SIAM J. MATH. DATA SCI.
Vol. 2, No. 3, pp. 840–871

Analysis of Asymptotic Escape of Strict Saddle Sets in Manifold Optimization*

Thomas Y. Hou†, Zhenzhen Li†, and Ziyun Zhang†

Abstract. In this paper, we provide some analysis on the asymptotic escape of strict saddles in manifold optimization using the projected gradient descent algorithm (PGD). One of our main contributions is that we extend the current analysis to include nonisolated and possibly continuous saddle sets with complicated geometry. We prove that the PGD is able to escape strict critical submanifolds under certain conditions on the geometry and the distribution of the saddle point sets. We also show that the PGD may fail to escape strict saddles under weaker assumptions even if the saddle point set has zero measure. We apply this saddle analysis to the phase retrieval problem on the low-rank matrix manifold, prove that there are only a finite number of saddles, and that in a specific region, they are strict saddles with high probability. We also show the potential application of our analysis for a broader range of manifold optimization problems.

Key words. manifold optimization, projected gradient descent, strict saddles, stable manifold theorem, phase retrieval

AMS subject classifications. 58D17, 37D10, 65F10, 90C26

DOI. 10.1137/19M129437X

1. Introduction.

Manifold optimization has long been studied in mathematics. From signal and imaging science [43], [46], to computer vision [59], [60] and quantum information [34], [48], [65], it finds applications in various disciplines. However, it is the field of machine learning that sees the most diverse applications. Examples include matrix sensing [8] and matrix completion [61] on the low-rank matrix manifold $\{\boldsymbol{Z} : \boldsymbol{Z} \in \mathbb{R}^{n\times n}, rank(\boldsymbol{Z}) = r\}$; independent component analysis on the Stiefel manifold [1]; covariance estimation on the low-rank elliptope [20]; kernel learning, feature selection, and dimension reduction on the Grassmannian manifold [62]; dictionary learning on the special orthogonal group manifold [23]; blind deconvolution [28], [53]; point cloud denoising [43], [66]; tensor completion [11], [32], [35]; metric learning [51]; and Gaussian mixture [25], just to name a few. Among them, low-rank matrix recovery problems have drawn the most attention in recent years, mainly because of the intrinsic low-rank structure that naturally occurs in real-world applications.

Convergence analysis of manifold optimization is similar to that of classical optimization [55], except that spurious local minima and saddle points might occur due to the possible nonconvex structure of the manifold. However, numerous previous works on asymptotic convergence analysis reveal that many applications are actually void of spurious local minima [14],

*Received by the editors October 22, 2019; accepted for publication (in revised form) June 18, 2020; published electronically September 23, 2020.

https://doi.org/10.1137/19M129437X

Funding: This research was supported in part by NSF grants DMS-1613861, DMS-1907977, and DMS-1912654.

†Applied and Computational Mathematics, California Institute of Technology, Pasadena, CA 91125 USA (hou@cms.caltech.edu, zhenzhen@caltech.edu, zyzhang@caltech.edu).

[18], [41]. More importantly, they point out that the saddle points do not interfere with the convergence of various first-order algorithms towards a minimum, as long as they are "strict saddles," even though saddles are first-order stationary points [17], [30]. Intuitively speaking, a strict saddle point is a stationary point whose tangent space has a direction with negative curvature. The point series generated by the algorithm tend to "slip away" or escape from strict saddles because of this negative curvature direction. This "escape and convergence" property is observed not only for stochastic algorithms but also for deterministic ones, such as the simplest fixed-stepsize gradient descent.

The current analysis on strict saddles mostly focuses on isolated saddle points in the Euclidean space [37], [38]. There is a lack of systematic analysis on the Riemannian manifolds, and even less for the explicit treatment for nonisolated continuous saddle sets, despite their prevalence. One example is the variational linear eigenvalue problem on the sphere $\mathbb{S}^{n-1}$. It can have infinitely many saddle points if a nonminimal eigenvalue is duplicated. Another is the Gaussian phase retrieval problem, where the critical points are even more elusive due to the stochasticity of the model. This motivates us to provide some general analytic tools to study these algorithms.

One of the main contributions of this paper is that we provide a systematic analysis for the asymptotic escape of nonisolated and possibly continuous saddle sets with complicated geometry. Based on the analytic tools that we develop to study optimization on low-rank manifolds, we prove that the projected gradient descent algorithm (PGD) is able to escape strict critical submanifolds under certain conditions. These conditions are concerned with some geometric property of the saddle point set or the distribution of these saddle points. We argue that these conditions are necessary to guarantee the asymptotic escape of the strict saddles by the PGD in manifold optimization. However, these conditions are not stringent and are usually satisfied by familiar applications. We compare our conditions with those of the recent work [45] and point out that these two are consistent. We also give some examples that violate the conditions and result in failures of asymptotic escape, for the purpose of theoretical interest.

What lies at the core of this asymptotic analysis is an interesting interplay of dynamical systems and nonconvex optimization, and a translation of languages from the Morse theory [3], [4], [9] into gradient flow lines and further into gradient descents. Although these tools were initially developed to study homology, they have provided invaluable insight into the converging/escaping sets of strict saddle points with nontrivial geometry. We draw inspirations from them and propose a new unified tool to analyze asymptotic convergence and escape from saddle points.

We are aware that there is a parallel line of research on the stochastic/perturbed version of gradient descent, as well as its variant on the Riemannian manifold. Recent works including [10], [17], [30], [58] show that the stochastic/perturbed gradient descent is a powerful tool to get rid of saddles and does not impose any constraint on the geometry of saddle sets as we do. The reason our analysis focuses on the unperturbed gradient descent is that only by eliminating the perturbation effect can we single out the essential property of gradient descent itself. The development of a thorough asymptotic theory for this simple PGD algorithm is crucial to the understanding of why simple PGD works sufficiently well in many applications.

As an application of our asymptotic escape analysis for strict saddles, we consider the

phase retrieval problem [15], [19], [29], [52] that has received considerable attention in recent years. We combine the perspectives of Riemannian manifold optimization [7] and landscape analysis [41], [57] and derive new results. We analyze the saddle points of the phase retrieval problem on the low-rank matrix manifold. Surprisingly, we are able to prove that there are only a finite number of saddles and they are all strict saddles with very high probability. Our analysis provides a rigorous explanation for the robust performance of the PGD in solving the phase retrieval problem as a low-rank optimization problem, a phenomenon that has been observed in previous applications reported in the literature [41].

Although our primary focus is the low-rank matrix manifold, the asymptotic convergence to the minimum and the escape of strict saddles (strict critical submanifolds) are also valid on any arbitrary finite dimensional Riemannian manifold. In particular, the properties of the PGD are well preserved if the manifold is embedded in a Banach space and inherits its metric. Popular examples of manifold optimization include optimization problems on the sphere, the Stiefel manifold, the Grassmann manifold, the Wasserstein statistical manifold [40], the flag manifold for multiscale analysis [64], and the low-rank tensor manifold [24], [42]. They find applications in many fields, including physics, statistics, quantum information, and machine learning.

To illustrate this, we consider the optimization on the unit sphere and the Stiefel manifold as two examples of applications. In the first example, we consider a variational eigenvalue problem on a sphere. Both linear and nonlinear eigenvalue problems are considered. In the case of a linear eigenvalue problem, we show that the first eigenvector is the unique global minimum with a positive Hessian, and all the subsequent eigenvectors are all strict saddles whose Hessian has at least one negative curvature direction given by the eigenvector itself. Thus, our asymptotic escape analysis applies to this manifold optimization problem. In the case of the nonlinear eigenvalue problem, we cannot guarantee by our analysis that the saddles are all strict saddles. But we can verify this numerically. Our numerical results show that the PGD gives good convergence for both linear and nonlinear eigenvalue problems. In the case when we have a cluster of eigenvalues, we further propose an acceleration method by formulating a simultaneous eigensolver on the Stiefel manifold. We observe a considerable speedup in the convergence of the PGD method on the Stiefel manifold.

We point out that the analysis of this paper is purely asymptotic. The quantitative study of convergence rate will be the focus of our future work [27]. Notably, in contrast to the caveat that gradient descent could take exponential time to escape saddles in the worst-case scenario [12], empirical evidence in the figures of sections 4 and 5 and many examples reported in the literature demonstrate that we usually have almost linear convergence to minimizers.

The rest of the paper is organized as follows. Section 2 contains the main results on asymptotic escape of strict saddles (in particular nonisolated ones) on the Riemannian manifolds. In section 3, we explore the geometric structure of the closed low-rank matrix manifold $\overline{\mathcal{M}_r}$. In section 4, phase retrieval is analyzed as an example of asymptotic escape of saddles on $\overline{\mathcal{M}_r}$. We extend the application to other manifolds and a broader scope of problems in section 5. Finally, we make some concluding remarks and discuss future works in section 6.

2. The projected gradient descent and asymptotic escape of nonisolated strict saddles.

In this section, we discuss the optimization technique we use, namely the projected gradient

descent algorithm (PGD) with retraction onto the manifold. We will prove an important property of the proposed technique; namely it is able to escape any strict saddle points and always converge to minimizers. This may be obvious in the Euclidean spaces but not so obvious on manifolds.

We stress that although our motivations (and hence notation conventions) are based on the low-rank matrix manifold, what we discuss in this section can be easily generalized to arbitrary finite dimensional Riemannian manifolds, as is mentioned in section 1. This works as long as either (1) $\mathcal{M}$ is embedded in an ambient Banach space and inherits its metric (so that the *embedded* gradient ∇ is well defined), and there exists a well-defined first-order retraction R ; or (2) $\mathcal{M}$ is "flat" so that no retraction R is needed.

The asymptotic escape of strict saddle points combined with a good landscape of the objective function can lead to asymptotic convergence to minimizers. This could serve as a fundamental tool for various application tasks. More applications will be discussed section 5.

2.1. Projected gradient descent on the manifold. Assume we are given a function $f(\cdot) : \mathcal{M} \to \mathbf{R}$ where $\mathcal{M}$ can be a general manifold. We start from a proper initial guess $\boldsymbol{Z}_0 \in \mathcal{M}$. The iteration points $\{\boldsymbol{Z}_n\}_{n=0}^N$ are generated by

$$\boldsymbol{Z}_{n+1} = \mathcal{R}\left(\boldsymbol{Z}_n - \alpha_n P_{T_{Z_n}}(\nabla f(\boldsymbol{Z}_n))\right). \tag{2.1}$$

Here ∇f is the *embedded* gradient of f in its ambient Banach space , $P_{T_{Z_n}}$ is the projection onto the tangent space of $\mathcal{M}$ at point $\boldsymbol{Z}_n$, α_n is the nth stepsize, and $\mathcal{R} : T_Z \to \mathcal{M}$ is a first-order retraction. The retraction operation is necessary in that it makes sure the generated iteration point still stays on the manifold $\mathcal{M}$. Specifically, we define the retraction as follows.

Definition 2.1 (Retraction). *Let $\|\cdot\|$ be the norm of the embedded Banach space of $\mathcal{M}$. Let $T_{\boldsymbol{Z}}$ be the tangent space (or tangent cone) of $\mathcal{M}$ at $\boldsymbol{Z}$. We call $\mathcal{R}_{\boldsymbol{Z}} : T_{\boldsymbol{Z}} \to \mathcal{M}$ a* retraction *if, for any $\xi \in T_{\boldsymbol{Z}}$,*

$$\lim_{\alpha \to 0^+} \frac{\|\mathcal{R}_{\boldsymbol{Z}}(\alpha\xi) - (\boldsymbol{Z} + \alpha\xi)\|}{\alpha} = 0.$$

We also write $\mathcal{R}(\boldsymbol{Z} + \alpha\xi)$ equivalently for $\mathcal{R}_{\boldsymbol{Z}}(\alpha\xi)$.

We will refer to (2.1) as the *first-order retraction property*.

Remark 2.2. It is worth mentioning that there exist other manifold optimization techniques, e.g., without the projection step $T_{\boldsymbol{Z}_n}$. We choose the current PGD here mainly for two reasons:

(1) Most of the operations we list here rely heavily on tangent bundles. For example, the first-order retraction (2.1) only holds for $\xi \in T_{\boldsymbol{Z}}$. The embedded gradient ∇f is not always in $T_{\boldsymbol{Z}}$—only $P_{T_{\boldsymbol{Z}}}(\nabla f(\boldsymbol{Z}))$ is. In fact, only the projected embedded gradient corresponds to the true *Riemannian* gradient: $P_{T_{\boldsymbol{Z}}}(\nabla f(\boldsymbol{Z})) = \mathrm{grad} f(\boldsymbol{Z})$, where $\mathrm{grad} f(\boldsymbol{Z})$ is the Riemannian gradient of f on $\mathcal{M}$ at point $\boldsymbol{Z}$.

(2) Retraction $\mathcal{R}_{\boldsymbol{Z}}$ can be computed more efficiently when the increment lies in the tangent space. For example, on the low-rank matrix manifold it involves a smaller-scale SVD [50], [61], [63].

2.2. Asymptotic escape of isolated saddle points. Consider the PGD method with a fixed stepsize α. Let φ be the iteration operation and

$$\boldsymbol{Z}_{n+1} = \varphi(\boldsymbol{Z}_n) := R(\boldsymbol{Z}_n - \alpha P_{T_{\boldsymbol{Z}_n}}(\nabla f(\boldsymbol{Z}_n))).$$

Let $\boldsymbol{X}^*$ be a critical point of f. The most common nonminimizer critical point is a "strict saddle." Intuitively speaking, a strict saddle point $\boldsymbol{Z}^*$ is a point around which there is a strictly decreasing direction (Hess$f(\boldsymbol{Z}^*)$ has a negative eigenvalue) and no flat direction (Hess$f(\boldsymbol{Z}^*)$ has no zero eigenvalue). To rigorously define strict saddle points we need the following definitions.

Definition 2.3 (Levi-Civita connection). *The* Levi-Civita connection $\widetilde{\nabla}_\xi \eta$, *acting on two vectors or vector fields* η, ξ *in the tangent bundle* $T\mathcal{M}$, *is the unique affine connection on* $\mathcal{M}$ *that preserves the metric and is torsion-free.*

Note that this is not to be confused with the operator ∇ in (2.1), which is used to denote the gradient in the ambient space.

Definition 2.4 (Riemannian gradient). *Given* $f : \mathcal{M} \to \mathbb{R}$, *the* Riemannian gradient *of* f *is the vector field* $gradf$, *such that for any vector field* Y *on* $\mathcal{M}$,

$$\langle gradf, Y\rangle = Y(f),$$

where $\langle\cdot,\cdot\rangle$ *is the metric on* $\mathcal{M}$ *and* $Y(\cdot)$ *is the vector field action, i.e.,* $Y(f) = \sum_i Y_i \frac{\partial f}{\partial E_i}$ *for a basis* $\{E_i\}$.

Remark 2.5. The Riemannian gradient is equivalent to the tangent space projection of the embedded gradient in the ambient space, i.e.,

$$\text{grad } f(\boldsymbol{Z}) = P_{T_{\boldsymbol{Z}}}(\nabla f(\boldsymbol{Z})).$$

Further, if the metric of $\mathcal{M}$ is inherited from the ambient space, then the Levi-Civita connection on $\mathcal{M}$ is the tangent space projection of the Levi-Civita connection (natural gradient) of the ambient space. Specifically, for η, $\xi \in T\mathcal{M}$, we have

$$\widetilde{\nabla}_\xi \eta = P_{T_{\boldsymbol{Z}}}(\nabla \eta[\xi]).$$

Definition 2.6 (Hessian). *Given a function* $f : \mathcal{M} \to \mathbb{R}$, *the* Riemannian Hessian *of* f *at point* $\boldsymbol{Z}$ *is Hess* $f(\boldsymbol{Z}) : T_{\boldsymbol{Z}}\mathcal{M} \to T_{\boldsymbol{Z}}\mathcal{M}$ *defined by*

$$Hess\ f(\boldsymbol{Z})[\xi] = \widetilde{\nabla}_\xi grad\ f(\boldsymbol{Z}), \tag{2.2}$$

where $\widetilde{\nabla}_{(\cdot)}(\cdot)$ *is the Levi-Civita connection on* $\mathcal{M}$.

Proposition 2.7. *If the retraction is second-order, i.e.,*

$$P_{T_{\boldsymbol{Z}}}\left(\frac{d^2}{d\alpha^2}\mathcal{R}_{\boldsymbol{Z}}(\alpha\xi)\mid_{\alpha=0}\right) = 0,$$

then

$$Hess\ f(\boldsymbol{Z}) = Hess\ (f \circ \mathcal{R}_{\boldsymbol{Z}})(0). \tag{2.3}$$

In particular, (2.3) is true for the low-rank matrix manifold, and this gives a more computable definition of Hessian; see also [61]. It is proved in [2] that in the case of the low-rank matrix manifold, the above expression recovers Definition 2.6.

Definition 2.8 (Strict saddle point). *We call $\boldsymbol{Z} \in \mathcal{M}$ a* strict saddle point *of $f(\cdot) : \mathcal{M} \to \mathbb{R}$ if*

1. $P_{T_{\boldsymbol{Z}}}(\nabla f(\boldsymbol{Z})) = 0$;
2. *Hess $f(\boldsymbol{Z})$ as a linear operator has at least one positive eigenvalue, at least one negative eigenvalue, and no zero eigenvalue.*

In contrast to Definition 2.8, we call a point $\boldsymbol{Z} \in \mathcal{M}$ a *local minimizer* if Hess $f(\boldsymbol{Z})$ is positive semidefinite. Note that this is only legitimate in the broader sense, as the PGD does not distinguish between local minima and degenerate saddles, i.e., saddles that only have higher than second-order negative curvature.

Then we have the following main result.

Theorem 2.9 (PGD asymptotically only converges to a local minimum). *Let $f(\cdot) : \mathcal{M} \to \mathbb{R}$ be a C^2 function on $\mathcal{M}$. Suppose that f has either finitely many saddle points (denoted as set S), or countably many saddle points in a compact submanifold of $\mathcal{M}$, and all saddle points of f are strict saddles as defined in Definition 2.8. Let C be set of all local minimizers and $\{\boldsymbol{Z}_n\}$ be the series of points generated by the projected gradient descent algorithm on $\mathcal{M}$; then we have*

1. $Pr(\lim_{n\to\infty} \boldsymbol{Z}_n \in S) = 0$;
2. *if $\lim_{n\to\infty} \boldsymbol{Z}_n$ exists, then $Pr(\lim_{n\to\infty} \boldsymbol{Z}_n \in C) = 1$.*

In other words, the PGD with a random initialization is unlikely to converge to a saddle point $\boldsymbol{Z}^*$ as long as $\boldsymbol{Z}^*$ is a strict saddle.

To prove this result, the main tool is the stable manifold theorem on low-rank matrix manifolds, which is an extension of similar results in the Euclidean spaces.

Definition 2.10 (Definition 4.13 in [4]). *A fixed point $p \in \mathcal{M}$ of a smooth diffeomorphism $\varphi : \mathcal{M} \to \mathcal{M}$ is called* hyperbolic *iff none of the complex eigenvalues of $D\varphi_p : T_p\mathcal{M} \to T_p\mathcal{M}$ have length 1, where $D\varphi$ is the Riemannian gradient of φ at point p.*

For a fixed point p of φ, there is a splitting of $T_p\mathcal{M}$ that is preserved by φ:

$$D\varphi_p : T_p^s\mathcal{M} \oplus T_p^u\mathcal{M} \to T_p^s\mathcal{M} \oplus T_p^u\mathcal{M}, \tag{2.4}$$

where φ is contracting on $T_p^s\mathcal{M}$ and expanding on $T_p^u\mathcal{M}$. Since the manifold $\mathcal{M}$ is finite dimensional, for a hyperbolic fixed point p on $\mathcal{M}$, we can always find a $\lambda \in (0, 1)$ such that

$$\|D\varphi|_{T_p^s\mathcal{M}}\| < \lambda, \quad \|(D\varphi|_{T_p^u\mathcal{M}})^{-1}\| < \lambda.$$

Theorem 2.11 (Theorem 4.15 in [4]). *If $\varphi : \mathcal{M} \to \mathcal{M}$ is a smooth diffeomorphism of a finite dimensional smooth manifold $\mathcal{M}$, and p is a hyperbolic fixed point of φ, then*

$$W_p^s(\varphi) := \{x \in \mathcal{M} | \lim_{n\to\infty} \varphi^n(x) = p\}$$

is an immersed submanifold of $\mathcal{M}$ with $T_pW_p^s(\varphi) = T_p^s\mathcal{M}$. Moreover, $W_p^s(\varphi)$ is the surjective image of a smooth injective immersion

$$E^s : T_p^s\mathcal{M} \to W_p^s(\varphi) \subseteq \mathcal{M}.$$

Hence, $W_p^s(\varphi)$ is a smooth injectively immersed open disk in $\mathcal{M}$. We call it the stable manifold *of p with respect to φ. The* unstable manifold *$W_p^u(\varphi)$ of p can also be defined accordingly.*

Proof of Theorem 2.9. From Theorem 2.11, if φ is a diffeomorphism and $\boldsymbol{Z}^*$ is a hyperbolic point of φ, then $W^S_{\boldsymbol{Z}^*}(\varphi)$, the stable set of $\boldsymbol{Z}^*$, will be a lower dimensional submanifold in $\overline{\mathcal{M}}$. Then, the converging set of $\boldsymbol{Z}^*$ will have measure 0 with respect to the manifold, and any random initialization of PGD will escape such a strict saddle point almost surely.

We first prove that $\boldsymbol{Z}^*$ is a hyperbolic point. The diffeomorphic property of φ is actually contained in the proof of the former.

Given $\xi \in T_{\boldsymbol{Z}^*}\mathcal{M}$, for any $\tilde{\xi}$ that satisfies $\boldsymbol{Z}^* + \tilde{\xi} \in \mathcal{M}$, $P_{T_{\boldsymbol{Z}^*}}(\tilde{\xi}) = \xi$, we have

$$\begin{aligned}
\mathrm{Hess} f(\boldsymbol{Z}^*)[\xi] + \mathcal{O}(\|\xi\|^2) &= \widetilde{\nabla}_\xi \ \mathrm{grad} f(\boldsymbol{Z}^*) \\
&= P_{T_{\boldsymbol{Z}^*}}(\nabla \mathrm{grad} f(\boldsymbol{Z}^*)[\xi]) \\
&= P_{T_{\boldsymbol{Z}^*}}(\mathrm{grad} f(\boldsymbol{Z}^* + \tilde{\xi}) - \mathrm{grad} f(\boldsymbol{Z}^*)) + \mathcal{O}(\|\xi\|^2) \\
&= P_{T_{\boldsymbol{Z}^*}}(\mathrm{grad} f(\boldsymbol{Z}^* + \tilde{\xi})) + \mathcal{O}(\|\xi\|^2).
\end{aligned}$$

Note that $\mathrm{grad} f(\boldsymbol{Z}^*) = 0$ since Z^* is a critical point. Therefore, for $\varphi(\boldsymbol{Z}_n) = R(\boldsymbol{Z}_n - \alpha P_{T_{\boldsymbol{Z}_n}}(\nabla f(\boldsymbol{Z}_n)))$,

$$\begin{aligned}
D\varphi_{\boldsymbol{Z}^*}[\xi] &= P_{T_{\boldsymbol{Z}^*}}(\varphi(\boldsymbol{Z}^* + \tilde{\xi}) - \varphi(\boldsymbol{Z}^*)) + \mathrm{o}(\|\xi\|) \\
&= P_{T_{\boldsymbol{Z}^*}}(R(\boldsymbol{Z}^* + \tilde{\xi} - \alpha \ \mathrm{grad} f(\boldsymbol{Z}^* + \tilde{\xi})) - \boldsymbol{Z}^*) + \mathrm{o}(\|\xi\|) \\
&= P_{T_{\boldsymbol{Z}^*}}(\boldsymbol{Z}^* + \tilde{\xi} - \alpha \ \mathrm{grad} f(\boldsymbol{Z}^* + \tilde{\xi}) + \mathrm{o}(\|\xi\|) - \boldsymbol{Z}^*) \\
&= P_{T_{\boldsymbol{Z}^*}}(\tilde{\xi} - \alpha \ \mathrm{grad} f(\boldsymbol{Z}^* + \tilde{\xi})) + \mathrm{o}(\|\xi\|) \\
&= \xi - \alpha \ \mathrm{Hess} f(\boldsymbol{Z}^*)[\xi] + \mathrm{o}(\|\xi\|).
\end{aligned}$$

We have

$$D\varphi_{\boldsymbol{Z}^*}[\xi] = \xi - \alpha \mathrm{Hess}(f)(\boldsymbol{Z}^*)[\xi],$$

i.e.,

$$D\varphi_{\boldsymbol{Z}^*} = I - \alpha \ \mathrm{Hess}(f)(\boldsymbol{Z}^*).$$

Thus $\boldsymbol{Z}^*$ being strict saddle implies that, by choosing α sufficiently small (but only depending on the eigenvalues of $\mathrm{Hess}(f)(\boldsymbol{Z}^*)$), $\boldsymbol{Z}^*$ is hyperbolic with respect to φ.

Now, φ is a diffeomorphism at $\boldsymbol{Z}^*$ because, if we choose α small enough so that $\|\mathrm{Hess}(f)(\boldsymbol{Z}^*)\| < \frac{1}{\alpha}$, then $D\varphi$ is always invertible and bounded. If there are only finitely many strict saddle points, or there are a countably infinite number of them in a compact region, $\|\mathrm{Hess}(f)(\boldsymbol{Z}^*)\|$ shall be upper bounded, and such an α is always attainable.

Using Theorem 2.11, the set of points on $\mathcal{M}$ that converge to $\boldsymbol{Z}^*$ is a lower dimensional submanifold in $\mathcal{M}$, which has measure 0. We could safely deduce that, by randomly sampling a start point $\boldsymbol{Z}_0$ in $\mathcal{M}$, the probability of converging to a strict saddle point is 0, i.e.,

$$\text{Prob}(\lim_{k\to\infty} \boldsymbol{Z}_k = \boldsymbol{Z}^*) = 0.$$

Since there are only countably many strict saddle points, $\cup_{\boldsymbol{Z}^*\in S} W^S_{\boldsymbol{Z}^*}(\varphi)$ still has measure 0. So the PGD with a randomly sampled starting point converges to any point in S with probability 0. This proves the first argument.

As for the second argument, since the stepsize is a constant α, the only stationary points of the algorithm are first-order critical points of the loss function. The local maximizers are ruled out by the descent property of gradient descent. So if the limit point exists, it is a local minimizer with probability 1. ∎

Remark 2.12. As is mentioned at the beginning of this section, another case is when the manifold $\mathcal{M}$ is "flat" and no retraction is needed, i.e., $\boldsymbol{Z} + \xi \in \mathcal{M}$ for any $\boldsymbol{Z} \in \mathcal{M}$ and $\xi \in T_{\boldsymbol{Z}}\mathcal{M}$. If no embedding Banach space is present, the optimization technique might be reduced to using the Riemannian gradient directly. The above proof is also reduced to applying D directly to $\boldsymbol{Z} - \text{grad } f(\boldsymbol{Z})$ and is trivially true.

The natural question to ask is "What if $\text{Hess}(f)(X^*)$ has zero eigenvalue(s) in addition to negative eigenvalues?" In this case, $D\varphi$ has *center* directions that are neither expanding nor contracting. This can be taken into account by extending the stable manifold theorem to a slightly stronger version, which is the following.

Theorem 2.13. *Let $\varphi : \mathcal{M} \to \mathcal{M}$ be a smooth diffeomorphism of a finite dimensional smooth manifold $\mathcal{M}$, and p is a fixed point of φ. Assume that*

$$T_p\mathcal{M} = T^s_p\mathcal{M} \oplus T^c_p\mathcal{M} \oplus T^u_p\mathcal{M}, \tag{2.5}$$

which is the invariant splitting of $T_p\mathcal{M}$ into contracting, centering, and expanding subspaces corresponding to eigenvalues of magnitude less than, equal to, and greater than 1. *Let*

$$T^{cs}_p\mathcal{M} := T^s_p\mathcal{M} \oplus T^c_p\mathcal{M}.$$

Then we have that

$$W^s_p(\varphi) := \{x \in \mathcal{M} | \lim_{n\to\infty} \varphi^n(x) = p\}$$

is an immersed submanifold of $\mathcal{M}$ and $T_pW^s_p(\varphi) \subseteq T^{cs}_p\mathcal{M}$. We call it the (generalized) stable manifold *of p with respect to φ.*

For those who are interested in the proof, a detailed one for the Euclidean case can be found in Theorem III.7 in [54], and the extension to the manifold is similar to Theorem 2.11.

Definition 2.14 (Strict saddle point, the general definition). *We call $\boldsymbol{Z} \in \mathcal{M}$ a* strict saddle point *to $f(\cdot) : \mathcal{M} \to \mathbb{R}$ if*

1. $P_{T_{\boldsymbol{Z}}}(\nabla f(\boldsymbol{Z})) = 0$;

2. *Hess* $f(\boldsymbol{Z})$ *has at least one negative eigenvalue.*

Thus, for a strict saddle point $\boldsymbol{Z}^*$ where Hess(f) possibly has zero eigenvalue(s), we still have $\dim(W_p^s(\varphi)) \le \dim(T_p^{cs}\mathcal{M}) < \dim \mathcal{M}$, and we have $\text{Prob}(\lim_{k\to\infty} \boldsymbol{Z}_k = \boldsymbol{Z}^*) = 0$. Therefore, we have the following theorem.

Theorem 2.15. *Let* $f(\cdot) : \mathcal{M} \to \mathbb{R}$ *be a* C^2 *function on* $\mathcal{M}$. *Suppose that* $f(\cdot) : \mathcal{M} \to \mathbb{R}$ *has either finitely many saddle points, or countably many saddle points in a compact submanifold of* $\mathcal{M}$, *and all saddle points of* f *are strict saddles as defined in Definition* 2.14. *Then the results of Theorem* 2.9 *still hold, i.e.,*

1. $Pr(\lim_{n\to\infty} \boldsymbol{Z}_n \in S) = 0$;
2. *if* $\lim_{n\to\infty} \boldsymbol{Z}_n$ *exists, then* $Pr(\lim_{n\to\infty} \boldsymbol{Z}_n \in C) = 1$.

2.3. A closer look at the critical points: Nonisolated case. As is mentioned in section 1, it is very common that there are more than a countable number of strict saddle points, e.g., when they form a submanifold, or a more complicated set, with Lebesgue measure 0. Empirical evidence shows satisfactory convergence of the PGD to its minimum, which indicates successful escape from these strict saddles. But there is a lack of theoretical analysis to confirm this observation. In the following, we will use some further results from the Morse theory and its extensions to provide an analytical tool for this purpose.

Definition 2.16 (Critical submanifold). *For* $f : \mathcal{M} \mapsto \mathbb{R}$, *a connected submanifold* $\mathcal{N} \subset \mathcal{M}$ *is called a* critical submanifold *of* f *if every point* $\boldsymbol{Z}$ *in* $\mathcal{N}$ *is a critical point of* f, *i.e., grad* $f(\boldsymbol{Z}) = 0$ *for any* $\boldsymbol{Z} \in \mathcal{N}$.

Definition 2.17 (Strict critical submanifold). *A critical submanifold* $\mathcal{N}$ *of* f *is called a* strict critical submanifold *if, for all* $\boldsymbol{Z} \in \mathcal{N}$,

$$\lambda_{min}(Hess\ f(\boldsymbol{Z})) \le c < 0,$$

where $\lambda_{min}(\cdot)$ *takes the smallest eigenvalue, and* $c = c(\mathcal{N})$ *is a uniform constant for all* $\boldsymbol{Z} \in \mathcal{N}$ *depending only on* $\mathcal{N}$.

Analogous to the stable/unstable manifolds of critical points in Theorems 2.11 and 2.13, we may define stable/unstable manifolds of critical submanifolds.[1]

Definition 2.18 (Generalized stable/unstable manifold). *Let* $\varphi : \mathcal{M} \to \mathcal{M}$ *be a smooth diffeomorphism of* $\mathcal{M}$. *Then for a submanifold of* $\mathcal{N} \subset \mathcal{M}$, *the* stable manifold *and* unstable manifold *of* $\mathcal{N}$ *with respect to* φ *are defined as*

$$W_{\mathcal{N}}^s(\varphi) := \{x \in \mathcal{M} | \lim_{n\to\infty} \varphi^n(x) \in \mathcal{N}\},$$
$$W_{\mathcal{N}}^u(\varphi) := \{x \in \mathcal{M} | \lim_{n\to-\infty} \varphi^n(x) \in \mathcal{N}\}.$$

Given a nontrivial strict critical submanifold $\mathcal{N}$ of f, at every point $p \in \mathcal{N}$, the tangent space is split as

$$T_p\mathcal{M} = T_p\mathcal{N} \oplus \nu_p\mathcal{N},$$

[1]The reader shall be careful while distinguishing different "manifolds": the domain of the function f is a manifold, and the critical set of f is now a submanifold, but the names of stable/unstable manifolds are given regardless of the domain of f.

where $\nu_p\mathcal{N}$ is the normal space of $\mathcal{N}$ at p immersed in $\mathcal{M}$. Similar to (2.5), it is further split into

$$T_p\mathcal{M} = T_p\mathcal{N} \oplus (\nu_p^s\mathcal{M} \oplus \nu_p^c\mathcal{M} \oplus \nu_p^u\mathcal{M}).$$

To arrive at a result similar to that stated in Theorem 2.13, it suffices to define

$$T_p^{cs}\mathcal{M} := T_p\mathcal{N} \oplus (\nu_p^s\mathcal{M} \oplus \nu_p^c\mathcal{M})$$

and notice that $T_pW_{\mathcal{N}}^s(\varphi) \subseteq T_p^{cs}\mathcal{M}$ for any $p \in \mathcal{N}$. Since $T_p^{cs}\mathcal{M}$ is dimension deficient by the definition of a strict critical submanifold, any random initialization still falls into the converging set of $\mathcal{N}$ with probability 0. Because the union of a finite number of 0-measure sets still has measure 0, the above result works well with countably many critical submanifolds. To sum up, we have the following theorem.

Theorem 2.19. *Let $f(\cdot) : \mathcal{M} \to \mathbb{R}$ be a C^2 function on $\mathcal{M}$. Suppose that the saddle set S of $f(\cdot) : \mathcal{M} \to \mathbb{R}$ consists of finitely many critical submanifolds, or countably many critical submanifolds in a compact region of $\mathcal{M}$, and all of them are strict critical submanifolds as defined in Definition 2.18. Then the results of Theorem 2.9 still hold, i.e.,*

1. $Pr(\lim_{n\to\infty} \boldsymbol{Z}_n \in S) = 0$;
2. *if* $\lim_{n\to\infty} \boldsymbol{Z}_n$ *exists, then* $Pr(\lim_{n\to\infty} \boldsymbol{Z}_n \in C) = 1$.

For situations more complicated than those stated in the above theorem, it is conjectured that the transversality relationship of submanifolds can be exploited to find out the succession relationship of critical sets. We refer the reader to Appendix A for some useful tools and interesting insights in this direction.

Finally, we point out that the number of critical submanifolds being countable is an essential condition, but not a binding one. Of course, one reason for this statement is that it is often satisfied in practice. Namely, in the known applications with very complicated saddle geometries, e.g., matrix factorization and nonlinear eigenproblems, the saddles can still be grouped into countably many points or submanifolds. In those cases, Theorem 2.9, Theorem 2.15, or Theorem 2.19 is applicable.

But even from a purely theoretical point of view, the number of strict critical submanifolds being uncountable is unlikely. This is in accordance with the result of a recent work [45]. The result explicitly includes the case of "uncountably many critical points," but from the viewpoint of submanifolds, such a result belongs to the case of "countably many submanifolds" in our Theorem 2.19. (A submanifold can contain uncountably many points but is still a single object to escape.) This can also be inferred from the use of a countable subcover in Theorem 10 of [45] and the subsequent proof of the main theorem, where the convergence set to any saddle is categorized into a countable number of stable manifolds.

To further illustrate this point, here we give some interesting examples. The saddle sets in Example 2.20 occupy a zero measure set in the whole manifold. They cannot be assembled into countably many connected submanifolds. We will analyze and see why they cannot be almost surely avoided.

Example 2.20. Let $\mathcal{M} = [-1, 2] \times [-1, 1] \subset \mathbb{R}^2$ be a rectangular region, viewed as a 2-dimensional submanifold of $\mathbb{R}^2$. Then the tangent space $T\mathcal{M}$ equals $\mathcal{M}$. To construct the

function f on $\mathcal{M}$, we need the 1-dimensional Smith–Volterra–Cantor ("fat Cantor") set V in $[0,1]$. The construction is as follows:

(1) Remove the middle interval of length $\frac{1}{4}$ from $[0,1]$, and the remaining set is $[0,\frac{3}{8}]\cup[\frac{5}{8},1]$;
(2) Remove 2 middle subintervals of length $\frac{1}{4^2}$ from the 2 remaining intervals, and the remaining set is $[0,\frac{5}{32}]\cup[\frac{7}{32},\frac{3}{8}]\cup[\frac{5}{8},\frac{25}{32}]\cup[\frac{27}{32},1]$;
(3) Remove 4 middle subintervals of length $\frac{1}{4^3}$ from the 4 remaining intervals;
(4) ...

A visualization of the construction is given in Figure 1a. The construction is different from that of the classical Cantor set in that we remove proportionally shorter subintervals, instead of subintervals proportional to the mother interval. Therefore, V has positive measure in $\mathbb{R}$, $meas(V)=\frac{1}{2}$, while the classical Cantor set has zero measure. Still, V is nowhere dense.

We look for a synthetic objective function on $\mathcal{M}$ in the form

$$f:\mathcal{M}\to\mathbb{R},\quad f(x,y)=-p(x)+y^2,$$

where $p(x)$ is a function of certain regularity on the 1-dimensional interval $x\in[-1,2]$. Consider two examples:

(A) Define $p_A(x)=0$ for $x\in V$. As V is a closed set, write $V^c=[-1,2]\backslash V=(\bigcup_\alpha(a_\alpha,b_\alpha))\cup[-1,0)\cup(1,2]$ as the disjoint union of intervals. On each interval (a,b), let

$$p_A(x)=\begin{cases}(x-a)^2 & \text{for}\quad a<x\le a+\frac{(b-a)}{4},\\ C_1(x-\frac{a+b}{2})^4+C_2(x-\frac{a+b}{2})^2+C_3 & \text{for}\quad a+\frac{(b-a)}{4}<x\le b-\frac{(b-a)}{4},\\ (b-x)^2 & \text{for}\quad b-\frac{(b-a)}{4}<x<b,\end{cases}$$

where $C_1=\frac{8}{(b-a)^2}$, $C_2=-2$, $C_3=\frac{5(b-a)^2}{32}$. See a visualization in Figures 1b and 1c.

(B) Similar to (A), but on each interval (a,b), let

$$p_B(x)=\begin{cases}(x-a)^4 & \text{for}\quad a<x\le a+\frac{(b-a)}{4},\\ C_1(x-\frac{a+b}{2})^6+C_2(x-\frac{a+b}{2})^4+C_3 & \text{for}\quad a+\frac{(b-a)}{4}<x\le b-\frac{(b-a)}{4},\\ (b-x)^4 & \text{for}\quad b-\frac{(b-a)}{4}<x<b,\end{cases}$$

where $C_1=\frac{512}{3(b-a)^4}$, $C_2=-\frac{24}{(b-a)^2}$, $C_3=\frac{11(b-a)^2}{96}$.

It is easy to see that both functions $p_i(x)$, $i=A$ or B, satisfy $p(x)\ge 0$ and $p(x)=0$ iff $x\in V$. Thus for $f_i(x,y)=-p_i(x)+y^2$, the saddle set of f is $S=V\times[0]$. Viewed in the 2-dimensional manifold, it has zero measure. But the converging set of S is $W_S^s(\varphi)=V\times[-1,1]$. It has positive measure in $\mathcal{M}$: $meas(W_S^s(\varphi))=1$. If we start the PGD with a uniform random initialization, the probability that $\{\boldsymbol{Z}_n\}_{n=0}^\infty$ end up towards a saddle is

$$Pr(\lim_{n\to\infty}\boldsymbol{Z}_n\in S)=\frac{1}{6}>0.$$

So what happens? The reason that gradient descent fails to escape such a saddle set is well hidden. Specifically, in example (A), $p_A(x)$ is only C^1 but not C^2. For each $x\in V$, the

second derivatives of $p_A(x)$ on two sides are not equal. One side of x is an open interval in V^c, so the second derivative is 2; while on the other side x is the limit point of a sequence $\{x_j\}_{j=1}^{\infty} \subset V$, and the second derivative is not well defined. As for example (B), $p_B(x)$ is C^3 over $[-1, 2]$ and thus $f_B(x, y)$ satisfies the regularity requirements. However, $p_B''(x) = 0$ for all $x \in V$, so $x \in V$ are not *strict* saddles.

A loosely relevant discussion of the above constructions can be found in [49, Exercise 5.21]. This example is purely synthetic, but it raises a healthy warning as to how much the assumptions can be relaxed while the escape from saddle sets is still valid.

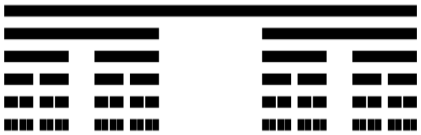

(a) Construction of Smith–Volterra–Cantor set for the first 5 steps.

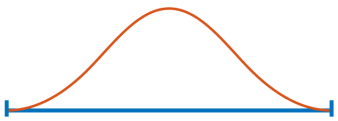

(b) Example of $p_A(x)$ on a single interval.

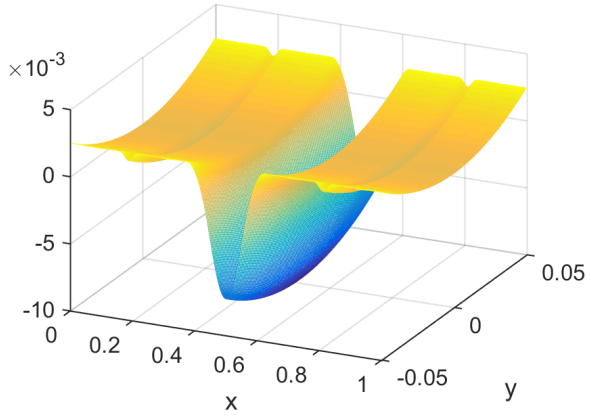

(c) Visualization of $f(x, y) = -p_A(x) + y^2$, where the saddle set is the middle of the "plateaus."

Figure 1. *Illustration of Example 2.20.*

3. Exploring the low-rank matrix manifold structure. The low-rank matrix manifold $\mathcal{M}_r := \{\boldsymbol{X} : rank(\boldsymbol{X}) = r\}$ has long been studied in low-rank recovery problems, and there has been a glossary of previous works on various applications. Some real-world applications carry an intrinsic low-rank matrix structure with them, e.g., matrix sensing and matrix completion, while others were originally formulated as the optimization in the Euclidean space, e.g., the phase retrieval problem to be discussed in section 4. One obvious advantage of manifold optimization in this case is that, instead of finding $\boldsymbol{U} \in \mathbb{R}^{n\times r}$ such that $\boldsymbol{X} = \boldsymbol{U}\boldsymbol{U}^\top$, solving $\boldsymbol{X}$ directly on the manifold $\mathcal{M}_r$ itself helps avoid the duplication of spurious critical points (i.e., $\boldsymbol{U}\boldsymbol{U}^\top$ versus $(\boldsymbol{U}\boldsymbol{R})(\boldsymbol{U}\boldsymbol{R})^\top$, $\boldsymbol{R}$ unitary). Another more important advantage is that replacing $\boldsymbol{U}\boldsymbol{U}^\top$ with $\boldsymbol{X}$ changes the form of the objective function, and it will drastically improve the convergence rate of a first-order method. The latter will be the topic of our upcoming paper [27].

This section is a self-contained discussion of some essential properties of the low-rank matrix manifold $\mathcal{M}_r$ and its closure $\overline{\mathcal{M}_r}$.

3.1. Manifold setting. First, it is necessary to point out that most previous works are based on the fixed-rank manifold $\mathcal{M}_r$, but the following reasons show that the fixed-rank manifold is not rigorous enough:

(1) The fixed-rank manifold $\mathcal{M}_r$ is *not* closed. Optimization techniques like the gradient descent generate a sequence of points towards the ground truth, and closedness is naturally necessary for asymptotic convergence analysis.
(2) It is possible that at some step $(\boldsymbol{Z}_n - \alpha P_{T_{\boldsymbol{Z}_n}}(\nabla f(\boldsymbol{Z}_n)))$ happens to have rank lower than r and falls out of $\mathcal{M}_r$.
(3) The fixed-rank manifold $\mathcal{M}_r$ also rules out the possibility that the ground truth matrix has (intrinsic) rank $\tilde{r} < r$. This prohibits *overapproximation*, i.e., when prior knowledge on $\tilde{r}$ is not available and people choose a larger guess attempting to capture more information.

Therefore, we propose considering $\overline{\mathcal{M}_r} = \{\boldsymbol{Z} : rank(\boldsymbol{Z}) \le r\}$ instead. This closure is obtained by taking the union of all lower-rank manifolds $\mathcal{M}_s$ $(s < r)$ together with the rank-r manifold. On the one hand, the essential definitions and properties of $\mathcal{M}_r$ only need to be slightly modified to accommodate $\overline{\mathcal{M}_r}$, for example by introducing the "tangent cone" to be defined below. On the other hand, this poses no actual challenge to numerical computation, as in practice the randomly generated points would fall into $\overline{\mathcal{M}_r}\backslash\mathcal{M}_r$ with probability 0, unless overapproximation is involved.

The following lemmas validate that the closure $\overline{\mathcal{M}_r}$ is a nice domain to consider. Note that the structure of the manifolds can be different on the real/complex field, or with/without symmetry, as they take different subsets of the Euclidean space $\mathbb{R}^{m\times n}$. So they are listed separately.

Lemma 3.1 (Real, asymmetric case). *Let* $\overline{\mathcal{M}_r} = \{\boldsymbol{Z} \in \mathbb{R}^{m\times n} : rank(\boldsymbol{Z}) \le r\}$ *and* $\mathcal{M}_s = \{\boldsymbol{Z} \in \mathbb{R}^{m\times n} : rank(\boldsymbol{Z}) = s\}$. *Then*

(1) $\mathcal{M}_r$ *is dense in* $\overline{\mathcal{M}_r}$;
(2) $\mathcal{M}_r$ *is connected;*
(3) *the local dimension of* $\mathcal{M}_s$ *is* $(m+n-s)s$;
(4) *the boundary of* $\mathcal{M}_r$ *is* $\overline{\mathcal{M}_r} \setminus \mathcal{M}_r = \cup_{0\le s<r}\mathcal{M}_s$.

Proof.

(1) For $\boldsymbol{Z} \in \overline{\mathcal{M}_r}\setminus\mathcal{M}_r$, it can be approached by a sequence of rank-r matrices $\{\boldsymbol{Z}_k\} \subset \mathcal{M}_r$ such that $\lim_{k\to\infty}\boldsymbol{Z}_k = \boldsymbol{Z}$.
(2) Consider

$$\widetilde{\Phi_r} : \mathbf{SO}(m,\mathbb{R}) \times \mathbb{R}^r_+ \times \mathbf{SO}(n,\mathbb{R}) \to \mathbb{R}^{m\times n},$$
$$(\widetilde{\boldsymbol{U}}, \boldsymbol{\sigma}_r, \widetilde{\boldsymbol{V}}) \mapsto \widetilde{\boldsymbol{U}}\widetilde{\boldsymbol{\Sigma}}\widetilde{\boldsymbol{V}}^\top,$$

where $\mathbf{SO}(m,\mathbb{R})$ is the real orthogonal group in dimension m, $\boldsymbol{X} = \widetilde{\boldsymbol{U}}\widetilde{\boldsymbol{\Sigma}}\widetilde{\boldsymbol{V}}^\top$ is the full SVD of $\boldsymbol{X}$, $\boldsymbol{\sigma}_r \in \mathbb{R}^r$, $\boldsymbol{\sigma}_r(i) \ne 0$, $i = 1,\dots,r$, and

$$\widetilde{\boldsymbol{\Sigma}} = \begin{pmatrix} \mathrm{diag}(\boldsymbol{\sigma}_r) & \\ & \mathbf{0}_{(m-r)\times(n-r)} \end{pmatrix}.$$

Since $\mathbf{SO}(m,\mathbb{R}) \times \mathbb{R}^r_+ \times \mathbf{SO}(n,\mathbb{R})$ is connected and $\widetilde{\Phi_r}$ is continuous, its orbit $\mathcal{M}_r$ is connected.

(3) Consider the compact SVD $\boldsymbol{X} = \boldsymbol{U}\mathrm{diag}(\boldsymbol{\sigma}_s)\boldsymbol{V}^\top$, where $\boldsymbol{U} \in \mathbb{R}^{m*s}$, $\boldsymbol{V} \in \mathbb{R}^{n*s}$, and $\boldsymbol{\sigma}_s$ is properly ordered. The local dimension is $s + \frac{(2m-s-1)s}{2} + \frac{(2n-s-1)s}{2} = (m+n-s)s$.
(4) It is obviously true. ∎

Lemma 3.2 (Complex, non-Hermitian case). *Let* $\overline{\mathcal{M}_r} = \{\boldsymbol{Z} \in \mathbb{C}^{m\times n} : rank(\boldsymbol{Z}) \le r\}$ *and* $\mathcal{M}_s = \{\boldsymbol{Z} \in \mathbb{C}^{m\times n} : rank(\boldsymbol{Z}) = s\}$. *Then*
(1) $\mathcal{M}_r$ *is dense in* $\overline{\mathcal{M}_r}$*;*
(2) $\mathcal{M}_r$ *is connected;*
(3) *the local dimension of* $\mathcal{M}_s$ *is* $(2m+2n-s)s$*;*
(4) *the boundary of* $\mathcal{M}_r$ *is* $\overline{\mathcal{M}_r} \setminus \mathcal{M}_r = \cup_{0\le s<r}\mathcal{M}_s$.

Proof.
(3) Consider the compact SVD $\boldsymbol{X} = \boldsymbol{U}\mathrm{diag}(\boldsymbol{\sigma}_r)\boldsymbol{V}^*$. The local dimension is $s + \frac{(4m-s-1)s}{2} + \frac{(4n-s-1)s}{2} = (2m+2n-s)s$. ∎

The Hermitian low-rank matrix manifold is also used very often as it fits the assumptions of many applications. But its structure is somewhat different from general non-Hermitian case because it has branches; see also [21]. The real symmetric case is very similar, and we omit the details.

Lemma 3.3 (Complex, Hermitian case). *Let* $\overline{\mathcal{M}_r} = \{\boldsymbol{Z} \in \mathbb{S}^n(\mathbb{C}) : rank(\boldsymbol{Z}) \le r\}$ *and* $\mathcal{M}_s = \{\boldsymbol{Z} \in \mathbb{S}^n(\mathbb{C}) : rank(\boldsymbol{Z}) = s\}$. *Then*
(1) $\mathcal{M}_r$ *is dense in* $\overline{\mathcal{M}_r}$*;*
(2) $\mathcal{M}_r$ *has* $r+1$ *disjoint branches and each branch is connected;*
(3) *the local dimension of* $\mathcal{M}_s$ *is* $\frac{(4m-s+1)s}{2}$*;*
(4) *the boundary of* $\mathcal{M}_r$ *is* $\overline{\mathcal{M}_r} \setminus \mathcal{M}_r = \cup_{0\le s<r}\mathcal{M}_s$.

Proof.
(2) Consider the set of matrices that has p positive eigenvalues and q negative eigenvalues, $p+q=r$. Define

$$\Psi_{p,q} : \mathbf{GL}^+(n,\mathbb{C}) \to \mathbb{S}^n,$$
$$\boldsymbol{P} \mapsto \boldsymbol{P}\boldsymbol{I}_{p,q}\boldsymbol{P}^*,$$

where $\mathbf{GL}^+(n,\mathbb{C})$ is the complex positive-determinant group in dimension n, and

$$\boldsymbol{I}_r = \begin{pmatrix} \boldsymbol{I}_p & & \\ & -\boldsymbol{I}_q & \\ & & \boldsymbol{0}_{(n-r)\times(n-r)} \end{pmatrix}.$$

Thus, the orbit of each $\Psi_{p,q}$ is connected. The tuple (p,q) is called the *signature* of the matrix. However, matrices with different signatures are not path-connected on $\mathcal{M}_r$ (they are path-connected only on $\overline{\mathcal{M}_r}$). So $\mathcal{M}_r$ has $r+1$ branches corresponding to the orbits of $\Psi_{r,0}, \Psi_{r-1,1}, \dots, \Psi_{0,r}$.
(3) Consider $\boldsymbol{X} \in \mathcal{M}_s$, and let $\boldsymbol{X} = \boldsymbol{U}\boldsymbol{D}_s\boldsymbol{U}^*$ be its compact eigenvalue decomposition with eigenvalues properly ordered. Consider the mapping

$$\Phi_s : (\boldsymbol{U}, \boldsymbol{\sigma}_s) \mapsto \boldsymbol{U}\boldsymbol{D}_s\boldsymbol{U}^*.$$

The local dimension is $s + \frac{(4m-s-1)s}{2} = \frac{(4m-s+1)s}{2}$. ∎

3.2. Geometric properties of the manifold. Recall that the PGD is defined as

$$\boldsymbol{Z}_{n+1} = \mathcal{R}\left(\boldsymbol{Z}_n - \alpha_n P_{T_{Z_n}}(\nabla f(\boldsymbol{Z}_n))\right), \tag{3.1}$$

which involves the embedded gradient ∇f, the tangent space projection $P_{T_{Z_n}}$, and the retraction $\mathcal{R}$ back onto the manifold.

For the low-rank matrix manifold, let it inherit the metric from its ambient Euclidean space, i.e., $\langle A, B\rangle = \text{trace}(A^\top B)$ and $\|A\| = \|A\|_F$. The tangent space is defined as follows.

Definition 3.4 (Tangent space, non-Hermitian case). *Let $\boldsymbol{X} \in \mathcal{M}_r$ and $\boldsymbol{X} = U\Sigma V^\top$ (or $\boldsymbol{X} = U\Sigma V^*$). Let $\mathcal{U} = Col(U)$, $\mathcal{V} = Col(V)$ be the column spaces of U and V, respectively. Then the* tangent space *of $\mathcal{M}_r$ at $\boldsymbol{X}$ is*

$$T_X\mathcal{M}_r = (\mathcal{U}\otimes\mathcal{V}) \oplus (\mathcal{U}\otimes\mathcal{V}^\perp) \oplus (\mathcal{U}^\perp\otimes\mathcal{V}).$$

The projection operator onto the tangent space is

$$P_{T_X} = P_{\mathcal{U}}\otimes I + I\otimes P_{\mathcal{V}} - P_{\mathcal{U}}\otimes P_{\mathcal{V}}.$$

Definition 3.5 (Tangent space, Hermitian case). *Let $\boldsymbol{X} \in \mathcal{M}_r$, $\boldsymbol{X} = UDU^\top$ (or $\boldsymbol{X} = UDU^*$), and $\mathcal{U} = Col(U)$. Then the* tangent space *of $\mathcal{M}_r$ at $\boldsymbol{X}$ is*

$$T_X\mathcal{M}_s = (\mathcal{U}\otimes\mathcal{U}) \oplus (\mathcal{U}\otimes\mathcal{U}^\perp) \oplus (\mathcal{U}^\perp\otimes\mathcal{U}).$$

The projection onto the tangent space is

$$P_{T_{\boldsymbol{X}}\mathcal{M}_r} = P_{\mathcal{U}}\otimes I + I\otimes P_{\mathcal{U}} - P_{\mathcal{U}}\otimes P_{\mathcal{U}}.$$

Since $\overline{\mathcal{M}_r}$ is constructed by "gluing together" all lower-rank matrix manifolds, it needs some special treatment at $\mathcal{M}_s$ $(s < r)$ in order to make up for the deficient dimension. In addition to the classical tangent space [61], [33], we need the *tangent cone* at these lower dimensional instances [50].

Definition 3.6 (Tangent cone, non-Hermitian case). *Let $\boldsymbol{X} \in \mathcal{M}_s \subset (\overline{\mathcal{M}_r})$ where $s < r$, $\boldsymbol{X} = U\Sigma V^\top$ (or $\boldsymbol{X} = U\Sigma V^*$), $\mathcal{U} = Col(U)$, and $\mathcal{V} = Col(V)$. Then the* tangent cone *of $\overline{M_r}$ at $\boldsymbol{X}$ is*

$$T_X\overline{\mathcal{M}_r} = T_X\mathcal{M}_s \oplus \{\eta : \eta \in \mathcal{U}^\perp\otimes\mathcal{V}^\perp,\ rank(\eta) = r-s\}.$$

The projection onto the tangent cone is the projection onto the tangent space plus a rank $(r-s)$ principal component, i.e.,

$$P_{T_{\boldsymbol{X}}\overline{\mathcal{M}_r}}(Y) = P_{T_{\boldsymbol{X}}\mathcal{M}_s}(Y) + Y_{r-s},$$

where Y_{r-s} is a best rank $(r-s)$ approximation of $Y - P_{T_{\boldsymbol{X}}\mathcal{M}_s}(Y)$ in the Frobenius norm.

Definition 3.7 (Tangent cone, Hermitian case). *Let* $\boldsymbol{X} \in \mathcal{M}_s \subset (\overline{\mathcal{M}_r})$ *where* $s < r$, $\boldsymbol{X} = UDU^\top$ *(or* $\boldsymbol{X} = UDU^*$*), and* $\mathcal{U} = Col(U)$. *Then the* tangent cone *of* $\overline{M_r}$ *at* $\boldsymbol{X}$ *is*

$$T_X\overline{M_r} = T_X\mathcal{M}_s \oplus \{\eta : \eta \in \mathcal{U}^\perp \otimes \mathcal{U}^\perp,\ rank(\eta) = r - s\}.$$

The projection onto the tangent cone is

$$P_{T_{\boldsymbol{X}}\overline{\mathcal{M}_r}}(Y) = P_{T_{\boldsymbol{X}}\mathcal{M}_s}(Y) + Y_{r-s},$$

where Y_{r-s} *is a best rank* $(r-s)$ *approximation of* $Y - P_{T_{\boldsymbol{X}}\mathcal{M}_s}(Y)$ *in the Frobenius norm.*

The retraction, or the projection onto the manifold as it appears in some literature, is defined as follows.

Definition 3.8 (Retraction). *The* natural retraction *on* $\mathcal{M}_r$ *is defined as*

$$\mathcal{R}_N(\boldsymbol{Z} + \xi) = \underset{\boldsymbol{Y}\in\mathcal{M}_r}{argmin} \|\boldsymbol{Z} + \xi - \boldsymbol{Y}\|_F.$$

In other words, $\mathcal{R}_N(\boldsymbol{Z} + \xi)$ is the best rank-r approximation of $\boldsymbol{Z} + \xi$. Such a retraction not only satisfies the first-order retraction property (2.1) but is actually second-order, i.e.,

$$\mathcal{R}(\boldsymbol{Z} + \alpha\xi) = \boldsymbol{Z} + \alpha\xi + \mathcal{O}(\alpha^2).$$

Vandereycken in [61] provides an explicit second-order approximation $\mathcal{R}_N^{(2)}$ to this second-order retraction $\mathcal{R}_N(\cdot)$ and shows that $\mathcal{R}_N(\boldsymbol{Z} + \xi) = \mathcal{R}_N^{(2)}(\boldsymbol{Z} + \xi) + \mathcal{O}(\|\xi\|^3)$.

It can also been seen from the above definition that the projected version of the gradient descent is also cheaper in computation. Namely, solving $\mathcal{R}_N(\boldsymbol{Z}_n - \alpha_n\nabla f(\boldsymbol{Z}_n))$ involves solving the SVD of a rank-n matrix, while $\mathcal{R}_N(\boldsymbol{Z}_n - \alpha_n P_{T_{\boldsymbol{Z}_n}}(\nabla f(\boldsymbol{Z}_n)))$ only involves that of a rank-$2r$ matrix. The latter is more favorable when $r \ll n$.

The above discussion of the low-rank matrix manifold lays the foundation for our analysis both in section 4 about phase retrieval and in the upcoming paper [27] on more general problems. While section 4 mainly focuses on $\mathcal{M}_1$ and studies the asymptotic convergence, the upcoming paper will make heavier use of $\overline{\mathcal{M}_r}$ and look at the convergence rate side of the problem. We will put multiple problems into the same framework and give a complete picture of the landscape of these optimization problems.

4. Asymptotic escape on the low-rank matrix manifold. In this section, we consider the phase retrieval problem [7], [41], [57] on the rank-1 matrix manifold. This serves both as an application of our asymptotic escape analysis for strict saddles and as a demonstration of the possibility of treating such problems rigorously on the manifold as opposed to the Euclidean space.

Since the phase retrieval problem involves a large number of stochastic measurements (i.e., random coefficient matrices $\{A_j(\omega)\}$ that constitute the objective function f_ω, ω indicating the random event), we will approach this problem in two steps. First, a crude analysis will be performed on its expectation $\mathbb{E}_\omega f$. In this case we will locate a strict critical submanifold in the shape of a "hyper ring." Then, for the nonexpectation case $f = f_{(\omega)}$, we will prove a

rather surprising result that it almost surely has only a finite number of saddle points. We will then show that with high probability, these saddles are strict saddles, and we know they are located near the above "hyper ring," so our asymptotic escape analysis is also applicable. The asymptotic escape is further supported by numerical experiments.

4.1. Phase retrieval on manifold: The expectation. The problem of phase retrieval in the case of real values aims to retrieve the information about $x \in \mathbb{R}^n$ from the phaseless measurements

$$y_j = |a_j^\top x|^2, \quad j = 1, \dots, m,$$

where the entries of $\{a_j(\omega)\}_{j=1}^m$ are drawn from i.i.d. Gaussian. Usually a large m is needed to ensure successful recovery of x.

Let $\boldsymbol{X} = xx^\top$, $A_j = a_j a_j^\top$; then $y_j = \langle A_j, \boldsymbol{X}\rangle$. The problem can be posed on the rank-1 matrix manifold $\mathcal{M}_1$ as

$$\min_{\boldsymbol{Z}\in\mathcal{M}_1} f(\boldsymbol{Z}) := \frac{1}{2m}\sum_{j=1}^m |\langle A_j, \boldsymbol{Z} - \boldsymbol{X}\rangle|^2.$$

We can apply the PGD to solve this problem on $\mathcal{M}_1$. We refer the reader to [7] in which the authors discussed the practical aspects of the PGD applied to phase retrieval. It is easily seen that $\boldsymbol{Z} = \boldsymbol{X}$ is the unique global minimizer. To ensure asymptotic convergence of the PGD to the global minimizer, it remains to rule out local minimizers and identify other critical points as strict saddles. Previous works [41], [57] have shown that phase retrieval has no spurious local minimum at least with high probability in the Euclidean setting. The analysis of saddles has been more complicated because of the stochasticity and Euclidean space parameterization.

It helps to take the expectation of $f(\boldsymbol{Z})$ and look into its landscape on the manifold. Note that

$$\bar{f}(\boldsymbol{Z}) := \mathbb{E}_\omega f(\boldsymbol{Z}) = \frac{3}{2}\|\boldsymbol{Z}\|_F^2 + \frac{3}{2}\|\boldsymbol{X}\|_F^2 - \|\boldsymbol{Z}\|_F\|\boldsymbol{X}\|_F - 2\langle \boldsymbol{Z}, \boldsymbol{X}\rangle,$$

and the Riemannian gradient (i.e., projected gradient) is

$$\operatorname{grad}\bar{f}(\boldsymbol{Z}) = P_{T_{\boldsymbol{Z}}}(\nabla\bar{f}(\boldsymbol{Z})) = P_{T_{\boldsymbol{Z}}}\left(\left(3 - \frac{\|\boldsymbol{X}\|_F}{\|\boldsymbol{Z}\|_F}\right)\boldsymbol{Z} - 2\boldsymbol{X}\right).$$

The first-order condition is satisfied if either $\boldsymbol{Z} = \boldsymbol{X}$ or

$$\|\boldsymbol{Z}\|_F = \frac{1}{3}\|\boldsymbol{X}\|_F, \quad \langle \boldsymbol{Z}, \boldsymbol{X}\rangle = 0.$$

The latter are spurious critical points, and they form an $(n-2)$-dimensional submanifold on $\mathcal{M}_1$. To see whether they are strict saddles, we look into their Hessian.

Let $\boldsymbol{Z} = zz^\top$, $u = z/\|z\|_2$; then $u \perp x$. Any element $\xi \in T_{\boldsymbol{Z}}$ can be represented as $\xi = wuu^\top + uv^\top + vu^\top$, where $w \in \mathbb{R}$, $v \in \mathbb{R}^n$, and $v \perp u$. From [61], $R_N(\boldsymbol{Z}+\xi) = \boldsymbol{Z}+\xi+\eta+\mathcal{O}(\|\xi\|^3)$ where $\eta = vv^\top/\|\boldsymbol{Z}\|_F$. Using the formula that $\operatorname{Hess} f(\boldsymbol{Z}) = \operatorname{Hess}(f\circ\mathcal{R}_{\boldsymbol{Z}})(t\xi)\,|_{t=0}$, we have

$$\begin{aligned} f\circ\mathcal{R}_{\boldsymbol{Z}}(\xi) &= f(\boldsymbol{Z}+\xi+\eta) + \mathcal{O}(\|\xi\|^3) \\ &= f(\boldsymbol{Z}) + \langle\nabla f(\boldsymbol{Z}), \xi\rangle + \langle\nabla f(\boldsymbol{Z}), \eta\rangle + \frac{1}{2}\langle\nabla^2 f(\boldsymbol{Z})[\xi], \xi\rangle + \mathcal{O}(\|\xi\|^3), \end{aligned}$$

and collecting second-order terms gives

$$\begin{aligned}
&\langle \mathrm{Hess}\bar{f}(\boldsymbol{Z})[\xi],\xi\rangle = 2\langle \nabla\bar{f}(\boldsymbol{Z}),\eta\rangle + \langle \nabla^2\bar{f}(\boldsymbol{Z})[\xi],\xi\rangle\\
&= \left(6 - 2\frac{\|\boldsymbol{X}\|_F}{\|\boldsymbol{Z}\|_F}\right)\langle \boldsymbol{Z},\eta\rangle - 4\langle \boldsymbol{X},\eta\rangle + \left(3 - \frac{\|\boldsymbol{X}\|_F}{\|\boldsymbol{Z}\|_F}\right)\|\xi\|_F^2 + \frac{\|\boldsymbol{X}\|_F}{\|\boldsymbol{Z}\|_F^3}\langle \boldsymbol{Z},\xi\rangle^2\\
&= -4\langle \boldsymbol{X},\eta\rangle + \frac{3}{\|\boldsymbol{Z}\|_F^2}\langle \boldsymbol{Z},\xi\rangle^2\\
&= -4\frac{|x^\top v|^2}{\|\boldsymbol{Z}\|_F} + 3w^2.
\end{aligned}$$

Let $\xi = ux^\top + xu^\top$; then $\langle \mathrm{Hess}\bar{f}(\boldsymbol{Z})[\xi],\xi\rangle = -12\|\boldsymbol{X}\|_F < 0$. Therefore, these spurious critical points are strict saddles. In fact they form a strict critical submanifold $\mathcal{N} = \{\boldsymbol{Z}\in\mathcal{M} \mid \|\boldsymbol{Z}\|_F = \frac{1}{3}\|\boldsymbol{X}\|_F,\ \langle \boldsymbol{Z},\boldsymbol{X}\rangle = 0\}$. For $p\in\mathcal{N}$, $T_p\mathcal{M} = \dim(T_p\mathcal{N}) = n-2$, $\dim(\nu_p^s\mathcal{M}) = \dim(\nu_p^u\mathcal{M}) = 1$, PGD will escape the strict critical submanifold and converge to the minimum of $\bar{f}$ almost surely by Theorem 2.19.

Note that although we focus on the real case (i.e., $\mathcal{M}_1(\mathbb{R})$) here, the above results can be generalized to the complex case easily, and the only change is in the constants concerning Gaussian moments.

4.2. Phase retrieval: Dive into specific realizations. Specific realizations of phase retrieval may have a much more complicated landscape than the expectation case. However, in the previous work [41] the authors have shown that for a slightly modified objective function, with high probability, the saddles of a specific realization of phase retrieval lie in the neighborhood of the above $\mathcal{N}$, the so-called *hyper ring*.

Consider

$$f(\boldsymbol{Z}) = \frac{1}{2m}\sum_{j=1}^m |\langle A_j, \boldsymbol{Z}-\boldsymbol{X}\rangle|^2$$

for a specific realization of $\{A_j(\omega)\}_{j=1}^m$. The Riemannian gradient is

$$\mathrm{grad} f(\boldsymbol{Z}) = P_{T_{\boldsymbol{Z}}}(\nabla f(\boldsymbol{Z})) = \frac{1}{m}\sum_{j=1}^m \langle A_j, \boldsymbol{Z}-\boldsymbol{X}\rangle P_{T_{\boldsymbol{Z}}}(A_j),$$

and the Riemannian Hessian is

$$\begin{aligned}
\langle \mathrm{Hess} f(\boldsymbol{Z})[\xi],\xi\rangle &= 2\langle \nabla f(\boldsymbol{Z}),\eta\rangle + \langle \nabla^2 f(\boldsymbol{Z})[\xi],\xi\rangle\\
&= \frac{1}{m}\sum_{j=1}^m (2\langle A_j, \boldsymbol{Z}-\boldsymbol{X}\rangle\langle A_j,\eta\rangle + \langle A_j,\xi\rangle^2).
\end{aligned}$$

The first result is a rather surprising one showing the finite number of critical points for phase retrieval.

Lemma 4.1. *When $m \ge n$, the above $f(\boldsymbol{Z})$ almost surely has only a finite number of critical points on the manifold $\mathcal{M}_1$.*

The proof of Lemma 4.1 is quite neat; it uses the following result from [16] that is restated in [39].

Lemma 4.2. *For a polynomial system $P(x) = (p_1(x), \dots, p_n(x))$ with $x = (x_1, \dots x_n)$ and $d_i =$ degree $p_i(x)$, let $p_i(x) = p_i^1(x) + p_i^2(x)$, where $p_i^1(x)$ consists of all the terms of $p_i(x)$ with degree d_i. If the homogeneous polynomial system $P^1(x) = (p_1^1(x), \dots, p_n^1(x)) = 0$ has only the trivial solution $x = 0$, then the original system $P(x) = 0$ only has a finite number of solutions. Moreover, the number of solutions is exactly $\Pi_{i=1}^n d_i$.*

Proof of Lemma 4.1. The first-order condition $\operatorname{grad} f(\boldsymbol{Z}) = 0$ is equivalent to

$$\frac{1}{m}\sum_{j=1}^m \langle A_j, \boldsymbol{Z} - \boldsymbol{X}\rangle P_{T_{\boldsymbol{Z}}}(A_j) = 0.$$

Let $\tilde{U} \in \mathbb{R}^{n\times(n-1)}$ be the orthonormal complement of u. Then we have

$$\begin{aligned} P_{T_{\boldsymbol{Z}}}(A_j) &= uu^\top A_j uu^\top + uu^\top A_j \tilde{U}\tilde{U}^\top + \tilde{U}\tilde{U}^\top A_j uu^\top \\ &= u(a_j^\top u)^2 u^\top + u(a_j^\top u \cdot a_j^\top \tilde{U})\tilde{U}^\top + \tilde{U}(a_j^\top u \cdot \tilde{U}^\top a_j)u^\top. \end{aligned}$$

Applying a basis transform $(u, \tilde{U})$ to the first-order condition, by symmetry, it is equivalent to

$$\begin{cases} \frac{1}{m}\sum_{j=1}^m \langle A_j, \boldsymbol{Z} - \boldsymbol{X}\rangle \cdot a_j^\top u \cdot a_j^\top u = 0, \\ \frac{1}{m}\sum_{j=1}^m \langle A_j, \boldsymbol{Z} - \boldsymbol{X}\rangle \cdot a_j^\top u \cdot a_j^\top \tilde{U} = 0, \end{cases}$$

which is equivalent to $\sum_{j=1}^m \langle A_j, \boldsymbol{Z} - \boldsymbol{X}\rangle (a_j^\top u) a_j = 0$, i.e., finding $z \in \mathbb{R}^n$ such that

$$\sum_{j=1}^m (|a_j^\top z|^2 - |a_j^\top x|^2)(a_j^\top z) a_j = 0. \tag{4.1}$$

This is a third-order heterogeneous polynomial system of n equations for n unknowns. The homogeneous part of the system is

$$\sum_{j=1}^m |a_j^\top z|^2 (a_j^\top z) a_j = 0.$$

This system almost surely only has the trivial solution $z = 0$. To see this, note that it requires $\sum_{j=1}^m |a_j^\top z|^4 = 0$, i.e.,

$$a_j^\top z = 0, \quad j = 1, \dots, m.$$

Since $\{a_j\}$ are i.i.d. Gaussian, when $m \geq n$ this linear system is almost surely nondegenerate. Now we can apply Lemma 4.2 and deduce that the original system only has a finite number of solutions, i.e., $f(\boldsymbol{Z})$ only has a finite number of critical points on the manifold. ∎

Remark 4.3. From (4.1), we can see that the first-order condition on the manifold $\mathcal{M}_1$ is equivalent to that in the parameterized Euclidean space. This means that their critical points match. Still, a critical point $\boldsymbol{Z} = zz^\top$ corresponds to at least two critical points $\pm z$ in the parameterized Euclidean space. Also, their Hessian can be very different.

Remark 4.4. The result of Lemma 4.1 only applies to the case $z \in \mathbb{R}^n$. In the case $z \in \mathbb{C}^n$, we conjecture that there would be a finite number of critical submanifolds instead. Each critical submanifold consists of $\{e^{i\theta} z_* : \theta \in [0, 2\pi)\}$, the family of *phaseless* vectors. To see this, we can impose the constraints $a_j^H z \in \mathbb{R}$ to the above equations (this is always possible by letting a_j absorb the phase information, which does not alter A_j). Now we can replace $|\cdot|$ with $(\cdot)$ and again get a polynomial system. Lemma 4.2 is still applicable, and we get the finiteness of solutions on this constrained subset. To remove the constraints, we put the phase information back and obtain the submanifolds.

The Hessians of saddle points in phase retrieval are treated in the next lemma. Note that the condition $m \geq n$ in Lemma 4.1 only ensures the finite number of saddles. To make sure that saddles are strict, we need $m \gtrsim n \log n$, which is consistent with recovery guarantees from previous works (see, e.g., [7] and references therein).

Theorem 4.5. *Given $\delta_0, \delta_1 > 0$. If $m \geq C(\delta_1) n \log n$, then with high probability no less than $1 - \frac{C_1}{m} - e^{-C_2 n}$, for all $\boldsymbol{Z}$ that satisfy the conditions*

$$\begin{cases} \langle \boldsymbol{Z}, \boldsymbol{X} \rangle \leq \delta_0 \|\boldsymbol{Z}\|_F \|\boldsymbol{X}\|_F, \\ \frac{1}{3} - \delta_0 \leq \frac{\|\boldsymbol{Z}\|_F}{\|\boldsymbol{X}\|_F} \leq \frac{1}{3} + \delta_0, \\ P_{T_{\boldsymbol{Z}}}(\nabla f(\boldsymbol{Z})) = 0, \end{cases}$$

we have

$$\lambda_{min}(Hess\ f(\boldsymbol{Z})) \leq \Lambda(\delta_0, \delta_1) < 0.$$

Here C_1, C_2 are absolute constants, $C(\delta_1)$ depends only on δ_1, and Λ depends only on δ_0 and δ_1. If we further require $\delta_0 < \frac{1}{6}$, $\delta_1 < \frac{5}{36}$, then $\lambda_{min}(Hessf(\boldsymbol{Z})) < -1$.

Proof of Theorem 4.5. The construction of a negative curvature direction is similar to that in the previous subsection. Let $\xi = xu^\top + ux^\top$; then $\xi \in T_{\boldsymbol{Z}}$. Since now x and z are not orthogonal, $\xi = wuu^\top + uv^\top + vu^\top$, where $w = 2u^\top x$ and $v = x - uu^\top x$. The Hessian is

$$\begin{aligned} \langle \mathrm{Hess} f(\boldsymbol{Z})[\xi], \xi \rangle &= \frac{1}{m} \sum_{j=1}^m (2\langle A_j, \boldsymbol{Z} - \boldsymbol{X} \rangle \langle A_j, \eta \rangle + \langle A_j, \xi \rangle^2) \\ &= \frac{1}{m} \sum_{j=1}^m \left(2\langle A_j, \boldsymbol{Z} - \boldsymbol{X} \rangle \left\langle A_j, \frac{xx^\top}{\|\boldsymbol{Z}\|_F} + \left(\eta - \frac{xx^\top}{\|\boldsymbol{Z}\|_F} \right) \right\rangle + \langle A_j, \xi \rangle^2 \right). \end{aligned}$$

An important observation is

$$\frac{1}{m} \sum_{j=1}^m \left(2\langle A_j, \boldsymbol{Z} - \boldsymbol{X} \rangle \left\langle A_j, \left(\eta - \frac{xx^\top}{\|\boldsymbol{Z}\|_F} \right) \right\rangle \right) = 0.$$

This is because

$$\begin{aligned}\eta \cdot \|\boldsymbol{Z}\|_F - xx^\top &= vv^\top - xx^\top = (x - uu^\top x)(x - uu^\top x)^\top - xx^\top \\ &= -uu^\top xx^\top - xx^\top uu^\top + uu^\top xx^\top uu^\top \in T_{\boldsymbol{Z}},\end{aligned}$$

and the first-order condition gives $\frac{1}{m}\sum_{j=1}^m \langle A_j, \boldsymbol{Z} - \boldsymbol{X}\rangle \langle A_j, \zeta\rangle = 0$ for any $\zeta \in T_{\boldsymbol{Z}}$.

Therefore, we have

$$\begin{aligned}\frac{\langle \mathrm{Hess} f(\boldsymbol{Z})[\xi], \xi\rangle}{\|\xi\|_F^2} &= \frac{\frac{1}{m}\sum_{j=1}^m (2\langle A_j, \boldsymbol{Z} - \boldsymbol{X}\rangle \langle A_j, \frac{xx^\top}{\|\boldsymbol{Z}\|_F}\rangle + \langle A_j, \xi\rangle^2)}{\|\xi\|_F^2} \\ &= \frac{\frac{1}{m}\sum_{j=1}^m (2(|a_j^\top z|^2 - |a_j^\top x|^2)|a_j^\top x|^2 + 4|a_j^\top z|^2 |a_j^\top x|^2)}{2(\|z\|^2\|x\|^2 + \langle x, z\rangle^2)} \\ &= \frac{\frac{1}{m}\sum_{j=1}^m (3|a_j^\top z|^2 |a_j^\top x|^2 - |a_j^\top x|^4)}{\|z\|^2\|x\|^2 + \langle x, z\rangle^2}.\end{aligned}$$

Using the concentration inequalities from section 4 in [27], with high probability no less than $1 - \frac{C_1}{m} - e^{-C_2 n}$, we have

$$\begin{aligned}\frac{\langle \mathrm{Hess} f(\boldsymbol{Z})[\xi], \xi\rangle}{\|\xi\|_F^2} &\le \frac{3(1+\delta_1)(\|\boldsymbol{Z}\|_F\|\boldsymbol{X}\|_F + 2\langle \boldsymbol{X}, \boldsymbol{Z}\rangle) - (3-\delta_1)\|\boldsymbol{X}\|_F^2}{\|\boldsymbol{Z}\|_F\|\boldsymbol{X}\|_F + \langle \boldsymbol{X}, \boldsymbol{Z}\rangle} \\ &\le \frac{3(1+\delta_1)(\frac{1}{3} + \delta_0 + 2\delta_0(\frac{1}{3} + \delta_0)) - (3-\delta_1)}{(\frac{1}{3} + \delta_0) + \delta_0(\frac{1}{3} + \delta_0)} := \Lambda(\delta_0, \delta_1).\end{aligned}$$

If $\delta_0 < \frac{1}{6}$, $\delta_1 < \frac{5}{36}$, then we get $\Lambda(\delta_0, \delta_1) < -1$. ∎

The above results give us a good idea of the critical points in the "hyper ring" region $\{\frac{1}{3} - \delta_0 \le \frac{\|\boldsymbol{Z}\|_F}{\|\boldsymbol{X}\|_F} \le \frac{1}{3} + \delta_0\}$ on the manifold. Specifically, Lemma 4.1 tells us that there are only a finite number of critical points, and Theorem 4.5 asserts that these critical points are all strict saddles on the manifold since they have a common negative curvature direction. We are particularly interested in the "hyper ring" region because Theorem 2.2 of [41] shows (with a slightly modified objective function) that all the critical points lie in this region with high probability, except the unique global minimum. From Theorem 2.9, we now know that the PGD will avoid saddles and converge to the global minimum.

Figure 2 shows the $log10$ error convergence of the PGD for phase retrieval on the manifold $\mathcal{M}_1$. The left figure is about the mean case, also called the population problem, while the right one is a specific case with a certain group of $\{A_j\}_{j=1}^m$, where $m = 12n$. In both experiments, we take $n = 256$, learning rate $\alpha = \frac{1}{3}$; draw 100 z_0 from i.i.d. Gaussian distribution ($\boldsymbol{Z}_0 = \boldsymbol{z}_0\boldsymbol{z}_0^\top$); and minimize $\mathbb{E}f(\boldsymbol{Z})$ or $f(\boldsymbol{Z})$ starting from these random initializations. The darker central line is the average, and the band shows the deviation. In general, it can be seen that the PGD is hardly affected by the possible existence of saddle points and converges to the minimum.

This experiment has also demonstrated the curious phenomenon mentioned at the beginning of section 3, namely a first-order method such as PGD converges exponentially quickly (i.e., linearly), even though in the Euclidean space it does not (i.e., only sublinearly). This will be explained in the upcoming work [27].

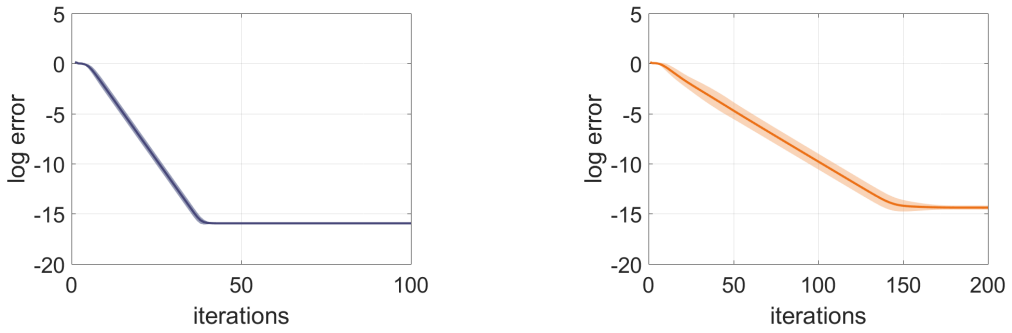

(a) $log10$ error for the mean case. (b) $log10$ error for a specific realization.

Figure 2. *Convergence (visualized as error band) of PGD for phase retrieval.*

5. Applications beyond the low-rank matrix manifold. As is mentioned in section 2, although our primary setting is the low-rank matrix manifold $\overline{\mathcal{M}_r}$, the asymptotic convergence to the minimum and escape of strict saddles (strict critical submanifolds) is valid on arbitrary finite dimensional Riemannian manifold $\mathcal{M}$. In particular, the properties of the PGD are well preserved if the manifold is embedded in a Banach space and inherits its metric. Below we discuss the optimization on the unit sphere and the Stiefel manifold as two examples of applications.

5.1. Variational eigenproblem on a sphere. Consider $\mathcal{M} = \mathbb{S}^{n-1}$, the sphere embedded in the Euclidean space $\mathbb{R}^n$. We consider the following eigenvalue problem:

$$g(z) = \lambda z, \quad z \in \mathbb{R}^n.$$

Note that $g(z)$ may or may not be linear in z. Assume that it admits eigenpairs $(\lambda_1, v_1), (\lambda_2, v_2), \ldots, (\lambda_k, v_k)$, $0 < \lambda_1 < \lambda_2 \leq \cdots \leq \lambda_k$. If $g(z) = \nabla f(z)$ for some function $f(z)$, then to find (λ_1, v_1) is to solve the following optimization problem:

$$\min_z f(z) \quad \text{s.t. } z \in \mathcal{M} = \mathbb{S}^{n-1}.$$

Viewed as an embedded Riemannian manifold, the tangent space, tangent space projection, and retraction on $\mathcal{M} = \mathbb{S}^{n-1}$ are given as follows:

$$\begin{aligned} T_z &= \{\xi \in \mathbb{R}^n : \xi^\top z = 0\}, \\ P_{T_z} &= I - zz^\top, \\ R(y) &= \frac{y}{\|y\|_2}. \end{aligned}$$

Note that $R(y)$ is a second-order retraction, because for any $z \in \mathcal{M}$, $\xi \in T_z$, we have

$$R(z + \alpha\xi) = \frac{z + \alpha\xi}{\|z + \alpha\xi\|_2} = (z + \alpha\xi)(1 + \alpha^2\|\xi\|_2^2)^{-\frac{1}{2}} = z + \alpha\xi + \mathcal{O}(\alpha^2).$$

The Levi-Civita connection on $\mathcal{M}$ is the projection of the Levi-Civita connection of the ambient space (which is the directional gradient in $\mathbb{R}^n$)

$$\widetilde{\nabla}_{\xi_z}\eta = P_{T_z}(\nabla_{\xi_z}\eta) = (I - zz^\top)(\nabla_{\xi_z}\eta), \quad \eta \in T_{\mathcal{M}}, \quad \xi_z \in T_z.$$

The Riemannian gradient on $\mathcal{M}$ is

$$\mathrm{grad} f(z) = P_{T_z}(\nabla f(z)).$$

So z is a critical point on $\mathcal{M}$ if and only if z is an eigenvector of the eigenproblem $g(z) = \lambda z$. The Riemannian Hessian on $\mathcal{M}$ is

$$\mathrm{Hess} f(z)[\xi] = P_{T_z}(\nabla^2 f(z)[\xi]) - (z^\top \nabla f(z))\xi.$$

If $g(z)$ is linear in z, then $f(z)$ is quadratic. With the positiveness assumption, we have $f(z) = z^\top Az$, where A is an SPD matrix. Then $f(x_i) = \lambda_i$, $\mathrm{grad} f(z) = Az - (z^\top Az)z$, and $\xi^\top \mathrm{Hess} f(z)[\xi] = \xi^\top A\xi - (z^\top Az)\xi^\top\xi$. It is easy to see that v_1 is the unique (up to sign) global minimum with a positive Hessian, and $v_s (s > 1)$ are all strict saddles whose Hessian has at least one negative curvature direction $\xi = v_s$.

It is interesting to look at the case where a nonminimal eigenvalue has multiplicity greater than 1. Assume that $\lambda_s = \lambda_{s+1} = \cdots = \lambda_{s+t}$; then the submanifold $\mathcal{N} = \{y \in \mathbb{R}^n \mid y = c_s v_s + \cdots + c_{s+t} v_{s+t},\ c_s^2 + \cdots + c_{s+t}^2 = 1\}$ is an immersed submanifold of $\mathcal{M}$, and it is a strict critical submanifold of f if $s \geq 2$. Since the number of such submanifolds is finite, escape from these submanifolds towards x_1 is ensured by the tools in section 2.3.

When $g(z)$ is not linear in z, as $f(z)$ now contains nonquadratic terms, it is not immediately clear from the algebraic expression whether $\mathrm{Hess} f(x_s), s > 1$, has negative curvature direction, though it can be verified numerically.

The first numerical example is from the discretized 1-dimensional Schrödinger eigenproblem $-\Delta u + V(x)u = \lambda u$ with periodic boundary condition, where $V(x)$ is taken to be the smoothed 1-dimensional Kronig–Penney (KP) potential describing free electrons in 1-dimensional crystal [36], [44]. Figure 3a shows the profile of the KP potential defined on $D = [0, 50]$ with 5 energy wells and periodic boundary condition. Figure 3b shows the first 30 eigenvalues of the operator $-\Delta + V(x)$. We can see that the first 5 eigenvalues are clustered (but not identical).

We discretize D into $n = 2^7$ grids and solve the discretized problem on $\mathcal{M} = \mathbb{S}^{n-1}$ with the PGD starting from a random initialization. The stepsize is $\alpha = 0.01$. In Figure 3c, we observe that the generated point series first seem to "stagnate" near a nonminimal eigenstate but then escape and converge towards the minimum. Figure 3d shows the profile of the computed ground energy state v_1, which is quite close to the true ground state but slightly deformed. An improvement will be proposed in the next subsection.

The second example is the nonlinear Schrödinger eigenproblem $-\Delta u + V(x)u + \beta|u|^2 u = \lambda u$, or the so-called Gross–Pitaevskii eigenvalue problem for the Bose–Einstein Condensate (BEC) [47]. It gives a more accurate description of the dynamics of Bosonic gases at ultra low temperature. With the presence of the nonlinear term $\beta|u|^2 u$, linear eigensolvers would fail, and the optimization of its variational form becomes the state-of-art solver; see, e.g., [22].

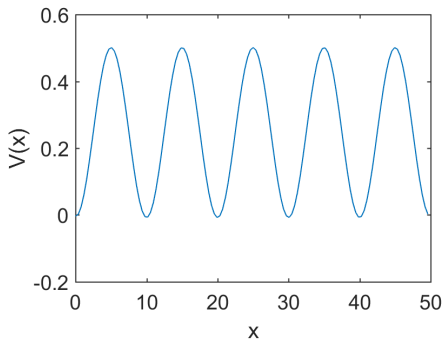

(a) Profile of $V(x)$.

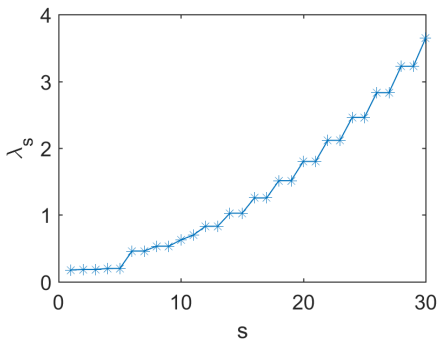

(b) First 30 eigenvalues of $-\Delta + V(x)$.

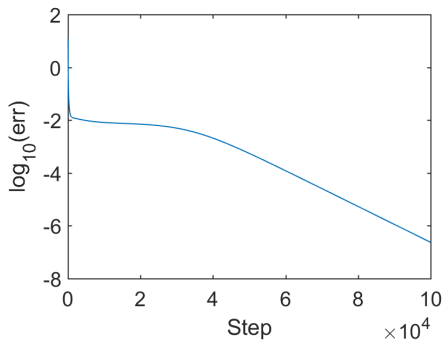

(c) Error decay in PGD.

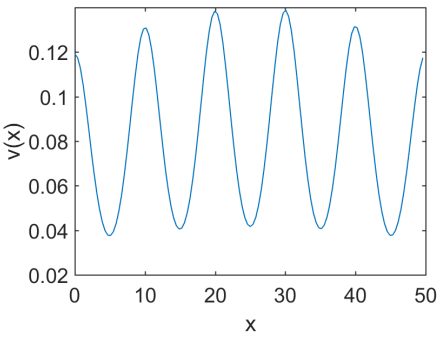

(d) Profile of the first eigenstate v_1.

Figure 3. *Solving the linear Schrödinger eigenproblem on the sphere.*

Apart from the PGD based on the L^2 metric, there can be other PGD algorithms based on other types of metrics and with different convergence theories, whose analysis is beyond the scope of this paper.

We use the same potential function $V(x)$ and discretization size as above. The nonlinear term has the weight $\beta = 1$. The objective function is

$$f(z) = \frac{1}{2} z^\top A z + \frac{\beta}{4} \sum_{j=1}^{n} z(j)^4, \quad A = -L + V.$$

For an eigenstate v_s, the eigenvalue associated to it is

$$\lambda_s = 2f(s) + \frac{\beta}{2} \sum_{j=1}^{n} z(j)^4.$$

We compute the first two eigenstates of the nonlinear Schrödinger problem using the PGD with stepsize $\alpha = 0.01$. Figure 4c shows their profiles. Figures 4a and 4b demonstrate the convergence of the PGD towards the computed eigenvalues.

To verify that v_2 is a strict saddle point, we numerically compute the smallest eigenvalue of $\mathrm{Hess} f(v_2)$ and $\lambda_{\min}(\mathrm{Hess} f(v_2)) = -0.0024 < 0$. Figure 4d shows a profile of the corresponding eigenvector $\xi_{\min}$, i.e., a negative curvature direction.

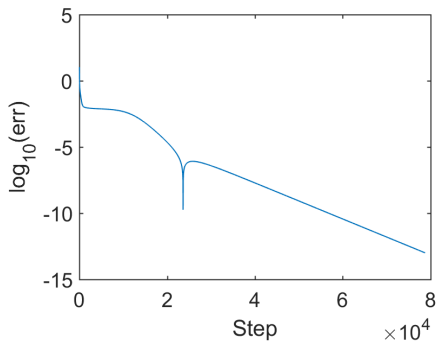

(a) Error decay of PGD when computing v_1.

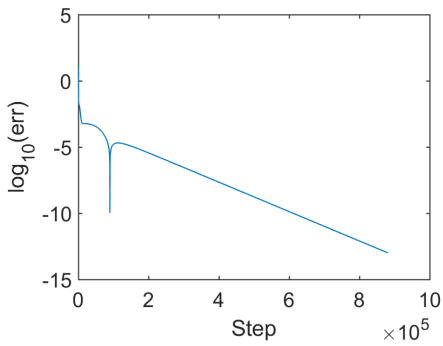

(b) Error decay of PGD when computing v_2.

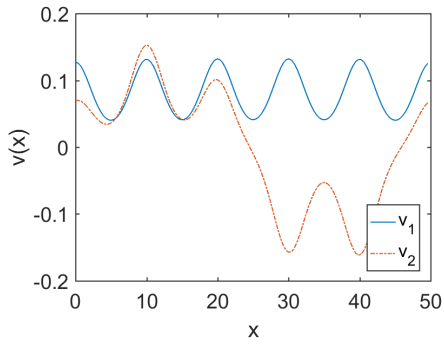

(c) Profile of eigenstates v_1 and v_2.

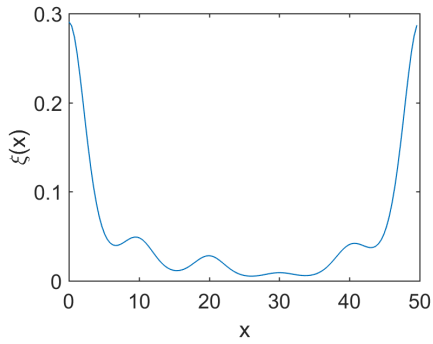

(d) A negative curvature direction.

Figure 4. *Solving the nonlinear Schrödinger eigenproblem on the sphere.*

Apart from physical problems like BEC eigenstates, linear and nonlinear eigenproblems also find applications in image processing and machine learning. For example, the Max-Cut problem corresponds to a linear eigenproblem, while the optimization of the Ginzburg–Landau-type functional in image segmentation and learning tasks corresponds to a nonlinear eigenproblem; see, e.g., [5], [26]. Although there are many algorithms tailored for linear eigenproblems, their nonlinear relatives often lack a rigorous convergence guarantee. Manifold optimization thus provides a more versatile point of view for them.

5.2. Simultaneous eigensolver on the Stiefel manifold. Subspace iteration is a common technique for accelerating the convergence of smallest eigenstates in linear eigenproblems, especially when the ground states are clustered, as in the previous examples.

From the viewpoint of manifold optimization, solving the first m eigenstates simultane-

ously can be posed as the optimization on the Stiefel manifold $\mathcal{M} = \{\boldsymbol{Z} \in \mathbb{R}^{n\times m} : \boldsymbol{Z}^\top \boldsymbol{Z} = I_m\}$:

$$\min_{\boldsymbol{Z}} \ \mathrm{trace}(f(\boldsymbol{Z})) \quad \text{s.t. } \boldsymbol{Z} \in \mathcal{M} = \{\boldsymbol{Z} \in \mathbb{R}^{n\times m} : \boldsymbol{Z}^\top \boldsymbol{Z} = I_m\}.$$

The Stiefel manifold [13], [31] is the set of all m-frames in $\mathbb{R}^n$. When $m = 1$, it reduces to the sphere $\mathbb{S}^{n-1}$. With the Euclidean metric, its tangent space, tangent space projection, and retraction are given as follows:

$$\begin{aligned} T_{\boldsymbol{Z}} &= \{\xi \in \mathbb{R}^{n\times m} : \xi^\top \boldsymbol{Z} + \boldsymbol{Z}^\top \xi = 0\}, \\ P_{T_{\boldsymbol{Z}}}(\mathbf{Y}) &= \mathbf{Y} - \mathbf{Z}\ \mathrm{sym}(\mathbf{Z}^\top \mathbf{Y}), \\ R(Y) &= \mathrm{qf}\ (Y), \end{aligned}$$

where sym takes the symmetric part and qf takes the Q factor of QR decomposition. Similar to the case of the sphere manifold, the Riemannian connection and gradient are defined by the projection onto the tangent space. When $f(\boldsymbol{Z}) = \boldsymbol{Z}^\top A \boldsymbol{Z}$, we have

$$\begin{aligned} \mathrm{grad} f(\boldsymbol{Z}) &= P_{T_{\boldsymbol{Z}}}(A\boldsymbol{Z}), \\ \langle \xi, \mathrm{Hess} f(\boldsymbol{Z})[\xi]\rangle &= \mathrm{tr}(\xi^\top A \xi - (\xi^\top \xi)(\boldsymbol{Z}^\top A \boldsymbol{Z})). \end{aligned}$$

It is easily verified that the minimum is achieved when $span\boldsymbol{Z} = span\{v_1, \ldots, v_m\}$, and all $\boldsymbol{Z}$ that span other eigensubspaces are strict saddles if all the eigenvalues are distinct.

We compute the first 5 eigenstates simultaneously for the linear Schrödinger eigenproblem with the same potential as in Figure 3a. The stepsize is $\alpha = 0.01$. Figures 5a and 5b compare the computed eigenstates extracted from $\boldsymbol{Z}$ and the true eigenstates, which are almost identical. In Figure 5c, we can see that the subspace iteration on the Stiefel manifold achieves much better convergence in fewer steps than the optimization on the sphere.

Application of the Stiefel manifold optimization can also be extended to data science, e.g., frame construction and dictionary learning [6], [56], if the frame/dictionary satisfies orthonormal assumptions.

6. Conclusion and future works. We have studied the asymptotic escape of strict saddle points of the PGD on the Riemannian manifolds. The first main contribution of this paper is that it pushes the boundary of current analysis to nonisolated saddle sets, proving when the PGD can escape and indicating when it cannot. As a general tool, it can be applied to various settings as long as the manifold of interest satisfies certain smoothness conditions. This is demonstrated by several representative examples from different fields.

The saddle analysis of phase retrieval on the low-rank matrix manifold serves as an application of the above asymptotic escape result, but it also stands as an insightful result by itself. We have shown that it always has a finite number of critical points, and the saddles are strict saddles with high probability. Essentially, the low-rank matrix manifold sheds light on the intrinsic quadratic (instead of quartic) structure of this problem.

In addition to the asymptotic convergence behavior of the PGD, the convergence rate is also an important issue. Empirical linear convergence rates in many low-rank matrix recovery problems are already observed but have yet to be explained. This will be the topic of our

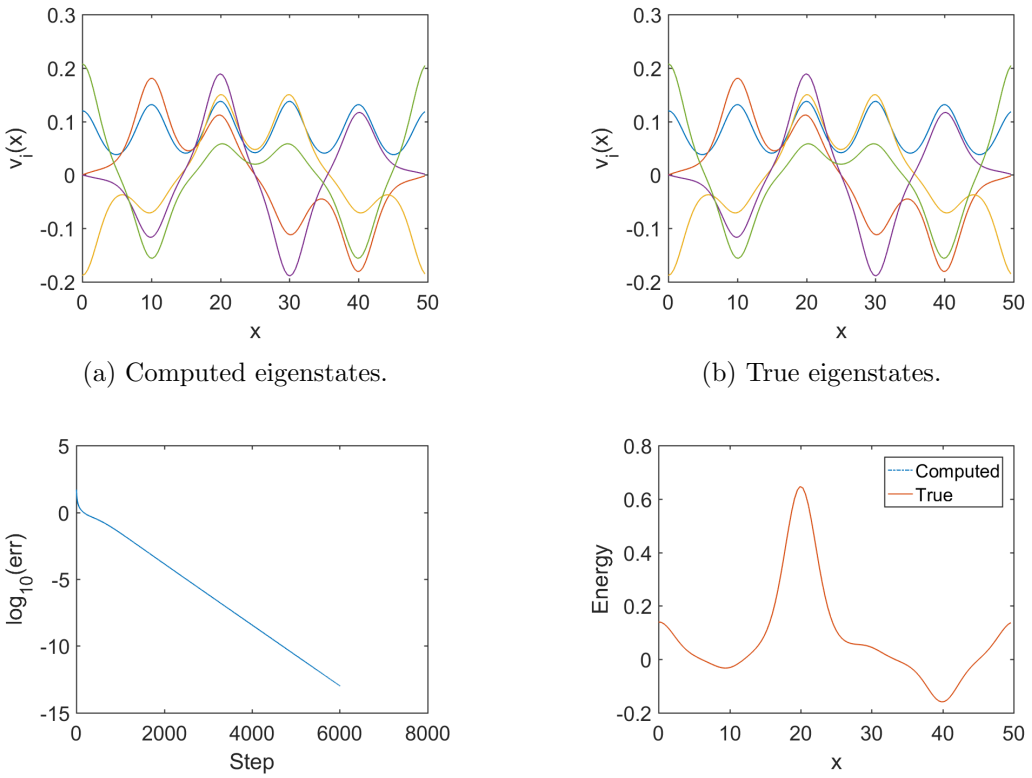

(a) Computed eigenstates.

(b) True eigenstates.

(c) Error decay of PGD.

(d) Accumulated energy of the first 5 eigenstates.

Figure 5. *Simultaneously solving the first* 5 *eigenstates of the linear Schrödinger problem on the Stiefel manifold.*

future work [27], where we prove the linear convergence rate using the quadratic nature of those problems on the manifold $\overline{\mathcal{M}_r}$.

Appendix A. As we mentioned in section 2.3, when there are a bunch of self-connected critical submanifolds (generalization of critical points), the escape of strict critical submanifolds (generalized strict saddles) and convergence to a minimum rely on the number or the structure of such critical submanifolds. When the number is uncountable, the situation can be quite complicated.

In this appendix, we discuss some structural properties of critical submanifolds that may help untangle their successive relations. We introduce the concepts of index and transversality, point out the transversality properties of certain functions and their consequences, and link the stable manifolds of the gradient flow to that of the gradient descent.

Definition A.1 (Index). *For $f : \mathcal{M} \mapsto \mathbb{R}$, let p be a critical point of f; then the* index *of p is*

$$\lambda_p := \dim T_p^u \mathcal{M}.$$

Remark A.2. All critical points in the same connected critical submanifold $\mathcal{N}$ have the same index, which is defined as the index $\lambda_{\mathcal{N}}$ of the submanifold $\mathcal{N}$. An equivalent way to define strict critical submanifold is $\lambda_{\mathcal{N}} > 0$.

Definition A.3 (Transversality). (1) *For smooth maps $f : \mathcal{N}_1 \mapsto \mathcal{M}$ and $g : \mathcal{N}_2 \mapsto \mathcal{M}$, we say that f is* transverse *to g iff for any X_1, X_2 such that $f(X_1) = g(X_2) = Y$,*

$$df(T_{X_1}\mathcal{N}_1) + dg(T_{X_2}\mathcal{N}_2) = T_Y\mathcal{M},$$

where df and dg are gradient vector fields of f and g.

(2) *If $\mathcal{N}_1$ and $\mathcal{N}_2$ are immersed submanifolds of $\mathcal{M}$, then $\mathcal{N}_1$ is* transverse *to $\mathcal{N}_2$ iff for any $X \in \mathcal{N}_1 \cap \mathcal{N}_2$,*

$$T_X\mathcal{N}_1 + T_X\mathcal{N}_2 = T_X\mathcal{M}.$$

Two immersed submanifolds vacuously transverse if they do not intersect.

Remark A.4. A function $f : \mathcal{M} \mapsto \mathbb{R}$ is called *Morse–Bott* if all its critical points lie in some disjoint union of connected and nondegenerate critical submanifolds; f is called *Morse–Smale* if it satisfies the Morse–Smale transversality condition, i.e., for any two critical submanifolds $\mathcal{N}_1, \mathcal{N}_2$, their stable and unstable manifolds intersect transversally.

The transversality condition for immersed manifolds simply means that two manifolds "cross" each other and do not "overlap." Figure 6 is a vivid illustration of transversality on a 2-dimensional manifold. If the objective function f is a Morse–Smale function, transversality implies more favorable properties.

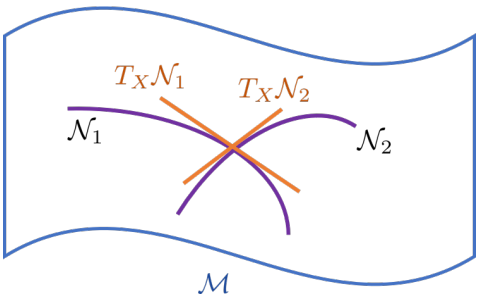

Figure 6. *An illustration of transversality.*

Theorem A.5 (Corollary 6.27 in [4]). *For a Morse–Smale function f, any critical point p of f satisfies*

$$\overline{W^u(p)} = \bigcup_{p \succeq q} W^u(q),$$
$$\overline{W^s(p)} = \bigcup_{r \succeq p} W^s(r),$$

where $W^s(p)$ (resp., $W^u(p)$) is the stable (resp., unstable) manifold of p defined by gradient flow line, and $p \succeq q$ means $W^u(p) \cap W^s(q) \neq \emptyset$.

Theorem A.6. *For a Morse–Smale function f, if two critical submanifolds $\mathcal{N}_1$ and $\mathcal{N}_2$ have the same index, then they vacuously transverse, i.e., $W^u(\mathcal{N}_1) \cap W^s(\mathcal{N}_2) = \emptyset$.*

Proof. By Proposition 6.2 in [4], if $W^u(\mathcal{N}_1) \cap W^s(\mathcal{N}_2) \neq \emptyset$, then their intersection is an embedded submanifold of dimension $(\lambda_{\mathcal{N}_1} - \lambda_{\mathcal{N}_2})$. But $\lambda_{\mathcal{N}_1} - \lambda_{\mathcal{N}_2} = 0$, which is a contradiction. ∎

Both Theorem A.5 and Theorem A.6 are helpful when taking the union of stable manifolds of infinitely many critical submanifolds. Theorem A.5 shows that the closure of the stable/unstable manifold of one critical set is the union of the stable/unstable manifolds of the sets that have successive relations with it. On the other hand, Theorem A.6 shows that the successive relations are strictly limited by the indices (i.e., negative curvature dimensions) of the critical sets. This successive relation simply cannot happen between sets of the same index.

It should be stressed that the above results are on the stable/unstable manifold of *gradient flows*, not *gradient descents*. Whether this can be generalized to gradient descents is still unclear. We know that with the first-order retraction property, as $\alpha \to 0$, the projected gradient descent on the manifold approximates the gradient flow line. It can be proved that the respective stable/unstable manifolds also converge as long as the retraction is at least first-order and the domain is compact. However, the transversality concerns the "angles" at the intersection of these submanifolds. Even the uniform convergence of submanifolds cannot ensure the preservation of their intersection angles along the convergence.

This discussion aims to draw interest to the vast possibilities that Morse theory has to offer. They point out a way to deal with complex geometries of critical point sets. We plan to conduct further studies along this direction to quantify the above relations.

Acknowledgment. We would like to thank the anonymous referees for their insightful comments, which improved the quality of our manuscript.

REFERENCES

[1] P.-A. Absil and K. A. Gallivan, *Joint diagonalization on the oblique manifold for independent component analysis*, in Proceedings of the 2006 IEEE International Conference on Acoustics Speech and Signal Processing, Vol. 5, IEEE, Washington, DC, 2006, pp. V–V.

[2] P.-A. Absil, R. Mahony, and J. Trumpf, *An extrinsic look at the Riemannian Hessian*, in International Conference on Geometric Science of Information, Springer, New York, 2013, pp. 361–368.

[3] A. Banyaga and D. Hurtubise, *Morse-Bott homology*, Trans. Amer. Math. Soc., 362 (2010), pp. 3997–4043.

[4] A. Banyaga and D. Hurtubise, *Lectures on Morse Homology*, Texts Math. Sci. 29, Springer, Dordrecht, The Netherlands, 2004.

[5] A. L. Bertozzi and A. Flenner, *Diffuse interface models on graphs for classification of high dimensional data*, Multiscale Model. Simul., 10 (2012), pp. 1090–1118, https://doi.org/10.1137/11083109X.

[6] J.-F. Cai, H. Ji, Z. Shen, and G.-B. Ye, *Data-driven tight frame construction and image denoising*, Appl. Comput. Harmon. Anal., 37 (2014), pp. 89–105.

[7] J.-F. Cai and K. Wei, *Solving Systems of Phaseless Equations via Riemannian Optimization with Optimal Sampling Complexity*, preprint, https://arxiv.org/abs/1809.02773, 2018.

[8] Y. Chi, Y. M. Lu, and Y. Chen, *Nonconvex optimization meets low-rank matrix factorization: An overview*, IEEE Trans. Signal Process., 67 (2019), pp. 5239–5269.

[9] R. Cohen, *Topics in Morse Theory*, Department of Mathematics, Stanford University, Stanford, CA, 1991.
[10] C. Criscitiello and N. Boumal, *Efficiently escaping saddle points on manifolds*, in Advances in Neural Information Processing Systems, NeurIPS, San Diego, CA, 2019, pp. 5985–5995.
[11] C. Da Silva and F. J. Herrmann, *Optimization on the hierarchical Tucker manifold–applications to tensor completion*, Linear Algebra Appl., 481 (2015), pp. 131–173.
[12] S. S. Du, C. Jin, J. D. Lee, M. I. Jordan, A. Singh, and B. Poczos, *Gradient descent can take exponential time to escape saddle points*, in Advances in Neural Information Processing Systems, NeurIPS, San Diego, CA, 2017, pp. 1067–1077.
[13] A. Edelman, T. A. Arias, and S. T. Smith, *The geometry of algorithms with orthogonality constraints*, SIAM J. Matrix Anal. Appl., 20 (1998), pp. 303–353, https://doi.org/10.1137/S0895479895290954.
[14] S. Fattahi and S. Sojoudi, *Exact Guarantees on the Absence of Spurious Local Minima for Non-negative Robust Principal Component Analysis*, preprint, https://arxiv.org/abs/1812.11466v1, 2018.
[15] J. R. Fienup, *Phase retrieval algorithms: A comparison*, Appl. Optics, 21 (1982), pp. 2758–2769.
[16] C. B. Garcia and T. Y. Li, *On the number of solutions to polynomial systems of equations*, SIAM J. Numer. Anal., 17 (1980), pp. 540–546, https://doi.org/10.1137/0717046.
[17] R. Ge, F. Huang, C. Jin, and Y. Yuan, *Escaping from saddle points—online stochastic gradient for tensor decomposition*, in Proceedings of the 28th Conference on Learning Theory, PMLR, 2015, pp. 797–842.
[18] R. Ge, J. D. Lee, and T. Ma, *Matrix completion has no spurious local minimum*, in Advances in Neural Information Processing Systems, NeurIPS, San Diego, CA, 2016, pp. 2973–2981.
[19] R. W. Gerchberg, *A practical algorithm for the determination of phase from image and diffraction plane pictures*, Optik, 35 (1972), pp. 237–246.
[20] I. Grubišić and R. Pietersz, *Efficient rank reduction of correlation matrices*, Linear Algebra Appl., 422 (2007), pp. 629–653.
[21] U. Helmke and M. A. Shayman, *Critical points of matrix least squares distance functions*, Linear Algebra Appl., 215 (1995), pp. 1–19.
[22] P. Henning and D. Peterseim, *Sobolev gradient flow for the Gross–Pitaevskii eigenvalue problem: Global convergence and computational efficiency*, SIAM J. Numer. Anal., 58 (2020), pp. 1744–1772, https://doi.org/10.1137/18M1230463.
[23] J. Ho, Y. Xie, and B. Vemuri, *On a nonlinear generalization of sparse coding and dictionary learning*, in Proceedings of the 30th International Conference on Machine Learning, Atlanta, GA, 2013, pp. 1480–1488.
[24] S. Holtz, T. Rohwedder, and R. Schneider, *On manifolds of tensors of fixed TT-rank*, Numer. Math., 120 (2012), pp. 701–731.
[25] R. Hosseini and S. Sra, *Matrix manifold optimization for Gaussian mixtures*, in Advances in Neural Information Processing Systems, NeurIPS, San Diego, CA, 2015, pp. 910–918.
[26] T. Y. Hou, D. Huang, K. C. Lam, and Z. Zhang, *A fast hierarchically preconditioned eigensolver based on multiresolution matrix decomposition*, Multiscale Model. Simul., 17 (2019), pp. 260–306, https://doi.org/10.1137/18M1180827.
[27] T. Y. Hou, Z. Li, and Z. Zhang, *Fast Global Convergence for Low-rank Matrix Recovery via Riemannian Gradient Descent with Random Initialization*, arxiv preprint, in preparation, 2020.
[28] W. Huang and P. Hand, *Blind deconvolution by a steepest descent algorithm on a quotient manifold*, SIAM J. Imaging Sci., 11 (2018), pp. 2757–2785, https://doi.org/10.1137/17M1151390.
[29] K. Jaganathan, Y. C. Eldar, and B. Hassibi, *Phase Retrieval: An Overview of Recent Developments*, preprint, https://arxiv.org/abs/1510.07713, 2015.
[30] C. Jin, R. Ge, P. Netrapalli, S. M. Kakade, and M. I. Jordan, *How to escape saddle points efficiently*, in Proceedings of the 34th International Conference on Machine Learning (Sydney, Australia), PMLR 70, 2017, pp. 1724–1732.
[31] T. Kaneko, S. Fiori, and T. Tanaka, *Empirical arithmetic averaging over the compact Stiefel manifold*, IEEE Trans. Signal Process., 61 (2012), pp. 883–894.
[32] H. Kasai and B. Mishra, *Low-rank tensor completion: A Riemannian manifold preconditioning approach*, in Proceedings of the International Conference on Machine Learning (New York, NY), PMLR 48, 2016, pp. 1012–1021.

[33] O. Koch and C. Lubich, *Dynamical low-rank approximation*, SIAM J. Matrix Anal. Appl., 29 (2007), pp. 434–454, https://doi.org/10.1137/050639703.
[34] M. Kolodrubetz, V. Gritsev, and A. Polkovnikov, *Classifying and measuring geometry of a quantum ground state manifold*, Phys. Rev. B, 88 (2013), 064304.
[35] D. Kressner, M. Steinlechner, and B. Vandereycken, *Low-rank tensor completion by Riemannian optimization*, BIT, 54 (2014), pp. 447–468.
[36] R. D. L. Kronig and W. G. Penney, *Quantum mechanics of electrons in crystal lattices*, Proc. Roy. Soc. London Ser. A, 130 (1931), pp. 499–513.
[37] J. D. Lee, I. Panageas, G. Piliouras, M. Simchowitz, M. I. Jordan, and B. Recht, *First-Order Methods Almost Always Avoid Saddle Points*, preprint, https://arxiv.org/abs/1710.07406, 2017.
[38] J. D. Lee, M. Simchowitz, M. I. Jordan, and B. Recht, *Gradient Descent Converges to Minimizers*, preprint, https://arxiv.org/abs/1602.04915, 2016.
[39] T.-Y. Li, *Solving polynomial systems*, Mathematical Intelligencer, 9 (1987), pp. 33–39.
[40] W. Li and G. Montúfar, *Natural gradient via optimal transport*, Inform. Geom., 1 (2018), pp. 181–214.
[41] Z. Li, J.-F. Cai, and K. Wei, *Towards the optimal construction of a loss function without spurious local minima for solving quadratic equations*, IEEE Trans. Inform. Theory, 66 (2020), pp. 3242–3260.
[42] C. Lubich, I. V. Oseledets, and B. Vandereycken, *Time integration of tensor trains*, SIAM J. Numer. Anal., 53 (2015), pp. 917–941, https://doi.org/10.1137/140976546.
[43] S. Osher, Z. Shi, and W. Zhu, *Low dimensional manifold model for image processing*, SIAM J. Imaging Sci., 10 (2017), pp. 1669–1690, https://doi.org/10.1137/16M1058686.
[44] V. Ozoliņš, R. Lai, R. Caflisch, and S. Osher, *Compressed modes for variational problems in mathematics and physics*, Proc. Natl. Acad. Sci. USA, 110 (2013), pp. 18368–18373.
[45] I. Panageas and G. Piliouras, *Gradient Descent Only Converges to Minimizers: Non-Isolated Critical Points and Invariant Regions*, preprint, https://arxiv.org/abs/1605.00405, 2016.
[46] G. Peyré, *Manifold models for signals and images*, Comput. Vis. Image Understanding, 113 (2009), pp. 249–260.
[47] L. Pitaevskii and S. Stringari, *Bose-Einstein Condensation and Superfluidity*, Internat. Ser. Monogr. Phys. 164, Oxford University Press, Oxford, UK, 2016.
[48] J. Provost and G. Vallee, *Riemannian structure on manifolds of quantum states*, Comm. Math. Phys., 76 (1980), pp. 289–301.
[49] W. Rudin, *Principles of Mathematical Analysis*, 2nd ed., McGraw-Hill, New York, 1964.
[50] R. Schneider and A. Uschmajew, *Convergence results for projected line-search methods on varieties of low-rank matrices via Łojasiewicz inequality*, SIAM J. Optim., 25 (2015), pp. 622–646, https://doi.org/10.1137/140957822.
[51] U. Shalit, D. Weinshall, and G. Chechik, *Online learning in the embedded manifold of low-rank matrices*, J. Mach. Learn. Res., 13 (2012), pp. 429–458.
[52] Y. Shechtman, Y. C. Eldar, O. Cohen, H. N. Chapman, J. Miao, and M. Segev, *Phase retrieval with application to optical imaging: A contemporary overview*, IEEE Signal Process. Mag., 32 (2015), pp. 87–109.
[53] L. Shi and Y. Chi, *Manifold gradient descent solves multi-channel sparse blind deconvolution provably and efficiently*, in Proceedings of the 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2020), IEEE, Washington, DC, 2020, pp. 5730–5734.
[54] M. Shub, *Global Stability of Dynamical Systems*, Springer, New York, 1987.
[55] S. T. Smith, *Optimization techniques on Riemannian manifolds*, Fields Inst. Commun., 3 (1994), pp. 113–135.
[56] J. Sun, Q. Qu, and J. Wright, *Complete dictionary recovery over the sphere* I*: Overview and the geometric picture*, IEEE Trans. Inform. Theory, 63 (2016), pp. 853–884.
[57] J. Sun, Q. Qu, and J. Wright, *A geometric analysis of phase retrieval*, Found. Comput. Math., 18 (2018), pp. 1131–1198.
[58] Y. Sun, N. Flammarion, and M. Fazel, *Escaping from saddle points on Riemannian manifolds*, in Advances in Neural Information Processing Systems, NeurIPS, San Diego, CA, 2019, pp. 7274–7284.
[59] P. Turaga, A. Veeraraghavan, and R. Chellappa, *Statistical analysis on Stiefel and Grassmann manifolds with applications in computer vision*, in Proceedings of the 2008 IEEE Conference on Computer Vision and Pattern Recognition, IEEE, Washington, DC, 2008, pp. 1–8.

[60] P. K. Turaga and A. Srivastava, *Riemannian Computing in Computer Vision*, Springer, New York, 2016.

[61] B. Vandereycken, *Low-rank matrix completion by Riemannian optimization*, SIAM J. Optim., 23 (2013), pp. 1214–1236, https://doi.org/10.1137/110845768.

[62] J. J.-Y. Wang, H. Bensmail, and X. Gao, *Feature selection and multi-kernel learning for sparse representation on a manifold*, Neural Networks, 51 (2014), pp. 9–16.

[63] K. Wei, J.-F. Cai, T. F. Chan, and S. Leung, *Guarantees of Riemannian optimization for low rank matrix recovery*, SIAM J. Matrix Anal. Appl., 37 (2016), pp. 1198–1222, https://doi.org/10.1137/15M1050525.

[64] K. Ye, K. S.-W. Wong, and L.-H. Lim, *Optimization on Flag Manifolds*, preprint, https://arxiv.org/abs/1907.00949, 2019.

[65] P. Zanardi, P. Giorda, and M. Cozzini, *Information-theoretic differential geometry of quantum phase transitions*, Phys. Rev. Lett., 99 (2007), 100603.

[66] J. Zeng, G. Cheung, M. Ng, J. Pang, and C. Yang, *3D point cloud denoising using graph Laplacian regularization of a low dimensional manifold model*, IEEE Trans. Image Process., 29 (2020), pp. 3474–3489.