SPIN: Iterative 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 *Successive Projection onto INcoherent manifolds* (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 signal recovery models and algorithms for low-dimensional linear inverse problems, and matches (or exceeds) the current state of the art in terms of performance.**

I. Introduction

Estimation of an unknown signal from linear observations is a core problem in signal processing, statistics, and information theory. Particular attention 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. Examples of such instances include morphological component analysis [1]; sparse approximation and compressive sensing [2, 3]; and affine rank minimization [4].

In this work, we will 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 *arbitrary nonlinear sub-manifolds* of the signal space $\mathbb{R}^N$. We are given access to noisy linear measurements of $\mathbf{x}^*$:

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

where $\mathbf{\Phi} \in \mathbb{R}^{M \times N}$ 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 that might appear to be insurmountable:

Identifiability I. Consider even 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 $\mathbf{x} = \mathbf{a}^* + \mathbf{b}^*$. In this case, $\mathbf{x}$ 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.

Identifiability II. To complicate matters, in more general situations the linear operator $\mathbf{\Phi}$ in (1) might have fewer rows that columns, so that $M < N$. Thus, $\mathbf{\Phi}$ possesses a nontrivial nullspace. Indeed, we are particularly interested in cases where $M \ll N$, i.e., the nullspace of Φ is extremely large relative to the ambient space. This further emphasizes the problem of identifiability of the ordered pair $(\mathbf{a}^*, \mathbf{b}^*)$, given the available observations $\mathbf{z}$.

Nonconvexity. Even if the above 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 applied; neither can the litany of convex optimization methods that have been designed for linear inverse problems with certain types of signal priors [3, 4].

In this paper, we propose a simple method called *Successive Projection onto INcoherent manifolds* (SPIN) to recover the component signals $(\mathbf{a}^*, \mathbf{b}^*)$ from $\mathbf{z}$. Despite the highly nonconvex nature of the problem and the possibility of underdetermined measurements, SPIN *provably* recovers the signal components $\mathbf{a}^*$ and $\mathbf{b}^*$. 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 $\mathbf{\Phi}$ 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 II.

SPIN is iterative in nature. Each iteration consists of three steps: computation of the gradient of the error function $\psi(\mathbf{a}, \mathbf{b}) = \frac{1}{2}\|\mathbf{z} - \mathbf{\Phi}(\mathbf{a} + \mathbf{b})\|^2$, 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. Interestingly, SPIN exhibits a convergence rate comparable to several state-of-the-art algorithms [5, 6], despite the nonlinear and nonconvex nature of the above reconstruction problem.

The core essence of our proposed approach has been extensively studied in a number of different contexts. In particular, SPIN is an iterative projected gradient method with the same basic approach as two recent signal recovery algorithms — Gradient Descent with Sparsification (GraDeS) [5] and Manifold Iterative Pursuit (MIP) [7]. Our method generalizes 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 complicated. An appealing feature of our method is its conceptual simplicity plus its ability to generalize to mixtures of arbitrary nonlinear manifolds.

Owing to space constraints, we state our theoretical claims and only provide a brief outline of the proofs following each statement. Refer to the extended version of this paper [8] for a thorough discussion of our approach, with complete proofs and additional applications.

II. Geometric Assumptions

The analysis and proof of accuracy of SPIN (Algorithm 1) involves 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 $\mathbf{\Phi}$ 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 $\mathbb{R}^N$.

A. Manifold Incoherence

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 $\mathcal{A}$.

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 via the secant manifold.

Definition 1: Suppose $\mathcal{A}, \mathcal{B}$ are submanifolds are $\mathbb{R}^N$. Let

$$\sup_{\mathbf{u} \in \mathcal{S}(\mathcal{A}),\ \mathbf{u}' \in \mathcal{S}(\mathcal{B})} |\langle \mathbf{u}, \mathbf{u}' \rangle| = \epsilon, \tag{2}$$

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 $\epsilon < 1$. 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 sketch. Let $\mathbf{x}_1 = \mathbf{a} + \mathbf{b}$, and $\mathbf{x}_2 = \mathbf{a}' + \mathbf{b}'$. Expand the relation $\|\mathbf{x}_1 - \mathbf{x}_2\|^2 = 0$, and apply the inequality between the arithmetic and geometric means to obtain the uniqueness result. □

In Section III, we will see that the condition for *exact recovery* of $(\mathbf{a}, \mathbf{b})$ from $\mathbf{x}$ will require a mild tightening of the upper bound on ϵ.

B. Restricted Isometry

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

Definition 2: Let $\mathcal{C}$ be a submanifold of $\mathbb{R}^N$. Then, the matrix $\mathbf{\Phi} \in \mathbb{R}^{M \times N}$ 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

$$1 - \delta \leq \|\Phi \mathbf{u}\|^2 \leq 1 + \delta. \tag{3}$$

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 [3, 4]. While the RIP has traditionally been studied in the context of sparse signal models, (3) generalizes of this notion to *arbitrary* nonlinear manifolds. A key result [9] states that under certain upper bounds on the curvature of the manifold $\mathcal{C}$, there exist probabilistic constructions of matrices $\mathbf{\Phi} \in \mathbb{R}^{M \times N}$ that satisfy the RIP on $\mathcal{C}$ such that the number of rows of $\mathbf{\Phi}$ is proportional to the intrinsic dimension of $\mathcal{C}$, rather than the ambient dimension N of the signal space.

Algorithm 1 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 $\widehat{\mathbf{a}} \in \mathcal{A}, \widehat{\mathbf{b}} \in \mathcal{B}$

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

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

$\quad \mathbf{g}_k \leftarrow \Phi^T \mathbf{r}$

$\quad \widetilde{\mathbf{a}}_k \leftarrow \mathbf{a}_k + \eta \mathbf{g}_k$, $\widetilde{\mathbf{b}}_k \leftarrow \mathbf{b}_k + \eta \mathbf{g}_k$

$\quad \mathbf{a}_{k+1} \leftarrow \mathcal{P}_{\mathcal{A}}(\widetilde{\mathbf{a}}_k)$, $\mathbf{b}_{k+1} \leftarrow \mathcal{P}_{\mathcal{B}}(\widetilde{\mathbf{b}}_k)$

$\quad \mathbf{r} \leftarrow \mathbf{z} - \Phi(\mathbf{a}_{k+1} + \mathbf{b}_{k+1})$

$\quad k \leftarrow k + 1$

end while

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

C. 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 $\mathcal{A}$:

$$\mathcal{P}_{\mathcal{A}}(\mathbf{x}) = \arg \min_{\mathbf{x}' \in \mathcal{A}} \|\mathbf{x}' - \mathbf{x}\|^2 . \qquad (4)$$

We observe that for arbitrary nonconvex manifolds $\mathcal{A}$, the minimization problem (4) 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 (4).

III. The SPIN Recovery Algorithm

We propose an efficient algorithm to solve the linear inverse problem (1) and detail it in pseudocode form in Algorithm 1. The algorithm can be viewed as a generalization of first order recovery methods that have been developed for a number of different models [5, 7]. The key innovation in SPIN is that we formulate *two* proxy vectors for the signal components, $\widetilde{\mathbf{a}}_k$ and $\widetilde{\mathbf{b}}_k$, and project these onto the corresponding manifolds $\mathcal{A}$ and $\mathcal{B}$.

We show that SPIN possesses strong *uniform recovery guarantees*, even under the presence of limited and highly inaccurate measurements. These guarantees are comparable to those exhibited by existing state-of-the-art signal recovery algorithms for sparse approximation and compressive sensing, while encompassing a very broad range of nonlinear signal models.

Theorem 1 (Stable recovery): Suppose $\mathcal{A}, \mathcal{B}$ are ϵ-incoherent manifolds in $\mathbb{R}^N$. Let $\mathbf{\Phi}$ 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 $\mathbf{z} = \mathbf{\Phi}(\mathbf{a}^* + \mathbf{b}^*) + \mathbf{e}$, where $\mathbf{a}^* \in \mathcal{A}, \mathbf{b}^* \in \mathcal{B}$. Suppose $0 \leq \delta < (1 - 11\epsilon)/(3 + 7\epsilon)$. Then, SPIN (Algorithm 1) with step size $\eta = 1/(1+\delta)$ and projection operators $\mathcal{P}_{\mathcal{A}}, \mathcal{P}_{\mathcal{B}}$ outputs $\mathbf{a}_T \in \mathcal{A}, \mathbf{b}_T \in B$ such that $\|\mathbf{z} - \mathbf{\Phi}(\mathbf{a}_T + \mathbf{b}_T)\|^2 \leq \beta \|\mathbf{e}\|^2 + \nu$ 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.

Proof sketch. For a given set of measurements $\mathbf{z}$ obeying (1), define the error function $\psi : \mathcal{A} \times B \rightarrow \mathbb{R}$ as: $\psi(\mathbf{a}, \mathbf{b}) = \frac{1}{2} \|\mathbf{z} - \mathbf{\Phi}(\mathbf{a} + \mathbf{b})\|^2$. In the extended version of this paper [8], we show that at any iteration k, the SPIN signal estimates satisfy the relation

$$\psi(\mathbf{a}_{k+1}, \mathbf{b}_{k+1}) \leq \alpha \psi(\mathbf{a}_k, \mathbf{b}_k) + C \|\mathbf{e}\|^2 , \qquad (5)$$

where

$$\alpha = \frac{\frac{2\delta}{1-\delta} + 6 \frac{1+\delta}{1-\delta} \frac{\epsilon}{1-\epsilon}}{1 - 4 \frac{1+\delta}{1-\delta} \frac{\epsilon}{1-\epsilon}}, \quad C = \frac{\frac{1}{2} + 5 \frac{1+\delta}{1-\delta} \frac{\epsilon}{1-\epsilon}}{1 - 4 \frac{1+\delta}{1-\delta} \frac{\epsilon}{1-\epsilon}}.$$

Equation (5) describes a linear recurrence relation for the sequence of positive real numbers $\psi(\mathbf{a}_k, \mathbf{b}_k)$, $k = 0, 1, 2, \ldots$ with leading coefficient $\alpha < 1$. By the 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 $T = \lceil \frac{1}{\log(1/\alpha)} \log \frac{\|\mathbf{z}\|^2}{2\nu} \rceil$, the result follows. □

Some implications of Theorem 1 are as follows. For the special case where there is no measurement noise (i.e., $\mathbf{e} = 0$), Theorem 1 states that after a finite number of iterations, SPIN outputs signal component estimates $(\widehat{\mathbf{a}}, \widehat{\mathbf{b}})$ such that $\left\|\mathbf{z} - \mathbf{\Phi}(\widehat{\mathbf{a}} + \widehat{\mathbf{b}})\right\| < \nu$ for any desired precision parameter ν. Since we can set ν to an arbitarily small value, we have that the SPIN estimates $(\widehat{\mathbf{a}}, \widehat{\mathbf{b}})$ converges to the true signals $(\mathbf{a}^*, \mathbf{b}^*)$.

Further, 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 [7]. 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 in Theorem 1 on the inequalities linking the isometry constant δ and the manifold incoherence parameter ϵ implies that $\epsilon < 1/11$. This represents a mild tightening of the condition on ϵ required for a unique decomposition (Lemma 1), even with full measurements (i.e., when $\mathbf{\Phi}$ is the identity operator, or more generally, when $\delta = 0$).

IV. Applications

The two-manifold signal model described in this paper is applicable to a wide variety of problems. We discuss two representative instances.

A. Articulation Manifolds

Consider the set 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: 1D acoustic chirps of varying frequencies (where θ represents the chirp rate); 2D images of a white disk translating on a black background (where θ represents the planar location of the disk center); and 2D images of a 3D 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 manifolds, it has been shown [9] that there exist *randomized* constructions of measurement operators $\mathbf{\Phi} \in \mathbb{R}^{M \times N}$ that satisfy the RIP on the secants of $\mathcal{M}$ with constant δ, and with probability at least ρ, provided

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

for some constant $C_{\mathcal{M}}$ that depends only on the smoothness and volume of the manifold $\mathcal{M}$. Therefore, the range space of $\mathbf{\Phi}$ is proportional merely to the *number of articulation parameters* K, and only logarithmic in the ambient dimension N. Moreover, given such a measurement matrix $\mathbf{\Phi}$ 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 $\mathbf{y} = \mathbf{\Phi}\mathbf{x}$ using Manifold Iterative Pursuit (MIP) [7].

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 2D images, where the manifolds $\mathcal{A}, \mathcal{B}$ comprise of translations of fixed template images $I_A(\mathbf{t}), I_B(\mathbf{t})$, where $\mathbf{t}$ denotes the 2D domain over which the image is defined. Then, the signal of interest is a 2D image of the form

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

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

We demonstrate that SPIN offers an easy, efficient technique to recover the component images. Figure 1(a) displays the results of SPIN recovery of a 64×64 image from merely $M = 50$ random Gaussian measurements. The unknown image consists of the linear sum of arbitrary translations of template images $I_A(\mathbf{t}), I_B(\mathbf{t})$. Here, the templates $I_A(\mathbf{t}), I_B(\mathbf{t})$ are assumed to be smoothed 0/1 images on a black background of a white disk and a white square respectively. From Fig. 1 we observe that SPIN is able to perfectly recover the component signals from very limited observations. Here, the operator $\mathcal{P}_{\mathcal{A}}(\mathbf{x})$ onto the manifold $\mathcal{A}$ consists of running a matched filter between the template $I(\mathbf{t})$ and the input signal $\mathbf{x}$, and returning $I(\mathbf{t} + \widehat{\theta})$, where the parameter value $\widehat{\theta}$ corresponds to the 2D location of the maximum of the matched filter response. This can be very efficiently carried out in $\mathcal{O}(N \log N)$ operations using the Fast Fourier Transform (FFT).

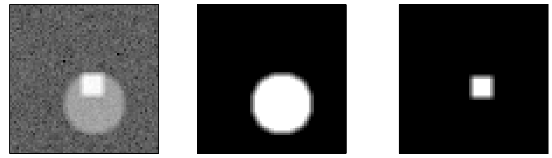

Original image Disk component Square component

Fig. 1. *SPIN recovery of a 64×64 image from compressive measurements. The 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. Here, the signal length $N = 64 \times 64 = 4096$, and the number of linear measurements $M = 50$. SPIN perfectly reconstructs both components from just M/N = 1.2% measurements.*

B. 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 correspond to the undesirable noise. In this case, the linear observations are more accurately modeled as

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

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 $\mathbf{\Phi}$ satisfies the RIP on the direct sum $\mathcal{M} \oplus \Sigma_{K'}$.

Figure 2 displays the results of a numerical experiment that illustrates the utility of SPIN in this setting. We consider a manifold of 1D signals of length $N = 10000$ comprising 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

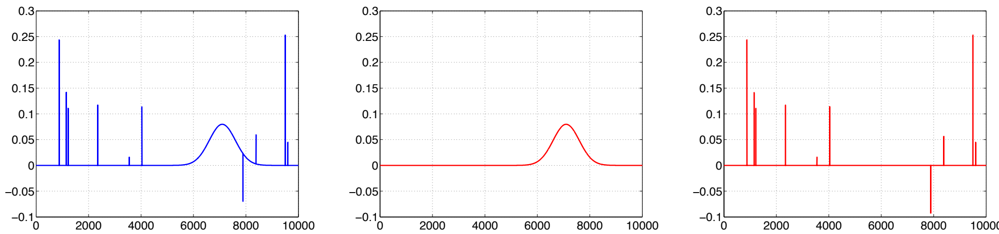

(a) Original signal. (b) Manifold component. (c) Noise component.

Fig. 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. $N = 10,000$, $K' = 10$, $M = 150$. (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.*

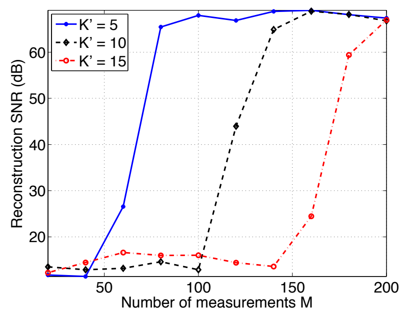

Fig. 3. *Monte Carlo simulation of SPIN recovery, averaged over 100 trials. In each trial, the measured signal is the sum of a randomly shifted Gaussian pulse and a random K'-sparse signal. SPIN can tolerate a higher K' with increasing M.*

is sampled using $M = 150$ random measurements to obtain an observation vector $\mathbf{z}$. For this problem, 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.

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, we can tolerate an increased number K' of nuisance spikes; we can show that this relationship between M and K' is in fact *linear* [8]. 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.

V. Summary

We have proposed a projected gradient-descent algorithm (SPIN) for the recovery of signals originating from a pair of incoherent manifolds, given limited and noisy measurements of their linear sum. 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 finite number Q of manifolds. Refer to the extended version of this paper [8] for full proofs, detailed discussions, and a comparison of SPIN to classical approaches.

Acknowledgements

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 Texas Instruments Leadership University Program.

References

[1] M. Elad, J.-L. Starck, P. Querre, and D. Donoho, "Simultaneous cartoon and texture image inpainting using morphological component analysis," *Appl. Comput. Harmon. Anal.*, vol. 19, no. 3, pp. 340–358, 2005.

[2] J. Tropp, "Greed is good: Algorithmic results for sparse approximation," *IEEE Trans. Inform. Theory*, vol. 50, no. 10, pp. 2231–2242, 2004.

[3] E. Candès, "Compressive sampling," in *Proc. Int. Congress of Math.*, Madrid, Spain, Aug. 2006.

[4] 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.

[5] 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.

[6] 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.

[7] 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.

[8] C. Hegde and R. Baraniuk, "Signal recovery on incoherent manifolds," *Arxiv preprint arXiv:1202.1595*, 2012, Preprint.

[9] R. Baraniuk and M. Wakin, "Random projections of smooth manifolds," *Found. Comput. Math.*, vol. 9, no. 1, pp. 51–77, 2009.