



First-order methods almost always avoid strict saddle points

Jason D. Lee¹ · Ioannis Panageas² · Georgios Piliouras³ · Max Simchowitz⁴ · Michael I. Jordan⁴ · Benjamin Recht⁴

Received: 19 October 2017 / Accepted: 5 February 2019 / Published online: 18 February 2019
© Springer-Verlag GmbH Germany, part of Springer Nature and Mathematical Optimization Society 2019

Abstract

We establish that first-order methods avoid strict saddle points for almost all initializations. Our results apply to a wide variety of first-order methods, including (manifold) gradient descent, block coordinate descent, mirror descent and variants thereof. The connecting thread is that such algorithms can be studied from a dynamical systems perspective in which appropriate instantiations of the Stable Manifold Theorem allow for a global stability analysis. Thus, neither access to second-order derivative information nor randomness beyond initialization is necessary to provably avoid strict saddle points.

Keywords Gradient descent · Smooth optimization · Saddle points · Local minimum · Dynamical systems

Mathematics Subject Classification 90C26

1 Introduction

Saddle points have long been regarded as a major obstacle for machine learning methodology that require optimizing a non-convex objective [21,43]. It is well understood that in many applications of interest, the number of saddle points significantly outnumber the number of local minima, which is especially problematic when the

This paper extends upon the special case of gradient descent dynamics developed in the conference proceedings of the authors [32,42].

Electronic supplementary material The online version of this article (<https://doi.org/10.1007/s10107-019-01374-3>) contains supplementary material, which is available to authorized users.

✉ Jason D. Lee
jasonlee@marshall.usc.edu

Extended author information available on the last page of the article

solutions associated with worst-case saddle points are considerably worse than those associated with worst-case local minima [19].

Moreover, it is not hard to construct examples where a worst-case initialization of gradient descent (or other first-order methods) provably converge to saddle points [39, Sect. 1.2.3].

The main message of our paper is that, under very mild regularity conditions, saddle points have little effect on the asymptotic behavior of first-order methods. Building on tools from the theory of dynamical systems, we generalize recent analysis of gradient descent [32,42] to establish that a wide variety of first-order methods—including gradient descent, proximal point algorithm, block coordinate descent, mirror descent—avoid so-called “strict” saddle points for almost all initializations; that is, saddle points where the Hessian of the objective function admits at least one direction of negative curvature (see Definition 1).

Our results provide a unified theoretical framework for analyzing the asymptotic behavior of a wide variety of classic optimization algorithms in non-convex optimization. Furthermore, we believe that furthering our understanding of the behavior and geometry of deterministic optimization techniques with random initialization can serve in the development of stochastic algorithms which improve upon their deterministic counterparts and achieve strong convergence-rate results; indeed, such insights have already led to significant improves in modifying gradient descent to navigate saddle-point geometry [22,28].

1.1 Related work

In recent years, the optimization and machine learning communities have dedicated much effort to understanding the geometry of non-convex landscapes by searching for unified geometric properties which could be leverage by general-purpose optimization techniques. The strict saddle property (Definition 1) is one such property which has been shown to hold in a wide and diverse range of salient objective functions: PCA, a fourth-order tensor factorization [24], formulations of dictionary learning [52,53], phase retrieval [51], low-rank matrix factorizations [12,25,26], and simple neural networks [16,23,49]. It is also known that, in the worst case, the strict saddle property is unavoidable as finding descent-directions at critical points with degenerate Hessians is NP-hard in general [38].

Earlier work had shown that first-order descent methods can circumvent strict saddle points, provided that they are augmented with unbiased noise whose variance is sufficiently large in each direction. For example, Pemantle [44] establishes convergence of the Robbins–Monro stochastic approximation method to local minimizers for strict saddle functions. More recently, Ge et al. [24] give quantitative rates on the convergence of noisy gradient descent to local minimizers, for strict saddle functions. O'Neill and Wright [41] use the stable manifold theorem to prove that the accelerated gradient method converges to a single strict saddle point with probability zero. Theorem 2 is complementary and strengthens their result to guarantee the probability of converging to the set of strict saddles points is zero, rather than a single (or countably many) strict saddle point.

To obtain provable guarantees without the addition of stochastic noise, Sun et al. [51–53] adopt trust-region methods which leverage Hessian information in order to circumvent saddle points. This approach represents a refinement of a long tradition of related, “second-order” strategies, including: a modified Newton’s method with curvilinear line search [37], the modified Cholesky method [27], trust-region methods [20], and the related cubic regularized Newton’s method [40]. Specialized to deep learning applications Dauphin et al. [21] and Pascanu et al. [43] have introduced a saddle-free Newton method.

However, such curvature-based optimization algorithms have a per-iteration computational complexity which scales quadratically or even cubically in the dimension d , rendering them unsuitable for optimization of high-dimensional functions. In more recent work, several works have presented faster curvature-based methods including Liu and Yang [34], Reddi et al. [45], Royer and Wright [47] by combining fast first-order methods with fast eigenvector algorithms, to obtain lower per-iteration complexity.

Fortunately, it appears that neither the addition of isotropic noise, nor the use of second-order methods are necessary for circumventing saddle points. For example, recent work by Jin et al. [28] showed that by carefully perturbing the iterates of gradient descent in the vicinity of possible saddles results in a first-order method which converges to local minimizers in a number of iterations with only poly-logarithmic dimension dependence. Moreover, many recent works have shown that, even without any random perturbations, a combination of gradient descent and a smart-initialization provably converges to the global minimum for a variety of non-convex problems: such settings include matrix factorization [29,55], phase retrieval [17,18], dictionary learning [6], and latent-variable models [11,54]. While our results only guarantee convergence to local minimizers, they eschew the need for complex and often computationally prohibitive initialization procedures.

In addition to what has been established theoretically, there is a broadly-accepted folklore in the field that running gradient descent with a random initialization is sufficient to identify a local optima. For example, the authors of Sun et al. [51] empirically observe gradient descent with 100 random initializations on the phase retrieval problem always converges to a local minimizer, one whose quality matches that of the solution found using more costly trust-region techniques. It is the purpose of this work to place these intuitions on firm mathematical footing.

Finally, we emphasize that there are many settings in which all local optima (but not saddles!) have objective values which are nearly as small as those of the global minima; see for example Ge et al. [25,26], Soltanolkotabi et al. [49], Sun et al. [50,52]. Some preliminary results have suggested that this may be a quite general phenomenon. For example, Choromanska et al. [19] study the loss surface of a particular Gaussian random field as a proxy for understanding the objective landscape of deep neural nets. The results leverage the Kac–Rice Theorem [5,9], and establish that critical points with more positive eigenvalues have lower expected function value, often close to that of the global minimizer. We remark that functions drawn from this Gaussian random field model share the strict saddle property defined above, and so our results apply in this setting. On the other hand, our results are considerably more

general, as they do not place stringent generative assumptions¹ on the objective function f .

1.2 Organization

The rest of the paper is organized as follows. Section 2 introduces the notation and definitions used throughout the paper. Section 3 provides an intuitive explanation for why it is unlikely that gradient descent converges to a saddle point, by studying a non-convex quadratic and emphasizing the analogy with power iteration. Section 4 develops the main technical theorem, which uses the stable manifold theorem to show that the stable set of unstable fixed points has measure zero. Section 5 applies the main theorem to show that gradient descent, block coordinate descent, proximal point, manifold gradient descent, and mirror descent all avoid strict saddle points. Finally, we conclude in Sect. 6 by suggesting several directions of future work.

2 Preliminaries

Throughout the paper, we will use $f : \mathcal{X} \rightarrow \mathbf{R}$ to denote a real-valued function in C^2 , the space of twice-continuously differentiable functions.

Definition 1 (*Strict saddle*) When $\mathcal{X} = \mathbf{R}^n$,

1. A point x^* is a critical point of f if $\nabla f(x^*) = 0$.
2. A point x^* is a strict saddle² of f if x^* is a critical point and $\lambda_{\min}(\nabla^2 f(x^*)) < 0$, where $\lambda_{\min}(H)$ is the smallest eigenvalue of H . Let \mathcal{X}^* denote the set of strict saddle points.

When \mathcal{X} is a manifold, the same definition applies, but with gradient and Hessian replaced by the Riemannian gradient $\nabla_R f(x)$ and Riemannian Hessian $\nabla_R^2 f(x)$. See Sect. 5.5 for details, and Chapter 5.5 of Absil et al. [2].

Our interest is in the attraction region of an optimization algorithm g , viewed as a mapping from $\mathcal{X} \rightarrow \mathcal{X}$. The iterates of the algorithm are generated by the sequence

$$x_k = g(x_{k-1}) = g^k(x_0),$$

where g^k is the k -fold composition of g . As an example, gradient descent corresponds to $g(x) = x - \alpha \nabla f(x)$.

Since we are interested in the region of attraction of a critical point, we provide the definition of the stable set.

Definition 2 (*Global stable set*) The global stable set of the strict saddles is the set of initial conditions where iteration of the mapping g converges to a strict saddle. This is defined as

¹ This line of work assumes that f is a random function with a specific distribution.

² For the purposes of this paper, strict saddle points include local maximizers.

$$W_g = \left\{ x_0 : \lim_k g^k(x_0) \in \mathcal{X}^* \right\}.$$

3 Intuition

To illustrate why gradient descent and related first-order methods do not converge to saddle points, consider the case of a non-convex quadratic, $f(x) = \frac{1}{2}x^T Hx$. Without loss of generality, assume $H = \mathbf{diag}(\lambda_1, \dots, \lambda_n)$ with $\lambda_1, \dots, \lambda_k > 0$ and $\lambda_{k+1}, \dots, \lambda_n < 0$. $x^* = 0$ is the unique critical point of this function and the Hessian at x^* is H . Gradient descent initialized from x_0 has iterates

$$x_{t+1} = g(x_t) = \sum_{i=1}^n (1 - \alpha\lambda_i)^{t+1} \langle e_i, x_0 \rangle e_i.$$

where e_i denote the standard basis vectors. This iteration resembles power iteration with the matrix $I - \alpha H$.

Let $L = \max |\lambda_i|$, and suppose $\alpha < 1/L$. Thus we have $(1 - \alpha\lambda_i) < 1$ for $i \leq k$ and $(1 - \alpha\lambda_i) > 1$ for $i > k$. If $x_0 \in E_s := \text{span}(e_1, \dots, e_k)$, then x_t converges to the saddle point at zero since $(1 - \alpha\lambda_i)^{t+1} \rightarrow 0$. However, if x_0 has a component outside E_s then gradient descent diverges to ∞ . For this simple quadratic function, we see that the global stable set (attractive set) of zero is the subspace E_s . Now, if we choose our initial point at random, the probability of that point landing in E_s is zero as long as $k < n$ (i.e., E_s is not full dimensional).

As an example of this phenomenon for a non-quadratic function, consider the following example from Nesterov [39, Sect. 1.2.3]. Letting $f(x, y) = \frac{1}{2}x^2 + \frac{1}{4}y^4 - \frac{1}{2}y^2$, the corresponding gradient mapping is

$$g(x) = \begin{bmatrix} (1 - \alpha)x \\ (1 + \alpha)y - \alpha y^3 \end{bmatrix}.$$

The critical points are

$$z_1 = \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \quad z_2 = \begin{bmatrix} 0 \\ -1 \end{bmatrix}, \quad z_3 = \begin{bmatrix} 0 \\ 1 \end{bmatrix}.$$

The points z_2 and z_3 are isolated local minima, and z_1 is a saddle point.

Gradient descent initialized from any point of the form $\begin{bmatrix} x \\ 0 \end{bmatrix}$ converges to the saddle point z_1 . Any other initial point either diverges, or converges to a local minimum, so the stable set of z_1 is the x -axis, which is a zero-measure set in \mathbf{R}^2 . By computing the Hessian,

$$\nabla^2 f(x) = \begin{bmatrix} 1 & 0 \\ 0 & 3y^2 - 1 \end{bmatrix},$$

we find that $\nabla^2 f(z_1)$ has one positive eigenvalue with eigenvector that spans the x -axis, thus agreeing with our above characterization of the stable set. If the initial point is chosen randomly, there is zero probability of initializing on the x -axis and thus zero probability of converging to the saddle point z_1 .

For gradient descent, the local attractive set of a critical point x^* is well-approximated by the span of the eigenvectors corresponding to positive eigenvalues of the Hessian. By an application of Taylor's theorem, one can see that if the initial point x_0 is chosen uniformly random in a small neighborhood around x^* , then the probability of initializing in the span of these eigenvectors is zero whenever there is a negative eigenvalue. Thus, gradient descent initialized at x_0 will leave the neighborhood of x^* . Although this argument provides valuable intuition, there are several difficulties with formalizing this argument: (1) x_0 is randomly distributed over the entire domain, not a small neighborhood around x^* , and Taylor's theorem does not provide any global guarantees, and (2) it does not rule out converging to a different saddle point.

4 Stable manifold theorem and unstable fixed points

4.1 Setup

For the rest of this paper, g is a mapping from \mathcal{X} to itself, and \mathcal{X} is a n -dimensional manifold without boundary. Recall that a C^k -smooth, n -dimensional manifold is a space \mathcal{X} , together with a collection of *charts* $\{(U_\alpha, \phi_\alpha)\}$, called an *atlas*, where each ϕ_α is a homeomorphism from an open subset $U_\alpha \subset \mathcal{X}$ to \mathbb{R}^n . The charts are required to be compatible in the sense that, whenever $U_\alpha \cap U_\beta \neq \emptyset$, then the transition map $\phi_\alpha \circ \phi_\beta^{-1}$ is a C^k map from $\phi_\beta(U_\beta \cap U_\alpha) \rightarrow \mathbb{R}^n$. We also require that $\bigcup_{\alpha \in \mathcal{I}} U_\alpha = \mathcal{X}$, and \mathcal{X} is second countable, which means that for any set U contained in $\bigcup_{\alpha \in \mathcal{I}} U_\alpha$ for some index set \mathcal{I} , there exists a countable set $\mathcal{J} \subset \mathcal{I}$ such that $U \subset \bigcup_{\alpha \in \mathcal{J}} U_\alpha$. We can now recall the definition of a measure zero subset of a manifold:

Definition 3 (Section 5.4 of Mikusinski and Taylor [36]) Given a n -dimensional manifold \mathcal{X} , we say that a set $S \subset \mathcal{X}$ is measure zero if there is an atlas $\{U_i, \phi_i\}_{i \geq 1}$ such that $\phi_i(S \cap U_i)$ has Lebesgue-measure zero as a subset of \mathbb{R}^n . In this case, we use the shorthand $\mu(S) = 0$. The measure zero property is independent of the choice of atlas [36, Chapter 5].

Definition 4 (Chapter 3 of Absil et al. [2]) The differential of g , denoted as $Dg(x)$, is a linear operator from $T(x) \rightarrow T(g(x))$, where $T(x)$ is the tangent space of \mathcal{X} at point x . Given a curve $\gamma(t)$ in \mathcal{X} with $\gamma(0) = x$ and $\frac{d\gamma}{dt}(0) = v \in T(x)$, the linear operator is defined as $Dg(x)v = \frac{d(g \circ \gamma)}{dt}(0) \in T(g(x))$. The determinant of the linear operator $\det(Dg(x))$ is the determinant of the matrix representing $Dg(x)$ with respect to an arbitrary basis.³

Lemma 1 Let $E \subset \mathcal{X}$ be a measure zero subset. If $\det(Dg(x)) \neq 0$ for all $x \in \mathcal{X}$, then $\mu(g^{-1}(E)) = 0$.

³ The determinant is invariant under similarity transformations, so is independent of the choice of basis.

Proof Let $h = g^{-1}$, and $\{(V_i, \psi_i)\}$ be a countable collection of charts of the co-domain of h . By countable additivity of measure, it suffices to show that each $h(E) \cap V_i$ is measure zero. Without loss of generality, we may assume that $h(E)$ is contained in a chart (V, ψ) , else we could repeat the same argument for each element of the chart.

We wish to show that $\mu(\psi \circ h(E)) = 0$. Let (U_j, ϕ_j) be another countable collection of charts of the domain of g . Define $E_j = E \cap U_j$, and note that $E = \cup_{i=1}^\infty E_i = \cup \phi_i^{-1} \circ \phi_i(E_i)$. Thus

$$\begin{aligned} \mu(\psi \circ h(E)) &= \mu(\psi \circ h(\cup_i \phi_i^{-1} \circ \phi_i(E_i))) \\ &\leq \sum_{i=1}^\infty \mu(\psi \circ h \circ \phi_i^{-1}(\phi_i(E_i))). \end{aligned}$$

By assumption, $\phi_i(E_i)$ is measure zero. The function $\psi \circ h \circ \phi_i^{-1} = \psi \circ g^{-1} \circ \phi_i^{-1}$ is C^1 if $\det(Dg) \neq 0$, and thus locally Lipschitz, so preserves measure zero sets. By countable additivity and the displayed equation above, E has measure zero. \square

4.2 Unstable fixed points

Definition 5 (Set of unstable fixed points) Let

$$\mathcal{A}_g^* = \left\{ x : g(x) = x, \max_i |\lambda_i(Dg(x))| > 1 \right\}$$

be the set of fixed points where the differential has at least a single eigenvalue with magnitude greater than one. These are the unstable fixed points. \square

Theorem 1 (Theorem III.7, Shub [48]) Let x^* be a fixed point for the C^r local diffeomorphism $g : \mathcal{X} \rightarrow \mathcal{X}$. Suppose that $E = E_{cs} \oplus E_u$, where E_{cs} is the span of the eigenvectors corresponding to eigenvalues of magnitude less than or equal to one of $Dg(x^*)$, and E_u is the span of the eigenvectors corresponding to eigenvalues of magnitude greater than one of $Dg(x^*)$. Then there exists a C^r embedded disk W_{loc}^{cs} that is tangent to E_{cs} at x^* called the local stable center manifold. Moreover, there exists a neighborhood B of x^* , such that $g(W_{loc}^{cs}) \cap B \subset W_{loc}^{cs}$ and $\cap_{k=0}^\infty g^{-k}(B) \subset W_{loc}^{cs}$.

Theorem 2 Let g be a C^1 mapping from $\mathcal{X} \rightarrow \mathcal{X}$ and $\det(Dg(x)) \neq 0$ for all $x \in \mathcal{X}$. Then the set of initial points that converge to an unstable fixed point has measure zero, $\mu(\{x_0 : \lim_{k \rightarrow \infty} x_k \in \mathcal{A}_g^*\}) = 0$.

Proof For each $x^* \in \mathcal{A}_g^*$, there is an associated open neighborhood B_{x^*} promised by the Stable Manifold Theorem 1. $\cup_{x^* \in \mathcal{A}_g^*} B_{x^*}$ forms an open cover, and since \mathcal{X} is second-countable we can extract a countable subcover, so that $\cup_{x^* \in \mathcal{A}_g^*} B_{x^*} = \cup_{i=1}^\infty B_{x_i^*}$.

Define $W = \{x_0 : \lim_{k \rightarrow \infty} x_k \in \mathcal{A}_g^*\}$. Fix a point $x_0 \in W$. Since $x_k \rightarrow x^* \in \mathcal{A}_g^*$, then for some non-negative integer T and all $t \geq T$, $g^t(x_0) \in \cup_{x^* \in \mathcal{A}_g^*} B_{x^*}$. Since we have a countable sub-cover, $g^t(x_0) \in B_{x_i^*}$ for some $x_i^* \in \mathcal{A}_g^*$ and all $t \geq T$. Since

$g^{t+k}(x_0) \in B_{x_i^*}$, then $g^t(x_0) \in g^{-k}(B_{x_i^*})$, and thus $g^t(x_0) \in \bigcap_{k=0}^\infty g^{-k}(B_{x_i^*})$ for all $t \geq T$. By Theorem 1, $S_i \triangleq \bigcap_{k=0}^\infty g^{-k}(B_{x_i^*})$ is a subset of the local center stable manifold which has co-dimension at least one, and S_i is thus measure zero.

Finally, $g^T(x_0) \in S_i$ implies that $x_0 \in g^{-T}(S_i)$. Since T is unknown we union over all non-negative integers, to obtain $x_0 \in \bigcup_{j=0}^\infty g^{-j}(S_i)$. Since x_0 was arbitrary, we have shown that $W \subset \bigcup_{i=1}^\infty \bigcup_{j=0}^\infty g^{-j}(S_i)$. Using Lemma 1 and that countable union of measure zero sets is measure zero, W has measure zero. \square

Next, we state a simple corollary that only requires verifying $\det(Dg(x)) \neq 0$, and $\mathcal{X}^* \subset \mathcal{A}_g^*$.

Corollary 1 *Under the same conditions as Theorem 2, and in addition assume $\mathcal{X}^* \subset \mathcal{A}_g^*$, then the global stable set (defined in Definition 2) is of measure zero, $\mu(W_g) = 0$.*

Proof Since $\mathcal{X}^* \subset \mathcal{A}_g^*$, then the global stable set satisfies $W_g \subset \{x_0 : \lim_{k \rightarrow \infty} g^k(x_0) \in \mathcal{A}_g^*\}$. Using Theorem 2, $\mu(W_g) = 0$. \square

5 Application to optimization

5.1 Gradient descent

As an application of Theorem 2, we show that gradient descent avoids saddle points. Consider the gradient descent algorithm with step-size α :

$$x_{k+1} = g(x_k) \triangleq x_k - \alpha \nabla f(x_k). \tag{1}$$

Assumption 1 (Bounded Hessian) Let $f \in C^2$, and $\|\nabla^2 f(x)\|_2 \leq L$ for all x .

Proposition 1 *Every strict saddle point x^* is an unstable fixed point of gradient descent, meaning $\mathcal{X}^* \subset \mathcal{A}_g^*$.*

Proof First we verify that critical points of f are fixed points of g . Since $\nabla f(x) = 0$, then $g(x) = x - \alpha \nabla f(x) = x$ and is a fixed point.

At a strict saddle $x^* \in \mathcal{X}^*$, $Dg(x^*) = I - \alpha \nabla^2 f(x^*)$ with eigenvalues $1 - \alpha \lambda_i$, where λ_i are eigenvalues of $\nabla^2 f(x^*)$. Since x^* is a strict saddle, then there is at least one eigenvalue $\lambda < 0$, and $1 - \alpha \lambda_i > 1$. Thus $x^* \in \mathcal{A}_g^*$. \square

Proposition 2 *Under Assumption 1 and $\alpha < \frac{1}{L}$, then $\det(Dg(x)) \neq 0$ for all x .*

Proof By a straightforward calculation

$$Dg(x) = I - \alpha \nabla^2 f(x) = I - \alpha V D V^T,$$

where $\nabla^2 f(x) = V D V^T$. The eigenvalues of $Dg(x)$ are $1 - \alpha \lambda_i$, and so

$$\det(Dg(x)) = \prod_i (1 - \alpha \lambda_i).$$

Using the bounded Hessian assumption, $\alpha < 1/|\lambda_i|$ and each term in the product is positive, so $\det(Dg(x)) > 0$. □

Corollary 2 *Let g be the gradient descent algorithm as defined in Eq. (1). Under Assumption 1 and $\alpha < \frac{1}{L}$, the stable set of the strict saddle points has measure zero, meaning $\mu(W_g) = 0$.*

Proof The proof is a straightforward application of the previous two Propositions and Corollary 1. Proposition 1 shows that $\mathcal{X}^* \subset \mathcal{A}_g^*$, and Proposition 2 shows that $\det(Dg(x)) \neq 0$. By applying Corollary 1, we conclude that $\mu(\{x_0 : \lim_k g^k(x_0) \in \mathcal{X}^*\}) = 0$. □

5.2 Proximal point

The proximal point algorithm is given by the iteration

$$x_{k+1} = g(x) \triangleq \arg \min_z f(z) + \frac{1}{2\alpha} \|x_k - z\|^2. \tag{2}$$

Proposition 3 *Under Assumption 1 and $\alpha < \frac{1}{L}$, then*

1. $\det(Dg(x)) \neq 0$.
2. Every strict saddle point x^* is an unstable fixed point of proximal point, meaning $\mathcal{X}^* \subset \mathcal{A}_g^*$.

Proof Since ∇f is L -Lipschitz, $f(z) + \frac{1}{2\alpha} \|x - z\|^2$ is strongly convex for $\alpha < \frac{1}{L}$, and the arg min is well-defined and unique. By the optimality conditions, $g(x) + \alpha \nabla f(g(x)) = x$. By implicit differentiation, $Dg(x) + \alpha \nabla^2 f(g(x))Dg(x) = I$, and so

$$Dg(x) = (I + \alpha \nabla^2 f(g(x)))^{-1}.$$

At a strict saddle x^* , $Dg(x^*) = (I + \alpha \nabla^2 f(x^*))^{-1}$, and thus has an eigenvalue greater than one. For $\alpha < \frac{1}{L}$, $Dg(x)$ is invertible, and thus $\det(Dg(x)) \neq 0$. □

By combining Proposition 3 and Corollary 1, we have the following:

Corollary 3 (Proximal Point) *Let g be the proximal point algorithm as defined in Eq. (2). Under Assumption 1 and $\alpha < \frac{1}{L}$, the stable set of the strict saddle points has measure zero, meaning $\mu(W_g) = 0$.*

5.3 Coordinate descent

We define $g_i(x) = x - \alpha(0, \dots, 0, \frac{\partial f(x)}{\partial x_i}, 0, \dots, 0)$ to be the coordinate descent update of index i in Algorithm 1. One iteration of coordinate gradient descent corresponds to the update

$$x_{k+1} = g(x_k) = g_n \circ g_{n-1} \circ \dots \circ g_1(x_k). \tag{4}$$

Algorithm 1: Coordinate Descent

- 1 **Input:** Function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, step size α , initial point $x_0 \in \mathbb{R}^n$
- 2 **For** $k = 0, 1, \dots,$
- 3 **For index** $i = 1, \dots, n$
- 4 $x_{k+1}^i \leftarrow x_k^i - \alpha \frac{\partial f(y_k^{i-1})}{\partial x_i}$, where

$$y_k^0 = x_k \text{ and } y_k^i = (x_{k+1}^1, \dots, x_{k+1}^i, x_k^{i+1}, \dots, x_k^n) \tag{3}$$

Assumption 2 (*Lipschitz coordinate gradient*) Let $f \in C^2$, and $|\nabla^2 f(x)_{ii}| \leq L_{\max}$ for all $i \in [n]$ and x .

Lemma 2 *The differential is*

$$Dg(x_k) = \prod_{j=1}^n (I - \alpha e_{n-j+1} e_{n-j+1}^T \nabla^2 f(y_k^{n-j})), \tag{5}$$

where e_i is the i th standard basis vector.

Proof This is an application of the chain rule. The differential of the composition of two functions $f \circ h$ is just $Df(h(x)) \cdot Dh(x)$. By repeatedly applying this and observing that $Dg_i(x) = I - \alpha e_i e_i^T \nabla^2 f(x)$, we have the result. \square

Proposition 4 *Under Assumption 2 and $\alpha < \frac{1}{L_{\max}}$, then $\det(Dg(x)) \neq 0$.*

Proof It suffices to prove that every term of Eq. (5) is an invertible matrix. Using the matrix determinant lemma, the characteristic polynomial of the matrix $I - \alpha e_i e_i^T \nabla^2 f(x)$ is equal to $(\lambda - 1)^{n-1} (\lambda - 1 + \alpha \frac{\partial^2 f(x)}{\partial x_i^2})$. For $\alpha < \frac{1}{\frac{\partial^2 f(x)}{\partial x_i^2}}$, the eigenvalues of Dg_i are all positive, and thus Dg_i is invertible. \square

Proposition 5 (*Instability at saddle points*) *Under Assumption 2 and $\alpha < \frac{1}{L_{\max}}$, every strict saddle point x^* is an unstable fixed point of coordinate descent, meaning $\mathcal{X}^* \subset \mathcal{A}_g^*$.*

Proof Let $H = \nabla^2 f(x^*)$, $J = Dg(x^*) = \prod_{j=1}^n (I - \alpha e_{n-j+1} e_{n-j+1}^T H)$, and y_0 be the eigenvector corresponding to the smallest (leftmost) eigenvalue of H .

We shall prove that $\|J^t y_0\|_2 \geq c(1 + \frac{\epsilon}{4})^t$ for some ϵ, c which depend on H , but not on t . By proving the aforementioned fact, we get that $\|J^t\|_2 \geq \frac{\|J^t y_0\|_2}{\|y_0\|_2} \geq \frac{c}{\|y_0\|_2} (1 + \frac{\epsilon}{4})^t$. Gelfand’s theorem states that the spectral radius ρ satisfies

$$\rho(J) = \lim_{t \rightarrow \infty} \|J^t\|_2^{1/t},$$

and thus

$$\rho(J) = \lim_{t \rightarrow \infty} \|J^t\|_2^{1/t} \geq \lim_{t \rightarrow \infty} \left(\frac{c}{\|y_0\|_2} \right)^{1/t} \left(1 + \frac{\epsilon}{4} \right) = \left(1 + \frac{\epsilon}{4} \right),$$

i.e., J must have an eigenvalue greater than one.

We fix some arbitrary iteration t and let $y_t = J^t y_0$. We will first show that there exists an $\epsilon > 0$ so that

$$y_{t+1}^T H y_{t+1} \leq (1 + \epsilon) y_t^T H y_t, \tag{6}$$

for all $t \in \mathbb{N}$. Let $z_1 = y_t$ and $z_{i+1} = (I - \alpha e_i e_i^T H) z_i = z_i - \alpha (e_i^T H z_i) e_i$, so that $y_{t+1} = J y_t = z_{n+1}$. We see that the sequence $z_{i+1}^T H z_{i+1}$ is decreasing (non-increasing),

$$\begin{aligned} z_{i+1}^T H z_{i+1} &= [z_i^T - \alpha (e_i^T H z_i) e_i^T] H [z_i - \alpha (e_i^T H z_i) e_i] \\ &= z_i^T H z_i - \alpha (z_i^T H e_i) (e_i^T H z_i) - \alpha (e_i^T H z_i) e_i^T H z_i + \alpha^2 (e_i^T H z_i)^2 e_i^T H e_i \\ &= z_i^T H z_i - \alpha (z_i^T H e_i)^2 (2 - \alpha e_i^T H e_i) \\ &< z_i^T H z_i - \alpha (z_i^T H e_i)^2, \end{aligned} \tag{7}$$

where the last inequality uses that $\alpha < \frac{1}{L_{\max}}$.

Next, we use the following claim to show sufficient decrease by lower bounding $(z_i^T H e_i)^2$.

Claim Let y_t be in the range of H . There exists a $j \in [n]$ so that $\alpha |e_j^T H z_j| \geq \delta \|z_j\|_2$ for some global constant $\delta > 0$ that depends on H, n .

The proof of the claim is provided in ‘‘Appendix A’’.

Decompose $y_t = y_{\mathcal{N}} + y_{\mathcal{R}}$ into the orthogonal components defined by the nullspace $\mathcal{N}(H)$ and range space $\mathcal{R}(H)$. Notice that J acts as the identity on $\mathcal{N}(H)$, so

$$y_{t+1} = J y_t = y_{\mathcal{N}} + J y_{\mathcal{R}}.$$

Define an auxiliary sequence $\bar{y}_{t+1} = J \bar{y}_t$, and $\bar{y}_t = y_{\mathcal{R}}$. Similarly, $\bar{z}_1 = y_{\mathcal{R}}$, $\bar{z}_{i+1} = (I - \alpha e_i e_i^T H) \bar{z}_i$, and $\bar{z}_{n+1} = \bar{y}_{t+1}$. It follows that

$$\begin{aligned} y_{t+1}^T H y_{t+1} &= (y_{\mathcal{N}} + J y_{\mathcal{R}})^T H (y_{\mathcal{N}} + J y_{\mathcal{R}}) \\ &= (J y_{\mathcal{R}})^T H (J y_{\mathcal{R}}) \\ &= \bar{z}_{n+1}^T H \bar{z}_{n+1} \\ &\leq \bar{z}_{j+1}^T H \bar{z}_{j+1} \quad (\text{non-increasing property in Eq. (7)}) \\ &< \bar{z}_j^T H \bar{z}_j - \alpha (\bar{z}_j^T H e_j)^2 \quad (\text{using Eq. (7)}) \\ &< \bar{z}_j^T H \bar{z}_j - \frac{\delta^2}{\alpha} \|\bar{z}_j\|_2^2 \quad (\text{using Claim 4}) \end{aligned}$$

$$\begin{aligned}
 &\leq \bar{z}_j^T H \bar{z}_j + \frac{\delta^2}{\alpha L} \bar{z}_j^T H \bar{z}_j \quad \left(\text{since } L \|\bar{z}_j\|^2 \geq \bar{z}_j^T H \bar{z}_j\right) \\
 &= \left(1 + \frac{\delta^2}{\alpha L}\right) \bar{z}_j^T H \bar{z}_j \\
 &\leq \left(1 + \frac{\delta^2}{\alpha L}\right) \bar{z}_1^T H \bar{z}_1 \quad (\text{non-increasing property}) \\
 &= \left(1 + \frac{\delta^2}{\alpha L}\right) \bar{y}_i^T H \bar{y}_i \\
 &= \left(1 + \frac{\delta^2}{\alpha L}\right) y_i^T H y_i,
 \end{aligned}$$

where we used that Eq. (7) applies also the sequence \bar{z}_j . Let $\epsilon = \frac{\delta^2}{\alpha L}$. By inducting, and noting that $y_0^T H y_0 = -\lambda$,

$$\begin{aligned}
 y_i^T H y_i &\leq (1 + \epsilon)^i y_0^T H y_0 \\
 &= -\lambda(1 + \epsilon)^i.
 \end{aligned}$$

Using $-\lambda \|y_i\|_2^2 \leq y_i^T H y_i$,

$$\begin{aligned}
 -\lambda \|y_i\|_2^2 &\leq -\lambda(1 + \epsilon)^i \\
 \|y_i\|_2^2 &\geq (1 + \epsilon)^i \\
 \|J^i y_0\|_2 &\geq (1 + \epsilon)^{i/2} \\
 &\geq \left(1 + \frac{\epsilon}{4}\right)^i,
 \end{aligned}$$

where the last inequality uses that $\epsilon \leq \frac{1}{2}$. By Gelfand’s theorem, we have established

$$\rho(J) \geq 1 + \frac{\epsilon}{4},$$

and thus J has an eigenvalue of magnitude greater than one. Thus $x^* \in \mathcal{A}_g^*$. □

By combining Propositions 5, 4, and Corollary 1, we have the following:

Corollary 4 (Coordinate Descent) *Let g be the coordinate descent algorithm as defined in Eq. (4). Under Assumption 2 and $\alpha < \frac{1}{L_{\max}}$, the stable set of the strict saddle points has measure zero, meaning $\mu(W_g) = 0$.*

Remark 1 In the worst-case, $L_{\max} = L$, but in many instances $L_{\max} \ll L$, so coordinate descent can use more aggressive step-sizes. The step-size choice $\alpha < \frac{1}{L_{\max}}$ is standard for coordinate-descent methods [46].

Algorithm 2: Block Coordinate Descent

- 1 **Input:** Function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, step size α , initial point $x_0 \in \mathbb{R}^n$
- 2 **For** $k = 0, 1, \dots$,
- 3 **For block** $i = 1, \dots, b$
- 4 **For index** j **in block** i
- 5 $x_{k+1}^j \leftarrow x_k^j - \alpha \frac{\partial f(y_k^{S_i-1})}{\partial x_j}$, where

$$y_k^{S_0} = x_k \text{ and } y_k^{S_i} = (x_{k+1}^{S_1}, \dots, x_{k+1}^{S_i}, x_k^{S_{i+1}}, \dots, x_k^{S_b}) \tag{8}$$

5.4 Block coordinate descent

The results of this section are a generalization of the previous section, but we present the coordinate descent case separately, since the proofs are simpler.

We partition the set $[n] = \{1, 2, \dots, n\}$ to b blocks $\{S_1, S_2, \dots, S_b\}$ such that $[n] = \cup_i S_i$. For ease of notation, we define $S_0 = \emptyset$.

We define $g_i(x)$ to be the block coordinate descent update of block i in Algorithm 2. Block coordinate gradient descent is a dynamical system

$$x_{k+1} = g(x_k) = g_b \circ g_{b-1} \circ \dots \circ g_1, \tag{9}$$

where $g_i(x) = x - \alpha \sum_{j \in S_i} e_j^T \nabla f(x)$. We define the matrix $P_S = \sum_{i \in S} e_i e_i^T$, i.e., the projector onto the entries in S .

Lemma 3 *The differential is*

$$Dg(x_k) = \prod_{i=1}^b (I - \alpha P_{S_{b-i+1}} \nabla^2 f(y_k^{S_{b-i}})). \tag{10}$$

Proof This is an application of the chain rule. The differential of the composition of two functions $f \circ h$ is just $Df(h(x)) \cdot Dh(x)$. By repeatedly applying this and observing that $Dg_i(x) = I - \alpha P_{S_i} \nabla^2 f(x)$, we obtain the result. □

Assumption 3 Let $f \in C^2$, and $\nabla^2 f(x)_S$ be the submatrix of $\nabla^2 f(x)$ by extracting the rows and columns indexed by S . Let $L_b = \max_{i \in [b]} \|\nabla^2 f(x)_{S_i}\|_2$

Proposition 6 *Under Assumption 3 and $\alpha < \frac{1}{L_b}$, then $\det(Dg(x)) \neq 0$.*

Proof It suffices to prove that every term of the product (10) is an invertible matrix. Every matrix of the form $I - \alpha P_{S_i} \nabla^2 f(x)$ has $n - |S_i|$ eigenvalues equal to one and the rest of its eigenvalues correspond to eigenvalues of $I_{S_i} - \alpha \nabla^2 f(x)_{S_i, S_i}$. Since $\alpha < \frac{1}{L_b}$, then the eigenvalues of $I_{S_i} - \alpha \nabla^2 f(x)_{S_i, S_i}$ are all greater than zero. Thus each $I - \alpha P_{S_i} \nabla^2 f(x)$ is invertible, and Dg is also invertible. □

Proposition 7 (Stability at fixed points) *Let x^* be a strict saddle point of f . The Jacobian of the update rule of block coordinate descent computed at point x^* has an eigenvalue of modulus greater than one.*

The proof is provided in ‘‘Appendix B’’. By combining Propositions 7, 6, and Corollary 1, we have the following:

Corollary 5 (Block Coordinate Descent) *Let g be the block coordinate descent algorithm as defined in Eq. (9). Under Assumption 3 and $\alpha < \frac{1}{L_b}$, the stable set of the strict saddle points has measure zero, meaning $\mu(W_g) = 0$.*

Remark 2 The step-size choice $\alpha < \frac{1}{L_b}$ is standard for block coordinate descent methods [46].

5.5 Manifold gradient descent

Let \mathcal{M} be a submanifold of \mathbf{R}^n , and $\mathcal{T}(x)$ be the tangent space of \mathcal{M} at x . $P_{\mathcal{X}}$ and $P_{\mathcal{T}(x)}$ are the orthogonal projectors onto \mathcal{M} and $\mathcal{T}(x)$ respectively, and $I_{\mathcal{T}}$ is the identity operator on \mathcal{T} . Let \bar{f} be a smooth extension of f to \mathbf{R}^n , and $f = \bar{f}|_{\mathcal{M}}$. The manifold gradient descent algorithm is:

$$x_{k+1} = g(x_k) \triangleq P_{\mathcal{M}}(x_k - \alpha P_{\mathcal{T}(x_k)} \nabla \bar{f}(x_k)). \tag{11}$$

Recall that the Riemannian gradient $\nabla_R f(x) = P_{\mathcal{T}(x)} \nabla \bar{f}(x)$, so the above iteration is precisely manifold gradient descent with $P_{\mathcal{M}}$ as retraction.

Proposition 8 *At a strict saddle point x^* , $Dg(x^*)$ has an eigenvalue of magnitude larger than 1.*

Proof Since x^* is a strict saddle, the Riemannian Hessian (see Sect. 5.5 of [2] for the definition) $\nabla_R^2 f(x^*)$ has a negative eigenvalue λ_v and eigenvector v , and $P_{\mathcal{T}(x^*)} \nabla \bar{f}(x^*) = 0$.

Using Lewis and Malick [33, Lemma 4], $DP_{\mathcal{M}}(x) = P_{\mathcal{T}(x)}$ for $x \in \mathcal{M}$,

$$Dg(x^*)v = P_{\mathcal{T}(x^*)}v - \alpha P_{\mathcal{T}(x^*)}D(P_{\mathcal{T}} \nabla \bar{f})(x^*)v.$$

Using Absil et al. [3, Eq. 4], $P_{\mathcal{T}(x)}D(P_{\mathcal{T}} \nabla \bar{f})(x)v = \nabla_R^2 f(x)v$, so

$$Dg(x^*)v = v - \lambda_v v.$$

Thus v is an eigenvector of $Dg(x^*)$ with eigenvalue $1 - \lambda_v > 1$. □

Proposition 9 *For a compact submanifold \mathcal{M} , there is a strictly positive α such that $\det(Dg) \neq 0$.*

Proof Since \mathcal{M} is a compact smooth manifold, $P_{\mathcal{M}}$ is unique and smooth in a neighborhood of radius r of the manifold [4]. Letting $\alpha < \frac{r}{\max_{x \in \mathcal{M}} \|\nabla f(x)\|}$, $P_{\mathcal{M}}(x - \alpha P_{\mathcal{T}(x)} \nabla f(x))$ and its derivatives exist.

We wish to show that $Dg(x) = DP_{\mathcal{M}}(x - \alpha P_{\mathcal{T}(x)} \nabla f(x))(I - \alpha D(P_{\mathcal{T}} \nabla f)(x))$ is invertible. Define

$$h_x(\alpha) = \det \left(DP_{\mathcal{M}}(x - \alpha P_{\mathcal{T}(x)} \nabla f(x))(I - \alpha D(P_{\mathcal{T}} \nabla f)(x)) \right).$$

Using $DP_{\mathcal{M}}(x) = P_{\mathcal{T}(x)}$ [4], $h_x(0) = 1$. Since

$$B := \max_{x \in \mathcal{M}, \alpha < \frac{r}{\max_{x \in \mathcal{M}} \|\nabla f(x)\|}} \left| \frac{nh_x}{n\alpha}(\alpha) \right| < \infty,$$

we see that for

$$\alpha < C_{\mathcal{M},f} \triangleq \min \left(\frac{r}{\max_{x \in \mathcal{M}} \|\nabla f(x)\|}, \frac{1}{B} \right)$$

that $h_x(\alpha)$ is positive, so Dg is invertible. □

Corollary 6 *Let g be the operator corresponding to the manifold gradient descent algorithm of Eq. (11), and \mathcal{M} be a compact sub-manifold of \mathbf{R}^n . Then there is a $C_{\mathcal{M},f} > 0$, that only depends on the properties of \mathcal{M} and f , such that for any step-size $\alpha < C_{\mathcal{M},f}$, the stable set of the strict saddle points has measure zero, meaning $\mu(W_g) = 0$.*

Proof The proof is a straightforward application of the previous two Propositions and Corollary 1. Proposition 8 shows that $\mathcal{X}^* \subset \mathcal{A}_g^*$, and Proposition 9 shows that $\det(Dg(x)) \neq 0$. By applying Corollary 1, we conclude that $\mu(\{x_0 : \lim_k g^k(x_0) \in \mathcal{X}^*\}) = 0$. The parameter $C_{\mathcal{M},f}$ is specified in the proof of Proposition 9. □

5.6 Mirror descent

In this section, we consider the mirror descent algorithm. Let \mathcal{D} be a convex open subset of \mathbf{R}^n , and $\mathcal{X} = \mathcal{D} \cap \mathcal{M}$ for some affine space \mathcal{M} . Given a mirror map Φ , we define the mirror descent algorithm in Algorithm 3.

Assumption 4 (Mirror map) We say that Φ is a *mirror* map if it satisfies the following properties:

1. Φ is a proper closed convex function with $\overline{\mathcal{D}} = \text{dom}(\Phi)$.
2. Φ is C^2 on $\mathcal{D} = \text{int}(\overline{\mathcal{D}})$.
3. The gradient of Φ is surjective onto \mathbf{R}^n , that is $\nabla \Phi(\mathcal{D}) = \mathbf{R}^n$.
4. The subgradient $\partial \Phi(x) = \{\nabla \Phi(x)\}$ for $x \in \mathcal{D}$ and $\partial \Phi(x) = \emptyset$ for $x \notin \mathcal{D}$.

Due to the affine constraint, \mathcal{X} may not be full-dimensional, so we define the appropriate notions of gradient and Hessian. Let \mathcal{T} be the tangent space of \mathcal{M} . The Riemannian gradient is $\nabla_R \Phi(x) = P_{\mathcal{T}} \nabla \Phi(x)$. Similarly the Riemannian Hessian is $\nabla_R^2 \Phi(x) = P_{\mathcal{T}} \nabla^2 \Phi(x) P_{\mathcal{T}}$ and is a linear mapping from $\mathcal{T} \rightarrow \mathcal{T}$. Finally, the mirror descent mapping is defined as

$$g(x) = h \circ F(x), \tag{12}$$

Algorithm 3: Mirror Descent

- 1 **Input:** Function $f : \mathcal{X} \rightarrow \mathbb{R}$, step size α , and initial point x_0 .
- 2 **For** $k = 0, 1, \dots$,
- 3 Update

$$x_{k+1} \leftarrow h(\nabla\Phi(x_k) - \alpha\nabla f(x_k)),$$

where $h(x) \triangleq \arg \max_{z \in \mathcal{X}} z^T x - \Phi(z)$.

with $F(x) = \nabla\Phi(x) - \alpha\nabla f(x)$ and $h(x) = \arg \max_{z \in \mathcal{X}} z^T x - \Phi(z)$.

Assumption 5 (*Strong convexity of Φ and Lipschitz gradient*) Let $I_{\mathcal{T}}$ be the identity mapping on \mathcal{T} . We assume that

1. Φ is μ -strongly convex, meaning $\nabla_R^2\Phi(x) \succeq \mu I_{\mathcal{T}}$ for $x \in \mathcal{X}$.
2. f has L -Lipschitz gradient, meaning $\nabla_R^2 f(x) \preceq L I_{\mathcal{T}}$ for $x \in \mathcal{X}$.

Example 1 (*Probability simplex*) Define the mirror map $\Phi(x) = \sum x_i \log x_i$, with \mathcal{D} being the positive orthant $\mathbf{R}_{>0}^n$, $\overline{\mathcal{D}}$ is the non-negative orthant, and affine space $\mathcal{M} = \{x : \sum_i x_i = 1\}$. The domain is $\overline{\mathcal{X}} = \overline{\mathcal{D}} \cap \mathcal{M}$ which is the probability simplex. The mirror descent algorithm corresponds to the update:

$$x_i \leftarrow \frac{x_i \exp\left(-\alpha \frac{\partial f}{\partial x_i}(x)\right)}{\sum_j x_j \exp\left(-\alpha \frac{\partial f}{\partial x_j}(x)\right)}.$$

The strong convexity parameter satisfies $\mu \geq 1$.

We first express the mapping g as a composition of simple mappings.

Lemma 4 *Assume that Φ is a μ -strongly convex mirror map. The mirror descent algorithm can be equivalently expressed as $g(x) = (\nabla_R\Phi)^{-1} \circ P_{\mathcal{T}} \circ F(x)$, and $(\nabla_R\Phi)^{-1} : \mathcal{T} \rightarrow \mathcal{X}$ is a local diffeomorphism.*

Proof Recall that $h(w) = \arg \max_{z \in \mathcal{X}} z^T w - \Phi(z)$. By a standard result [10, Theorem 9.8], the maximum is uniquely attained and $h(w) \in \mathcal{D}$.

By the first-order optimality conditions, $\nabla_R\Phi(h(w)) = P_{\mathcal{T}}w$, and thus

$$h(w) = (\nabla_R\Phi)^{-1} \circ P_{\mathcal{T}}(w).$$

As a shorthand, let $\Psi = (\nabla_R\Phi)^{-1}$. By existence and uniqueness of the maximizer, Ψ is a single-valued function from $\mathcal{T} \rightarrow \mathcal{X}$. Thus $g = h \circ F = \Psi \circ P_{\mathcal{T}} \circ F$.

Finally we verify that $\Psi = (\nabla_R\Phi)^{-1}$ is a local diffeomorphism. Since $D\Psi(\nabla_R\Phi(x)) = (\nabla_R^2\Phi(x))^{-1}$ and Φ is strongly convex over T , the inverse function theorem ensures that Ψ is a local diffeomorphism. \square

Proposition 10 *Under Assumptions 4, 5 and $\alpha < \frac{\mu}{L}$, then*

1. $\det(Dg(x)) \neq 0$.
2. Every strict saddle point of x^* is an unstable fixed point of mirror descent, meaning $x^* \in \mathcal{A}_g^*$.

Proof Again adopt the shorthand $\Psi = (\nabla_R \Phi)^{-1}$. Using Lemma 4, $g = \Psi \circ (P_{\mathcal{T}} \circ F)$, and Ψ is a local diffeomorphism. To show $\det(Dg) \neq 0$, it suffices to show that $P_{\mathcal{T}} \circ F$ is a local diffeomorphism, or equivalently verify that $D(P_{\mathcal{T}} \circ F)(x) : \mathcal{T} \rightarrow \mathcal{T}$ is an invertible linear transformation for every $x \in \mathcal{X}$.

First, $D(P_{\mathcal{T}} \circ F)(x) = \nabla_R^2 \Phi(x) - \alpha \nabla_R^2 f(x)$. Then using that $\nabla_R^2 \Phi(x) \succeq \mu I_{\mathcal{T}}$, $\nabla_R^2 f(x) \preceq LI_{\mathcal{T}}$, and $\alpha \leq \frac{\mu}{L}$,

$$\alpha \nabla_R^2 f(x) \preceq \mu I_{\mathcal{T}} \preceq \nabla_R^2 \Phi(x).$$

Thus $\nabla_R^2 \Phi(x) - \alpha \nabla_R^2 f(x) \succ 0$ and is invertible. This completes our proof of the first part.

Let $x^* \in \mathcal{X}$ be a strict saddle point. First we verify that it is a fixed point of g . Using that x^* is a critical point,

$$\begin{aligned} g(x^*) &= \Psi(P_{\mathcal{T}} \nabla \Phi(x^*) - P_{\mathcal{T}} \nabla f(x^*)) \\ &= \Psi(P_{\mathcal{T}} \nabla \Phi(x^*)) \\ &= \Psi(\nabla_R \Phi(x^*)) \\ &= x^*. \end{aligned}$$

Next we verify that $Dg(x^*)$ has an eigenvalue of magnitude greater than one. Using the chain rule and then inverse function theorem,

$$\begin{aligned} Dg(x^*) &= D\Psi(P_{\mathcal{T}} \circ F(x^*))(\nabla_R^2 \Phi(x^*) - \alpha \nabla_R^2 f(x^*)) \quad (\text{chain rule}) \\ &= \nabla_R^2 \Phi(x^*)^{-1}(\nabla_R^2 \Phi(x^*) - \alpha \nabla_R^2 f(x^*)) \\ &\quad (\text{inverse function theorem and } \Psi \circ P_{\mathcal{T}} \circ F(x^*) = x^*) \\ &= I_{\mathcal{T}} - \alpha \nabla_R^2 \Phi(x^*)^{-1} \nabla_R^2 f(x^*). \end{aligned}$$

Define $A = \nabla_R^2 \Phi(x^*)$ and $H = \nabla_R^2 f(x^*)$. By similarity transformation under $A^{1/2}$,

$$A^{1/2} Dg(x^*) A^{-1/2} = I_{\mathcal{T}} - \alpha A^{-1/2} H A^{-1/2},$$

which is a symmetric linear operator. Define $\bar{v} = A^{1/2} v$, where v is an eigenvector of $\nabla_R^2 f(x^*)$ corresponding to a strictly negative eigenvalue λ , then $\bar{v}^T A^{-1/2} H A^{-1/2} \bar{v} < 0$, so $\lambda_{\min}(A^{-1/2} H A^{-1/2}) < 0$. Thus $1 - \alpha \lambda_{\min}(A^{-1/2} H A^{-1/2})$ is an eigenvalue of $I_{\mathcal{T}} - \alpha A^{-1/2} H A^{-1/2}$ that is greater than one. Since similarity transformations preserve eigenvalues, $Dg(x^*)$ also has an eigenvalue greater than one, and so $x^* \in \mathcal{A}_g^*$. \square

By combining Proposition 10 with Corollary 1, we have the following:

Corollary 7 *Let g be the mirror descent algorithm defined in Eq. (12). Under Assumptions 4, 5, and $\alpha < \frac{\mu}{L}$, then the stable set of the strict saddles in \mathcal{X} is measure zero, meaning $\mu(W_g) = 0$.*

Remark 3 This corollary does not guarantee that the stable set of saddles on $\overline{\mathcal{X}} \setminus \mathcal{X}$ is measure zero. For example in the Multiplicative Weights algorithm, there are fixed points on $\overline{\mathcal{X}} \setminus \mathcal{X}$ (e.g. all the vectors x with support size 1).

5.7 Lojasiewicz inequality and existence of a single limit point

In the previous parts of this section, we have implicitly assumed that $\lim_k x_k$ exists, otherwise the theorems are vacuously true. As we will show with specific citations to recent papers, the existence of the limit is ensured when the function satisfies the Lojasiewicz inequality, and the iterates are bounded. Both assumptions are very weak, as we can explain below.

Assumption 6 (Lojasiewicz inequality) Let x^* be a critical point of f . f is said to satisfy the Lojasiewicz inequality if there are constants $c > 0$ and $\mu \in [0, 1)$ such that

$$\|\nabla f(x)\| \geq c|f(x) - f(x^*)|^\mu,$$

for all x in a neighborhood of x^* .

The Lojasiewicz inequality is very general as discussed in [7,8,13], and applies to all real analytic functions.

Assumption 7 (Bounded iterates) The iterates $\{x_k\}$ remain in a bounded set.

Remark 4 The bounded iterates assumption is necessary for $\lim x_k$ to exist. For example, gradient descent on $f(x) = e^x$ has iterates that minimize f by generating a sequence that $x_k \downarrow -\infty$, and so $\lim x_k$ does not exist. For algorithms that generate non-increasing objective values, bounded iterates can be ensured when the function has compact sub-level sets.

For gradient descent, Absil et al. [1] showed that when f satisfies the Lojasiewicz inequality that $\lim x_k$ exists. We recall the results below.

Proposition 11 (Theorems 3.2 and 4.1 of [1]) *Let f satisfy Assumptions 1, 6, and 7. Assume that the stepsize $\alpha < \frac{1}{L}$, then $\lim x_k$ exists.*

Proof Theorem 3.2 of [1] shows that $\lim x_k$ exists whenever the function satisfies the Lojasiewicz inequality and the algorithm satisfies the sufficient descent condition. An immediate consequence of Theorem 4.1 of [1] is that gradient descent with stepsize $\alpha < \frac{1}{L}$ satisfies the sufficient descent condition. \square

For block coordinate descent and manifold gradient descent, Bolte et al. [14] showed a similar result.

Proposition 12 (Block coordinate descent limit exists [14]) *Let f satisfy Assumptions 3, 6, and 7. Assume that $\alpha < \frac{1}{L_b}$, then the block coordinate descent iterates attain a limit, $\lim x_k$ exists.*

Proof This is proved in Theorem 1 and Sect. 3.6 of [14]. □

Proposition 13 (Mirror Descent limit exists [15]) *Under Assumptions 4, 5, 7, and 6 with $\alpha < \frac{\mu}{L}$, then the mirror descent iterates attain a limit, $\lim x_k$ exists.*

Proof This follows from a minor modification of Bolte et al. [15], which studies mirror descent on composite functions of the form $\Psi(x) = f(x) + g(x)$, where f is smooth and differentiable and g is proper and lower semicontinuous. For our setting, we choose g as the indicator function of M , $g(x) = \mathbf{1}_M(x)$. Bolte et al. assume that $\text{dom}(\Phi) = \mathbf{R}^n$, but this is only used in Eq. (4.7) of [15]. The same result can be obtained by the strong convexity of Φ over the domain $\mathcal{X} = \mathcal{D} \cap \mathcal{M}$ and using that $x_{k+1} - x_k \in \mathcal{T}$. Equation (4.7) with

$$\begin{aligned} f(x_k) - f(x_{k+1}) &\geq \left(\frac{1}{\alpha} - L\right) (\Phi(x_{k+1}) - \Phi(x_k) - \nabla\Phi(x_k)^T(x_{k+1} - x_k)) \\ &= \left(\frac{1}{\alpha} - L\right) \int dt \frac{1}{2}(x_{k+1} - x_k)^T \nabla^2\Phi(x_k + t(x_{k+1} - x_k))(x_{k+1} - x_k) \\ &\geq \left(\frac{1}{\alpha} - L\right) \frac{\mu}{2} \|x_{k+1} - x_k\|^2. \end{aligned}$$

The right-hand side is the same as Eq. (4.7) of [15], so the rest of the argument proceeds exactly as in Theorem 4.1 of Bolte et al. □

Remark 5 For functions f that do not satisfy the Lojasiewicz inequality, it is possible to prove that $\lim_k x_k$ exists under the assumption that the critical points are isolated and Assumption 7. The strategy is standard in the literature and can be found in [31,35], so we omit the details in the sketch below.

Let $\{x_k\}_{k \in \mathbb{N}}$ be the trajectory of any of the algorithms discussed and Ω be the set of limit points. First, Ω is non-empty due to Assumption 7. Since the algorithms are sufficient descent methods, it can be shown Ω consists of critical points, and is a compact, connected set.

Finally, by the assumption that the critical points are isolated, it follows that Ω is a singleton critical point (since Ω is *connected*), i.e., $\Omega = \{x^*\}$ and hence $\lim_k x_k = x^*$.

6 Conclusion

We have shown that first-order methods with random initialization and appropriate constant step-size do not converge to a strict saddle point. Our results apply to gradient descent, proximal point algorithm, coordinate descent, block coordinate descent, manifold gradient descent and mirror descent. The key common insight in analyzing all these optimization methods is to treat these algorithms as dynamical systems. Every strict saddle point is shown to be locally unstable for these first-order methods

and applications of the center-stable manifold theorem suffice to characterize the local behavior. As long as the mapping induced by the optimization method is sufficiently well behaved, e.g. local diffeomorphism, these local arguments can be extended to the whole domain. Proving the instability of saddle points as well as the smoothness and invertibility of the corresponding maps depends upon careful instantiations of these generic arguments (e.g. choice of step-size) on a case-by-case basis. The global instability of saddle points for first-order methods is many times informally invoked without careful discussion about the necessary technical conditions needed to formalize these arguments. We hope that this work will help ground these arguments on a unified formal foundation. We end this paper with a brief discussion of some open directions:

Step-size It is not clear if the step size restrictions are necessary to avoid saddle points (e.g. $\alpha < 1/L$ for gradient descent; see Panageas and Piliouras [42] in which examples are provided where $\alpha < 2/L$ is necessary for gradient descent). Most of the constructions where the gradient method converges to saddle points require fragile initial conditions as discussed in Sect. 3. It remains a possibility that adaptive choice of step-size by Wolfe Line Search or backtracking, may still avoid saddle points provided the initial point is chosen at random.

Strict saddles It is also important to understand how stringent the strict saddle assumption is. Will a perturbation of a function always satisfy the strict saddle property? Adler and Taylor [5] provide very general sufficient conditions for a random function to be Morse, meaning the eigenvalues at critical points are non-zero, which implies the strict saddle condition. These conditions rely on checking that the density of $\nabla^2 f(x)$ has full support conditioned on the event that $\nabla f(x) = 0$. This can be explicitly verified for functions f that arise from learning problems. Similar arguments for applications that arise in game theory are developed in Kleinberg et al. [30].

However, we note that there are very difficult unconstrained optimization problems where the strict saddle condition fails. Perhaps the simplest is optimization of quartic polynomials. Indeed, checking if zero is a local minimizer of the quartic

$$f(x) = \sum_{i,j=1}^n q_{ij} x_i^2 x_j^2$$

is equivalent to checking whether the matrix $Q = [q_{ij}]$ is co-positive, a co-NP complete problem. For this f , the Hessian at $x = 0$ is zero, so $x = 0$ is a second-order KKT point, but not necessarily a local minimizer. By the change of variables $z_i = x_i^2$, we see that checking local minimality in a problem with quadratic objective and non-negative inequality constraints is also co-NP complete.

Speed of convergence Although gradient descent can take exponential amount of time to escape from saddle points at least for some carefully constructed non-convex functions [22], its stochastic counterparts perform much better [28]. It would be interesting to characterize these hard instances to the extent possible and to understand whether they are indeed prevalent in applications of interest (e.g. deep learning). In the other direction, it would be rather useful to show that all first-order methods can be sped up by switching to carefully chosen stochastic variants.

Beyond saddle points Even if saddle points are provably avoided, there can be multiple local minima of widely different objective value. The performance of first-order methods would depend crucially on whether they converge for most initial conditions to nearly optimal global minima. Panageas and Piliouras panageas2016average analyze such a game theoretic application and show that indeed the size of the region of attraction of the good local optima dominates that of the bad local optima implying nearly optimal average case performance. Such arguments depend crucially both on the setting as well as on the chosen optimization method and it would be interesting to explore their applicability in other settings.

A Proof of Claim 4

Proof We assume that $\alpha|e_j^T H z_j| < \delta \|z_j\|_2$ for all $j \in \{1, \dots, n\}$, for some δ to be chosen later. For the base case $j = 2$, it holds that $\|y_t - z_2\|_2 = \|z_1 - z_2\|_2 = \alpha|e_1^T H z_1| < \delta \|z_1\|_2 < 2\delta \|y_t\|_2$ and $\|z_2\|_2 < (1 + 2\delta) \|y_t\|_2$. Suppose for $j \geq 2$ that $\|y_t - z_j\|_2 < 2(j - 1)\delta \|y_t\|_2$ and thus $\|z_j\|_2 < [1 + 2(j - 1)\delta] \|y_t\|_2$. Using induction and triangle inequality we get

$$\begin{aligned} \|y_t - z_{j+1}\|_2 &\leq \|y_t - z_j\|_2 + \|z_j - z_{j+1}\|_2 \\ &= 2(j - 1)\delta \|y_t\|_2 + \alpha|e_j^T H z_j| \\ &< 2(j - 1)\delta \|y_t\|_2 + \delta \|z_j\|_2 \\ &< 2(j - 1)\delta \|y_t\|_2 + \delta[1 + 2(j - 1)\delta] \|y_t\|_2 \\ &\leq 2j\delta \|y_t\|_2, \end{aligned}$$

where we assume $\delta < \frac{1}{2n}$ so that $2(j - 1)\delta < 1$ for all $j \in [n]$. Using the above calculation,

$$\begin{aligned} \alpha|e_i^T H y_t| &< \alpha|e_i^T H z_i| + \alpha|e_i^T H (y_t - z_i)| \\ &< \delta \|z_i\|_2 + \alpha \|H e_i\|_2 \|y_t - z_i\|_2 \\ &< \delta(1 + 2(i - 1)\delta) \|y_t\|_2 + \alpha \|H e_i\|_2 (2(i - 1)\delta) \|y_t\|_2 \\ &\leq \delta(1 + 2n\delta + 2n\alpha L) \|y_t\|. \end{aligned}$$

Thus $\alpha \|H y_t\|_2 < \sqrt{n}\delta(1 + 2n\delta + 2n\alpha L) \|y_t\|_2$, and

$$\sigma_{\min^+}(H) \|y_t\|_2 \leq \|H y_t\|_2 < \frac{\sqrt{n}}{\alpha} \delta(1 + 2n\delta + 2n\alpha L) \|y_t\|_2,$$

where σ_{\min^+} is the smallest non-zero singular value of H . Thus by choosing δ small enough such that

$$\sigma_{\min^+}(H) \geq \frac{\sqrt{n}}{\alpha} \delta(1 + 2n\delta + 2n\alpha L),$$

we have obtained a contradiction. □

B Proof of Proposition 7

Proof Let $H = \nabla^2 f(x^*)$, $J = \text{Dg}(x^*) = \prod_{i=1}^b (I - \alpha P_{S_{b-i+1}} H)$, and y_0 be an eigenvector of the Hessian at x^* .

The proof technique is very similar to that of the proof of Proposition 5. We shall prove that $\|J^t y_0\|_2 \geq c(1 + \eta)^t$. Hence by Gelfand’s theorem J must have at least one eigenvalue with magnitude greater than one.

We fix some arbitrary iteration t and let $y_t = J^t y_0$. We will first show that there exists an $\epsilon > 0$,

$$y_{t+1}^T H y_{t+1} \leq (1 + \epsilon) y_t^T H y_t, \tag{13}$$

for all $t \in \mathbb{N}$. Let $z_1 = y_t$ and $z_{i+1} = (I - \alpha P_{S_i} H) z_i = z_i - \alpha \sum_{j \in S_i} (e_j^T H z_i) e_j$, so that $y_{t+1} = J y_t = z_{b+1}$. We get that

$$\begin{aligned} z_{i+1}^T H z_{i+1} &= \left(z_i^T - 2\alpha \sum_{j \in S_i} (e_j^T H z_i) e_j^T \right) H \left(z_i - \alpha \sum_{j \in S_i} (e_j^T H z_i) e_j \right) \\ &= z_i^T H z_i - 2\alpha \sum_{j \in S_i} (e_j^T H z_i)^2 + \alpha^2 \left(\sum_{j \in S_i} (e_j^T H z_i) e_j \right)^T H \left(\sum_{j \in S_i} (e_j^T H z_i) e_j \right) \\ &< z_i^T H z_i - 2\alpha \sum_{j \in S_i} (e_j^T H z_i)^2 + \alpha^2 L_b \left\| \sum_{j \in S_i} (e_j^T H z_i) e_j \right\|_2^2 \quad (\text{using } \|H_{S_i}\|_2 \leq L_b) \\ &= z_i^T H z_i - \alpha(2 - \alpha L_b) \left\| \sum_{j \in S_i} (e_j^T H z_i) e_j \right\|_2^2 \\ &\leq z_i^T H z_i - \alpha \left\| \sum_{j \in S_i} (e_j^T H z_i) e_j \right\|_2^2 \quad (\text{using } \alpha L_b < 1). \end{aligned}$$

Thus $z_i^T H z_i$ is a decreasing (non-increasing) sequence.

We shall prove that there exists an $i \in [b]$ so that $z_{i+1}^T H z_{i+1} \leq (1 + \delta) z_i^T H z_i$ for some global constant δ to be chosen later.

Claim Let y_t be in the range of H . There exists an $i \in [b]$ so that $\alpha \sum_{j \in S_i} |e_j^T H z_i| \geq \delta \|z_i\|_2$ for some $\delta > 0$.

To finish the proof of the lemma, suppose that Claim B applies. Then by Cauchy–Schwarz, there exists an index i such that

$$\begin{aligned}
 z_{i+1}^T H z_{i+1} &< z_i^T H z_i - \alpha \sum_{j \in S_i} (z_i^T H e_j)^2 \\
 &< z_i^T H z_i - \frac{\alpha}{n} \left(\sum_{j \in S_i} |z_i^T H e_j| \right)^2 < z_i^T H z_i - \frac{\delta^2}{n\alpha} \|z_i\|_2^2.
 \end{aligned}$$

However, $w^T H w \geq \lambda_{\min}(H) \|w\|_2^2 \geq -L \|w\|_2^2$, hence we get that

$$z_{i+1}^T H z_{i+1} < \left(1 + \frac{\delta^2}{\alpha L n} \right) z_i^T H z_i. \tag{14}$$

By choosing $\epsilon = \frac{\delta^2}{\alpha L n}$ we showed that $y_{t+1}^T H y_{t+1} \leq (1 + \epsilon) y_t^T H y_t$ as long as y_t is in the range of H .

Assume that $y_t = y_{\mathcal{N}} + y_{\mathcal{R}}$. It is easy to see $y_t^T H y_t = y_{\mathcal{R}}^T H y_{\mathcal{R}}$ and also $y_{t+1} = J y_t = y_{\mathcal{N}} + J y_{\mathcal{R}}$, hence $y_{t+1}^T H y_{t+1} = (J y_{\mathcal{R}})^T H (J y_{\mathcal{R}})$. Therefore from Inequality (14) proved above, if the starting vector is $y_{\mathcal{R}}$, which Claim B applies too, then $(J y_{\mathcal{R}})^T H J y_{\mathcal{R}} \leq (1 + \epsilon) y_{\mathcal{R}}^T H y_{\mathcal{R}} = (1 + \epsilon) y_t^T H y_t$.

To sum up, we showed that $y_t^T H y_t \leq (1 + \epsilon)^t y_0^T H y_0$ and since y_0 is an eigenvector of H (of norm one) with corresponding negative eigenvalue λ , it follows that $y_t^T H y_t \leq \lambda(1 + \epsilon)^t$. Finally using $y_t^T H y_t \geq \lambda_{\min}(H) \|y_t\|_2^2$, we get $\|y_t\|_2 \geq (1 + \epsilon)^{t/2} \frac{\lambda}{\lambda_{\min}(H)}$. Observe that $\frac{\lambda}{\lambda_{\min}(H)} > 0$ is a positive constant, $(1 + \epsilon)^{t/2} \geq (1 + \epsilon/4)^t$ (since $\epsilon \leq 1/2$) and the proof follows (the parameters as claimed in the beginning will be $c = \frac{\lambda}{\lambda_{\min}(H)}$ and $\eta = \epsilon/4$). \square

Proof of Claim B We assume that $\alpha \sum_{j \in S_i} |e_j^T H z_i| < \delta \|z_i\|_2$ for all $i \in [b]$. For base case $i = 2$, it holds that $\|y_t - z_2\|_2 = \|z_1 - z_2\|_2 = \alpha |\sum_{j \in S_1} e_j^T H z_1| < \alpha \sum_{j \in S_1} |e_j^T H z_1| < \delta \|z_1\|_2 < 2\delta \|y_t\|_2$ and $\|z_2\|_2 < (1 + 2\delta) \|y_t\|_2$. Suppose for $i \geq 2$ that $\|y_t - z_i\|_2 < 2(i - 1)\delta \|y_t\|_2$ and thus $\|z_i\|_2 < [1 + 2(i - 1)\delta] \|y_t\|_2$. Using induction and triangle inequality we obtain

$$\begin{aligned}
 \|y_t - z_{i+1}\|_2 &\leq \|y_t - z_i\|_2 + \|z_i - z_{i+1}\|_2 \\
 &= 2(i - 1)\delta \|y_t\|_2 + \alpha \left| \sum_{j \in S_i} e_j^T H z_i \right| \\
 &\leq 2(i - 1)\delta \|y_t\|_2 + \alpha \sum_{j \in S_i} |e_j^T H z_i| \\
 &< 2(i - 1)\delta \|y_t\|_2 + \delta \|z_i\|_2 \\
 &< 2(i - 1)\delta \|y_t\|_2 + \delta [1 + 2(i - 1)\delta] \|y_t\|_2 \\
 &\leq 2i\delta \|y_t\|_2,
 \end{aligned}$$

where we assume $\delta < \frac{1}{2b}$ so that $2(i - 1)\delta < 1$ for all $i \in [b]$. Using the above,

$$\begin{aligned}
\alpha \sum_{j \in S_i} \left| e_j^T H y_t \right| &< \alpha \sum_{j \in S_i} \left| e_j^T H z_i \right| + \alpha \sum_{j \in S_i} \left| e_j^T H (y_t - z_i) \right| \\
&< \delta \|z_i\|_2 + \alpha \left(\sum_{j \in S_i} \|H e_j\|_2 \right) \|y_t - z_i\|_2 \\
&< \delta [1 + 2(i - 1)\delta] \|y_t\|_2 + \alpha [2(i - 1)\delta] \|y_t\|_2 \left(\sum_{j \in S_i} \|H e_j\|_2 \right).
\end{aligned}$$

Since $\|H e_i\|_2 < \sigma_{\max}(H) \leq L$, we get that $\alpha \sum_{j \in S_i} \|H e_j\|_2 < |S_i| \leq n$ and we conclude

$$\alpha \sum_{j \in S_i} \left| e_j^T H y_t \right| < 2n^2 \delta \|y_t\|_2. \quad (15)$$

Finally, using Inequality (15) it follows that $\alpha \|H y_t\|_2 < 2n^2 \delta \sqrt{n} \|y_t\|_2$. Let $w \in \text{Im}(H)$ be a vector that is orthogonal to $\text{null}(H)$ (since H is symmetric). Then it holds that $\|H w\|_2 \geq \sigma_{\min^+}(H) \|w\|_2$ where $\sigma_{\min^+}(H)$ denotes the smallest positive singular value of H (greater than zero). Assume that $y_t \in \text{Im}(H)$ and we get $\|H y_t\|_2 < \frac{2n^2 \delta \sqrt{n}}{\alpha} \|y_t\|_2$. However, $\|H y_t\|_2 \geq \sigma_{\min^+}(H) \|y_t\|_2$ thus by choosing $\frac{2n^2 \sqrt{n} \delta}{\alpha} < \sigma_{\min^+}(H)$ we reach a contradiction. \square

References


1. Absil, P.A., Mahony, R., Andrews, B.: Convergence of the iterates of descent methods for analytic cost functions. *SIAM J. Optim.* **16**(2), 531–547 (2005)
2. Absil, P.A., Mahony, R., Sepulchre, R.: *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, Princeton (2009)
3. Absil, P.A., Mahony, R., Trunpf, J.: An extrinsic look at the Riemannian Hessian. In: Nielsen, F., Barbaresco, F. (eds.) *Geometric Science of Information*, pp. 361–368. Springer, Berlin (2013)
4. Absil, P.A., Malick, J.: Projection-like retractions on matrix manifolds. *SIAM J. Optim.* **22**(1), 135–158 (2012)
5. Adler, R.J., Taylor, J.E.: *Random Fields and Geometry*. Springer, Berlin (2009)
6. Arora, S., Ge, R., Ma, T., Moitra, A.: Simple, efficient, and neural algorithms for sparse coding. In: *Proceedings of The 28th Conference on Learning Theory*, pp. 113–149 (2015)
7. Attouch, H., Bolte, J., Redont, P., Soubeyran, A.: Proximal alternating minimization and projection methods for nonconvex problems: an approach based on the Kurdyka–Lojasiewicz inequality. *Math. Oper. Res.* **35**(2), 438–457 (2010)
8. Attouch, H., Bolte, J., Svaiter, B.F.: Convergence of descent methods for semi-algebraic and tame problems: proximal algorithms, forward–backward splitting, and regularized Gauss–Seidel methods. *Math. Program.* **137**(1–2), 91–129 (2013)
9. Auffinger, A., Arous, G.B., Černý, J.: Random matrices and complexity of spin glasses. *Commun. Pure Appl. Math.* **66**(2), 165–201 (2013)
10. Beck, A.: *First-Order Methods in Optimization*, vol. 25. SIAM, New Delhi (2017)
11. Belkin, M., Rademacher, L., Voss, J.: Basis learning as an algorithmic primitive. In: *Conference on Learning Theory*, pp. 446–487 (2016)
12. Bhojanapalli, S., Neyshabur, B., Srebro, N.: Global optimality of local search for low rank matrix recovery. In: *Advances in Neural Information Processing Systems*, pp. 3873–3881 (2016)

13. Bolte, J., Daniilidis, A., Ley, O., Mazet, L., et al.: Characterizations of Lojasiewicz inequalities: sub-gradient flows, talweg, convexity. *Trans. Am. Math. Soc.* **362**(6), 3319–3363 (2010)
14. Bolte, J., Sabach, S., Teboulle, M.: Proximal alternating linearized minimization or nonconvex and nonsmooth problems. *Math. Program.* **146**(1–2), 459–494 (2014)
15. Bolte, J., Sabach, S., Teboulle, M., Vaisbourd, Y.: First order methods beyond convexity and lipschitz gradient continuity with applications to quadratic inverse problems. *SIAM J. Optim.* **28**(3), 2131–2151 (2018)
16. Brutzkus, A., Globerson, A.: Globally optimal gradient descent for a convnet with Gaussian inputs. arXiv preprint [arXiv:1702.07966](https://arxiv.org/abs/1702.07966) (2017)
17. Cai, T.T., Li, X., Ma, Z., et al.: Optimal rates of convergence for noisy sparse phase retrieval via thresholded wirtinger flow. *Ann. Stat.* **44**(5), 2221–2251 (2016)
18. Candes, E.J., Li, X., Soltanolkotabi, M.: Phase retrieval via Wirtinger flow: theory and algorithms. *IEEE Trans. Inf. Theory* **61**(4), 1985–2007 (2015)
19. Choromanska, A., Henaff, M., Mathieu, M., Arous, G.B., LeCun, Y.: The loss surfaces of multilayer networks. In: *Artificial Intelligence and Statistics*, pp. 192–204 (2015)
20. Conn, A.R., Gould, N.I., Toint, P.L.: *Trust Region Methods*, vol. 1. SIAM, New Delhi (2000)
21. Dauphin, Y.N., Pascanu, R., Gulcehre, C., Cho, K., Ganguli, S., Bengio, Y.: Identifying and attacking the saddle point problem in high-dimensional non-convex optimization. In: *Advances in Neural Information Processing Systems*, pp. 2933–2941 (2014)
22. Du, S.S., Jin, C., Lee, J.D., Jordan, M.I., Póczos, B., Singh, A.: Gradient descent can take exponential time to escape saddle points. arXiv preprint [arXiv:1705.10412](https://arxiv.org/abs/1705.10412) (2017)
23. Du, S.S., Lee, J.D., Tian, Y.: When is a convolutional filter easy to learn? arXiv preprint [arXiv:1709.06129](https://arxiv.org/abs/1709.06129) (2017)
24. Ge, R., Huang, F., Jin, C., Yuan, Y.: Escaping from saddle points online stochastic gradient for tensor decomposition. In: *Conference on Learning Theory*, pp. 797–842 (2015)
25. Ge, R., Jin, C., Zheng, Y.: No spurious local minima in nonconvex low rank problems: a unified geometric analysis. arXiv preprint [arXiv:1704.00708](https://arxiv.org/abs/1704.00708) (2017)
26. Ge, R., Lee, J.D., Ma, T.: Matrix completion has no spurious local minimum. In: *Advances in Neural Information Processing Systems*, pp. 2973–2981 (2016)
27. Gill, P.E., Murray, W.: Newton-type methods for unconstrained and linearly constrained optimization. *Math. Program.* **7**(1), 311–350 (1974)
28. Jin, C., Ge, R., Netrapalli, P., Kakade, S.M., Jordan, M.I.: How to escape saddle points efficiently. arXiv preprint [arXiv:1703.00887](https://arxiv.org/abs/1703.00887) (2017)
29. Keshavan, R.H., Montanari, A., Oh, S.: Matrix completion from a few entries. *IEEE Trans. Inf. Theory* **56**(6), 2980–2998 (2009)
30. Kleinberg, R., Piliouras, G., Tardos, E.: Multiplicative updates outperform generic no-regret learning in congestion games. In: *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*, pp. 533–542. ACM (2009)
31. Lange, K.: *Optimization*, vol. 95. Springer (2013)
32. Lee, J.D., Simchowitz, M., Jordan, M.I., Recht, B.: Gradient descent only converges to minimizers. In: *Conference on Learning Theory*, pp. 1246–1257 (2016)
33. Lewis, A.S., Malick, J.: Alternating projections on manifolds. *Math. Oper. Res.* **33**(1), 216–234 (2008)
34. Liu, M., Yang, T.: On noisy negative curvature descent: competing with gradient descent for faster non-convex optimization. arXiv preprint [arXiv:1709.08571](https://arxiv.org/abs/1709.08571) (2017)
35. Losert, V., Akin, E.: Dynamics of games and genes: discrete versus continuous time. *J. Math. Biol.* **17**, 241–251 (1983)
36. Mikusinski, P., Taylor, M.: *An Introduction to Multivariable Analysis From Vector to Manifold*. Springer, Berlin (2012)
37. Moré, J.J., Sorensen, D.C.: On the use of directions of negative curvature in a modified Newton method. *Math. Program.* **16**(1), 1–20 (1979)
38. Murty, K.G., Kabadi, S.N.: Some NP-complete problems in quadratic and nonlinear programming. *Math. Program.* **39**(2), 117–129 (1987)
39. Nesterov, Y.: *Introductory Lectures on Convex Optimization*, vol. 87. Springer, Berlin (2004)
40. Nesterov, Y., Polyak, B.T.: Cubic regularization of Newton method and its global performance. *Math. Program.* **108**(1), 177–205 (2006)
41. O'Neill, M., Wright, S.: Behavior of accelerated gradient methods near critical points of nonconvex problems. arXiv preprint [arXiv:1706.07993](https://arxiv.org/abs/1706.07993) (2017)

42. Panageas, I., Piliouras, G.: Gradient descent only converges to minimizers: non-isolated critical points and invariant regions. In: *Innovations of Theoretical Computer Science (ITCS)* (2017)
43. Pascanu, R., Dauphin, Y.N., Ganguli, S., Bengio, Y.: On the saddle point problem for non-convex optimization. [arXiv:1405.4604](https://arxiv.org/abs/1405.4604) (2014)
44. Pemantle, R.: Nonconvergence to unstable points in urn models and stochastic approximations. *Ann. Probab.* **18**, 698–712 (1990)
45. Reddi, S.J., Zaheer, M., Sra, S., Póczos, B., Bach, F., Salakhutdinov, R., Smola, A.J.: A generic approach for escaping saddle points. [arXiv:1709.01434](https://arxiv.org/abs/1709.01434) (2017)
46. Richtárik, P., Takáč, M.: Iteration complexity of randomized block-coordinate descent methods for minimizing a composite function. [arXiv:1107.2848](https://arxiv.org/abs/1107.2848) (2011)
47. Royer, C.W., Wright, S.J.: Complexity analysis of second-order line-search algorithms for smooth nonconvex optimization. [arXiv:1706.03131](https://arxiv.org/abs/1706.03131) (2017)
48. Shub, M.: *Global Stability of Dynamical Systems*. Springer, Berlin (1987)
49. Soltanolkotabi, M., Javanmard, A., Lee, J.D.: Theoretical insights into the optimization landscape of over-parameterized shallow neural networks. [arXiv:1707.04926](https://arxiv.org/abs/1707.04926) (2017)
50. Sun, J., Qu, Q., Wright, J.: When are nonconvex problems not scary? [arXiv:1510.06096](https://arxiv.org/abs/1510.06096) (2015)
51. Sun, J., Qu, Q., Wright, J.: A geometric analysis of phase retrieval. In: *2016 IEEE International Symposium on Information Theory (ISIT)*, pp. 2379–2383. IEEE (2016)
52. Sun, J., Qu, Q., Wright, J.: Complete dictionary recovery over the sphere I: overview and the geometric picture. *IEEE Trans. Inf. Theory* **63**(2), 853–884 (2017)
53. Sun, J., Qu, Q., Wright, J.: Complete dictionary recovery over the sphere II: recovery by Riemannian trust-region method. *IEEE Tran. Inf. Theory* **63**(2), 885–914 (2017)
54. Zhang, Y., Chen, X., Zhou, D., Jordan, M.I.: Spectral methods meet EM: a provably optimal algorithm for crowdsourcing. In: *Advances in Neural Information Processing Systems*, pp. 1260–1268 (2014)
55. Zhao, T., Wang, Z., Liu, H.: Nonconvex low rank matrix factorization via inexact first order oracle. *Adv. Neural Inf. Process. Syst.* 559–567 (2015)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Affiliations

Jason D. Lee¹  · Ioannis Panageas² · Georgios Piliouras³ · Max Simchowitz⁴ · Michael I. Jordan⁴ · Benjamin Recht⁴

Ioannis Panageas
ioannis@sutd.edu.sg

Georgios Piliouras
georgios@sutd.edu.sg

Max Simchowitz
msimchow@berkeley.edu

Michael I. Jordan
jordan@cs.berkeley.edu

Benjamin Recht
brecht@berkeley.edu

¹ Data Sciences and Operations, University of Southern California, Los Angeles, CA, USA

² Department of Information Systems, Singapore University of Technology, Tampines, Singapore

-
- ³ Engineering Systems and Design Pillar, Singapore University of Technology and Design, Tampines, Singapore
 - ⁴ Department of Electrical Engineering and Computer Science, UC Berkeley, Berkeley, USA