that $\|\mathbf{a} + \mathbf{b} - (\mathbf{a}' + \mathbf{b}')\|^2 = 0$, i.e.,

$$\|\mathbf{a} - \mathbf{a}'\|^2 + \|\mathbf{b} - \mathbf{b}'\|^2 = -2\langle \mathbf{a} - \mathbf{a}', \mathbf{b} - \mathbf{b}' \rangle \leq 2|\langle \mathbf{a} - \mathbf{a}', \mathbf{b} - \mathbf{b}' \rangle|.$$

However, due to the manifold incoherence assumption, the (unnormalized) secants $\mathbf{a} - \mathbf{a}'$, $\mathbf{b} - \mathbf{b}'$ obey the relation:

$$|\langle \mathbf{a} - \mathbf{a}', \mathbf{b} - \mathbf{b}' \rangle| \leq \epsilon \|\mathbf{a} - \mathbf{a}'\| \|\mathbf{b} - \mathbf{b}'\| \leq \frac{1}{2}\epsilon(\|\mathbf{a} - \mathbf{a}'\|^2 + \|\mathbf{b} - \mathbf{b}'\|^2), \quad (4)$$

where the last inequality follows from the relation between arithmetic and geometric means (henceforth referred to as the *AM-GM inequality*). Therefore, we have that

$$\|\mathbf{a} - \mathbf{a}'\|^2 + \|\mathbf{b} - \mathbf{b}'\|^2 \leq \epsilon \left(\|\mathbf{a} - \mathbf{a}'\|^2 + \|\mathbf{b} - \mathbf{b}'\|^2 \right),$$

for $\epsilon < 1$, which is impossible unless $\mathbf{a} = \mathbf{a}'$, $\mathbf{b} = \mathbf{b}'$. \square

We can also prove the following relation between secants and direct sums of signals lying on incoherent manifolds.

Lemma 2 *Suppose that \mathcal{A}, \mathcal{B} are ϵ -incoherent with $0 < \epsilon < 1$. Consider $\mathbf{x}_1 = \mathbf{a}_1 + \mathbf{b}_1$, $\mathbf{x}_2 = \mathbf{a}_2 + \mathbf{b}_2$ where $\mathbf{a}_1, \mathbf{a}_2 \in \mathcal{A}$ and $\mathbf{b}_1, \mathbf{b}_2 \in \mathcal{B}$. Then*

$$|\langle \mathbf{a}_1 - \mathbf{a}_2, \mathbf{b}_1 - \mathbf{b}_2 \rangle| \leq \frac{\epsilon}{2(1-\epsilon)} \|\mathbf{x}_1 - \mathbf{x}_2\|^2.$$

Proof. From (4), we have

$$\begin{aligned} |\langle \mathbf{a}_1 - \mathbf{a}_2, \mathbf{b}_1 - \mathbf{b}_2 \rangle| &\leq \frac{\epsilon}{2} (\|\mathbf{a}_1 - \mathbf{a}_2\|^2 + \|\mathbf{b}_1 - \mathbf{b}_2\|^2) \\ &= \frac{\epsilon}{2} \|\mathbf{a}_1 + \mathbf{b}_1 - \mathbf{a}_2 - \mathbf{b}_2\|^2 - \epsilon \langle \mathbf{a}_1 - \mathbf{a}_2, \mathbf{b}_1 - \mathbf{b}_2 \rangle \\ &\leq \frac{\epsilon}{2} \|\mathbf{x}_1 - \mathbf{x}_2\|^2 + \epsilon |\langle \mathbf{a}_1 - \mathbf{a}_2, \mathbf{b}_1 - \mathbf{b}_2 \rangle|. \end{aligned}$$

Rearranging terms, we obtain the desired result. \square

2.2 Restricted isometry

Next, we address the situation where the measurement operator $\Phi \in \mathbb{R}^{M \times N}$ contains a nontrivial nullspace, i.e., when $M < N$. We will require that Φ satisfies a *restricted isometry* criterion on the secants of the direct sum manifold $\mathcal{C} = \mathcal{A} \oplus \mathcal{B}$.

Definition 3 *Let \mathcal{C} be a submanifold of \mathbb{R}^N . Then, the matrix $\Phi \in \mathbb{R}^{M \times N}$ satisfies the restricted isometry property (RIP) on \mathcal{C} with constant $\delta \in [0, 1)$, if for every normalized secant \mathbf{u} belonging to the secant manifold $\mathcal{S}(\mathcal{C})$, we have that*

$$1 - \delta \leq \|\Phi \mathbf{u}\|^2 \leq 1 + \delta. \quad (5)$$

The notion of restricted isometry (and its generalizations) is an important component in the analysis of many algorithms in sparse approximation, compressive sensing, and low-rank matrix recovery [4, 6]. While the RIP has traditionally been studied in the context of sparse signal models, (5) generalizes of this notion to *arbitrary* nonlinear manifolds. The restricted isometry condition is of particular interest when the range space of the matrix Φ is low-dimensional. A key result [16] states that, under certain upper bounds on the curvature of the manifold \mathcal{C} , there exist probabilistic constructions of matrices Φ that satisfy the RIP on \mathcal{C} such that the number of rows of Φ is proportional to the intrinsic dimension of \mathcal{C} , rather than the ambient dimension N of the signal space. We will discuss this further in Section 5.

2.3 Projections onto manifolds

Given an arbitrary nonlinear manifold $\mathcal{A} \in \mathbb{R}^N$, we define the operator $\mathcal{P}_{\mathcal{A}}(\cdot)$ as the Euclidean projection projector onto \mathcal{A} :

$$\mathcal{P}_{\mathcal{A}}(\mathbf{x}) = \arg \min_{\mathbf{x}' \in \mathcal{A}} \|\mathbf{x}' - \mathbf{x}\|^2. \quad (6)$$

We observe that for arbitrary nonconvex manifolds \mathcal{A} , the above minimization problem (6) may not yield a unique optimum. Technically, therefore, $\mathcal{P}_{\mathcal{A}}(\mathbf{x})$ is a set-valued operator. For ease of exposition, $\mathcal{P}_{\mathcal{A}}(\mathbf{x})$ will henceforth refer to *any* arbitrarily chosen element of the set of signals that minimize the ℓ_2 -error in (6).

The projection operator $\mathcal{P}_{\mathcal{A}}$ plays a crucial role in the development of our proposed signal recovery algorithm in Section 3. Note that in the case of general nonlinear manifolds, $\mathcal{P}_{\mathcal{A}}$ may be quite difficult to compute exactly. Therefore, following the lead of [13], we also define a γ -approximate projection operator onto \mathcal{A} :

$$\mathbf{x}' = \mathcal{P}_{\mathcal{A}}^{\gamma}(\mathbf{x}) \implies \mathbf{x}' \in \mathcal{A}, \text{ and } \|\mathbf{x}' - \mathbf{x}\|^2 \leq \|\mathcal{P}_{\mathcal{A}}(\mathbf{x}) - \mathbf{x}\|^2 + \gamma, \quad (7)$$

so that $\mathcal{P}_{\mathcal{A}}^{\gamma}(\mathbf{x})$ yields a vector $\mathbf{x}' \in \mathcal{A}$ that approximately minimizes the squared distance from \mathbf{x} to \mathcal{A} . Again, $\mathcal{P}_{\mathcal{A}}^{\gamma}(\mathbf{x})$ need not be uniquely defined for a particular \mathbf{x} .

Certain specific instances of nonconvex manifolds do admit efficient exact projection operators. For example, consider the space of all K -sparse signals of length N ; this can be viewed as the union of the $\binom{N}{K}$ canonical subspaces in \mathbb{R}^N or, alternately, a K -dimensional submanifold of \mathbb{R}^N . Then, the projection of an arbitrary vector $\mathbf{x} \in \mathbb{R}^N$ onto this manifold is merely the best K -sparse approximation to \mathbf{x} , which can be very efficiently computed via simple thresholding. We discuss additional examples in Section 5.

3 The SPIN Algorithm

We now describe an algorithm to solve the linear inverse problem (1). Our proposed algorithm, *Successive Projections onto INcoherent manifolds* (SPIN), can be viewed as a generalization of several first order methods for signal recovery for a variety of different models [9, 10, 13]. SPIN is described in pseudocode form in Algorithm 1.

The key innovation in SPIN is that we formulate *two* proxy vectors for the signal components $\tilde{\mathbf{a}}_k$ and $\tilde{\mathbf{b}}_k$ and project these onto the corresponding manifolds \mathcal{A} and \mathcal{B} . We demonstrate that SPIN possesses strong *uniform recovery* guarantees comparable to existing state-of-the-art algorithms for sparse approximation and compressive sensing, while encompassing a very broad range

Algorithm 1 Successive Projections onto Incoherent Manifolds (SPIN)

Inputs: Observation matrix Φ , measurements \mathbf{z} , projection operators $\mathcal{P}_{\mathcal{A}}(\cdot), \mathcal{P}_{\mathcal{B}}(\cdot)$,
number of iterations T , step size η

Outputs: Estimated signal components $\hat{\mathbf{a}} \in \mathcal{A}, \hat{\mathbf{b}} \in \mathcal{B}$

Initialize: $\mathbf{a}_0 = 0, \mathbf{b}_0 = 0, \mathbf{r} = \mathbf{z}, k = 0$

while $k \leq T$ **do**

$\mathbf{g}_k \leftarrow \Phi^T \mathbf{r}$	{form gradient}
$\tilde{\mathbf{a}}_k \leftarrow \mathbf{a}_k + \eta \mathbf{g}_k, \tilde{\mathbf{b}}_k \leftarrow \mathbf{b}_k + \eta \mathbf{g}_k$	{form signal proxies}
$\mathbf{a}_{k+1} \leftarrow \mathcal{P}_{\mathcal{A}}(\tilde{\mathbf{a}}_k), \mathbf{b}_{k+1} \leftarrow \mathcal{P}_{\mathcal{B}}(\tilde{\mathbf{b}}_k)$	{apply projection operators}
$\mathbf{r} \leftarrow \mathbf{z} - \Phi(\mathbf{a}_{k+1} + \mathbf{b}_{k+1})$	{update residual}
$k \leftarrow k + 1$	

end while

return $(\hat{\mathbf{a}}, \hat{\mathbf{b}}) \leftarrow (\mathbf{a}_T, \mathbf{b}_T)$

of nonlinear signal models. The following theoretical result describes the performance of SPIN for signal recovery.

Theorem 2 (Main result) *Suppose \mathcal{A}, \mathcal{B} are ϵ -incoherent manifolds in \mathbb{R}^N . Let Φ be a measurement matrix with restricted isometry constant δ over the direct sum manifold $\mathcal{C} = \mathcal{A} \oplus \mathcal{B}$. Suppose we observe noisy linear measurements $\mathbf{z} = \Phi(\mathbf{a}^* + \mathbf{b}^*) + \mathbf{e}$, where $\mathbf{a}^* \in \mathcal{A}$ and $\mathbf{b}^* \in \mathcal{B}$. If*

$$0 \leq \delta < \frac{1 - 9\epsilon}{3 + 5\epsilon}, \quad (8)$$

then SPIN (Algorithm 1) with step size $\eta = 1/(1 + \delta)$ with exact projections $\mathcal{P}_{\mathcal{A}}, \mathcal{P}_{\mathcal{B}}$ outputs $\mathbf{a}_T \in \mathcal{A}$ and $\mathbf{b}_T \in \mathcal{B}$, such that $\|\mathbf{z} - \Phi(\mathbf{a}_T + \mathbf{b}_T)\|^2 \leq \beta \|\mathbf{e}\|^2 + \nu$ in no more than $T = \lceil \frac{1}{\log(1/\alpha)} \log \frac{\|\mathbf{z}\|^2}{2\nu} \rceil$ iterations for any $\nu > 0$. Here, α, β are moderately sized positive constants.

For the special case when there is no measurement noise (i.e., $\mathbf{e} = 0$), Theorem 2 states that after a finite number of iterations, SPIN outputs signal component estimates $(\hat{\mathbf{a}}, \hat{\mathbf{b}})$ such that $\|\mathbf{z} - \Phi(\hat{\mathbf{a}} + \hat{\mathbf{b}})\| < \nu$ for any desired precision parameter ν . From the restricted isometry assumption on Φ and Lemma 1, we immediately obtain Theorem 1. Since we can set ν to an arbitrarily small value, we have that the SPIN estimate $(\hat{\mathbf{a}}, \hat{\mathbf{b}})$ converges to the true signal pair $(\mathbf{a}^*, \mathbf{b}^*)$. Exact convergence of the algorithm might potentially take a very large number of iterations, but convergence to any desired positive precision constant β takes only a finite number of iterations. For the rest of the paper, we will informally denote signal “recovery” to imply convergence to a sufficiently fine precision.

SPIN assumes the availability of the exact projection operators $\mathcal{P}_{\mathcal{A}}, \mathcal{P}_{\mathcal{B}}$. In certain cases, it might be feasible to numerically compute only γ -approximate projections, as in (7). In this case, the bound on the norm of the error $\mathbf{z} - \Phi(\mathbf{a}_T + \mathbf{b}_T)$ is only guaranteed to be upper bounded by a positive multiple of the approximation parameter γ . The following theoretical guarantee (with a near-identical proof mechanism as Theorem 2) captures this behavior.

Theorem 3 (Approximate projections) *Under the same suppositions of Theorem 2, SPIN (Algorithm 1) with γ -approximate projections and step size $\eta = 1/(1 + \delta)$ outputs $\mathbf{a}_T \in \mathcal{A}$ and $\mathbf{b}_T \in \mathcal{B}$ such that $\|\mathbf{z} - \Phi(\mathbf{a}_T + \mathbf{b}_T)\|^2 \leq \beta \|\mathbf{e}\|^2 + \frac{1+\delta}{1-\alpha} \gamma + \nu$, in no more than $T = \lceil \frac{1}{\log(1/\alpha)} \log \frac{\|\mathbf{z}\|^2}{2\nu} \rceil$ iterations.*

We note some implications of Theorem 2. First, suppose that Φ is the identity operator, i.e., we have full measurements of the signal $\mathbf{x}^* = \mathbf{a}^* + \mathbf{b}^*$. Then, $\delta = 0$ and the lower bound on the restricted isometry constant holds with equality. However, we still require that $\epsilon < 1/9$ for guaranteed recovery using SPIN. We will discuss this condition further in Section 5.

Second, suppose that the one of the components manifolds is the trivial (zero) manifold; then, we have that $\epsilon = 0$. In this case, SPIN reduces to the Manifold Iterative Pursuit (MIP) algorithm for recovering signals from a single manifold [13]. Moreover, the condition on δ reduces to $0 \leq \delta < 1/3$, which exactly matches the condition required for guaranteed recovery using MIP.

Lastly, the condition (8) in Theorem 2 automatically implies that $\epsilon < 1/9$. This represents a mild tightening of the condition on ϵ required for a unique decomposition (Lemma 1), even with full measurements (i.e., when Φ is the identity operator or, more generally, when $\delta = 0$).

4 Analysis

The analysis of SPIN is based on the proof technique developed by [12] used in analyzing the iterative hard thresholding (IHT) algorithm for sparse recovery and further extended in [10, 13]. For a given set of measurements \mathbf{z} obeying (1), define the error function $\psi : \mathcal{A} \times \mathcal{B} \rightarrow \mathbb{R}$ as

$$\psi(\mathbf{a}, \mathbf{b}) = \frac{1}{2} \|\mathbf{z} - \Phi(\mathbf{a} + \mathbf{b})\|^2.$$

It is clear that $\psi(\mathbf{a}^*, \mathbf{b}^*) = \frac{1}{2} \|\mathbf{e}\|^2$. The following lemma bounds the error of the estimated signals output by SPIN at the $(k+1)$ -th iteration in terms of the error incurred at the k -th iteration, and the norm of the measurement error.

Lemma 3 *Define $(\mathbf{a}_k, \mathbf{b}_k)$ as the intermediate estimates obtained by SPIN at the k -th iteration. Let δ, ϵ be as defined in Theorem 2. Then,*

$$\psi(\mathbf{a}_{k+1}, \mathbf{b}_{k+1}) \leq \alpha \psi(\mathbf{a}_k, \mathbf{b}_k) + C \|\mathbf{e}\|^2, \quad (9)$$

where

$$\alpha = \frac{\frac{2\delta}{1-\delta} + 6\frac{1+\delta}{1-\delta}\frac{\epsilon}{1-\epsilon}}{1 - 2\frac{1+\delta}{1-\delta}\frac{\epsilon}{1-\epsilon}}, \quad C = \frac{\frac{1}{2} + 4\frac{1+\delta}{1-\delta}\frac{\epsilon}{1-\epsilon}}{1 - 2\frac{1+\delta}{1-\delta}\frac{\epsilon}{1-\epsilon}}.$$

Proof. Fix a current estimate of the signal components $(\mathbf{a}_k, \mathbf{b}_k)$ at iteration k . Define $\mathbf{x}_k \doteq \mathbf{a}_k + \mathbf{b}_k$. Then, for any other pair of signals $(\mathbf{a}, \mathbf{b}) \in \mathcal{A} \times \mathcal{B}$, we have

$$\begin{aligned} \psi(\mathbf{a}, \mathbf{b}) - \psi(\mathbf{a}_k, \mathbf{b}_k) &= \frac{1}{2} \left(\|\mathbf{z} - \Phi(\mathbf{a} + \mathbf{b})\|^2 - \|\mathbf{z} - \Phi(\mathbf{a}_k + \mathbf{b}_k)\|^2 \right) \\ &= \frac{1}{2} \|\mathbf{z} - \Phi(\mathbf{a} + \mathbf{b})\|^2 - \frac{1}{2} \|\mathbf{z} - \Phi(\mathbf{a} + \mathbf{b})\|^2 + \langle \Phi(\mathbf{a}_k + \mathbf{b}_k) - \Phi(\mathbf{a} + \mathbf{b}), \mathbf{z} \rangle \\ &= \frac{1}{2} \|\Phi\mathbf{x} - \Phi\mathbf{x}_k\|^2 + \langle \mathbf{z} - \Phi\mathbf{x}_k, \Phi\mathbf{x}_k - \Phi\mathbf{x} \rangle, \end{aligned}$$

where $\mathbf{x} \doteq \mathbf{a} + \mathbf{b}$, and the last equation is obtained by completing the square. Since Φ is a linear operator, we can take the adjoint within the inner product to obtain

$$\begin{aligned} \psi(\mathbf{a}, \mathbf{b}) - \psi(\mathbf{a}_k, \mathbf{b}_k) &= \frac{1}{2} \|\Phi\mathbf{x} - \Phi\mathbf{x}_k\|^2 + \langle \Phi^T(\mathbf{z} - \Phi\mathbf{x}_k), \mathbf{x}_k - \mathbf{x} \rangle \\ &\leq \frac{1}{2} (1 + \delta) \|\mathbf{x} - \mathbf{x}_k\|^2 + \langle \Phi^T(\mathbf{z} - \Phi\mathbf{x}_k), \mathbf{x}_k - \mathbf{x} \rangle. \end{aligned} \quad (10)$$

The last inequality occurs due to the RIP of Φ applied to the secant vector $\mathbf{x} - \mathbf{x}_k \in \mathcal{S}(\mathcal{C})$. To the right hand side of (10), we further add and subtract $\frac{1}{2(1+\delta)} \|\Phi^T(\mathbf{z} - \Phi\mathbf{x}_k)\|^2$ to complete the square:

$$\psi(\mathbf{a}, \mathbf{b}) - \psi(\mathbf{a}_k, \mathbf{b}_k) \leq \frac{1}{2}(1+\delta) \left\| c - c_k + \frac{1}{1+\delta} \Phi^T(\mathbf{z} - \Phi\mathbf{x}_k) \right\|^2 - \frac{1}{2(1+\delta)} \|\Phi^T(\mathbf{z} - \Phi\mathbf{x}_k)\|^2.$$

Define $\mathbf{g}_k \doteq \frac{1}{1+\delta} \Phi^T(\mathbf{z} - \Phi(\mathbf{a}_k + \mathbf{b}_k))$ as the gradient of the error function at $(\mathbf{a}_k, \mathbf{b}_k)$. Then,

$$\psi(\mathbf{a}, \mathbf{b}) - \psi(\mathbf{a}_k, \mathbf{b}_k) \leq \frac{1}{2}(1+\delta) \left(\|\mathbf{a} + \mathbf{b} - (\mathbf{a}_k + \mathbf{b}_k + \mathbf{g}_k)\|^2 - \|\mathbf{g}_k\|^2 \right). \quad (11)$$

Next, define the function ζ on $\mathcal{A} \times \mathcal{B}$ as $\zeta(\mathbf{a}, \mathbf{b}) \doteq \|\mathbf{a} + \mathbf{b} - (\mathbf{a}_k + \mathbf{b}_k + \mathbf{g}_k)\|^2$. Then, we have

$$\begin{aligned} \zeta(\mathbf{a}_{k+1}, \mathbf{b}_{k+1}) &= \|\mathbf{a}_{k+1} - (\mathbf{a}_k + \mathbf{g}_k) + \mathbf{b}_{k+1} - (\mathbf{b}_k + \mathbf{g}_k) + \mathbf{g}_k\|^2 \\ &= \|\mathbf{a}_{k+1} - (\mathbf{a}_k + \mathbf{g}_k)\|^2 + \|\mathbf{b}_{k+1} - (\mathbf{b}_k + \mathbf{g}_k)\|^2 + \|\mathbf{g}_k\|^2 \\ &\quad + 2\langle \mathbf{a}_{k+1} - (\mathbf{a}_k + \mathbf{g}_k), \mathbf{b}_{k+1} - (\mathbf{b}_k + \mathbf{g}_k) \rangle + 2\langle \mathbf{g}_k, \mathbf{a}_{k+1} + \mathbf{b}_{k+1} - (\mathbf{a}_k + \mathbf{b}_k + 2\mathbf{g}_k) \rangle. \end{aligned}$$

But, by definition, $\mathbf{a}_{k+1} = \mathcal{P}_{\mathcal{A}}(\mathbf{a}_k + \mathbf{g}_k)$, and hence $\|\mathbf{a}_{k+1} - (\mathbf{a}_k + \mathbf{g}_k)\| \leq \|\mathbf{a} - (\mathbf{a}_k + \mathbf{g}_k)\|$ for any $\mathbf{a} \in \mathcal{A}$. An analogous relation can be formed between \mathbf{b}_{k+1} and \mathbf{b}^* . Hence, we have

$$\begin{aligned} \|\mathbf{a}_{k+1} - (\mathbf{a}_k + \mathbf{g}_k)\| &\leq \|\mathbf{a}^* - (\mathbf{a}_k + \mathbf{g}_k)\|, \quad \text{and} \\ \|\mathbf{b}_{k+1} - (\mathbf{b}_k + \mathbf{g}_k)\| &\leq \|\mathbf{b}^* - (\mathbf{b}_k + \mathbf{g}_k)\|. \end{aligned}$$

Substituting for $(\mathbf{a}_{k+1}, \mathbf{b}_{k+1})$, we obtain

$$\begin{aligned} \zeta(\mathbf{a}_{k+1}, \mathbf{b}_{k+1}) &\leq \|\mathbf{a}^* - (\mathbf{a}_k + \mathbf{g}_k)\|^2 + \|\mathbf{b}^* - (\mathbf{b}_k + \mathbf{g}_k)\|^2 + \|\mathbf{g}_k\|^2 \\ &\quad + 2\langle \mathbf{a}_{k+1} - (\mathbf{a}_k + \mathbf{g}_k), \mathbf{b}_{k+1} - (\mathbf{b}_k + \mathbf{g}_k) \rangle + 2\langle \mathbf{g}_k, \mathbf{a}_{k+1} + \mathbf{b}_{k+1} - (\mathbf{a}_k + \mathbf{b}_k + 2\mathbf{g}_k) \rangle \\ &= \|\mathbf{a}^* - (\mathbf{a}_k + \mathbf{g}_k)\|^2 + \|\mathbf{b}^* - (\mathbf{b}_k + \mathbf{g}_k)\|^2 + \|\mathbf{g}_k\|^2 \\ &\quad + 2\langle \mathbf{a}^* - (\mathbf{a}_k + \mathbf{g}_k), \mathbf{b}^* - (\mathbf{b}_k + \mathbf{g}_k) \rangle + 2\langle \mathbf{g}_k, \mathbf{a}^* + \mathbf{b}^* - (\mathbf{a}_k + \mathbf{b}_k + 2\mathbf{g}_k) \rangle \\ &\quad + 2\langle \mathbf{a}_{k+1} - (\mathbf{a}_k + \mathbf{g}_k), \mathbf{b}_{k+1} - (\mathbf{b}_k + \mathbf{g}_k) \rangle - 2\langle \mathbf{a}^* - (\mathbf{a}_k + \mathbf{g}_k), \mathbf{b}_{k+1}^* - (\mathbf{b}_k + \mathbf{g}_k) \rangle \\ &\quad + 2\langle \mathbf{g}_k, \mathbf{a}_{k+1} + \mathbf{b}_{k+1} - (\mathbf{a}^* + \mathbf{b}^*) \rangle. \end{aligned}$$

Completing the squares, we have:

$$\begin{aligned} \zeta(\mathbf{a}_{k+1}, \mathbf{b}_{k+1}) &\leq \|\mathbf{a}^* + \mathbf{b}^* - (\mathbf{a}_k + \mathbf{b}_k + \mathbf{g}_k)\|^2 + 2\langle \mathbf{a}_{k+1} - \mathbf{a}_k, \mathbf{b}_{k+1} - \mathbf{b}_k \rangle - 2\langle \mathbf{a}^* - \mathbf{a}_k, \mathbf{b}^* - \mathbf{b}_k \rangle \\ &\quad + 2\langle \mathbf{g}_k, -\mathbf{a}_{k+1} + \mathbf{a}_k - \mathbf{b}_{k+1} + \mathbf{b}_k + \mathbf{a}^* - \mathbf{a}_k + \mathbf{b}^* - \mathbf{b}_k + \mathbf{a}_{k+1} + \mathbf{b}_{k+1} - (\mathbf{a}^* + \mathbf{b}^*) \rangle. \end{aligned}$$

The last term on the right hand side equals zero, and so we obtain

$$\zeta(\mathbf{a}_{k+1}, \mathbf{b}_{k+1}) \leq \zeta(\mathbf{a}^*, \mathbf{b}^*) + 2\langle \mathbf{a}_{k+1} - \mathbf{a}_k, \mathbf{b}_{k+1} - \mathbf{b}_k \rangle - 2\langle \mathbf{a}^* - \mathbf{a}_k, \mathbf{b}^* - \mathbf{b}_k \rangle.$$

Combining this inequality with (11), we obtain the series of inequalities

$$\begin{aligned} \psi(\mathbf{a}_{k+1}, \mathbf{b}_{k+1}) - \psi(\mathbf{a}_k, \mathbf{b}_k) &\leq \frac{1}{2}(1+\delta) \left(\zeta(\mathbf{a}_{k+1}, \mathbf{b}_{k+1}) - \|\mathbf{g}_k\|^2 \right) \\ &\leq \overbrace{\frac{1}{2}(1+\delta) \left(\zeta(\mathbf{a}^*, \mathbf{b}^*) - \|\mathbf{g}_k\|^2 \right)}^{\mathbb{T}_1} \\ &\quad + \overbrace{(1+\delta) \left(\langle \mathbf{a}_{k+1} - \mathbf{a}_k, \mathbf{b}_{k+1} - \mathbf{b}_k \rangle - \langle \mathbf{a}^* - \mathbf{a}_k, \mathbf{b}^* - \mathbf{b}_k \rangle \right)}^{\mathbb{T}_2} \quad (12) \\ &= \mathbb{T}_1 + \mathbb{T}_2. \end{aligned}$$

We can further bound the right hand side of (12) as follows. First, we expand \mathbb{T}_1 to obtain

$$\begin{aligned}
\mathbb{T}_1 &= \frac{1}{2}(1+\delta) \left(\|\mathbf{a}^* + \mathbf{b}^* - (\mathbf{a}_k + \mathbf{b}_k + \mathbf{g}_k)\|^2 - \|\mathbf{g}_k\|^2 \right) \\
&= \frac{1}{2}(1+\delta) \left(\|\mathbf{a}^* + \mathbf{b}^* - (\mathbf{a}_k + \mathbf{b}_k)\|^2 + 2\langle \mathbf{g}_k, \mathbf{a}^* + \mathbf{b}^* - (\mathbf{a}_k + \mathbf{b}_k) \rangle \right) \\
&= \frac{1}{2}(1+\delta) \|\mathbf{x}^* - \mathbf{x}_k\|^2 + \langle \Phi^T(\mathbf{z} - \Phi \mathbf{x}_k), \mathbf{x}_k - \mathbf{x}^* \rangle.
\end{aligned}$$

Again, $\mathbf{x}^* - \mathbf{x}_k$ is a secant on the direct sum manifold \mathcal{C} . By the RIP property of Φ , we have

$$\begin{aligned}
\mathbb{T}_1 &\leq \frac{1}{2} \frac{1+\delta}{1-\delta} \|\Phi \mathbf{x}^* - \Phi \mathbf{x}_k\|^2 + \langle \Phi^T(\mathbf{z} - \Phi \mathbf{x}_k), \mathbf{x}_k - \mathbf{x}^* \rangle \\
&= \frac{1}{2} \|\Phi \mathbf{x}^* - \Phi \mathbf{x}_k\|^2 + \langle \Phi^T(\mathbf{z} - \Phi \mathbf{x}_k), \mathbf{x}_k - \mathbf{x}^* \rangle + \frac{\delta}{1-\delta} \|\mathbf{z} - \Phi \mathbf{x}_k\|^2 \\
&= \psi(\mathbf{a}^*, \mathbf{b}^*) - \psi(\mathbf{a}_k, \mathbf{b}_k) + \frac{2\delta}{1-\delta} \psi(\mathbf{a}_k, \mathbf{b}_k),
\end{aligned} \tag{13}$$

where the last expression follows easily from (10) and the definition of ψ .

The term \mathbb{T}_2 can be bounded using Lemma 2 as follows. We have

$$\begin{aligned}
\langle \mathbf{a}^* - \mathbf{a}_k, \mathbf{b}^* - \mathbf{b}_k \rangle &\leq \frac{\epsilon}{2(1-\epsilon)} \|\mathbf{x}^* - \mathbf{x}_k\|^2, \quad \text{and} \\
\langle \mathbf{a}_{k+1} - \mathbf{a}_k, \mathbf{b}_{k+1} - \mathbf{b}_k \rangle &\leq \frac{\epsilon}{2(1-\epsilon)} \|\mathbf{x}_{k+1} - \mathbf{x}_k\|^2 \\
&\leq \frac{\epsilon}{(1-\epsilon)} \left(\|\mathbf{x}_{k+1} - \mathbf{x}^*\|^2 + \|\mathbf{x}_k - \mathbf{x}^*\|^2 \right),
\end{aligned} \tag{14}$$

where (14) follows from the the triangle inequality and the AM-GM inequality. Further, \mathbb{T}_2 can be simplified as

$$\begin{aligned}
\mathbb{T}_2 &\leq (1+\delta) \frac{\epsilon}{1-\epsilon} \left(\frac{3}{2} \|\mathbf{x}^* - \mathbf{x}_k\|^2 + \|\mathbf{x}^* - \mathbf{x}_{k+1}\|^2 \right) \\
&\leq \frac{1+\delta}{1-\delta} \frac{\epsilon}{1-\epsilon} \left(\frac{3}{2} \|\Phi \mathbf{x}^* - \Phi \mathbf{x}_k\|^2 + \|\Phi \mathbf{x}^* - \Phi \mathbf{x}_{k+1}\|^2 \right).
\end{aligned}$$

But,

$$\begin{aligned}
\|\Phi \mathbf{x}^* - \Phi \mathbf{x}_k\|^2 &= \|\Phi \mathbf{x}^* - \Phi \mathbf{x}_k + \mathbf{e} - \mathbf{e}\|^2 \\
&\leq 2\|\Phi \mathbf{x}^* - \Phi \mathbf{x}_k + \mathbf{e}\|^2 + 2\|\mathbf{e}\|^2 = 4\psi(\mathbf{a}_k, \mathbf{b}_k) + 4\psi(\mathbf{a}^*, \mathbf{b}^*),
\end{aligned}$$

via the same technique used to obtain (14). Similarly,

$$\|\Phi \mathbf{x}^* - \Phi \mathbf{x}_{k+1}\|^2 \leq 4\psi(\mathbf{a}_{k+1}, \mathbf{b}_{k+1}) + 4\psi(\mathbf{a}^*, \mathbf{b}^*).$$

Hence, we obtain

$$\mathbb{T}_2 \leq \frac{1+\delta}{1-\delta} \frac{2\epsilon}{1-\epsilon} (3\psi(\mathbf{a}_k, \mathbf{b}_k) + \psi(\mathbf{a}_{k+1}, \mathbf{b}_{k+1}) + 4\psi(\mathbf{a}^*, \mathbf{b}^*)). \tag{15}$$

Combining (12), (13), and (15), we obtain

$$\begin{aligned} \psi(\mathbf{a}_{k+1}, \mathbf{b}_{k+1}) &\leq \psi(\mathbf{a}^*, \mathbf{b}^*) + \frac{2\delta}{1-\delta} \psi(\mathbf{a}^k, \mathbf{b}^k) \\ &\quad + \frac{1+\delta}{1-\delta} \frac{2\epsilon}{1-\epsilon} (3\psi(\mathbf{a}_k, \mathbf{b}_k) + \psi(\mathbf{a}_{k+1}, \mathbf{b}_{k+1}) + 4\psi(\mathbf{a}^*, \mathbf{b}^*)). \end{aligned}$$

Rearranging, we obtain Lemma 3. \square

Proof of Theorem 2. Equation 9 describes a linear recurrence relation for the sequence of positive real numbers $\psi(\mathbf{a}_k, \mathbf{b}_k)$, $k = 0, 1, 2, \dots$ with leading coefficient $\alpha < 1$. By choice of initialization, $\psi(\mathbf{a}_0, \mathbf{b}_0) = \frac{\|\mathbf{z}\|^2}{2}$. Therefore, for all $k \in \mathbb{N}$, we have the relation

$$\psi(\mathbf{a}_k, \mathbf{b}_k) \leq \alpha^k \psi(\mathbf{a}_0, \mathbf{b}_0) + \frac{C}{1-\alpha} \|\mathbf{e}\|^2.$$

Choosing $\beta = \frac{C}{1-\alpha}$, and $k \geq T$ such that $T = \lceil \frac{1}{\log(1/\alpha)} \log \frac{\|\mathbf{z}\|^2}{2\beta} \rceil$, the result follows. \square

5 Applications

The two-manifold signal model described in this paper is applicable to a wide variety of problems that have attracted considerable interest in the literature over the last several years. We discuss a few representative instances and show how SPIN can be utilized for efficient signal recovery in each of these instances. We also present several numerical experiments that indicate the kind of gains that SPIN can offer in practice.

5.1 Sparse representations in pairs of bases

We revisit the classical problem of decomposing signals in an overcomplete dictionary that is the union of a pair of orthonormal bases and show how SPIN can be used to efficiently solve this problem. Let $\mathbf{\Lambda}, \mathbf{\Lambda}'$ be orthonormal bases of \mathbb{R}^N . Let \mathcal{A} be the set of all K_1 -sparse signals in \mathbb{R}^N in the basis expansion of $\mathbf{\Lambda}$, and let \mathcal{B} be the set of all K_2 -sparse signals in the basis expansion of $\mathbf{\Lambda}'$. Then, \mathcal{A} and \mathcal{B} can be viewed as K_1 - and K_2 -dimensional submanifolds of \mathbb{R}^N , respectively. Consider a signal $\mathbf{x}^* = \mathbf{a}^* + \mathbf{b}^*$, where $\mathbf{a}^* \in \mathcal{A}, \mathbf{b}^* \in \mathcal{B}$ so that

$$\mathbf{a}^* = \sum_{i=1}^{K_1} a_i \lambda_i, \quad \mathbf{b}^* = \sum_{i=1}^{K_2} b_i \lambda'_i.$$

The problem is to recover $(\mathbf{a}^*, \mathbf{b}^*)$ given \mathbf{x}^* . This problem has been studied in many different forms in the literature, and several algorithms have been proposed in order to solve it efficiently [1, 3, 17]. See [18] for an in-depth study of the various state-of-the-art methods. All these methods assume a certain notion of incoherence between the two bases, most commonly referred to as the *mutual coherence* μ , which is defined as

$$\mu(\mathbf{\Lambda}, \mathbf{\Lambda}') \doteq \max_{i,j} |\langle \lambda_i, \lambda'_j \rangle|. \quad (16)$$

It is clear that $\mu \leq 1$. It can be shown that the mutual coherence also obeys the lower bound $\mu \geq 1/\sqrt{N}$ for any basis of \mathbb{R}^N . This lower bound is in fact tight; equality is achieved, for example,

when $\mathbf{\Lambda}$ is the canonical basis in \mathbb{R}^N and $\mathbf{\Lambda}'$ is the basis defining the Walsh-Hadamard transform or the Discrete Fourier Basis [3, 17].

We establish the following simple relation between μ and the manifold incoherence between \mathcal{A} and \mathcal{B} .

Lemma 4 *Let \mathcal{A} be the set of all K_1 -sparse signals in $\mathbf{\Lambda}$, and \mathcal{B} be the set of all K_2 -sparse signals in $\mathbf{\Lambda}'$. Let ϵ denote the manifold incoherence between \mathcal{A} and \mathcal{B} . Then,*

$$\epsilon \leq \mu(\mathbf{\Lambda}, \mathbf{\Lambda}')(K_1 + K_2).$$

Proof. The secant manifold $\mathcal{S}(\mathcal{A})$ is equivalent to the set of signals in \mathbb{R}^N that are $2K_1$ -sparse in \mathcal{A} ; similarly, $\mathcal{S}(\mathcal{B})$ is equivalent to the set of $2K_2$ -sparse signals in \mathcal{B} . Therefore, if one considers unit norm vectors $\mathbf{u} \in \mathcal{S}(\mathcal{A}), \mathbf{u}' \in \mathcal{S}(\mathcal{B})$, we obtain

$$\begin{aligned} |\langle \mathbf{u}, \mathbf{u}' \rangle| &= \left| \left\langle \sum_{i=1}^{2K_1} a_i \lambda_i, \sum_{j=1}^{2K_2} b_j \lambda'_j \right\rangle \right| = \left| \sum_i \sum_j \langle \lambda_i, \lambda'_j \rangle a_i b_j \right|, \\ &\leq \mu \sum_i \sum_j |a_i b_j| = \mu \left(\sum_{i=1}^{2K_1} |a_i| \right) \left(\sum_{j=1}^{2K_2} |b_j| \right), \end{aligned} \quad (17)$$

where the last relation follows from the triangle inequality. From Theorem 1 of [17], we have that the right hand side is maximized when $a_i = \frac{1}{\sqrt{2K_1}}, b_j = \frac{1}{\sqrt{2K_2}}$ for all $i = 1, \dots, 2K_1$ and $j = 1, \dots, 2K_2$. The result of Lemma 4 follows from inserting these values into (17) and applying the AM-GM inequality. \square

We show how SPIN can be used to solve the linear inverse problem of recovering $(\mathbf{a}^*, \mathbf{b}^*)$ from \mathbf{x}^* . The restricted isometry assumption is not relevant in this case, since we assume that we have full measurements of the signal; therefore $\delta = 0$. An upper bound for the manifold incoherence parameter ϵ is specified in Lemma 4. The (exact) projection operators $\mathcal{P}_{\mathcal{A}}, \mathcal{P}_{\mathcal{B}}$ can be easily implemented; we simply perform a coefficient expansion in the corresponding orthonormal basis and retain the coefficients of largest magnitude. Mixing these ingredients together, we can guarantee that, given any signal \mathbf{x}^* , SPIN will return the true components $\mathbf{a}^*, \mathbf{b}^*$. This guarantee is summarized in the following result.

Corollary 1 (SPIN for pairs of bases) *Let $\mathbf{x}^* = \mathbf{a}^* + \mathbf{b}^*$, where \mathbf{a}^* is K_1 -sparse in $\mathbf{\Lambda}$ and \mathbf{b}^* is K_2 -sparse in $\mathbf{\Lambda}'$. Let μ denote the mutual coherence between $\mathbf{\Lambda}$ and $\mathbf{\Lambda}'$. Then, SPIN exactly recovers $(\mathbf{a}^*, \mathbf{b}^*)$ from \mathbf{x}^* provided*

$$K_1 + K_2 < \frac{1}{9\mu}. \quad (18)$$

Proof. If (18) holds, then from Lemma 4 we know that the manifold incoherence ϵ between \mathcal{A} and \mathcal{B} is smaller than $1/9$. But this is exactly the condition required for guaranteed convergence of SPIN to the true signal components $(\mathbf{a}^*, \mathbf{b}^*)$. \square

SPIN thus offers a conceptually simple method to separate mixtures of signals that are sparse in incoherent bases. The best known approach for this problem is an ℓ_1 -minimization formulation that features a similar guarantee that is known to be tight [17, 19]:

$$K_1 + K_2 < \frac{\sqrt{2} - 0.5}{\mu}.$$

Therefore, SPIN yields a recovery guarantee that is equivalent to the best possible method, modulo a small constant. It is likely that the constant $1/9$ in Corollary 1 can be sharpened, but we will not pursue this direction in this paper.

5.2 Articulation manifolds

Articulation manifolds provide a powerful, flexible conceptual tool for modeling signals and image ensembles in a number of applications [20, 21]. Consider an ensemble of signals $\mathcal{M} \subset \mathbb{R}^N$ that are generated by varying K parameters $\theta \in \Theta$, $\Theta \subset \mathbb{R}^K$. Then, we say that the signals trace out a nonlinear K -dimensional articulation manifold in \mathbb{R}^N , where θ is called the *articulation parameter vector*. Examples of articulation manifolds include: acoustic chirp signals of varying frequencies (where θ represents the chirp rate); images of a white disk translating on a black background (where θ represents the planar location of the disk center); and images of a solid object with variable pose (where θ represents the 6D pose parameters, three corresponding to spatial location and three corresponding to orientation).

We consider the class of *compact, smooth, K -dimensional articulation manifolds* $\mathcal{M} \subset \mathbb{R}^N$. For such manifold classes, it is possible to construct linear measurement operators Φ that preserve the pairwise secant geometry of \mathcal{M} . Specifically, it has been shown [16] that there exist *randomized* constructions of measurement operators $\Phi \in \mathbb{R}^{M \times N}$ that satisfy the RIP on the secants of \mathcal{M} with constant δ , and with probability at least ρ , provided

$$M \geq \mathcal{O} \left(K \frac{\log(CN\delta^{-1}) \log(\rho^{-1})}{\delta^2} \right),$$

for some constant C that depends only on the smoothness and volume of the manifold \mathcal{M} . Therefore, the dimension of the range space of Φ is proportional to the *number of degrees of freedom* K , but is only logarithmic in the ambient dimension N . Moreover, given such a measurement matrix Φ with isometry constant $\delta < 1/3$ and a projection operator $\mathcal{P}_{\mathcal{M}}(\cdot)$ onto \mathcal{M} , any signal $\mathbf{x} \in \mathcal{M}$ can be reconstructed from its compressive measurements $\mathbf{y} = \Phi \mathbf{x}$ using Manifold Iterative Pursuit (MIP) [13].

We generalize this setting to the case where the unknown signal of interest arises as a mixture of signals from two manifolds \mathcal{A} and \mathcal{B} . For instance, suppose we are interested in the space of images, where \mathcal{A} and \mathcal{B} comprise of translations of fixed template images $f(\mathbf{t})$ and $g(\mathbf{t})$, where \mathbf{t} denotes the 2D domain over which the image is defined. Then, the signal of interest is an image of the form

$$\mathbf{x}^* = \mathbf{a}^* + \mathbf{b}^* = f(\mathbf{t} + \theta_1) + g(\mathbf{t} + \theta_2),$$

where θ_1 and θ_2 denote the unknown translation parameters. The problem is to recover $(\mathbf{a}^*, \mathbf{b}^*)$, or equivalently (θ_1, θ_2) , given compressive measurements $\mathbf{z} = \Phi(\mathbf{a}^* + \mathbf{b}^*)$.

We demonstrate that SPIN offers an easy, efficient technique to recover the component images. This example also demonstrates that SPIN is robust to practical considerations such as noise. Figure 1 displays the results of SPIN recovery of a 64×64 image from very limited measurements. The unknown image consists of the linear sum of arbitrary translations of template images $f(\mathbf{t})$ and $g(\mathbf{t})$, that are smoothed binary images on a black background of a white disk and a white square, respectively. Further, the image has been contaminated with significant Gaussian noise (SNR = 14dB) prior to measurement (Fig. 1(a)). From Figs. 1(b) and 1(c), we observe that SPIN is able to perfectly recover the original component signals from merely $M = 50$ random linear measurements.

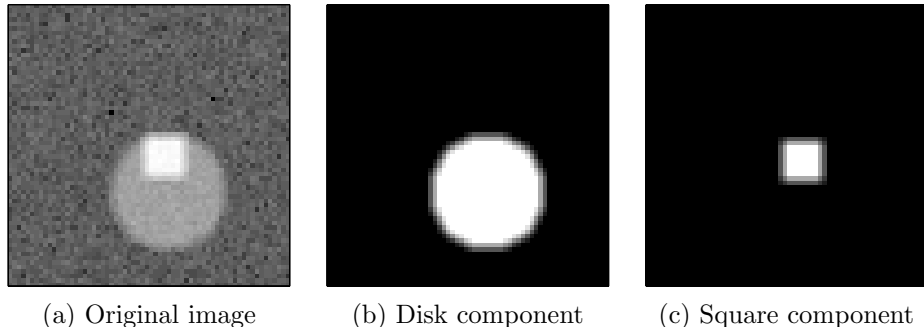


Figure 1: *SPIN* recovery of a noisy 64×64 image from compressive measurements. The clean image consists of the linear superposition of a disk and a square of fixed pre-specified sizes, but the locations of the centers of the disk and the square are unknown. Additive Gaussian noise ($\text{SNR} = 14\text{dB}$) is added to the image prior to measurement. Signal length $N = 64 \times 64 = 4096$, number of compressive measurements $M = 50$. (a) Original noisy image. (b) Reconstructed disk. (c) Reconstructed square. *SPIN* perfectly reconstructs both components from just $M/N = 1.2\%$ measurements.

For guaranteed *SPIN* convergence, we require that the manifolds \mathcal{A}, \mathcal{B} are incoherent. Informally, the condition of incoherence on the secants of \mathcal{A} and \mathcal{B} is always valid when the template images $f(\mathbf{t}), g(\mathbf{t})$ are “sufficiently” distinct. This intuition is made precise using the bound in (3). More generally, we can state the following theoretical guarantee for *SPIN* performance in the case of general higher-dimensional manifolds.

Corollary 2 (SPIN for pairs of manifolds) *Let the entries of $\Phi \in \mathbb{R}^{M \times N}$ be chosen from a standard Gaussian probability distribution. Let \mathcal{A} and \mathcal{B} be ϵ -incoherent compact submanifolds of \mathbb{R}^N of dimensions K and K' respectively. Let $\mathbf{z} = \Phi(\mathbf{a}^* + \mathbf{b}^*)$, where $\mathbf{a}^* \in \mathcal{A}$ and $\mathbf{b}^* \in \mathcal{B}$. Then, with probability greater than $1 - \rho$, *SPIN* exactly recovers $(\mathbf{a}^*, \mathbf{b}^*)$ from \mathbf{z} , provided*

$$M \geq \mathcal{O}((K \log(C_{\mathcal{A}}N) + K' \log(C_{\mathcal{B}}N)) \log(\rho^{-1})). \quad (19)$$

Here, $C_{\mathcal{A}}, C_{\mathcal{B}}$ are constants that depend only on certain intrinsic geometric parameters (such as volume) of \mathcal{A}, \mathcal{B} respectively.

Proof. It is easy to show that the matrix Φ satisfies the RIP on the secants of the direct sum $\mathcal{C} = \mathcal{A} \oplus \mathcal{B}$ if (19) is true. This can be proved using a simple union bound argument in conjunction with the techniques used in Section 3.2.2 in [16]. Once the RIP of Φ is established, *SPIN* recovery follows from Theorem 2. \square

An important consideration in *SPIN* is the tractable computation of the projection $\mathcal{P}_{\mathcal{M}}(\mathbf{x})$ given any $\mathbf{x} \in \mathbb{R}^N$. For example, in the numerical example in Fig. 1, the operator $\mathcal{P}_{\mathcal{A}}(\mathbf{x})$ onto the manifold \mathcal{A} consists of running a matched filter between the template $f(\mathbf{t})$ and the input signal \mathbf{x} and returning $f(\mathbf{t} + \hat{\theta})$, where the parameter value $\hat{\theta}$ corresponds to the 2D location of the maximum of the matched filter response. This is very efficiently carried out in $\mathcal{O}(N \log N)$ operations using the Fast Fourier Transform (FFT). However, for more complex articulation manifolds, the exact projection operation might not be tractable, whereas only approximate numerical projections can be efficiently computed. In this case also, *SPIN* can recover the signal components $(\mathbf{a}^*, \mathbf{b}^*)$, but with weaker convergence guarantees (Theorem 3).




Figure 2: *SPIN recovery of a shifted Gaussian pulse from compressive measurements. The shift parameter of the pulse is unknown, and the signal is corrupted with K' -sparse, impulsive noise of unknown amplitudes and locations. $N = 10000, K' = 10, M = 150$. (a) Original signal. (b) Reconstructed Gaussian pulse (Recovery SNR = 80.09 dB). (c) Estimated noise component. SPIN perfectly reconstructs both components from just $M/N = 1.5\%$ measurements.*

5.3 Signals in impulsive noise

In some situations, the signal of interest \mathbf{x} might be corrupted with *impulsive noise* (or shot noise) prior to signal acquisition via linear measurements. For example, consider Fig. 2(a), where the Gaussian pulse is the signal of interest, and the spikes indicate the undesirable noise. In this case, the linear observations are more accurately modeled as:

$$\mathbf{z} = \Phi(\mathbf{x} + \mathbf{n}), \quad \text{such that } \mathbf{x} \in \mathcal{M},$$

and n is a K' -sparse signal in the canonical basis. Therefore, SPIN can be used to recover \mathbf{x} from \mathbf{z} , provided that the manifold \mathcal{M} is incoherent with the set of sparse signals $\Sigma_{K'}$ and Φ satisfies the RIP on the direct sum $\mathcal{M} + \Sigma_{K'}$.

Figure 2 displays the results of a numerical experiment that illustrates the utility of SPIN in this setting. We consider a manifold of signals of length $N = 10000$ that consist of shifts of a Gaussian pulse of fixed width $g_0 \in \mathbb{R}^N$. The unknown signal \mathbf{x} is an element of this manifold \mathcal{M} , and is corrupted by $K' = 10$ spikes of unknown magnitudes and locations. This degraded signal is sampled using $M = 150$ random linear measurements to obtain an observation vector \mathbf{z} .

We apply SPIN to recover \mathbf{x} from \mathbf{z} . The projection operator $\mathcal{P}_{\mathcal{M}}(\cdot)$ consists of a matched filter with the template pulse g_0 , while the projection operator $\mathcal{P}_{\Sigma_{K'}}(\cdot)$ simply returns the best K' -term approximation in the canonical basis. Assuming that we have knowledge of the number of nonzeros in the noise vector \mathbf{n} , we can use SPIN to reconstruct both \mathbf{x} and \mathbf{n} . We observe from Fig. 2(b) that SPIN recovers the true signal \mathbf{x} with near-perfect accuracy. Further, this recovery is possible with only a small number $M = 150$ linear measurements of \mathbf{x} , which constitutes but a fraction of the ambient dimension of the signal space.

Figure 3 plots the number of measurements M vs. the signal reconstruction error (normalized relative to the signal energy, and plotted in dB). We observe that, by increasing M , SPIN can tolerate an increased number K' of nuisance spikes. Further, by Corollary 2, we observe that this relationship between M and K' is in fact *linear*. This result can be extended to any situation where the signals of interest obey a “hybrid” model that is a mixture of a nonlinear manifold and the set of sparse signals.

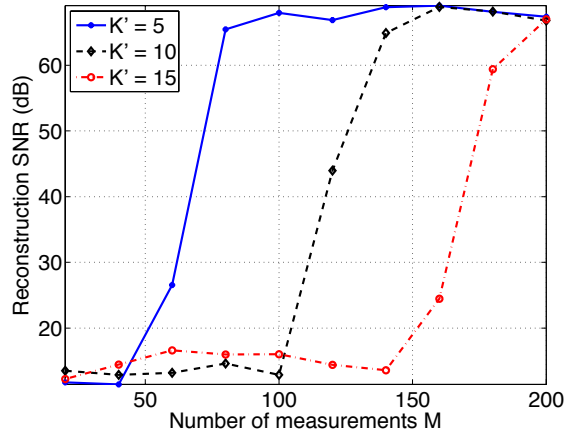


Figure 3: Monte Carlo simulation of SPIN signal recovery in impulsive noise, averaged over 100 trials. In each trial, the measured signal is the sum of a randomly shifted Gaussian pulse and a random K' -sparse signal. SPIN can tolerate a higher number K' nuisance impulses by increasing the number of measurements M . By Corollary 2 this dependence of M on K' can be shown to be linear.

6 Discussion

We have proposed and rigorously analyzed an algorithm, which we dub Successive Projections onto INcoherent Manifolds (SPIN), for the recovery of a pair of signals given a small number of measurements of their linear sum. For SPIN to guarantee signal recovery, we require two main geometric criteria to hold: (i) the component signals should arise from two disjoint manifolds that are in a specific sense *incoherent*, and (ii) the linear measurement operator should satisfy a *restricted isometry* criterion on the secants of the direct sum of the two manifolds. The computational efficiency of SPIN is determined by the tractability of the *projection operators* onto either component manifold. We have presented indicative numerical experiments demonstrating the utility of SPIN, but defer a thorough experimental study of SPIN to future work.

Practical considerations. SPIN is an iterative gradient projection algorithm and requires as input parameters the number of iterations T and the gradient step size η . The iteration count T can be chosen using one of many commonly-used stopping criteria. For example, convergence can be declared if the norm of the error $\psi(\mathbf{a}_k, \mathbf{b}_k)$ at the $(k+1)$ -th time step does not differ significantly from the error at k -th time step. The choice of optimal step size η is more delicate. Theorem 2 relates the step size to the restricted isometry constant δ of Φ , but this constant is not easy to calculate. In our preliminary findings, a step size in the range $0.5 \leq \eta \leq 0.7$ consistently gave good results. See [22] for a discussion on the choice of step size for hard thresholding methods.

In practical scenarios, the signal of interest rarely belongs exactly to a low-dimensional submanifold \mathcal{M} of the ambient space, but only is well-approximated by \mathcal{M} . Interestingly, in such situations the effect of this mismatch can be studied using the concept of γ -approximate projections (7). Theorem 3 rigorously demonstrates that SPIN is robust to such approximations. Further, our main result (Theorem 2) indicates that SPIN is stable with respect to inaccurate measurements, owing to the fact that the reconstruction error is bounded by a constant times the norm of the measurement noise vector \mathbf{e} .

More than two manifolds. For clarity and brevity, we have focused our attention on signals

belonging to the direct sum of two signal manifolds. However, SPIN (and its accompanying proof mechanism) can be conceptually extended to sums of any Q manifolds. In such a scenario, the conditions of convergence of SPIN would require that the component manifolds are Q -wise incoherent, and the measurement operator Φ satisfies a restricted isometry on the Q -wise direct sum of the component manifolds.

Connections to matrix recovery. An intriguing open question is whether SPIN (or a similar first-order projected gradient algorithm) is applicable to situations where either of the component manifolds is the set of low-rank matrices. The problem of reconstructing, from affine measurements, matrices that are a sum of low-rank and sparse matrices has attracted significant attention in the recent literature [7, 8, 23]. The key stumbling block is that the manifold of low-rank matrices is *not incoherent* with the manifold of sparse matrices; indeed, the two manifolds share a nontrivial intersection (i.e., there exist low rank matrices that are also sparse, and vice versa). Phenomena such as these make the analysis of SPIN (or similar algorithms) quite challenging, and it may be that higher-order techniques for signal reconstruction will be needed.

References

- [1] S. Chen, D. Donoho, and M. Saunders, “Atomic decomposition by basis pursuit,” *SIAM J. Sci. Comp.*, vol. 20, no. 1, pp. 33–61, 1998.
- [2] M. Elad, J.L. Starck, P. Querre, and D.L. Donoho, “Simultaneous cartoon and texture image inpainting using morphological component analysis,” *Appl. Comput. Harmon. Anal.*, vol. 19, no. 3, pp. 340–358, 2005.
- [3] J. Tropp, “Greed is good: Algorithmic results for sparse approximation,” *IEEE Trans. Inform. Theory*, vol. 50, no. 10, pp. 2231–2242, 2004.
- [4] E. Candès, “Compressive sampling,” in *Proc. Int. Congress of Math.*, Madrid, Spain, Aug. 2006.
- [5] D. Donoho, “Compressed sensing,” *IEEE Trans. Inform. Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.
- [6] M. Fazel, E. Candès, B. Recht, and P. Parrilo, “Compressed sensing and robust recovery of low rank matrices,” in *Proc. 40th Asilomar Conf. Signals, Systems and Computers*, Pacific Grove, CA, Nov. 2008.
- [7] E. Candès, X. Li, Y. Ma, and J. Wright, “Robust principal component analysis?,” *J. ACM*, vol. 58, no. 3, pp. 1–37, May 2011.
- [8] V. Chandrasekaran, S. Sanghavi, P. Parrilo, and A. Willsky, “Sparse and low-rank matrix decompositions,” in *Proc. Allerton Conf. on Comm., Contr., and Comp.*, Monticello, IL, Sep. 2009.
- [9] T. Blumensath and M. Davies, “Iterative hard thresholding for compressive sensing,” *Appl. Comput. Harmon. Anal.*, vol. 27, no. 3, pp. 265–274, 2009.
- [10] R. Meka, P. Jain, and I. Dhillon, “Guaranteed rank minimization via singular value projection,” in *Proc. Adv. in Neural Processing Systems (NIPS)*, Vancouver, BC, Dec. 2010.

- [11] I. Daubechies, M. Defrise, and C. De Mol, “An iterative thresholding algorithm for linear inverse problems with a sparsity constraint,” *Comm. Pure Appl. Math.*, vol. 57, no. 11, pp. 1413–1457, 2004.
- [12] R. Garg and R. Khandekar, “Gradient descent with sparsification: an iterative algorithm for sparse recovery with restricted isometry property,” in *Proc. Int. Conf. Machine Learning*, Montreal, Canada, Jun. 2009.
- [13] P. Shah and V. Chandrasekharan, “Iterative projections for signal identification on manifolds,” in *Proc. Allerton Conf. on Comm., Contr., and Comp.*, Monticello, IL, Sept. 2011.
- [14] D. Donoho, A. Maleki, and A. Montanari, “Message passing algorithms for compressed sensing,” *Proc. Natl. Acad. Sci.*, vol. 106, no. 45, pp. 18914–18919, 2009.
- [15] Z. Lin, M. Chen, L. Wu, and Y. Ma, “The augmented Lagrange multiplier method for exact recovery of corrupted low-rank matrices,” *Arxiv preprint arXiv:1009.5055*, 2010.
- [16] R. Baraniuk and M. Wakin, “Random projections of smooth manifolds,” *Found. Comput. Math.*, vol. 9, no. 1, pp. 51–77, 2009.
- [17] M. Elad and A. Bruckstein, “A generalized uncertainty principle and sparse representation in pairs of bases,” *IEEE Trans. Inform. Theory*, vol. 48, no. 9, pp. 2558–2567, 2002.
- [18] C. Studer, P. Kuppinger, G. Pope, and H. Bölcskei, “Recovery of sparsely corrupted signals,” *IEEE Trans. Inform. Theory*, 2012, To appear.
- [19] A. Feuer and A. Nemirovski, “On sparse representation in pairs of bases,” *IEEE Trans. Inform. Theory*, vol. 49, no. 6, pp. 1579–1581, 2003.
- [20] C. Grimes and D. Donoho, “Image manifolds which are isometric to euclidean space,” *J. Math. Imag. and Vision*, vol. 23, no. 1, pp. 5–24, 2005.
- [21] M. Wakin, D. Donoho, H. Choi, and R. Baraniuk, “High-resolution navigation on non-differentiable image manifolds,” in *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Processing (ICASSP)*, Philadelphia, PA, Mar. 2005.
- [22] A. Kyrillidis and V. Cevher, “Recipes for hard thresholding methods,” Tech. Rep., EPFL, Oct. 2011.
- [23] A. Waters, A. Sankaranarayanan, and R. Baraniuk, “SpaRCS: Recovering low-rank and sparse matrices from compressive measurements,” in *Proc. Adv. in Neural Processing Systems (NIPS)*, Granada, Spain, Dec. 2011.