



Fast convergence to non-isolated minima: four equivalent conditions for C^2 functions

Quentin Rebjock¹ · Nicolas Boumal¹

Received: 26 December 2023 / Accepted: 5 August 2024
© The Author(s) 2024

Abstract

Optimization algorithms can see their local convergence rates deteriorate when the Hessian at the optimum is singular. These singularities are inescapable when the optima are non-isolated. Yet, under the right circumstances, several algorithms preserve their favorable rates even when optima form a continuum (e.g., due to over-parameterization). This has been explained under various structural assumptions, including the Polyak–Łojasiewicz condition, Quadratic Growth and the Error Bound. We show that, for cost functions which are twice continuously differentiable (C^2), those three (local) properties are equivalent. Moreover, we show they are equivalent to the Morse–Bott property, that is, local minima form differentiable submanifolds, and the Hessian of the cost function is positive definite along its normal directions. We leverage this insight to improve local convergence guarantees for safe-guarded Newton-type methods under any (hence all) of the above assumptions. First, for adaptive cubic regularization, we secure quadratic convergence even with inexact subproblem solvers. Second, for trust-region methods, we argue capture can fail with an exact subproblem solver, then proceed to show linear convergence with an inexact one (Cauchy steps).

Mathematics Subject Classification 90C26 · 90C30

1 Introduction

We consider local convergence of algorithms for unconstrained optimization problems

$$\min_{x \in \mathcal{M}} f(x),$$

This work was supported by the Swiss State Secretariat for Education, Research and Innovation (SERI) under contract number MB22.00027.

✉ Quentin Rebjock
quentin.rebjock@epfl.ch

¹ Institute of Mathematics, Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland

where \mathcal{M} is a Riemannian manifold¹ and $f: \mathcal{M} \rightarrow \mathbb{R}$ is at least C^1 (continuously differentiable).

When f is C^2 (twice continuously differentiable), the most classical local convergence results ensure favorable rates for standard algorithms *provided* they converge to a non-singular local minimum \bar{x} , that is, one such that the Hessian $\nabla^2 f(\bar{x})$ is positive definite. And indeed, those rates can degrade if the Hessian is merely positive semidefinite. For example, with $f(x) = x^4$, gradient descent (with an appropriate step-size) converges only sublinearly to the minimum, and Newton's method converges only linearly.

This is problematic if the minimizers of f are not isolated, because in that case the Hessian cannot be positive definite there. This situation arises commonly in applications for structural reasons such as over-parameterization, redundant parameterizations and symmetry—see Sect. 1.2.

Notwithstanding, algorithms often exhibit good local behavior near non-isolated minimizers. As early as the 1960s, this has prompted investigations into properties that such cost functions may satisfy and which lead to fast local rates despite singular Hessians. We study four such properties.

In all that follows, we are concerned with the behavior of algorithms in the vicinity of its local minima. Since we do not assume that they are isolated, rather than selecting one local minimum \bar{x} , we select all local minima of the same value. Formally, given a local minimum \bar{x} , let

$$\mathcal{S} = \{x \in \mathcal{M} : x \text{ is a local minimum of } f \text{ and } f(x) = f_{\bar{x}}\} \quad (1)$$

denote the set of all local minima with a given value $f_{\mathcal{S}} = f(\bar{x})$.

For f of class C^2 , it is particularly favorable if \mathcal{S} is a differentiable submanifold of \mathcal{M} around \bar{x} . In that case the set \mathcal{S} has a *tangent space* $T_{\bar{x}}\mathcal{S}$ at \bar{x} . It is easy to see that each vector $v \in T_{\bar{x}}\mathcal{S}$ must be in the kernel of the Hessian $\nabla^2 f$ at \bar{x} because the gradient ∇f is constant (zero) on \mathcal{S} . Thus, $T_{\bar{x}}\mathcal{S} \subseteq \ker \nabla^2 f(\bar{x})$. Since \bar{x} is a local minimum, we also know that $\nabla^2 f(\bar{x})$ is positive semidefinite. Then, in the spirit of asking the Hessian to be “as positive definite as possible”, the best we can hope for is that the kernel of $\nabla^2 f(\bar{x})$ is exactly $T_{\bar{x}}\mathcal{S}$, in which case the restriction of $\nabla^2 f(\bar{x})$ to the *normal space* $N_{\bar{x}}\mathcal{S}$, that is, the orthogonal complement of $T_{\bar{x}}\mathcal{S}$ in $T_{\bar{x}}\mathcal{M}$, is positive definite.

We call this the *Morse–Bott* property (MB), and we write μ -MB to indicate that the positive eigenvalues are at least $\mu > 0$. The definition requires f to be twice differentiable.

Definition 1.1 Let \bar{x} be a local minimum of f with associated set \mathcal{S} (1). We say f satisfies the *Morse–Bott* property at \bar{x} if

$$\mathcal{S} \text{ is a } C^1 \text{ submanifold around } \bar{x} \quad \text{and} \quad \ker \nabla^2 f(\bar{x}) = T_{\bar{x}}\mathcal{S}. \quad (\text{MB})$$

¹ The contributions are relevant for $\mathcal{M} = \mathbb{R}^n$ too. We treat the more general manifold case as it involves only mild overhead in notation, summarized in Table 2.

If also $\langle v, \nabla^2 f(\bar{x})[v] \rangle \geq \mu \|v\|^2$ for some $\mu > 0$ and all $v \in N_{\bar{x}}\mathcal{S}$ then we say f satisfies μ -MB at \bar{x} .

At first, a reasonable objection to the above is that one may not want to assume that \mathcal{S} is a submanifold. Perhaps for that reason, it is far more common to encounter other assumptions in the optimization literature. We focus on three: *Polyak–Lojasiewicz* (PŁ), *error bound* (EB) and *quadratic growth* (QG). The first goes back to the 1960s [81]. The latter two go back at least to the 1990s [22, 70].

Below, the first two definitions (as stated) require f to be differentiable. The distance to a set is defined as usual: $\text{dist}(x, \mathcal{S}) = \inf_{y \in \mathcal{S}} \text{dist}(x, y)$ where $\text{dist}(x, y)$ is the Riemannian distance on \mathcal{M} .

Definition 1.2 Let \bar{x} be a local minimum of f with associated set \mathcal{S} (1). We say f satisfies

- the *Polyak–Lojasiewicz* condition with constant $\mu > 0$ (μ -PŁ) around \bar{x} if

$$f(x) - f_{\mathcal{S}} \leq \frac{1}{2\mu} \|\nabla f(x)\|^2; \quad (\text{PŁ})$$

- the *error bound* with constant $\mu > 0$ (μ -EB) around \bar{x} if

$$\mu \text{dist}(x, \mathcal{S}) \leq \|\nabla f(x)\|; \quad (\text{EB})$$

- *quadratic growth* with constant $\mu > 0$ (μ -QG) around \bar{x} if

$$f(x) - f_{\mathcal{S}} \geq \frac{\mu}{2} \text{dist}(x, \mathcal{S})^2; \quad (\text{QG})$$

all understood to hold for all x in some neighborhood of \bar{x} .

Note that all the definitions are *local* around a *point* \bar{x} . Two observations are immediate: (i) QG implies that \bar{x} is a strict minimum relatively to \mathcal{S} (1), meaning that $f(x) > f_{\mathcal{S}}$ for all $x \notin \mathcal{S}$ close enough to \bar{x} , and (ii) both EB and PŁ imply that critical points and \mathcal{S} coincide around \bar{x} . Thus, both EB and PŁ rule out existence of saddle points near \bar{x} . By extension, we say that f satisfies any of these four properties around a *set* of local minima if it holds around each point of that set.

1.1 Contributions

A number of relationships between PŁ, EB and QG are well known already for f of class C^1 : see Table 1 and Sect. 1.3. Our first main contribution in this paper is to show that:

If f is of class C^2 , then PŁ, EB, QG and MB are essentially equivalent.

Here, “essentially” means that the constant μ may degrade (arbitrarily little) and the neighborhoods where properties hold may shrink. Notably, we show that if f is C^p with $p \geq 2$, then PŁ, EB and QG all imply that the set of local minima \mathcal{S} is locally smooth

(at least C^{p-1}). We also give counter-examples when f is only C^1 . Explicitly, in Sect. 2 we summarize known results for f of class C^1 and we contribute the following:

- Theorem 2.16 shows PŁ around \bar{x} implies \mathcal{S} is a C^{p-1} submanifold around \bar{x} if f is C^p with $p \geq 2$. Remark 2.19 provides counter-examples if f is only C^1 . If f is analytic, so is \mathcal{S} , as already shown by Feehan [44].
- Lemma 2.14 is instrumental to prove Theorem 2.16 and to analyze algorithms. It states that under PŁ the gradient locally aligns with the dominant eigenvectors of the Hessian.
- Corollary 2.17 deduces that μ -PŁ implies μ -MB if f is C^2 .
- Proposition 2.8 shows μ -QG implies μ' -EB with $\mu' < \mu$ arbitrarily close if f is C^2 . Remark 2.9 provides a counter-example if f is only C^1 .
- Proposition 2.4 shows μ -EB implies μ' -PŁ with $\mu' < \mu$ arbitrarily close if f is C^2 . If f is only C^1 with L -Lipschitz continuous ∇f , Karimi et al. [52] showed the same with $\mu' = \mu^2/L$.

In Sect. 4, we study the classical globalized versions of Newton's method. We strengthen their local convergence guarantees when minimizers are not isolated but the C^2 cost function satisfies any (hence all) of the above. The key observation that enables those improvements is the fact that we may use all four conditions in the analysis without loss of generality. Specifically:

- Cubic regularization enjoys superlinear convergence under PŁ, as shown by Nesterov and Polyak [76]. Yue et al. [97] further showed quadratic convergence under EB. Both references assume an exact subproblem solver. Leveraging our results above, we show that quadratic convergence still holds with inexact subproblem solvers (Theorem 4.1).
- For the trust-region method with exact subproblem solver, we were surprised to find that even basic capture-type convergence properties can fail in the presence of non-isolated local minima (Sect. 4.2). Notwithstanding, common implementations of the trust-region method use a truncated conjugate gradient (tCG) subproblem solver, and those do, empirically, exhibit superlinear convergence under the favorable conditions discussed above. We discuss this further in Remark 4.14, and we show a partial result in Theorem 4.9, namely, that using the Cauchy step (i.e., the first iterate of tCG) yields linear convergence.

As is classical, to prove the latter results, we rely on capture theorems and Lyapunov stability. Those hold under assumptions of vanishing step-sizes and bounded path length. In Sect. 3, we state those building blocks succinctly, adapted to accommodate non-isolated local minima.

1.2 Non-isolated minima in applications

We now illustrate how optimization problems with continuous sets of minima occur in applications.

In all three scenarios below, we can cast the cost function $f: \mathcal{M} \rightarrow \mathbb{R}$ as a composition of some other function $g: \mathcal{N} \rightarrow \mathbb{R}$ through a map $\varphi: \mathcal{M} \rightarrow \mathcal{N}$, where \mathcal{N} is a smooth

Table 1 Summary of implications

Statement	f is C^p	Constants	Global?	Comments
$PL \Rightarrow QG$	$p \geq 1$	$\mu' = \mu$	Yes	Otto and Villani [79], Ioffe [50]
$MB \Rightarrow QG$	$p \geq 2$	$\mu' < \mu$	n/a	Taylor expansion
$EB \Rightarrow PL$	$p \geq 2$	$\mu' < \mu$	No	
$Rmrk 2.5$	$p \geq 1$	$\mu' = \frac{\mu^2}{L}$	Yes	History in Sect. 1.3
$QG \Rightarrow EB$	$p \geq 2$	$\mu' < \mu$	No	Fails for $p = 1$, Rmrk 2.9
$PL \Rightarrow EB$	$p \geq 1$	$\mu' = \mu$	Yes	History in Sect. 1.3
$PL \Rightarrow MB$	$p \geq 2$	$\mu' = \mu$	n/a	Fails for $p = 1$, Rmrk 2.19
$EB \Rightarrow QG$	$p \geq 1$	$\mu' = \mu$	Yes	Corvellec and Motreanu [35]

For a row $A \Rightarrow B$, we mean that A with constant μ implies B (in a possibly different neighborhood) with constant μ' . The condition $\mu' < \mu$ means μ' can be taken arbitrarily close to μ by shrinking the neighborhood. For $EB \Rightarrow PL$ with $p = 1$, we let L denote the Lipschitz constant of ∇f . If B holds globally when A holds globally, we write “Yes” in the column “Global?”. Implications we found in the literature are marked with a reference




Fig. 1 Optimization through the map φ




Fig. 2 (Left) Submanifold of minima where MB holds. (Middle) A C^1 function that satisfies QG but not PL nor EB. (Right) A C^1 function that satisfies PL, yet whose set of minima is a cross

manifold (see Fig. 1 and [60]). If g and φ are twice differentiable, then the Morse–Bott property (MB) for $f = g \circ \varphi$ can come about as follows. Consider a local minimum \bar{y} for g . The set $\mathfrak{X} = \varphi^{-1}(\bar{y})$ consists of local minima for f . Pick a point $\bar{x} \in \mathfrak{X}$. Assume $x \mapsto \text{rank } D\varphi(x)$ is constant in a neighborhood of \bar{x} . Then, the set \mathfrak{X} is an embedded submanifold of \mathcal{M} around \bar{x} with tangent space $\ker D\varphi(\bar{x})$. Moreover, the Hessians of f and g at \bar{x} are related by

$$\nabla^2 f(\bar{x}) = D\varphi(\bar{x})^* \circ \nabla^2 g(\varphi(\bar{x})) \circ D\varphi(\bar{x}). \tag{2}$$

Therefore, if $\nabla^2 g(\varphi(\bar{x}))$ is positive definite, then $\ker \nabla^2 f(\bar{x}) = T_{\bar{x}}\mathfrak{X}$ and $\nabla^2 f(\bar{x})$ is positive definite along the orthogonal complement. In other words: f satisfies the Morse–Bott property (MB) at \bar{x} . We present below a few concrete examples of optimization problems where this can happen.

Over-parameterization and nonlinear regression. Consider minimizing $f(x) = \frac{1}{2} \|F(x) - b\|^2$ with $F: \mathbb{R}^m \rightarrow \mathbb{R}^n$ a C^2 function. We cast this as above with $g(y) = \frac{1}{2} \|y - b\|^2$ and $\varphi = F$. Suppose $\mathfrak{X} = \varphi^{-1}(b) = \{x : F(x) = b\}$ is non-empty (interpolation regime), which is typical in deep learning. This is the set of global minimizers of f . If $\text{rank } DF(x)$ is equal to a constant r in a neighborhood of \mathfrak{X} , then \mathfrak{X} is a smooth submanifold of \mathbb{R}^m of dimension $m - r$. If additionally the problem is over-parameterized, that is, $m > n \geq r$, then \mathfrak{X} has positive dimension (see Fig. 2 for an illustration). The discussion above immediately implies that f satisfies MB on \mathfrak{X} . See also Nesterov and Polyak [76, §4.2] who argue that PL holds in this setting.

Redundant parameterizations and submersions. Say we want to minimize $g: \mathcal{N} \rightarrow \mathbb{R}$ constrained to $\mathcal{C} \subseteq \mathcal{N}$. If \mathcal{C} is complicated, and if we have access to a parameteriza-

tion φ for that set (so that $\varphi(\mathcal{M}) = \mathcal{C}$), it may be advantageous to minimize $f = g \circ \varphi$ instead. If the parameterization is redundant, this can cause f to have non-isolated minima, even if the minima of g are isolated.

As an example, consider minimizing $g: \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ over the bounded-rank matrices $\mathcal{C} = \{Y \in \mathbb{R}^{m \times n} : \text{rank } Y \leq r\}$. A popular approach consists in lifting the search space to $\mathcal{M} = \mathbb{R}^{m \times r} \times \mathbb{R}^{n \times r}$ and minimizing $f = g \circ \varphi$, where $\varphi: \mathcal{M} \rightarrow \mathcal{N}$ is defined as $\varphi(L, R) = LR^\top$. The parameterization is redundant because $\varphi(LJ^{-1}, RJ^\top) = \varphi(L, R)$ for all invertible J . In particular, given a local minimum $Y \in \mathcal{C}$ of g , the fiber $\varphi^{-1}(Y)$ is unbounded, which hinders convergence analyses (see [59]). However, if Y is of *maximal rank* r then $D\varphi$ has constant rank in a neighborhood of $\varphi^{-1}(Y)$. From the discussion above, it follows that f satisfies MB on $\varphi^{-1}(Y)$ if the (Riemannian) Hessian of g (restricted to the manifold of rank- r matrices) is positive definite.

Similarly, Burer and Monteiro [25, 26] introduced a popular approach to minimize a function g over the set of positive semidefinite matrices of bounded rank through the map $\varphi: Y \mapsto YY^\top$. The resulting function $f = g \circ \varphi$ can have non-isolated minima. However, the same arguments as above ensure that MB holds at minimizers of maximal rank when g is strongly convex (this setting is for example considered in [103]). This further extends to tensors [63].

Symmetries and quotients. Some optimization problems have intrinsic symmetries. For example, in estimation problems, if the measurements are invariant under particular transformations of the signal, then the signal can only be retrieved up to those transformations. The likelihood function then has symmetries, and possibly a continuous set of optima as a result. Sometimes, factoring these symmetries out (that is, passing to the quotient) yields a quotient manifold, and we can investigate optimization on that manifold [4]. In the notation of our general framework above, φ is then the quotient map. In particular, φ is a submersion, so that if $\bar{y} \in \mathcal{N}$ is a non-singular minimum of g then f satisfies MB on $\varphi^{-1}(\bar{y})$ (which is a submanifold of dimension $\dim \mathcal{M} - \dim \mathcal{N}$). See also [24, §9.9] for the case where \mathcal{N} is a *Riemannian* quotient of \mathcal{M} .

1.3 Related work

Historical note. Discussions about convergence to singular minima appear in the literature at least as early as [82, §6.1]. Luo and Tseng [70] introduced the EB condition explicitly to study gradient methods around singular minima. The QG property is arguably as old as optimization, though the earliest work we could locate is by Bonnans and Ioffe [22]. They employed QG to understand complicated landscapes with non-isolated minima. Łojasiewicz [68, 69] introduced his inequalities and used them subsequently to analyze gradient flow trajectories. Specifically, he proved that for analytic functions the trajectories either converge to a point or diverge. Concurrently, Polyak [81] introduced what became known as the Polyak–Łojasiewicz (PŁ) variant (also “gradient dominance”) to study both gradient flows and discrete gradient methods. Later, Kurdyka [55] developed generalizations now known as Kurdyka–Łojasiewicz (KŁ) inequalities. They are satisfied by most functions encountered in

practice, as discussed in [9, §4]. In contrast, the Morse–Bott property has received little attention in the optimization literature. Early work by Shapiro [87] analyzes perturbations of optimization problems assuming a property similar to MB. There is also a mention of gradient flow under MB in [49, Prop. 12.3].

Relationships between properties. Several articles have explored the interplay between MB, PŁ, EB and QG in the last decades. The implication $\text{PŁ} \Rightarrow \text{QG}$ has a rich history. It can be obtained as a corollary from [50, Basic lemma] (based on Ekeland’s variational principle; see also [41, Lem. 2.5]). It also follows from Łojasiewicz-type arguments that consist in bounding the length of gradient flow trajectories [79, Prop. 1]. Likewise, Bolte et al. [19] study growth under KŁ inequalities, with $\text{PŁ} \Rightarrow \text{QG}$ as a special case. For lower-semicontinuous functions, Corvellec and Motreanu [35, Thm. 4.2] show that $\mu\text{-EB} \Rightarrow \mu\text{-QG}$. They also find that $\mu\text{-QG} \Rightarrow \frac{\mu}{2}\text{-EB}$ for convex functions [35, Prop. 5.3]. In a somewhat different setting, Drusvyatskiy et al. [40, Cor. 3.2] prove that $\mu\text{-EB} \Rightarrow \mu'\text{-QG}$ with arbitrary $\mu' < \mu$. With extra assumptions, they also show $\mu\text{-QG} \Rightarrow \mu'\text{-EB}$ but without control of μ' . Later, Bolte et al. [21] proved an equivalence between KŁ inequalities and function growth for convex and potentially non-smooth functions. Their results seem to generalize to semi-convex functions. See also [100] and [39] for equivalences between EB and QG. Karimi et al. [52] established implications between several properties encountered in the optimization literature. In particular, they also show that PŁ and EB are equivalent, and that they both imply QG. Liao et al. [65] extended this to a non-smooth and weakly convex setting. Li and Pong [62] showed that EB implies PŁ for non-smooth functions under a level set separation assumption, though with no control on the PŁ constant. Implications between PŁ and EB are also reported in [42, Thm. 3.7, Prop. 3.8] for non-smooth functions under broad conditions. In the context of functional analysis, Feehan [44] proved that MB and PŁ are equivalent for *analytic* functions defined on Banach spaces. The work of [94, Ex. 2.9] also mentions that MB implies PŁ for C^2 functions. A more general implication is given by Arbel and Mairal [8, Prop. 1] for parameterized optimization problems. Previously, Bonnans and Ioffe [22] had exhibited sufficient conditions (similar to MB) for QG to hold. As a side note, Marteau-Ferey et al. [71] proved that MB is a sufficient condition to ensure that a non-negative function is globally decomposable as a sum of squares of smooth functions.

Convergence guarantees. The error bound approach of Luo and Tseng [70] has proven to be fruitful as multiple analyses based on this condition followed. Notably, Tseng [91] proved local superlinear convergence rates for some Newton-type methods applied to systems of nonlinear equations. They relied specifically on EB and did not assume isolated minima. Later, Yamashita and Fukushima [96] employed EB to establish capture theorems and quadratic convergence rates for the Levenberg–Marquardt method. Fan and Yuan [43], Behling et al. [12] and Boos et al. [23] generalized their results (in particular to solutions with a non-zero residual). More recently, Bellavia and Morini [14] found that two adaptive regularized methods converge quadratically for nonlinear least-squares problems (assuming that EB holds).

An early work of Anitescu [6] combines the QG property with other conditions to ensure isolated minima in constrained optimization, then deducing convergence results. Later, QG has found applications mainly in the context of convex optimization: see [67] for coordinate descent, [75] for various gradient methods, and [39] for the

proximal gradient method. It is also worth mentioning that the definition of QG does not require differentiability of the function. For this reason, QG is valuable to study algorithms in non-smooth optimization too [36, 61].

The literature about convergence results based on Łojasiewicz inequalities is vast, and we touch here on some particularly relevant references. Absil et al. [2] discretized the arguments from [69] and obtained capture results for a broad class of optimization algorithms. Lageman [56, 57] provided generalizations to broader classes of functions. Later, such arguments have been used in many contexts to prove algorithmic convergence guarantees, among which [9, 20] are particularly influential works. Moreover, Attouch et al. [10] proposed a general abstract framework based on KŁ to derive capture results and convergence rates, and Frankel et al. [46] extended their statements. See also [74] for a framework that encompasses higher-order methods. Li and Pong [62] studied the preservation of Łojasiewicz inequalities under function transformations (such as sums and compositions). See also [11, 90] for the preservation of PŁ through function compositions. Interestingly, the PŁ condition is known to be a necessary condition for gradient descent to converge linearly [1, Thm. 5]. Recently, Yue et al. [98] proved that acceleration is impossible to minimize globally PŁ functions: gradient descent is optimal absent further structure. Assuming PŁ, Stonyakin et al. [88] formulated stopping criteria for gradient methods when the gradient is corrupted with noise. Using KŁ inequalities, Noll and Rondepierre [78] and Khanh et al. [53] analyzed the convergence of line-search gradient descent and trust-region methods. Łojasiewicz inequalities have also proved relevant for the study of second-order algorithms and superlinear convergence rates. A prominent example is the regularized Newton algorithm that converges superlinearly when PŁ holds, as shown in [76]. More recently, Zhou et al. [105], Yue et al. [97] provided finer analyses of this algorithm, respectively assuming Łojasiewicz inequalities and EB. Qian and Pan [84] extended the abstract framework of Attouch et al. [10] to establish superlinear convergence rates.

Stochastic algorithms have also been extensively studied through Łojasiewicz inequalities, and we briefly mention a few references here. Dereich and Kassing [37, 38] analyzed stochastic gradient descent (SGD) in the presence of non-isolated minima, using Łojasiewicz inequalities among other things. Local analyses of SGD using PŁ inequalities are given by Li et al. [64] and Wojtowytsch [94]. Ko and Li [54] studied the local stability and convergence of SGD in the presence of a compact set of minima with a condition that is weaker than PŁ. As for second-order algorithms, Masiha et al. [72] proved that a stochastic version of regularized Newton has fast convergence under PŁ.

Łojasiewicz inequalities are also particularly suited to analyze the convergence of flows. Notably, Łojasiewicz [69] bounded the path length of gradient flow trajectories and Polyak [81] derived linear convergence of flows assuming PŁ. Related results for flows but under MB are claimed in [49, Prop. 12.3]. More recently, Apidopoulos et al. [7] considered the Heavy-Ball differential equation and deduced convergence guarantees from PŁ. Wojtowytsch [95] studied a continuous model for SGD and the impact of the noise on the trajectory.

Finally, we found only few convergence results based on MB in the optimization literature. Fehrman et al. [45] derive capture theorems and asymptotic sublinear convergence rates for gradient descent assuming that MB holds on the set of minima. They

Table 2 Simplifications in the case where \mathcal{M} is a Euclidean space

	$\text{Exp}_x(s)$	$\text{Log}_x(y)$	Γ_x^y, Γ_v	$\text{dist}(x, y)$	$\text{inj}(x)$
$\mathcal{M} = \mathbb{R}^n$	$x + s$	$y - x$	identity	$\ x - y\ $	$+\infty$

Here, Γ_x^y denotes parallel transport along the minimizing geodesic between x and y (assuming the points are close enough)

also provide probabilistic bounds for stochastic variants. Usevich et al. [92] consider optimization problems over unitary matrices. They propose sufficient conditions for MB to hold at a local optimum and then exploit the induced PŁ condition to obtain convergence rates. In order to solve systems of nonlinear equations, Zeng [99] proposes a Newton-type method that is robust to non-isolated solutions. It enjoys local quadratic convergence assuming a MB-type property. The algorithm requires the knowledge of the dimension of the set of solutions.

Applications. Non-isolated minima arise in all sorts of optimization problems. It is common for non-convex inverse problems to have continuous symmetries, hence non-isolated minima (see [104]). In the context of deep learning, Cooper [34] proved that the set of global minima of a sufficiently over-parameterized neural network is a smooth manifold. In the last decade, there has been a renewed interest in Łojasiewicz inequalities because they are compatible with these complicated non-convex landscapes. In particular, a whole line of research exploits them to understand deep learning problems specifically. As an example, Oymak and Soltanolkotabi [80] employed PŁ to analyze the path taken by (stochastic) gradient descent in the vicinity of minimizers. Several other works suggested that non-convex machine learning loss landscapes can be understood in over-parameterized regimes through the lens of Łojasiewicz inequalities [11, 13, 66, 90]. Specifically, they argue that PŁ holds on a significant part of the search space and analyze (stochastic) gradient methods. Chatterjee [30] also establishes local convergence results for a large class of neural networks with PŁ inequalities.

1.4 Notation and geometric preliminaries

This section anchors notation and some basic geometric facts. In the important case where $\mathcal{M} = \mathbb{R}^n$, several objects reduce as summarized in Table 2.

We let $\langle \cdot, \cdot \rangle$ denote the inner product on $T_x\mathcal{M}$ —it may depend on $x \in \mathcal{M}$, but the base point is always clear from context. The associated norm is $\|v\| = \sqrt{\langle v, v \rangle}$. The map $\text{dist} : \mathcal{M} \times \mathcal{M} \rightarrow \mathbb{R}_+$ is the Riemannian distance on \mathcal{M} . We let $B(x, \delta)$ denote the open ball of radius δ around $x \in \mathcal{M}$. The tangent bundle is $T\mathcal{M} = \{(x, v) : x \in \mathcal{M} \text{ and } v \in T_x\mathcal{M}\}$.

Moving away from $x \in \mathcal{M}$ along the geodesic with (sufficiently small) initial velocity $v \in T_x\mathcal{M}$ for unit time produces the point $\text{Exp}_x(v) \in \mathcal{M}$ (Riemannian exponential). The injectivity radius at x is $\text{inj}(x) > 0$. It is defined such that, given $y \in B(x, \text{inj}(x))$, there exists a unique smallest vector $v \in T_x\mathcal{M}$ for which $\text{Exp}_x(v) = y$. We denote this v by $\text{Log}_x(y)$ (Riemannian logarithm). Additionally, given $x \in \mathcal{M}$

and $y \in B(x, \text{inj}(x))$, we let $\Gamma_x^y: T_x\mathcal{M} \rightarrow T_y\mathcal{M}$ denote parallel transport along the unique minimizing geodesic between x and y . If $v = \text{Log}_x(y)$, we also let $\Gamma_v = \Gamma_x^y$.

Let \mathfrak{X} be a subset of \mathcal{M} . We need the notions of tangent and normal cones to \mathfrak{X} , defined below.

Definition 1.3 The *tangent cone* to a set \mathfrak{X} at $x \in \mathfrak{X}$ is the closed set

$$T_x\mathfrak{X} = \left\{ \lim_{k \rightarrow +\infty} \frac{1}{t_k} \text{Log}_x(x_k) \mid x_k \in \mathfrak{X}, t_k > 0 \text{ for all } k, x_k \rightarrow x, t_k \rightarrow 0 \right\}.$$

We also let $N_x\mathfrak{X} = \{w \in T_x\mathcal{M} : \langle w, v \rangle \leq 0 \text{ for all } v \in T_x\mathfrak{X}\}$ denote the normal cone to \mathfrak{X} at x .

When \mathfrak{X} is a submanifold of \mathcal{M} around x , the cones $T_x\mathfrak{X}$ and $N_x\mathfrak{X}$ reduce to the tangent and normal spaces of \mathfrak{X} at x . (By ‘‘submanifold’’, we always mean *embedded* submanifold.)

Given $x \in \mathcal{M}$, we let $\text{dist}(x, \mathfrak{X}) = \inf_{y \in \mathfrak{X}} \text{dist}(x, y)$ denote the distance of x to \mathfrak{X} . We further let $\text{proj}_{\mathfrak{X}}(x)$ denote the set of minima of the optimization problem $\min_{y \in \mathfrak{X}} \text{dist}(x, y)$. If this set is non-empty (which is the case in particular if \mathfrak{X} is closed), then we have:

$$\forall \bar{x} \in \mathfrak{X}, y \in \text{proj}_{\mathfrak{X}}(x), \quad \text{dist}(y, \bar{x}) \leq 2 \text{dist}(x, \bar{x}). \tag{3}$$

Indeed, the triangle inequality yields $\text{dist}(y, \bar{x}) \leq \text{dist}(x, y) + \text{dist}(x, \bar{x})$, and $\text{dist}(x, y) = \text{dist}(x, \mathfrak{X}) \leq \text{dist}(x, \bar{x})$. Moreover, if $y \in \text{proj}_{\mathfrak{X}}(x)$ with $\text{dist}(x, y) < \text{inj}(y)$ then $\text{Log}_y(x) \in N_y\mathfrak{X}$.

The set of local minima \mathcal{S} defined in (1) may not be closed: consider for example the function $f(x) = \text{sgn}(x) \exp(-\frac{1}{x^2})(1 + \sin(\frac{1}{x^2}))$ with $f_{\mathcal{S}} = 0$. It follows that the projection onto \mathcal{S} may be empty. Notwithstanding, the following holds:


Lemma 1.4 *Around each $\bar{x} \in \mathcal{S}$ there exists a neighborhood in which $\text{proj}_{\mathcal{S}}$ is non-empty.*

Proof Let \mathcal{U} be an open neighborhood of \bar{x} such that $f(x) \geq f_{\mathcal{S}}$ for all $x \in \mathcal{U}$. Let $\mathcal{V}_1, \mathcal{V}_2 \subset \mathcal{U}$ be two closed balls around \bar{x} of radii $\delta > 0$ and $\frac{1}{4}\delta$ respectively. Then $\mathcal{S} \cap \mathcal{V}_1 = f^{-1}(f_{\mathcal{S}}) \cap \mathcal{V}_1$, showing that $\mathcal{S} \cap \mathcal{V}_1$ is closed and the projection onto this set is non-empty. Let $x \in \mathcal{V}_2$ and $y \in \text{proj}_{\mathcal{S} \cap \mathcal{V}_1}(x)$. Then $\text{dist}(x, y) \leq \frac{1}{4}\delta$. Moreover, for all $y' \in \mathcal{S} \setminus \mathcal{V}_1$ we have $\text{dist}(x, y') \geq \frac{3}{4}\delta$. It follows that $\text{proj}_{\mathcal{S}}(x) = \text{proj}_{\mathcal{S} \cap \mathcal{V}_1}(x)$, and this is non-empty. \square

From these considerations we deduce that the projection onto \mathcal{S} is always locally well behaved.

Lemma 1.5 *Let $\bar{x} \in \mathcal{S}$. There exists a neighborhood \mathcal{U} of \bar{x} such that for all $x \in \mathcal{U}$ the set $\text{proj}_{\mathcal{S}}(x)$ is non-empty, and for all $y \in \text{proj}_{\mathcal{S}}(x)$ we have $\text{dist}(x, y) < \text{inj}(y)$. In particular, $v = \text{Log}_y(x)$ is well defined and $v \in N_y\mathcal{S}$.*

Fig. 3 Implication graph when f is C^2 . The main missing pieces were $P\mathbb{L} \Rightarrow MB$ and $QG \Rightarrow EB$, both secured with the right constants and under the right regularity assumptions



Proof Let \mathcal{U} be a neighborhood of \bar{x} where $\text{proj}_{\mathcal{S}}$ is non-empty (given by Lemma 1.4). Given $\bar{\delta} < \text{inj}(\bar{x})$, the ball $\bar{B}(\bar{x}, \delta)$ is compact for all $\delta < \bar{\delta}$. Define $h(\delta) = \inf_{x \in \bar{B}(\bar{x}, 2\delta)} \text{inj}(x)$ on the interval $[0, \frac{\bar{\delta}}{2}]$. The function h is continuous with $h(0) = \text{inj}(\bar{x}) > 0$ so we can pick $\delta > 0$ such that $\delta \leq h(\delta)$. Let $x \in B(\bar{x}, \delta) \cap \mathcal{U}$ and $y \in \text{proj}_{\mathcal{S}}(x)$. By definition of the projection we have $\text{dist}(x, y) \leq \text{dist}(x, \bar{x}) < \delta$. Moreover, inequality (3) yields $\text{dist}(y, \bar{x}) \leq 2 \text{dist}(x, \bar{x}) \leq 2\delta$ so $h(\delta) \leq \text{inj}(y)$. It follows that $\text{dist}(x, y) < \delta \leq h(\delta) \leq \text{inj}(y)$ and $v = \text{Log}_y(x)$ is well defined. The fact that v is in the normal cone follows from optimality conditions of projections. \square

Given a self-adjoint linear map H , we let $\lambda_i(H)$ denote the i th largest eigenvalue of H , and $\lambda_{\min}(H)$ and $\lambda_{\max}(H)$ denote the minimum and maximum eigenvalues respectively.

2 Four equivalent properties

In this section, we establish that MB, P \mathbb{L} , EB and QG (see Definitions 1.1 and 1.2) are equivalent around a local minimum \bar{x} when f is C^2 . Specifically, we show the implication graph in Fig. 3.

It is well known that P \mathbb{L} implies QG around minima: see references in Sect. 1.3. Perhaps the most popular argument relies on the bounded length of gradient flow trajectories under the more general Łojasiewicz inequality [2, 19, 68, 69, 79].

Definition 2.1 Let \bar{x} be a local minimum of f with associated set \mathcal{S} (1). We say f satisfies the Łojasiewicz inequality with constants $\theta \in [0, 1[$ and $\mu > 0$ around \bar{x} if

$$|f(x) - f_{\mathcal{S}}|^{2\theta} \leq \frac{1}{2\mu} \|\nabla f(x)\|^2 \tag{L}$$

for all x in some neighborhood of \bar{x} .

Notice that if f is Łojasiewicz with exponent θ then it is Łojasiewicz with exponent θ' for all $\theta \leq \theta' < 1$ (though possibly in a different neighborhood). The case $\theta = \frac{1}{2}$ is exactly the (P \mathbb{L}) condition.

Proposition 2.2 ($P\mathbb{L} \Rightarrow QG$) Suppose that f satisfies (L) around $\bar{x} \in \mathcal{S}$. Then f satisfies

$$f(x) - f_{\mathcal{S}} \geq ((1 - \theta)\sqrt{2\mu})^{\frac{1}{1-\theta}} \text{dist}(x, \mathcal{S})^{\frac{1}{1-\theta}}$$

for all x sufficiently close to \bar{x} . In particular, if $\theta = \frac{1}{2}$, this shows μ -(PŁ) \Rightarrow μ -(QG).

We include a classical proof in Appendix A for completeness, with care regarding neighborhoods.

2.1 Two straightforward implications

In this section we show that MB \Rightarrow QG and EB \Rightarrow PŁ. These implications are known and direct. We give succinct proofs for completeness. The first one follows immediately from a Taylor expansion.

Proposition 2.3 (MB \Rightarrow QG) *Suppose that f is C^2 and satisfies μ -(MB) at $\bar{x} \in \mathcal{S}$. Then f satisfies μ' -(QG) around \bar{x} for all $\mu' < \mu$.*

Proof Let \mathcal{U} be a neighborhood of \bar{x} as in Lemma 1.5. Let d be the codimension of \mathcal{S} (around \bar{x}). Given $\mu' < \mu$, pick $\varepsilon \in]0, \mu - \mu'[$ and shrink \mathcal{U} so that for all $x \in \mathcal{U}$ and $y \in \text{proj}_{\mathcal{S}}(x)$ we have $\lambda_d(\nabla^2 f(y)) \geq \mu' + \varepsilon$. Given $x \in \mathcal{U}$ and $y \in \text{proj}_{\mathcal{S}}(x)$, a Taylor expansion around y gives

$$f(x) - f_{\mathcal{S}} = \frac{1}{2} \langle v, \nabla^2 f(y)[v] \rangle + o(\|v\|^2) \geq \frac{\mu' + \varepsilon}{2} \text{dist}(x, \mathcal{S})^2 + o(\text{dist}(x, \mathcal{S})^2),$$

where $v = \text{Log}_y(x)$ is normal to \mathcal{S} . We get the inequality μ' -(QG) for all x sufficiently close to \bar{x} . □

Proposition 2.4 (EB \Rightarrow PŁ) *Suppose that f is C^2 and satisfies μ -(EB) around $\bar{x} \in \mathcal{S}$. Then f satisfies μ' -(PŁ) around \bar{x} for all $\mu' < \mu$.*

Proof Let \mathcal{U} be the intersection of two neighborhoods of \bar{x} : one where μ -(EB) holds, and the other provided by Lemma 1.5. Given $x \in \mathcal{U}$ and $y \in \text{proj}_{\mathcal{S}}(x)$, a Taylor expansion around y yields

$$f(x) - f_{\mathcal{S}} = \frac{1}{2} \langle v, \nabla^2 f(y)[v] \rangle + o(\|v\|^2) \quad \text{and} \quad \nabla f(x) = \Gamma_v \nabla^2 f(y)[v] + o(\|v\|),$$

where $v = \text{Log}_y(x)$. Using the Cauchy–Schwarz inequality and the triangle inequality, it follows that

$$f(x) - f_{\mathcal{S}} \leq \frac{1}{2} \|v\| \|\nabla^2 f(y)[v]\| + o(\|v\|^2) \leq \frac{1}{2} \|v\| \|\nabla f(x)\| + o(\|v\|^2).$$

Finally, EB gives that $\|v\| \leq \frac{1}{\mu} \|\nabla f(x)\|$ so $f(x) - f_{\mathcal{S}} \leq \frac{1}{2\mu} \|\nabla f(x)\|^2 + o(\|\nabla f(x)\|^2)$. We get the inequality μ' -(PŁ) for all x sufficiently close to \bar{x} . □

Remark 2.5 Suppose that f is only C^1 and ∇f is locally L -Lipschitz continuous around \bar{x} . If μ -EB holds around \bar{x} then f satisfies PŁ with constant $\frac{\mu^2}{L}$ around \bar{x} [52, Thm. 2]. (This still holds locally on manifolds with the same proof.) The constant worsens but that is inevitable: see the example in Remark 2.19.

2.2 Quadratic growth implies error bound

In this section, we show that QG implies EB for C^2 functions. Other works proving this implication either assume that f is convex (see [35, Prop. 5.3], [52], and [39, Cor. 3.6]) or do not provide control on the constants (see [40, Cor. 3.2]). For this, we first characterize a distance growth rate when we move from \mathcal{S} in a normal direction (see Definition 1.3). Recall that for now \mathcal{S} is not necessarily smooth, and therefore $N_{\bar{x}}\mathcal{S}$ is a priori only a cone.

Lemma 2.6 *Let $\bar{x} \in \mathcal{S}$ and $v \in N_{\bar{x}}\mathcal{S}$ unitary. Then $\text{dist}(\text{Exp}_{\bar{x}}(tv), \mathcal{S}) = t + o(t)$ as $t \rightarrow 0, t \geq 0$.*

Proof Let \mathcal{U} be a neighborhood of \bar{x} as in Lemma 1.5. Shrink \mathcal{U} so that for all $x \in \mathcal{U}$ and $y \in \text{proj}_{\mathcal{S}}(x)$ we have $\text{dist}(y, \bar{x}) < \text{inj}(\bar{x})$. Given a small parameter $t > 0$, define $x(t) = \text{Exp}_{\bar{x}}(tv)$ and let $y(t) \in \text{proj}_{\mathcal{S}}(x(t))$. From (3) we have $\text{dist}(y(t), \bar{x}) \leq 2t$, and it follows that $y(t) \rightarrow \bar{x}$ as $t \rightarrow 0$. Define $u(t) = \text{Log}_{y(t)}(x(t))$ and $w(t) = \text{Log}_{\bar{x}}(y(t))$. Then

$$\text{dist}(x(t), \mathcal{S})^2 = \|u(t)\|^2 = \|tv - w(t) + \Gamma_{y(t)}^{\bar{x}}u(t) - tv + w(t)\|^2 = \|tv - w(t)\|^2 + o(t^2)$$

as $t \rightarrow 0$ because $\|\Gamma_{y(t)}^{\bar{x}}u(t) - tv + w(t)\| = o(t)$ as $t \rightarrow 0$ [89, Eq. (6)]. In particular, for all t sufficiently small we have

$$\text{dist}(x(t), \mathcal{S})^2 = t^2 - 2t\langle w(t), v \rangle + \|w(t)\|^2 + o(t^2) \geq t^2 - 2t\langle w(t), v \rangle + o(t^2).$$

Let $I \subseteq \mathbb{R}_{>0}$ be the times when $y(t) \neq \bar{x}$. If $\inf I > 0$ then the final claim holds because $y(t) = \bar{x}$ for small enough t . Suppose now that $\inf I = 0$. Define $r(t) = \frac{w(t)}{\|w(t)\|}$ on I , and let \mathcal{A} be the set of accumulation points of $r(t)$ as $t \rightarrow 0, t \in I$. Then \mathcal{A} is included in the unit sphere and $\mathcal{A} \subseteq T_{\bar{x}}\mathcal{S}$ by definition. Given $a \in \mathcal{A}$ and $t \in I$, we can use $\langle a, v \rangle \leq 0$ to find that

$$\text{dist}(x(t), \mathcal{S})^2 \geq t^2 - 2t\|w(t)\|\langle r(t) - a + a, v \rangle + o(t^2) \geq t^2 - 4t^2\|r(t) - a\| + o(t^2).$$

It follows that $\text{dist}(x(t), \mathcal{S})^2 \geq t^2 - 4t^2 \text{dist}(r(t), \mathcal{A}) + o(t^2) = t^2 + o(t^2)$ because $\text{dist}(r(t), \mathcal{A}) \rightarrow 0$ as $t \rightarrow 0$. □

Using this rate, we now find that QG implies the following bounds for $\nabla^2 f$ in the normal cones of \mathcal{S} . Note that the following proposition does not yet show $\text{QG} \Rightarrow \text{MB}$, because to establish MB we still need to argue that \mathcal{S} is a smooth set.

Proposition 2.7 *Suppose f is C^2 and satisfies (QG) with constant μ around $\bar{x} \in \mathcal{S}$. Then for all $y \in \mathcal{S}$ sufficiently close to \bar{x} and all $v \in N_y\mathcal{S}$ we have*

$$\langle v, \nabla^2 f(y)[v] \rangle \geq \mu\|v\|^2 \quad \text{and} \quad \|\nabla^2 f(y)[v]\| \geq \mu\|v\|.$$

Proof Let \mathcal{U} be an open neighborhood of \bar{x} where μ -QG holds. Let $y \in \mathcal{S} \cap \mathcal{U}$ and $v \in N_y \mathcal{S}$. For all small enough $t > 0$, we obtain from a Taylor expansion that

$$\begin{aligned} f(\text{Exp}_y(tv)) - f_S &= \frac{t^2}{2} \langle v, \nabla^2 f(y)[v] \rangle + o(t^2) \geq \frac{\mu}{2} \text{dist}(\text{Exp}_y(tv), \mathcal{S})^2 \\ &= \frac{\mu}{2} t^2 \|v\|^2 + o(t^2), \end{aligned}$$

where the inequality comes from QG and the following equality comes from Lemma 2.6. Take $t \rightarrow 0$ to get the first inequality. The other inequality follows by Cauchy–Schwarz. \square

We deduce from this that, under QG, the gradient norm is locally bounded from below and from above by the distance to \mathcal{S} , up to some constant factors. This notably secures EB.

Proposition 2.8 (QG \Rightarrow EB) *Suppose f is C^2 and satisfies μ -(QG) around $\bar{x} \in \mathcal{S}$. For all $\mu^b < \mu$ and $\lambda^\sharp > \lambda_{\max}(\nabla^2 f(\bar{x}))$ there exists a neighborhood \mathcal{U} of \bar{x} such that for all $x \in \mathcal{U}$ we have*

$$\mu^b \text{dist}(x, \mathcal{S}) \leq \|\nabla f(x)\| \leq \lambda^\sharp \text{dist}(x, \mathcal{S}).$$

Proof Let \mathcal{U} be a neighborhood of \bar{x} as in Lemma 1.5. Shrink \mathcal{U} so that for all $x \in \mathcal{U}$ and $y \in \text{proj}_{\mathcal{S}}(x)$ the inequalities of Proposition 2.7 hold. Now let $x \in \mathcal{U}$ and $y \in \text{proj}_{\mathcal{S}}(x)$. Define $v = \text{Log}_y(x)$ and $\gamma(t) = \text{Exp}_y(tv)$, so that $y = \gamma(0)$ and $x = \gamma(1)$. Then a Taylor expansion of ∇f around y yields

$$\begin{aligned} \Gamma_v^{-1} \nabla f(x) &= \nabla^2 f(y)[v] + r(x) \text{ where} \\ r(x) &= \int_0^1 \left(\Gamma_{\tau v}^{-1} \circ \nabla^2 f(\gamma(\tau)) \circ \Gamma_{\tau v} - \nabla^2 f(y) \right) [v] d\tau. \end{aligned}$$

The Hessian is continuous so $r(x) = o(\|v\|)$ as $x \rightarrow \bar{x}$. Moreover, Proposition 2.7 provides that $\|\nabla^2 f(y)[v]\| \geq \mu \|v\|$ so using the triangle inequality and the reverse triangle inequality we get

$$\|v\|(\mu - o(1)) \leq \|\nabla f(x)\| \leq \|v\|(\lambda_{\max}(\nabla^2 f(y)) + o(1))$$

as $x \rightarrow \bar{x}$. We get the result if we choose x close enough to \bar{x} . \square

The first inequality in the previous proposition is exactly EB. In Proposition 2.4 we showed that EB implies PŁ when f is C^2 . Thus, it also holds that QG implies PŁ.

Remark 2.9 When f is only C^1 it is not true that QG implies EB or PŁ. To see this, consider the function $f(x) = 2x^2 + x^2 \sin(1/\sqrt{|x|})$ (see Fig. 2). It is C^1 and satisfies QG around the minimum $\bar{x} = 0$ because $f(x) \geq x^2$. However, there are other local minima arbitrarily close to \bar{x} . Those are critical points with function value strictly larger than $f(\bar{x})$, which disqualifies EB and PŁ around \bar{x} .

Remark 2.10 Combining Propositions 2.2 and 2.8 ($\mu\text{-P}\mathcal{L} \Rightarrow \mu\text{-Q}\mathcal{G}$ and $\mu\text{-Q}\mathcal{G} \Rightarrow \mu'\text{-E}\mathcal{B}$), one finds that $\mu\text{-P}\mathcal{L}$ implies $\mu'\text{-E}\mathcal{B}$ for C^2 functions. In fact, Karimi et al. [52] show that $\mu\text{-P}\mathcal{L}$ implies $\mu\text{-E}\mathcal{B}$ for C^1 functions, globally so if $\text{P}\mathcal{L}$ holds globally. Indeed, for x sufficiently close to \bar{x} , we have

$$\frac{\mu}{2} \text{dist}(x, \mathcal{S})^2 \leq f(x) - f_{\mathcal{S}} \leq \frac{1}{2\mu} \|\nabla f(x)\|^2,$$

where the second inequality comes from $\text{P}\mathcal{L}$, and the first inequality is $\text{Q}\mathcal{G}$ (implied by $\text{P}\mathcal{L}$).

Combining all previous implications, we obtain that $\text{M}\mathcal{B}$ implies $\text{P}\mathcal{L}$. This is also straightforward from Taylor expansion arguments.

Corollary 2.11 ($\text{M}\mathcal{B} \Rightarrow \text{P}\mathcal{L}$) *Suppose that f is C^2 and satisfies $\mu\text{-M}\mathcal{B}$ at $\bar{x} \in \mathcal{S}$. Let $0 < \mu' < \mu$. Then f satisfies $\mu'\text{-P}\mathcal{L}$ around \bar{x} .*

As a consequence, if f is $\mu\text{-M}\mathcal{B}$ on any subset of \mathcal{S} then for all $\mu' < \mu$ there exists a neighborhood of that subset where f is $\mu'\text{-P}\mathcal{L}$. When f is C^3 the size of the neighborhood where $\text{P}\mathcal{L}$ holds can be controlled (to some extent) with the third derivative. A version of this corollary appears in [94, Ex. 2.9] with a different trade-off between control of the neighborhood and the constant μ' . Feehan [44, Thm. 6] shows a similar result in Banach spaces assuming that the function is C^3 .

2.3 $\text{P}\mathcal{L}$ implies a smooth set of minima and $\text{M}\mathcal{B}$

The $\text{M}\mathcal{B}$ property is explicitly strong because it presupposes a smooth set of minima, and it clearly implies $\text{P}\mathcal{L}$, $\text{E}\mathcal{B}$, and $\text{Q}\mathcal{G}$. It raises a natural question: do the latter also enforce some structure on the set of minima? In this section we show that the answer is yes for C^2 functions: if $\text{P}\mathcal{L}$ (or $\text{E}\mathcal{B}$ or $\text{Q}\mathcal{G}$) holds around $\bar{x} \in \mathcal{S}$ then \mathcal{S} must be a submanifold around \bar{x} .

To get a sense of why \mathcal{S} cannot have singularities, suppose that $\mathcal{M} = \mathbb{R}^2$ and that around \bar{x} the set of minima is the union of two orthogonal lines (a cross) that intersect at \bar{x} . Then it must be that $\nabla^2 f(\bar{x}) = 0$ because the gradient is zero along both lines. However, if we assume $\text{P}\mathcal{L}$ then the spectrum of $\nabla^2 f(x)$ must contain at least one positive eigenvalue bigger than μ for all points $x \in \mathcal{S} \setminus \{\bar{x}\}$ close to \bar{x} , owing to the $\text{Q}\mathcal{G}$ property. We obtain a contradiction because the eigenvalues of $\nabla^2 f$ are continuous.

To generalize this intuition, we first show that $\text{P}\mathcal{L}$ induces a lower-bound on the positive eigenvalues of $\nabla^2 f$.

Proposition 2.12 *Suppose f is C^2 and $\mu\text{-P}\mathcal{L}$ around $\bar{x} \in \mathcal{S}$. If λ is a non-zero eigenvalue of $\nabla^2 f(\bar{x})$ then $\lambda \geq \mu$.*

Proof Let $\lambda > 0$ be an eigenvalue of $\nabla^2 f(\bar{x})$ with associated unit eigenvector v . Then

$$f(\text{Exp}_{\bar{x}}(tv)) - f_{\mathcal{S}} = \frac{\lambda}{2}t^2 + o(t^2) \quad \text{and} \quad \nabla f(\text{Exp}_{\bar{x}}(tv)) = \lambda t \Gamma_{tv} v + o(t).$$

The PL condition implies $\frac{\lambda}{2}t^2 + o(t^2) \leq \frac{1}{2\mu}\lambda^2t^2 + o(t^2)$, which gives the result as $t \rightarrow 0$. □

The latter argument is inconclusive when $\lambda = 0$. Still, we do get control of the Hessian’s rank.

Corollary 2.13 *Suppose f is C^2 and μ -(PL) around $\bar{x} \in \mathcal{S}$. Then $\text{rank}(\nabla^2 f(x)) = \text{rank}(\nabla^2 f(\bar{x}))$ for all $x \in \mathcal{S}$ close enough to \bar{x} .*

Proof Since f is C^2 the eigenvalues of $\nabla^2 f$ are continuous and the map $x \mapsto \text{rank}(\nabla^2 f(x))$ is lower semi-continuous, that is, if $x \in \mathcal{M}$ is close enough to \bar{x} then $\text{rank} \nabla^2 f(x) \geq d$, where $d = \text{rank} \nabla^2 f(\bar{x})$. Furthermore, if $y \in \mathcal{S}$ is sufficiently close to \bar{x} then $\lambda_{d+1}(\nabla^2 f(y)) < \mu$ by continuity of eigenvalues, and Proposition 2.12 then implies $\lambda_{d+1}(\nabla^2 f(y)) = 0$. □

This allows us to show that ∇f aligns locally in a special way with the eigenspaces of $\nabla^2 f$. This alignment will be particularly valuable to analyze second-order algorithms in Sect. 4.

Lemma 2.14 *Suppose f is C^2 and satisfies (PL) around $\bar{x} \in \mathcal{S}$. Let $d = \text{rank}(\nabla^2 f(\bar{x}))$. Then the orthogonal projector $P(x)$ onto the top d eigenspace of $\nabla^2 f(x)$ is well defined when x is sufficiently close to \bar{x} , and (with I denoting identity)*

$$\|(I - P(x))\nabla f(x)\| = o(\text{dist}(x, \mathcal{S})) = o(\|\nabla f(x)\|)$$

as $x \rightarrow \bar{x}$. Additionally, if $\nabla^2 f$ is locally Lipschitz continuous around \bar{x} then $\|(I - P(x))\nabla f(x)\| = O(\text{dist}(x, \mathcal{S})^2) = O(\|\nabla f(x)\|^2)$ as $x \rightarrow \bar{x}$.

Proof Given a point $x \in \mathcal{M}$, let $P(x): T_x\mathcal{M} \rightarrow T_x\mathcal{M}$ denote the orthogonal projector onto the top d eigenspace of $\nabla^2 f(x)$. This is well defined provided $\lambda_d(\nabla^2 f(x)) > \lambda_{d+1}(\nabla^2 f(x))$. Let \mathcal{U} be a neighborhood of \bar{x} as in Lemma 1.5. By continuity of the eigenvalues of $\nabla^2 f$, we can shrink \mathcal{U} so that for all $x \in \mathcal{U}$ and $y \in \text{proj}_{\mathcal{S}}(x)$ the projectors $P(x)$ and $P(y)$ are well defined. Given $x \in \mathcal{U}$, we let $y \in \text{proj}_{\mathcal{S}}(x)$ and $v = \text{Log}_y(x)$. Now define $\gamma(t) = \text{Exp}_y(tv)$. A Taylor expansion of the gradient around y gives that

$$\begin{aligned} \nabla f(x) &= \Gamma_v(\nabla^2 f(y)[v] + r(x)) \text{ where } r(x) \\ &= \int_0^1 (\Gamma_{\tau v}^{-1} \circ \nabla^2 f(\gamma(\tau)) \circ \Gamma_{\tau v} - \nabla^2 f(y))[v] d\tau. \end{aligned}$$

By Corollary 2.13, the rank of $\nabla^2 f$ is locally constant on \mathcal{S} (equal to d) so $\nabla^2 f(y) = P(y)\nabla^2 f(y)$ whenever x is sufficiently close to \bar{x} (using the bound $\text{dist}(y, \bar{x}) \leq 2 \text{dist}(x, \bar{x})$ from (3)). It follows that

$$(I - P(x))\nabla f(x) = (I - P(x))\Gamma_v(P(y)\nabla^2 f(y)[v] + r(x)).$$

Notice that $\Gamma_v \rightarrow I$ as $x \rightarrow \bar{x}$ and $P(y) \rightarrow P(\bar{x})$ as $x \rightarrow \bar{x}$ so $(I - P(x))\Gamma_v P(y) \rightarrow 0$ as $x \rightarrow \bar{x}$. It follows that $(I - P(x))\nabla f(x) = o(\|v\|)$ as $x \rightarrow \bar{x}$. The claim follows

by noting that $\|v\| = \text{dist}(x, \mathcal{S})$ and the fact that $\text{dist}(x, \mathcal{S})$ is commensurate $\|\nabla f(x)\|$ (as shown in Proposition 2.8).

Suppose now that $\nabla^2 f$ is locally Lipschitz continuous around \bar{x} . Then P is also locally Lipschitz continuous around \bar{x} [93, Thm. 1] and $(I - P(x))\Gamma_v P(y) = (I - P(x))\Gamma_v(P(y) - \Gamma_v^{-1}P(x)) = O(\|v\|)$. Moreover, we have $r(x) = O(\|v\|^2)$ so it follows that $(I - P(x))\nabla f(x) = O(\|v\|^2)$ as $x \rightarrow \bar{x}$. We conclude again with Proposition 2.8. \square

The lemma below exhibits a submanifold \mathcal{Z} that contains \bar{x} . It need *not* coincide with the set of minima \mathcal{S} . However, they *do* coincide if we assume that PL holds. This is the main argument to show that \mathcal{S} is locally a submanifold whenever f is PL and sufficiently regular.

Lemma 2.15 *Suppose f is C^p with $p \geq 2$. Let $\bar{x} \in \mathcal{S}$ and define $\mathcal{U} = \{x \in \mathcal{M} : \text{dist}(x, \bar{x}) < \text{inj}(\bar{x})\}$. Let $P(\bar{x})$ denote the orthogonal projector onto the image of $\nabla^2 f(\bar{x})$. Then the set*

$$\mathcal{Z} = \{x \in \mathcal{U} : P(\bar{x})\Gamma_{\bar{x}}\nabla f(x) = 0\}$$

is a C^{p-1} embedded submanifold locally around \bar{x} . If f is analytic then \mathcal{Z} is also analytic.

Proof We build a local defining function for \mathcal{Z} . Let d be the rank of $\nabla^2 f(\bar{x})$ and u_1, \dots, u_d be a set of orthonormal eigenvectors of $\nabla^2 f(\bar{x})$ with associated eigenvalues $\lambda_1, \dots, \lambda_d > 0$. We define $h : \mathcal{U} \rightarrow \mathbb{R}^d$ as $h_i(x) = \langle u_i, \Gamma_{\bar{x}}\nabla f(x) \rangle$. Clearly, $h(x) = 0$ if and only if $x \in \mathcal{Z}$. The function h is C^{p-1} if f is C^p , and analytic if f is analytic. For all $\dot{x} \in T_{\bar{x}}\mathcal{M}$ we have

$$Dh_i(\bar{x})[\dot{x}] = \langle u_i, \nabla^2 f(\bar{x})[\dot{x}] \rangle = \langle \nabla^2 f(\bar{x})[u_i], \dot{x} \rangle = \lambda_i \langle u_i, \dot{x} \rangle.$$

It follows that $Dh(\bar{x})$ has full rank. Thus, \mathcal{Z} is a submanifold around \bar{x} with the stated regularity. \square

A result similar to Lemma 2.15 is presented in [31, Lem. 1] for Banach spaces. We are now ready for one of our main theorems, regarding the regularity of the set of local minimizers \mathcal{S} (1).

Theorem 2.16 *Suppose f is C^p with $p \geq 2$ and satisfies (PL) around $\bar{x} \in \mathcal{S}$. Then \mathcal{S} is a C^{p-1} submanifold of \mathcal{M} locally around \bar{x} . If f is analytic then \mathcal{S} is also analytic.*

Proof We let d denote the rank of $\nabla^2 f(\bar{x})$. By Corollary 2.13, $\text{rank } \nabla^2 f(y) = d$ for all $y \in \mathcal{S}$ sufficiently close to \bar{x} . We let $\mathcal{U} \subseteq B(\bar{x}, \text{inj}(\bar{x}))$ be a neighborhood of \bar{x} such that for all $x \in \mathcal{U}$ the orthogonal projector $P(x) : T_x\mathcal{M} \rightarrow T_x\mathcal{M}$ onto the top d eigenspace of $\nabla^2 f(x)$ is well defined (it exists because eigenvalues of $\nabla^2 f$ are continuous). Lemma 2.15 ensures that

$$\mathcal{Z} = \{x \in \mathcal{U} : P(\bar{x})\Gamma_{\bar{x}}\nabla f(x) = 0\}$$

is a submanifold around \bar{x} . Clearly, $\mathcal{S} \cap \mathcal{U} \subseteq \mathcal{Z}$ holds. We now show the other inclusion to obtain that \mathcal{S} and \mathcal{Z} coincide around \bar{x} . From Lemma 2.14 we have $\|\nabla f(x)\| \leq \|P(x)\nabla f(x)\| + o(\|\nabla f(x)\|)$ as $x \rightarrow \bar{x}$. Moreover, the triangle inequality gives $\|P(x)\nabla f(x)\| \leq \|(P(x) - \Gamma_{\bar{x}}^x P(\bar{x})\Gamma_{\bar{x}}^x)\nabla f(x)\| + \|P(\bar{x})\Gamma_{\bar{x}}^x \nabla f(x)\|$. By continuity of P we have $P(x) - \Gamma_{\bar{x}}^x P(\bar{x})\Gamma_{\bar{x}}^x = o(1)$ as $x \rightarrow \bar{x}$ so it follows that

$$\|\nabla f(x)\| \leq \|P(\bar{x})\Gamma_{\bar{x}}^x \nabla f(x)\| + o(\|\nabla f(x)\|)$$

as $x \rightarrow \bar{x}$. We conclude that $P(\bar{x})\Gamma_{\bar{x}}^x \nabla f(x) = 0$ implies $\nabla f(x) = 0$ for all x sufficiently close to \bar{x} . This confirms that $\mathcal{Z} \subseteq \mathcal{S} \cap \mathcal{U}$ around \bar{x} because all critical points near \bar{x} are in \mathcal{S} by PL. \square

The codimension of \mathcal{S} is equal to the rank of $\nabla^2 f$ on \mathcal{S} , as expected. A similar result holds for Banach spaces when the function is assumed analytic [44, Thm. 1]. Around \bar{x} , the set of all minima of f and \mathcal{S} coincide when PL holds. Hence, Theorem 2.16 implies that the set of minima of f is a submanifold around \bar{x} . Using the QG property we now deduce that PL implies MB.

Corollary 2.17 *If f is C^2 and μ -PL around $\bar{x} \in \mathcal{S}$ then it satisfies μ -(MB) at \bar{x} . The same holds if f is μ -(EB) or μ -(QG) rather than μ -(PL).*

Proof Apply Theorem 2.16 to get that \mathcal{S} is locally a C^1 submanifold around \bar{x} . Proposition 2.2 gives (QG) around \bar{x} . Finally apply Proposition 2.7 to normal eigenvectors of $\nabla^2 f(\bar{x})$. This yields that the normal eigenvalues are at least μ . We obtain the same result if we suppose μ -EB or μ -QG instead of μ -PL. This is because they both imply μ' -PL for $\mu' < \mu$ arbitrarily close to μ . Taking the limit $\mu' \rightarrow \mu$ gives μ -MB. \square

Remark 2.18 (Connections to the distance function) Given a closed set $\mathfrak{X} \subseteq \mathcal{M}$, the function $f(x) = \frac{1}{2} \text{dist}_{\mathfrak{X}}(x)^2$ clearly satisfies QG. If f is C^p with $p \geq 2$ in a neighborhood of \mathfrak{X} then Theorem 2.16 applies, revealing that \mathfrak{X} is a C^{p-1} submanifold. This question is of independent interest, see for example [15] for a proof when assuming $p \geq 3$.

Remark 2.19 (Structure when f is only C^1) Theorem 2.16 requires f to be C^2 . And indeed, if f is only C^1 the set of minima may not be a submanifold. We provide two examples.

The function $f(x, y) = \frac{x^2 y^2}{x^2 + y^2}$ is C^1 and PL around the origin, yet its minimizers form a cross (see Fig. 2).² Incidentally, f is $\frac{1}{\sqrt{2}}$ -EB but only $\frac{1}{2}$ -PL around the origin, confirming that the constant worsens for the implication $\text{EB} \Rightarrow \text{PL}$ when f is only C^1 .

Additionally, let $\mathfrak{X} \subseteq \mathcal{M}$ be a closed set and suppose that the distance function $\text{dist}_{\mathfrak{X}}$ is C^1 around \mathfrak{X} (such a set is called proximally smooth [32]). For $f(x) = \frac{1}{2} \text{dist}_{\mathfrak{X}}(x)^2$, we find that $\nabla f(x) = x - \text{proj}_{\mathfrak{X}}(x)$, meaning that f is PL around \mathfrak{X} with constant $\mu = 1$. This holds in particular for all closed convex sets, yet many such sets fail to be C^0 submanifolds (e.g., consider a closed square in $\mathcal{M} = \mathbb{R}^2$). This provides further examples of C^1 functions satisfying PL yet for which the set of minima is not C^0 .

² We thank Christopher Criscitiello who found this function.

Remark 2.20 (Restricted secant inequality) It is possible to show equivalences with even more properties. For example, as in [101], we say f satisfies the *restricted secant inequality* (RSI) with constant μ around $\bar{x} \in \mathcal{S}$ if $\langle \nabla f(x), v \rangle \geq \mu \operatorname{dist}(x, \mathcal{S})^2$ for all x in a neighborhood of \bar{x} , where $v = -\operatorname{Log}_{\bar{x}}(y)$ and $y \in \operatorname{proj}_{\mathcal{S}}(x)$. From simple Taylor expansion arguments, we find that μ -MB implies μ' -RSI for all $\mu' < \mu$. By Cauchy–Schwarz, we also find that μ -RSI implies μ -EB for C^1 functions (see [52]). It follows that for C^2 functions RSI is also equivalent to the four properties that we consider.

Remark 2.21 (Other Łojasiewicz exponents) The PL condition is exactly (L) with exponent $\theta = \frac{1}{2}$. We comment here about other values of θ . First, suppose f is C^1 and ∇f locally Lipschitz continuous. If f is non-constant around $\bar{x} \in \mathcal{S}$ then it cannot satisfy (L) with an exponent $\theta < \frac{1}{2}$ around \bar{x} . This is because these assumptions are incompatible with the growth property from Proposition 2.2. See also [1, Thm. 4] for an algorithmic perspective on this. Now suppose f is C^2 and satisfies (L) with exponent θ around $\bar{x} \in \mathcal{S}$. If $\theta \in [\frac{1}{2}, \frac{2}{3}[$ and $\nabla^2 f$ is Lipschitz continuous around \bar{x} then PL holds around \bar{x} . Furthermore, if $\theta = \frac{2}{3}$ and f is additionally C^3 then PL also holds around \bar{x} . We include a proof of these observations in Appendix B.

3 Stability of minima and linear rates

In this section, we consider two types of algorithmic questions: the stability of minima and standard local convergence rates. We review some necessary classical arguments, taking this opportunity to generalize some of them to accommodate non-isolated minima. This will serve us well to analyze algorithms in Sect. 4.

3.1 Capture for sets of non-isolated minima

Typically, global convergence analyses of optimization algorithms merely guarantee that iterates accumulate only at critical points. The set of accumulation points may be empty (when the iterates diverge). Worse, it may even be infinite when minima are not isolated. See [2, §3.2.1] and [18, §5.3] for examples of pathological functions for which reasonable algorithms (such as gradient descent) produce iterates with continuous sets of accumulation points. The latter issue cannot occur when minima are isolated.

What kind of stability results still hold when minima are not isolated? Consider an algorithm generating iterates as $x_{k+1} = F_k(x_k, \dots, x_0)$, where F_k is a *descent mapping*: it satisfies $f(F_k(x_k, \dots, x_0)) \leq f(x_k)$ for all $x_k, \dots, x_0 \in \mathcal{M}$. Many deterministic algorithms fall in this category, including gradient descent and trust-region methods under suitable hypotheses. The standard capture theorem asserts that if the iterates generated by such descent mappings get sufficiently close to an *isolated* local minimum then the sequence eventually converges to it (under a few weak assumptions) [16, Prop. 1.2.5], [4, Thm. 4.4.2]. The result can be easily extended to a compact set of non-isolated local minima that satisfies several properties that we define now.

Definition 3.1 We say $\mathfrak{X} \subseteq \mathcal{M}$ is *isolated from critical points* if there exists a neighborhood \mathcal{U} of \mathfrak{X} such that $x \in \mathcal{U}$ and $\nabla f(x) = 0$ imply $x \in \mathfrak{X}$.

Note that the points in \mathfrak{X} do *not* need to be isolated: the set \mathfrak{X} may be a continuum of critical points. It is clear that if $\mathfrak{X} \subseteq \mathcal{S}(1)$ is isolated from critical points then there exists a neighborhood \mathcal{U} of \mathfrak{X} such that $f(x) > f_{\mathcal{S}}$ for all $x \in \mathcal{U} \setminus \mathfrak{X}$, where $f_{\mathcal{S}}$ is the value of f on \mathfrak{X} . The capture result below, based on [4, Thm. 4.4.2], states that if the set of minima $\mathfrak{X} \subseteq \mathcal{S}$ is both compact and isolated from critical points then it traps the iterates generated by all reasonable descent algorithms. A key hypothesis is that the steps have to be small around local minima: that is typically the case.

Definition 3.2 An algorithm which generates sequences on \mathcal{M} has the *vanishing steps* property on a set $\mathfrak{X} \subseteq \mathcal{M}$ if there exists a neighborhood \mathcal{U} of \mathfrak{X} and a continuous function $\eta: \mathcal{M} \rightarrow \mathbb{R}_+$ with $\eta(\mathfrak{X}) = 0$ such that, if x_k is an iterate in \mathcal{U} , then the next iterate x_{k+1} satisfies

$$\text{dist}(x_k, x_{k+1}) \leq \eta(x_k). \quad (\text{VS})$$

We say that the algorithm has the (VS) property at a point $\bar{x} \in \mathcal{S}$ if it holds on the set $\{\bar{x}\}$.

Proposition 3.3 (Capture of iterates) *Let \mathfrak{X} be a compact subset of \mathcal{S} isolated from critical points. Consider an algorithm that produces iterates as $x_{k+1} = F_k(x_k, \dots, x_0)$, where F_k is a descent mapping. Assume that it satisfies the (VS) property on \mathfrak{X} . Also suppose that the sequences generated by this algorithm accumulate only at critical points of f . Then there exists a neighborhood \mathcal{U} of \mathfrak{X} such that if a sequence enters \mathcal{U} then all subsequent iterates are in \mathcal{U} and $\text{dist}(x_k, \mathfrak{X}) \rightarrow 0$.*

Proof There exists a compact neighborhood \mathcal{V} of \mathfrak{X} such that $\mathcal{V} \setminus \mathfrak{X}$ does not contain any critical point and $f(x) > f_{\mathcal{S}}$ for all $x \in \mathcal{V} \setminus \mathfrak{X}$. The (VS) property implies that there exists an open neighborhood \mathcal{W} of \mathfrak{X} included in \mathcal{V} such that for all $k \in \mathbb{N}$, $x_k \in \mathcal{W}$ and $x_{k-1}, \dots, x_0 \in \mathcal{M}$ we have $F_k(x_k, \dots, x_0) \in \mathcal{V}$. The set $\mathcal{V} \setminus \mathcal{W}$ is compact, and we let $\alpha > f_{\mathcal{S}}$ denote the minimum of f on this set. We define $\mathcal{U} = \{x \in \mathcal{V} : f(x) < \alpha\}$, which is included in \mathcal{W} by minimality of α . Now let $K \in \mathbb{N}$, $x_K \in \mathcal{U}$ and $x_{K-1}, \dots, x_0 \in \mathcal{M}$. Then we have $F_K(x_K, \dots, x_0) \in \mathcal{V}$ by definition of \mathcal{W} . Moreover, F_K is a descent mapping so $f(F_K(x_K, \dots, x_0)) < \alpha$, and it implies that $F_K(x_K, \dots, x_0) \in \mathcal{U}$. It follows that x_{K+1} is in \mathcal{U} and all subsequent iterates are also in \mathcal{U} . Now we show that $\{x_k\}$ converges to \mathfrak{X} . The sequence $\{x_k\}$ eventually stays in a compact set (because $x_k \in \mathcal{U} \subseteq \mathcal{V}$ for all $k \geq K$) so it has a non-empty and compact set of accumulation points that we denote by \mathcal{A} . Then $\mathcal{A} \subseteq \mathfrak{X}$ because $\mathcal{A} \subseteq \mathcal{V}$ and the only critical points in \mathcal{V} are in \mathfrak{X} . The set of accumulation points of $\{\text{dist}(x_k, \mathcal{A})\}$ is $\{0\}$ (because $\{x_k\}$ is bounded). So we deduce that $\lim_{k \rightarrow +\infty} \text{dist}(x_k, \mathfrak{X}) \leq \lim_{k \rightarrow +\infty} \text{dist}(x_k, \mathcal{A}) = 0$. \square

This statement does *not* guarantee that the iterates converge to a specific point, but merely that $\text{dist}(x_k, \mathfrak{X}) \rightarrow 0$. Notice that we do not require any particular structure for \mathfrak{X} nor any form of function growth. The assumptions on the mappings F_k are also mild.

However, we need \mathfrak{X} to be compact and this cannot be relaxed. Indeed, consider the function $f: \mathbb{R}^2 \rightarrow \mathbb{R}$ defined as $f(x, y) = \exp(x)y^2$. The set of global minima is $\mathfrak{X} = \{(x, y) \in \mathbb{R}^2 : y = 0\}$, and it contains all the critical points of f . Consider the update rule $(x_{k+1}, y_{k+1}) = (x_k - y_k^2, y_k)$. It satisfies the descent condition because $f(x_{k+1}, y_{k+1}) = \exp(-y_k^2)f(x_k, y_k)$ for all k . The distance assumption (VS) also holds because $\text{dist}((x_{k+1}, y_{k+1}), (x_k, y_k)) = \text{dist}((x_k - y_k^2, y_k), (x_k, y_k)) = y_k^2$. However, the sequence $\{\text{dist}((x_k, y_k), \mathfrak{X})\}$ is constant (and not converging to zero) even if we initialize the algorithm arbitrarily close to \mathfrak{X} (but not exactly on \mathfrak{X}).

The vanishing steps property is a reasonable assumption. When $\mathcal{M} = \mathbb{R}^n$, many optimization algorithms have an update rule of the form $x_{k+1} = x_k + s_k$ and the vector s_k is small if x_k is close to a local minimum (e.g., $s_k = -\alpha_k \nabla f(x_k)$ with bounded α_k). For a general manifold \mathcal{M} , algorithms produce iterates as $x_{k+1} = R_{x_k}(s_k)$, where $R: \text{T}\mathcal{M} \rightarrow \mathcal{M}: (x, s) \mapsto R_x(s)$ is a retraction [24, Def. 3.47]. Specifically, for all $(x, s) \in \text{T}\mathcal{M}$ the curve $c(t) = R_x(ts)$ satisfies $c(0) = x$ and $c'(0) = s$.

To ensure vanishing steps, we must control the distance traveled by retractions. We let $c_r \geq 1$ be such that

$$\text{dist}(R_x(s), x) \leq c_r \|s\| \tag{RD}$$

for all $(x, s) \in \text{T}\mathcal{M}$ where $\|s\|$ is smaller than some fixed positive radius. For $R = \text{Exp}$ (including the Euclidean case) the choice $c_r = 1$ is valid for all s . More generally, c_r can be set arbitrarily close to 1 for all retractions as long as the radius is sufficiently small [86, Lem. 6].

3.2 Lyapunov stability and convergence to a single point

Pathological behavior such as a continuous set of accumulation points can be ruled out assuming Łojasiewicz inequalities. In this case, local minima are stable for a variety of algorithms, even without compactness hypothesis (in contrast to Proposition 3.3). This in turn ensures that the iterates converge to a single point. In this section, we review arguments from [2], [10, Lem. 2.6], [21, Thm. 14] and [81, Thm. 4].

A central property to obtain convergence to a single point is a bound on the path length of the iterates. We make this precise in the following definition.

Definition 3.4 An algorithm which generates sequences on \mathcal{M} has the *bounded path length* property on a set $\mathfrak{X} \subseteq \mathcal{M}$ if the following is true. There exist a neighborhood \mathcal{U} of \mathfrak{X} and a continuous function $\gamma: \mathcal{M} \rightarrow \mathbb{R}_+$ with $\gamma(\mathfrak{X}) = 0$ such that, if $x_L, \dots, x_K \in \mathcal{M}$ are consecutive points generated by the algorithm and which are all in \mathcal{U} , then

$$\sum_{k=L}^{K-1} \text{dist}(x_k, x_{k+1}) \leq \gamma(x_L). \tag{BPL}$$

We say that the algorithm has the (BPL) property at a point $\bar{x} \in \mathcal{M}$ if it does so on the set $\{\bar{x}\}$.

Here we think of the algorithm as an optimization method with some fixed hyper-parameters. The definition is given for a generic set \mathfrak{X} but the bounded path length property is usually only satisfied around local minima. Combined with the vanishing steps property, bounded path length ensures stability of local minima. For comparison, Absil et al. [2] and Attouch et al. [10] deduce (BPL) from a function decrease condition. Here, we factor out (BPL) to enable analysis of algorithms that do not satisfy that decrease condition.

Proposition 3.5 (Lyapunov stability) *Suppose that an algorithm satisfies the (VS) and (BPL) properties at $\bar{x} \in \mathcal{M}$. Given a neighborhood \mathcal{U} of \bar{x} , there exists a neighborhood \mathcal{V} of \bar{x} such that if a sequence generated by this algorithm enters \mathcal{V} then all subsequent iterates stay in \mathcal{U} .*

Proof The set \mathcal{U} contains an open ball of radius $\delta_u > 0$ around \bar{x} in which the (BPL) and (VS) properties are satisfied with some functions γ and η . By continuity of η there exists an open ball \mathcal{W} centered on \bar{x} of radius δ_w that satisfies $\delta_w + \eta(x) < \delta_u$ for all $x \in \mathcal{W}$. Likewise, by continuity of γ , there exists an open ball $\mathcal{V} \subset \mathcal{W}$ of radius $\delta_v > 0$ around \bar{x} such that for all $x \in \mathcal{V}$ we have $\delta_v + \gamma(x) < \delta_w$. Suppose that an iterate x_L is in \mathcal{V} . For contradiction, let $K \geq L$ be the first index such that $x_{K+1} \notin \mathcal{B}(\bar{x}, \delta_u)$. We deduce from the triangle inequality and the (BPL) property that

$$\text{dist}(x_K, \bar{x}) \leq \text{dist}(x_L, \bar{x}) + \sum_{k=L}^{K-1} \text{dist}(x_k, x_{k+1}) \leq \delta_v + \gamma(x_L) < \delta_w.$$

It follows that $x_K \in \mathcal{W}$. Using again the triangle inequality we find $\text{dist}(x_{K+1}, \bar{x}) \leq \text{dist}(x_K, \bar{x}) + \text{dist}(x_K, x_{K+1}) \leq \delta_w + \eta(x_K) < \delta_u$. This implies that x_{K+1} is in $\mathcal{B}(\bar{x}, \delta_u)$: a contradiction. \square

In particular, this excludes that the iterates diverge. We can also guarantee that accumulation points are actually limit points.

Corollary 3.6 *Suppose that an algorithm satisfies the (VS) and (BPL) properties at $\bar{x} \in \mathcal{M}$. If it generates a sequence that accumulates at \bar{x} then the sequence converges to \bar{x} .*

Proof Let \mathcal{U} be a neighborhood of \bar{x} . From Proposition 3.5 there is a neighborhood \mathcal{V} such that if an iterate is in \mathcal{V} then all subsequent iterates are in \mathcal{U} . Since \bar{x} is an accumulation point we know that such an iterate exists. Repeat with a sequence of smaller and smaller neighborhoods of \bar{x} . \square

Many optimization algorithms generate sequences that accumulate only at critical points. In that scenario, we can deduce that the sequence converges to a point, provided that it gets close enough to a set where (VS) and (BPL) hold.

Corollary 3.7 *Consider an algorithm that satisfies the (VS) and (BPL) properties on a set $\mathfrak{X} \subseteq \mathcal{M}$ and let $\bar{x} \in \mathfrak{X}$. Let \mathcal{U} be a neighborhood of \bar{x} such that if a sequence generated by this algorithm accumulates at a point $x_\infty \in \mathcal{U}$ then x_∞ is in \mathfrak{X} . There exists a neighborhood \mathcal{V} of \bar{x} such that if a sequence enters \mathcal{V} then all subsequent iterates stay in \mathcal{U} and converge to some $x_\infty \in \mathcal{U} \cap \mathfrak{X}$.*

Proof The set \mathcal{U} contains a compact neighborhood \mathcal{B} of \bar{x} such that the (VS) and (BPL) properties hold on $\mathcal{B} \cap \mathfrak{X}$. From Proposition 3.5 there exists a neighborhood \mathcal{V} of \bar{x} such that if $\{x_k\}$ enters \mathcal{V} then it stays in \mathcal{B} . The set \mathcal{B} is compact so $\{x_k\}$ has at least one accumulation point $x_\infty \in \mathcal{B}$. Our hypotheses ensure that x_∞ must be in \mathfrak{X} since $\mathcal{B} \subseteq \mathcal{U}$. We conclude with Corollary 3.6. \square

The conclusions of Corollary 3.7 are similar to the ones in [2, Prop. 3.3] and [10, Thm. 2.10].³

We describe below the argument that Absil et al. [2] used to show that many gradient descent algorithms satisfy (BPL) around points where f is Łojasiewicz. We say that the sequence $\{x_k\}$ satisfies the *strong decrease property* around $\bar{x} \in \mathcal{M}$ if there exists $\sigma > 0$ such that

$$f(x_k) - f(x_{k+1}) \geq \sigma \|\nabla f(x_k)\| \operatorname{dist}(x_k, x_{k+1}) \quad \text{and} \quad x_k \in \mathcal{S} \Rightarrow x_{k+1} = x_k \quad (4)$$

whenever x_k is sufficiently close to \bar{x} , as introduced by Absil et al. [2].

Lemma 3.8 *Suppose that f satisfies (L) around $\bar{x} \in \mathcal{S}$ with constants θ and μ . If an algorithm generates sequences $\{x_k\}$ that satisfy (4) around \bar{x} then it satisfies the (BPL) property at \bar{x} with*

$$\gamma(x) = \frac{1}{\sigma(1-\theta)\sqrt{2\mu}} |f(x) - f_{\mathcal{S}}|^{1-\theta}.$$

We include a proof of this statement in Appendix C for completeness. In fact, the algorithm would still satisfy the (BPL) property under the more general Kurdyka–Łojasiewicz assumption (see [2, §3.2.3]). In practice, many first-order algorithms (including gradient descent with constant step-sizes or with line-search) generate sequences with the strong decrease condition (4), as shown in [2, §4].

3.3 Asymptotic convergence rate

To conclude this section, we briefly review classical linear convergence results for gradient methods under the PŁ assumption, as needed for Sect. 4. Proofs are in Appendix D for completeness. It is well known that gradient descent with appropriate step-sizes converges linearly to a minimum when f satisfies PŁ globally and has a Lipschitz continuous gradient [81]. The same arguments lead to an asymptotic linear convergence rate when PŁ holds only locally. We say that the sequence $\{x_k\}$ satisfies the *sufficient decrease property* with constant $\omega > 0$ if

$$f(x_k) - f(x_{k+1}) \geq \omega \|\nabla f(x_k)\|^2. \quad (5)$$

whenever x_k is sufficiently close to a point \bar{x} . The classical result below follows from that inequality [81].

³ Notice that Absil et al. [2] do *not* assume vanishing steps, but in that case the property can fail when \bar{x} is not a *global* minimum. The statement of Corollary 3.7 accounts for this technical point.

Proposition 3.9 *Let $\{x_k\}$ be a sequence of iterates converging to some $\bar{x} \in \mathcal{S}$ and satisfying (5). Suppose f satisfies (PL) around \bar{x} with constant $\mu > 0$. Then the sequence $\{f(x_k)\}$ converges linearly to $f_{\mathcal{S}}$ with rate $1 - 2\omega\mu$. Moreover, $\{\|\nabla f(x_k)\|\}$ and $\{\text{dist}(x_k, \mathcal{S})\}$ converge linearly to zero with rate $\sqrt{1 - 2\omega\mu} \leq 1 - \omega\mu$.*

In the case where \mathcal{M} is a Euclidean space and $R_x(s) = x + s$, it is well known that the sufficient decrease condition (5) holds for many first-order algorithms when ∇f is Lipschitz continuous. This is also true for a general manifold \mathcal{M} and retraction R as we briefly describe now. We say that f and the retraction R locally satisfy a Lipschitz-type property around $\bar{x} \in \mathcal{S}$ if there exists $L > 0$ such that

$$f(R_x(s)) \leq f(x) + \langle \nabla f(x), s \rangle + \frac{L}{2} \|s\|^2 \tag{6}$$

for all x close enough to \bar{x} and s small enough. Note that if $f \circ R$ is C^2 then the inequality (6) is always (locally) satisfied. It is a classical result that (6) implies sufficient decrease for gradient descent with constant step-sizes. This yields the following statement.

Proposition 3.10 *Suppose f satisfies μ -(PL) around $\bar{x} \in \mathcal{S}$. Also assume that (6) holds around \bar{x} . Let $\{x_k\}$ be a sequence of iterates generated by gradient descent with constant step-size $\gamma \in]0, \frac{2}{L}[$, that is, $x_{k+1} = R_{x_k}(-\gamma \nabla f(x_k))$. Given a neighborhood \mathcal{U} of \bar{x} , there exists a neighborhood \mathcal{V} of \bar{x} such that if an iterate enters \mathcal{V} then the sequence converges linearly to some $x_\infty \in \mathcal{U} \cap \mathcal{S}$ with rate $\sqrt{1 - 2\mu(\gamma - \frac{L}{2}\gamma^2)}$.*

4 Aiming for superlinear convergence

Under fairly general assumptions, the PL condition (which is compatible with non-isolated minima) ensures stability of minima and linear convergence for first-order methods, as recalled in the previous section. We now assume f is C^2 and investigate superlinear convergence to non-isolated minima.

A natural starting point is Newton’s method which, in spite of terrible global behavior [51], enjoys quadratic convergence to a *non-singular* minimum, provided the method is initialized sufficiently close. Unfortunately, this does not extend to non-isolated minima.

We exhibit here an example showing that the MB property is in general not sufficient to ensure such a strong convergence behavior. The update rule is $x_{k+1} = x_k - \nabla^2 f(x_k)^{-1}[\nabla f(x_k)]$ (we may use the pseudo-inverse instead). Consider the cost function $f(x, y) = \frac{1}{2}(x^2 + 1)y^2$, whose set of minima is the line $\mathcal{S} = \{(x, y) \in \mathbb{R}^2 : y = 0\}$. The gradient and Hessian of f are

$$\nabla f(x, y) = \begin{bmatrix} xy^2 \\ (x^2 + 1)y \end{bmatrix} \quad \text{and} \quad \nabla^2 f(x, y) = \begin{bmatrix} y^2 & 2xy \\ 2xy & x^2 + 1 \end{bmatrix}.$$

One can check that f satisfies (MB). To see how Newton’s method behaves on f , notice that

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

whenever $3x^2 \neq 1$. Let $x(t) = \sqrt{\frac{1-t}{3}}$ and $y(t) = \sqrt{t}$. We can choose $t \in]0, 1[$ as small as desired to make the point $(x(t), y(t))$ arbitrarily close to \mathcal{S} . Yet computing the Newton step at $(x(t), y(t))$ results in a new point at a distance $\frac{2}{3} \frac{1-t}{\sqrt{t}}$ from the optimal set \mathcal{S} : that is arbitrarily far away. The failure of Newton’s method stems from a misalignment between the gradient and some eigenspaces of the Hessian.

The usual fix for Newton’s method is to regularize it. This yields two classes of algorithms in particular: regularized Newton with cubics and trust-region methods. We will show that cubic regularization enjoys satisfying local convergence properties, even in the presence of non-isolated minima. In contrast, the picture is less clear for trust-region methods.

Throughout, we make several local assumptions around a point $\bar{x} \in \mathcal{S}$ as stated below. The first two are Lipschitz-type properties.

A1 The Hessian $\nabla^2 f$ is locally L_H -Lipschitz continuous around \bar{x} for some $L_H \geq 0$.

A2 There exists a constant $L'_H \geq 0$ such that the Lipschitz-type inequality

$$f(\mathbf{R}_x(s)) - f(x) - \langle s, \nabla f(x) \rangle - \frac{1}{2} \langle s, \nabla^2 f(x)[s] \rangle \leq \frac{L'_H}{6} \|s\|^3$$

holds for all x close enough to \bar{x} and all s small enough.

These assumptions typically hold (locally). For the retraction $\mathbf{R} = \text{Exp}$, the first implies the other (with the same constant). This is true in particular when \mathcal{M} is a Euclidean space and $\mathbf{R}_x(s) = x + s$. The third assumption concerns the two classes of algorithms we consider. At every iterate x_k , they build a local model of the cost function around x_k based on a linear map H_k . We require it to be close to the Hessian $\nabla^2 f(x_k)$. That holds in particular if $H_k = \nabla^2 f(x_k)$.

A3 For all k the map H_k is linear, symmetric, and there is a constant $\beta_H \geq 0$ such that

$$\|H_k - \nabla^2 f(x_k)\| \leq \beta_H \|\nabla f(x_k)\|$$

whenever the iterate x_k is close enough to \bar{x} .

Finally, we let $c_r \geq 1$ be such that the retraction satisfies (RD) for sufficiently small tangent vectors (which is enough for the local analyses below).

4.1 Adaptive regularized Newton

The regularized Newton method using cubics was introduced by Griewank [48] and later revisited by Nesterov and Polyak [76]. An adaptive version of this algorithm was

proposed by Cartis et al. [27, 28], with extensions to manifolds by Qi [83], Zhang and Zhang [102] and Agarwal et al. [5]. The adaptive variants update the penalty weight automatically: they are called ARC. We consider those variants, and more specifically an algorithm that generates sequences $\{(x_k, \zeta_k)\}$, where x_k is the current iterate and ζ_k is the cubic penalty weight. The update rule is $x_{k+1} = R_{x_k}(s_k)$ for some step s_k . At each iteration k , we define a linear operator $H_k : T_{x_k}\mathcal{M} \rightarrow T_{x_k}\mathcal{M}$ and the step s_k is chosen to approximately minimize the regularized second-order model

$$m_k(s) = f(x_k) + \langle s, \nabla f(x_k) \rangle + \frac{1}{2} \langle s, H_k[s] \rangle + \frac{\zeta_k}{3} \|s\|^3 \quad (7)$$

in a way that we make precise below. We require H_k to be close to $\nabla^2 f(x_k)$ as prescribed in A3.

In the literature, there are a number of superlinear (but non-quadratic) convergence results for such algorithms. Assuming the PL condition, Nesterov and Polyak [76] showed that regularized Newton generates sequences that converge superlinearly, with exponent $4/3$. Later, assuming a Łojasiewicz inequality with exponent $\theta \in [0, 1]$, Zhou et al. [105] characterized the convergence speed of regularized Newton depending on θ . In particular, they also show that the PL condition implies superlinear convergence.⁴ More recently, Qian and Pan [84] developed an abstract framework that encompasses these superlinear convergence results, and Cartis et al. [29, §5.3] reviewed superlinear convergence rates of ARC under Łojasiewicz inequalities.

There is also a quadratic convergence result: Yue et al. [97] employed a local error bound assumption to show quadratic convergence for the regularized Newton method. As discussed in Sect. 2, this assumption is equivalent to local PL, making their result an improvement over the superlinear rates from the aforementioned references. This underlines one of the benefits of recognizing the equivalence of the four conditions MB, PL, EB and QG, as some may more readily lead to a sharp analysis than others.

Note that the results in [97] assume that the subproblem is solved exactly, meaning that $s_k \in \arg \min_s m_k(s)$. Several authors proposed weaker conditions on s_k (only requiring an approximate solution to the subproblem) to ensure convergence guarantees: this is important because we cannot find an exact solution in practice. Agarwal et al. [5] for example, following Birgin et al. [17], establish global convergence guarantees assuming only that $m_k(s_k) \leq m_k(0)$ and $\|\nabla m_k(s_k)\| \leq \kappa \|s_k\|^2$ for some $\kappa \geq 0$. For their results to hold, they require $H_k = \nabla^2 \hat{f}_k(0)$, where $\hat{f}_k = f \circ R_x$ is the pull-back of f at x . This choice of H_k is compatible with A3 for retractions with bounded initial acceleration (which is typical).

We revisit the results of Yue et al. [97] to obtain an asymptotic quadratic convergence rate for ARC under the PL assumption, even with approximate solutions to the subproblem. Specifically, throughout this section we suppose that the steps s_k satisfy

$$m_k(s_k) \leq m_k(0) \quad \text{and} \quad \|\nabla m_k(s_k)\| \leq \kappa \|s_k\| \|\nabla f(x_k)\| \quad (8)$$

⁴ In fact, Zhou et al. [105] show a superlinear convergence rate only assuming that $\theta \in]0, \frac{2}{3}[$ and that $\nabla^2 f$ is Lipschitz continuous. However, these assumptions imply PL as noted in Remark 2.21.

for some $\kappa \geq 0$. At each iteration k , we define the ratio

$$\varrho_k = \frac{f(x_k) - f(\mathbf{R}_{x_k}(s_k))}{m_k(0) - m_k(s_k) + \frac{\varsigma_k}{3} \|s_k\|^3} \tag{9}$$

(as do Birgin et al. [17]), which measures the adequacy of the local model. Iteration k is said to be *successful* when ϱ_k is larger than some fixed parameter $\varrho_c \in]0, 1[$. In this case, we set $x_{k+1} = \mathbf{R}_{x_k}(s_k)$ and decrease the penalty weight so that $\varsigma_{k+1} \leq \varsigma_k$. The update mechanism ensures that $\varsigma_k \geq \varsigma_{\min}$ for all k , where $\varsigma_{\min} > 0$ is a fixed parameter. Conversely, the step is *unsuccessful* when $\varrho_k < \varrho_c$: we set $x_{k+1} = x_k$ and increase the penalty so that $\varsigma_{k+1} > \varsigma_k$. The explicit updates for ς_{k+1} are stated in [5]. We prove the following result for this algorithm.

Theorem 4.1 *Suppose A1, A2, A3 and (PL) hold around $\bar{x} \in \mathcal{S}$. We run ARC with inexact subproblem solver satisfying (8). Given any neighborhood \mathcal{U} of \bar{x} , there exists a neighborhood \mathcal{V} of \bar{x} such that if an iterate enters \mathcal{V} then the sequence converges quadratically to some $x_\infty \in \mathcal{U} \cap \mathcal{S}$.*

We first adapt an argument from [5, Lem. 6] to show that ARC satisfies the vanishing steps property (VS) defined in Sect. 3.1.

Lemma 4.2 *Suppose A3 holds around a point \bar{x} . There exists a neighborhood \mathcal{U} of \bar{x} such that if an iterate x_k is in \mathcal{U} then the step-size has norm bounded as $\|s_k\| \leq \tilde{\eta}(x_k, \varsigma_k) \leq \tilde{\eta}(x_k, \varsigma_{\min})$, where*

$$\tilde{\eta}(x, \varsigma) = \sqrt{\frac{3\|\nabla f(x)\|}{\varsigma}} + \frac{3}{2\varsigma} \Lambda(x)$$

and $\Lambda(x) = \max\left(0, \beta_H \|\nabla f(x)\| - \lambda_{\min}(\nabla^2 f(x))\right)$.

In particular, ARC has the (VS) property around second-order critical points with $\eta(x) = c_r \tilde{\eta}(x, \varsigma_{\min})$, where c_r controls possible retraction distortion as in (RD).

Proof The model decrease in (8) and the model accuracy A3 ensure

$$\begin{aligned} \varsigma_k \|s_k\|^3 &\leq -3\left\langle s_k, \nabla f(x_k) + \frac{1}{2} \nabla^2 f(x_k)[s_k] + \frac{1}{2} (H_k - \nabla^2 f(x_k))[s_k] \right\rangle \\ &\leq 3\|s_k\| \left(\|\nabla f(x_k)\| + \frac{1}{2} \Lambda(x_k) \|s_k\| \right). \end{aligned}$$

Divide by $\|s_k\|$ and solve the quadratic inequality for $\|s_k\|$ to get the result, recalling $\varsigma_k \geq \varsigma_{\min}$. □

The function η is indeed continuous with value zero on \mathcal{S} , as required in (VS). To obtain the vanishing steps property we only relied on the decrease requirement $m_k(s_k) \leq m_k(0)$. Assuming a PL condition and a locally Lipschitz continuous Hessian, we now derive sharper bounds for the steps. We rely on the fact that the gradient of the model m_k (7)

$$\nabla m_k(s_k) = \nabla f(x_k) + H_k[s_k] + \varsigma_k \|s_k\| s_k \tag{10}$$

is small. We will exploit the particular alignment of ∇f given in Lemma 2.14. In the following statements, given a point $\bar{x} \in \mathcal{S}$ and $d = \text{rank}(\nabla^2 f(\bar{x}))$, we let $P(x) : T_x \mathcal{M} \rightarrow T_x \mathcal{M}$ denote the orthogonal projector onto the top d eigenspace of $\nabla^2 f(x)$. This is always well defined in a neighborhood of \bar{x} by continuity of eigenvalues. Additionally, we let $Q(x) = I - P(x)$ be the projector onto the orthogonal complement.

Lemma 4.3 *Suppose that A1, A3 and μ -(PL) hold around $\bar{x} \in \mathcal{S}$. Given $\varepsilon > 0$, $\mu^b < \mu$ and $\lambda^\sharp > \lambda_{\max}(\nabla^2 f(\bar{x}))$, there exists a neighborhood \mathcal{U} of \bar{x} and a constant $L_q \geq 0$ such that if x_k is an iterate in \mathcal{U} and s_k is a step satisfying (8) then*

$$(1 - \varepsilon) \frac{\|\nabla f(x_k)\|}{\lambda^\sharp + \varsigma_k \|s_k\|} \leq \|P(x_k)s_k\| \leq (1 + \varepsilon) \frac{\|\nabla f(x_k)\|}{\mu^b + \varsigma_k \|s_k\|} \quad \text{and} \tag{11}$$

$$\|Q(x_k)s_k\| \leq \frac{1}{\varsigma_k} \left((\kappa + \beta_H) \|\nabla f(x_k)\| + (L_H + L_q \sqrt{\varsigma_k}) \text{dist}(x_k, \mathcal{S}) \right). \tag{12}$$

Proof Let \mathcal{U} be a neighborhood of \bar{x} where the orthogonal projector $P(x)$ is well defined for all $x \in \mathcal{U}$. Let $x_k \in \mathcal{U}$ and s_k be a step that satisfies (8). We first bound the term $P(x_k)s_k$. Multiply (10) by $P(x_k)$ and use commutativity of $P(x_k)$ and $\nabla^2 f(x_k)$ to get

$$(\nabla^2 f(x_k) + \varsigma_k \|s_k\| I)P(x_k)s_k = P(x_k) \left(-\nabla f(x_k) + \nabla m_k(s_k) - (H_k - \nabla^2 f(x_k))[s_k] \right).$$

If we apply A3 and (8), we find that $\|\nabla m_k(s_k) - (H_k - \nabla^2 f(x_k))[s_k]\| \leq (\kappa + \beta_H) \|\nabla f(x_k)\| \|s_k\|$ when x_k is close enough to \bar{x} (shrink \mathcal{U} as needed). Consequently, the previous equality yields

$$\begin{aligned} \frac{\|P(x_k)\nabla f(x_k)\| - (\kappa + \beta_H)\|\nabla f(x_k)\|\|s_k\|}{\lambda_1(\nabla^2 f(x_k)) + \varsigma_k \|s_k\|} &\leq \|P(x_k)s_k\| \\ &\leq \frac{\|P(x_k)\nabla f(x_k)\| + (\kappa + \beta_H)\|\nabla f(x_k)\|\|s_k\|}{\lambda_d(\nabla^2 f(x_k)) + \varsigma_k \|s_k\|}. \end{aligned}$$

Lemma 2.14 gives that $\|P(x)\nabla f(x)\| = \|\nabla f(x)\| + o(\|\nabla f(x)\|)$ as $x \rightarrow \bar{x}$. Moreover, the steps s_k vanish (as shown in Lemma 4.2), so we obtain the bound (11) when

x_k is sufficiently close to \bar{x} . We now let $Q(x) = I - P(x)$ and consider the term $Q(x_k)s_k$. Multiply (10) by $Q(x_k)$ to obtain

$$Q(x_k)\nabla m_k(s_k) = Q(x_k)\nabla f(x_k) + \nabla^2 f(x_k)Q(x_k)s_k + Q(x_k)(H_k - \nabla^2 f(x_k))[s_k] + \varsigma_k \|s_k\| Q(x_k)s_k.$$

Taking the inner product of this expression with $Q(x_k)s_k$, applying the Cauchy–Schwarz inequality, dividing by $\|s_k\|$, and using A3 yields

$$\varsigma_k \|Q(x_k)s_k\|^2 \leq \|Q(x_k)\nabla f(x_k)\| + \left(\kappa \|\nabla f(x_k)\| + \Lambda(x_k)\right) \|Q(x_k)s_k\|,$$

where Λ is as in Lemma 4.2. Solving the quadratic inequality gives

$$\|Q(x_k)s_k\| \leq \frac{1}{\varsigma_k} \left(\sqrt{\varsigma_k \|Q(x_k)\nabla f(x_k)\|} + \kappa \|\nabla f(x_k)\| + \Lambda(x_k) \right).$$

Local Lipschitz continuity of $\nabla^2 f$ provides $\Lambda(x_k) \leq \beta_H \|\nabla f(x_k)\| + L_H \text{dist}(x_k, \mathcal{S})$ when x_k is close to \bar{x} . Via Lemma 2.14, it also provides a constant $L_q \geq 0$ such that $\sqrt{\|Q(x_k)\nabla f(x_k)\|} \leq L_q \text{dist}(x_k, \mathcal{S})$ when x_k is sufficiently close to \bar{x} . This is enough to secure (12). \square

Using these bounds, we now show that $\|P(x_k)s_k\|$ cannot be too small compared to $\|s_k\|$ when x_k is close to a minimum where PL holds.

Lemma 4.4 *Suppose that A1, A3 and μ -(PL) hold around $\bar{x} \in \mathcal{S}$. Given $\varepsilon > 0$ and $\lambda^\sharp > \lambda_{\max}(\nabla^2 f(\bar{x}))$, there exists a neighborhood \mathcal{U} of \bar{x} and a constant $L_q \geq 0$ such that if x_k is an iterate in \mathcal{U} and s_k is a step satisfying (8) then $\|P(x_k)s_k\| \geq r \|s_k\|$ where $r > 0$ is the constant*

$$r = \sqrt{1 - \frac{1}{1 + \tilde{r}^2}} \quad \text{with} \quad \tilde{r} = \frac{(1 - \varepsilon)\varsigma_{\min}}{(\lambda^\sharp + \varepsilon(1 + \sqrt{\varsigma_{\min}}))(\kappa + \beta_H + \frac{L_H + L_q \sqrt{\varsigma_{\min}}}{\mu})}.$$

Proof Assume x_k is sufficiently close to \bar{x} for the projectors $P(x_k)$ and $Q(x_k)$ to be well defined. Define $v_p = \|P(x_k)s_k\|$, $v_q = \|Q(x_k)s_k\|$ and $\xi = \frac{v_p}{v_q}$ (consider $v_q \neq 0$ as otherwise the claim is clear). We compute that $v_p^2 = (1 - \frac{1}{1 + \xi^2}) \|s_k\|^2$ and find a lower-bound for ξ . Remark 2.10 gives that $\text{dist}(x_k, \mathcal{S}) \leq \frac{1}{\mu} \|\nabla f(x_k)\|$ when x_k is sufficiently close to \bar{x} . Together with the bound on v_q in Lemma 4.3, this gives

$$v_q \leq \frac{\|\nabla f(x_k)\|}{\varsigma_k} \left(\kappa + \beta_H + \frac{L_H + L_q \sqrt{\varsigma_k}}{\mu} \right)$$

Combining this with the lower-bound on v_p in Lemma 4.3, we find

$$\xi \geq \frac{(1 - \varepsilon)\varsigma_k}{(\lambda^\sharp + \varsigma_k \|s_k\|)(\kappa + \beta_H + \frac{L_H + L_q \sqrt{\varsigma_k}}{\mu})}.$$

With Lemma 4.2 we can upper-bound $\varsigma_k \|s_k\| \leq \sqrt{3\varsigma_k} \|\nabla f(x_k)\| + \frac{3}{2} \Lambda(x_k)$, and consequently, $\varsigma_k \|s_k\| \leq \varepsilon(1 + \sqrt{\varsigma_k})$ whenever x_k is sufficiently close to \bar{x} . We finally notice that the resulting lower-bound on ξ is an increasing function of ς_k . Therefore, we get the desired inequality by using $\varsigma_k \geq \varsigma_{\min}$. \square

From this we deduce a lower-bound on the quadratic term of the model.

Lemma 4.5 *Suppose that A1, A3 and μ -(PŁ) hold around $\bar{x} \in \mathcal{S}$. Given $\mu^b < \mu$, there exists a constant $r > 0$ (provided by Lemma 4.4) and a neighborhood \mathcal{U} of \bar{x} such that if $x_k \in \mathcal{U}$ then the step satisfies*

$$\langle s_k, \nabla^2 f(x_k)[s_k] \rangle \geq \mu^b r^2 \|s_k\|^2.$$

Proof Let $r > 0$ be a constant and \mathcal{U} a neighborhood of \bar{x} as in Lemma 4.4. Shrink \mathcal{U} for the projectors P and Q to be well defined in \mathcal{U} . If x_k is in \mathcal{U} we compute

$$\begin{aligned} \langle s_k, \nabla^2 f(x_k)[s_k] \rangle &= \langle P(x_k)s_k, \nabla^2 f(x_k)[P(x_k)s_k] \rangle + \langle Q(x_k)s_k, \nabla^2 f(x_k)[Q(x_k)s_k] \rangle \\ &\geq \lambda_d(\nabla^2 f(x_k))r^2 \|s_k\|^2 + \lambda_{\min}(\nabla^2 f(x_k)) \|Q(x_k)s_k\|^2, \end{aligned}$$

where we used $\|P(x_k)s_k\| \geq r \|s_k\|$. We obtain the result by noticing that the second term is lower-bounded by $\max(0, -\lambda_{\min}(\nabla^2 f(x_k))) \|s_k\|^2$. \square

We now show that the ratio ϱ_k is large when x_k is close to a local minimum where PŁ holds. The upshot is that iterations near \mathcal{S} are successful.

Lemma 4.6 *Suppose that A1, A2, A3 and μ -(PŁ) hold around $\bar{x} \in \mathcal{S}$. For all $\varepsilon > 0$ there exists a neighborhood \mathcal{U} of \bar{x} such that if $x_k \in \mathcal{U}$ then $\varrho_k \geq 1 - \varepsilon$.*

Proof Let $\tau_k(s) = f(x_k) + \langle s, \nabla f(x_k) \rangle + \frac{1}{2} \langle s, H_k[s] \rangle$ be the second-order component of the model m_k . Then,

$$1 - \varrho_k = \frac{f(\mathbf{R}_{x_k}(s_k)) - \tau_k(s_k)}{\tau_k(0) - \tau_k(s_k)} \leq \frac{1}{6} \frac{L'_H \|s_k\|^3 + 3\beta_H \|s_k\|^2 \|\nabla f(x_k)\|}{\tau_k(0) - \tau_k(s_k)},$$

where the bound on the numerator comes from A2 and A3. The denominator is given by

$$\begin{aligned} \tau_k(0) - \tau_k(s_k) &= -\langle s_k, \nabla f(x_k) \rangle - \frac{1}{2} \langle s_k, H_k[s_k] \rangle \\ &= -\langle s_k, \nabla m_k(s_k) \rangle + \frac{1}{2} \langle s_k, (H_k - \nabla^2 f(x_k))[s_k] \rangle + \frac{1}{2} \langle s_k, \nabla^2 f(x_k)[s_k] \rangle \\ &\quad + \varsigma_k \|s_k\|^3, \end{aligned}$$

where we used identity (10) for the second equality. With the Cauchy–Schwarz inequality, A3 and (8), we bound the two first terms as

$$\begin{aligned} |\langle s_k, \nabla m_k(s_k) \rangle| &\leq \kappa \|\nabla f(x_k)\| \|s_k\|^2 \text{ and} \\ |\langle s_k, (H_k - \nabla^2 f(x_k))[s_k] \rangle| &\leq \beta_H \|\nabla f(x_k)\| \|s_k\|^2. \end{aligned}$$

If we combine these bounds with Lemma 4.5, we find that given $\mu^b < \mu$, there exists $r > 0$ and a neighborhood \mathcal{U} of the local minimum \bar{x} such that

$$x_k \in \mathcal{U} \implies \tau_k(0) - \tau_k(s_k) \geq \mu^b r^2 \|s_k\|^2,$$

and therefore $1 - \varrho_k \leq \frac{L'_H \|s_k\| + 3\beta_H \|\nabla f(x_k)\|}{6\mu^b r^2}$. Owing to Lemma 4.2 the steps s_k vanish around second-order critical points so we can guarantee $1 - \varrho_k$ becomes as small as desired. \square

We can now show that the iterates produced by ARC satisfy the strong decrease property (4) around minima where PL holds.

Proposition 4.7 *Suppose that A1, A2, A3 and μ -(PL) hold around $\bar{x} \in \mathcal{S}$. Given $\mu^b < \mu$ and $\lambda^\sharp > \lambda_{\max}(\nabla^2 f(\bar{x}))$, there exists a neighborhood of \bar{x} where ARC satisfies the strong decrease property (4) with constant $\sigma = \frac{r\mu^b}{2c_r\lambda^\sharp}$, where c_r is defined in (RD) and $r > 0$ is provided by Lemma 4.4.*

Proof From Lemma 4.6, there exists a neighborhood \mathcal{U} of \bar{x} in which all the steps are successful. Given an iterate x_k in \mathcal{U} , success implies $x_{k+1} = R_{x_k}(s_k)$ and therefore, by definition of ϱ_k (9),

$$\begin{aligned} f(x_k) - f(x_{k+1}) &= \varrho_k \left(m_k(0) - m_k(s_k) + \frac{5k}{3} \|s_k\|^3 \right) = -\varrho_k \left(\langle s_k, \nabla f(x_k) \rangle + \frac{1}{2} \langle s_k, H_k[s_k] \rangle \right). \end{aligned}$$

Also, taking the inner product of (10) with s_k and using (8) yields

$$\langle s_k, H_k[s_k] \rangle \leq -\langle s_k, \nabla f(x_k) \rangle + \kappa \|\nabla f(x_k)\| \|s_k\|^2.$$

Multiply the latter by $-\frac{\varrho_k}{2}$ and plug into the former to deduce

$$f(x_k) - f(x_{k+1}) \geq -\frac{\varrho_k}{2} \left(\langle s_k, \nabla f(x_k) \rangle + \kappa \|\nabla f(x_k)\| \|s_k\|^2 \right). \tag{13}$$

We now bound the inner product $\langle s_k, \nabla f(x_k) \rangle$. Let $d = \text{rank } \nabla^2 f(\bar{x})$ and restrict the neighborhood \mathcal{U} if need be to ensure $\lambda_d(\nabla^2 f(x_k)) > 0$ and $\lambda_d(\nabla^2 f(x_k)) > \lambda_{d+1}(\nabla^2 f(x_k))$. In particular, the orthogonal projector $P(x_k)$ onto the top d eigenspace of $\nabla^2 f(x_k)$ is well defined. Let $Q(x_k) = I - P(x_k)$. Decompose $s_k = P(x_k)s_k + Q(x_k)s_k$ and apply the Cauchy–Schwarz inequality to obtain

$$\langle s_k, \nabla f(x_k) \rangle \leq \langle P(x_k)s_k, \nabla f(x_k) \rangle + \|Q(x_k)\nabla f(x_k)\| \|s_k\|. \tag{14}$$

The second term is small owing to Lemma 2.14. Let us focus on the first term. To this end, multiply (10) by $P(x_k)$ to verify the following (recall that $\nabla^2 f(x_k)$ and $P(x_k)$

commute):

$$P(x_k)\nabla f(x_k) = -\nabla^2 f(x_k)P(x_k)s_k + P(x_k)\nabla m_k(s_k) - P(x_k)(H_k - \nabla^2 f(x_k))s_k - \varsigma_k \|s_k\| P(x_k)s_k.$$

On the one hand, we can use it with **A3** and **(8)** to lower-bound the norm of $P(x_k)s_k$, through:

$$\|P(x_k)\nabla f(x_k)\| \leq (\lambda_1(\nabla^2 f(x_k)) + \varsigma_k \|s_k\|) \|P(x_k)s_k\| + (\kappa + \beta_H)\|\nabla f(x_k)\|\|s_k\|.$$

On the other hand, we can use it to upper-bound $\langle s_k, P(x_k)\nabla f(x_k) \rangle$, also with **A3** and **(8)** and using the fact that $P(x_k)s_k$ lives in the top- d eigenspace of $\nabla^2 f(x_k)$, like so:

$$\begin{aligned} \langle s_k, P(x_k)\nabla f(x_k) \rangle &\leq -\left(\lambda_d(\nabla^2 f(x_k)) + \varsigma_k \|s_k\|\right) \|P(x_k)s_k\|^2 + (\kappa + \beta_H)\|\nabla f(x_k)\|\|s_k\|\|P(x_k)s_k\|. \end{aligned}$$

Combine the two inequalities above as follows: use the former to upper-bound one of the first factors $\|P(x_k)s_k\|$ in the latter. Also using $\frac{\lambda_d(\nabla^2 f(x_k))}{\lambda_1(\nabla^2 f(x_k))} \leq \frac{\lambda_d(\nabla^2 f(x_k)) + \varsigma_k \|s_k\|}{\lambda_1(\nabla^2 f(x_k)) + \varsigma_k \|s_k\|} \leq 1$, this yields:

$$\begin{aligned} \langle s_k, P(x_k)\nabla f(x_k) \rangle &\leq -\frac{\lambda_d(\nabla^2 f(x_k))}{\lambda_1(\nabla^2 f(x_k))} \|P(x_k)\nabla f(x_k)\|\|P(x_k)s_k\| \\ &\quad + 2(\kappa + \beta_H)\|\nabla f(x_k)\|\|s_k\|\|P(x_k)s_k\|. \end{aligned} \tag{15}$$

We now plug this back into **(14)**. Using Lemma 2.14 for its second term and also Lemma 4.2 which asserts $\|s_k\|$ is arbitrarily small for x_k near \bar{x} , we find that for all $\varepsilon > 0$ we can restrict the neighborhood \mathcal{U} in order to secure

$$\langle s_k, \nabla f(x_k) \rangle \leq -\frac{\lambda_d(\nabla^2 f(x_k))}{\lambda_1(\nabla^2 f(x_k))} \|P(x_k)\nabla f(x_k)\|\|P(x_k)s_k\| + \varepsilon\|\nabla f(x_k)\|\|s_k\|. \tag{16}$$

Lemma 4.4 provides a (possibly smaller) neighborhood and a positive r such that $\|P(x_k)s_k\| \geq r\|s_k\|$. Also, for any $\delta > 0$, Lemma 2.14 ensures $\|P(x_k)\nabla f(x_k)\| \geq (1 - \delta)\|\nabla f(x_k)\|$ upon appropriate neighborhood restriction. Thus,

$$\langle s_k, \nabla f(x_k) \rangle \leq -\left(\frac{\lambda_d(\nabla^2 f(x_k))}{\lambda_1(\nabla^2 f(x_k))}(1 - \delta)r - \varepsilon\right) \|\nabla f(x_k)\|\|s_k\|. \tag{17}$$

Now plugging back into (13) and possibly restricting the neighborhood again,

$$f(x_k) - f(x_{k+1}) \geq \frac{\varrho_k}{2} \left(\frac{\lambda_d(\nabla^2 f(x_k))}{\lambda_1(\nabla^2 f(x_k))} (1 - \delta)r - \varepsilon - \kappa \|s_k\| \right) \|\nabla f(x_k)\| \|s_k\|. \tag{18}$$

The result now follows since we can arrange to make ε , δ and $\|s_k\|$ arbitrarily small; to make $\lambda_d(\nabla^2 f(x_k))$ and $\lambda_1(\nabla^2 f(x_k))$ arbitrarily close to μ and $\lambda_{\max}(\nabla^2 f(\bar{x}))$ (respectively); and to make ϱ_k larger than a number arbitrarily close to 1 (Lemma 4.6). The final step is to account for potential distortion due to a retraction (RD): this adds the factor c_r . □

If we combine this result with Lemma 3.8 we obtain that ARC satisfies the bounded path length property (BPL). In Lemma 4.2 we found that it also satisfies the vanishing steps property (VS). Moreover, if the iterates of ARC stay in a compact region then they accumulate only at critical points [27, Cor. 2.6]. As a result, Corollary 3.7 applies to ARC: if an iterate gets close enough to a point where PL holds then the sequence has a limit. We conclude this section with the quadratic convergence rate of ARC.

Proposition 4.8 *Suppose that $\{x_k\}$ converges to some $\bar{x} \in \mathcal{S}$ and that f is (PL) around \bar{x} . Also assume that A1, A2 and A3 hold around \bar{x} . Then $\{\text{dist}(x_k, \mathcal{S})\}$ converges quadratically to zero.*

Proof From Lemma 4.6 all the steps are eventually successful. In particular, the penalty weights eventually stop increasing: there exists $\varsigma_{\max} > 0$ such that $\varsigma_k \leq \varsigma_{\max}$ for all k . Let $\mu^b < \mu$ (where μ is the PL constant) and $\lambda^\sharp > \lambda_{\max}(\nabla^2 f(\bar{x}))$. We first apply the Pythagorean theorem with the upper-bounds from Lemma 4.3. Together with the upper-bound in Proposition 2.8, it implies for all large enough k that

$$\|s_k\|^2 \leq c_1^2 \text{dist}(x_k, \mathcal{S})^2 \text{ where } c_1^2 = \left(\frac{\lambda^\sharp}{\mu^b} \right)^2 + \frac{1}{\varsigma_{\min}^2} \left((\kappa + \beta_H)\lambda^\sharp + L_H + L_q\sqrt{\varsigma_{\min}} \right)^2. \tag{19}$$

We now let $v_k = s_k - \text{Log}_{x_k}(x_{k+1})$, which is always well defined for large enough k . Using EB (given by Remark 2.10), we obtain that for all large enough k we have

$$\begin{aligned} \text{dist}(x_{k+1}, \mathcal{S}) &\leq \frac{1}{\mu} \|\nabla f(x_{k+1})\| \\ &= \frac{1}{\mu} \left\| \Gamma_{x_{k+1}}^{x_k} \nabla f(x_{k+1}) - \nabla f(x_k) - \nabla^2 f(x_k)[\text{Log}_{x_k}(x_{k+1})] \right. \\ &\quad \left. - \nabla^2 f(x_k)[v_k] - (H_k - \nabla^2 f(x_k))[s_k] - \varsigma_k \|s_k\| s_k + \nabla m_k(s_k) \right\|, \end{aligned}$$

where we used identity (10) for ∇m_k . Now the triangle inequality, A3 and (8) give

$$\text{dist}(x_{k+1}, \mathcal{S}) \leq \frac{1}{\mu} \left(\frac{L_H}{2} \text{dist}(x_k, x_{k+1})^2 + \lambda^\sharp \|v_k\| + \varsigma_k \|s_k\|^2 + (\kappa + \beta_H) \|\nabla f(x_k)\| \|s_k\| \right). \tag{20}$$

Notice that $\text{dist}(x_k, x_{k+1}) \leq c_r \|s_k\|$ using (RD). We now bound the quantity $\|v_k\|$. For all $x \in \mathcal{M}$, since $\text{DR}_x(0) = I$, the inverse function theorem implies that R_x is locally invertible and $\text{DR}_x^{-1}(x) = I$. It follows that there exists a neighborhood $\mathcal{U} \subseteq \text{B}(\bar{x}, \text{inj}(\bar{x}))$ of \bar{x} such that for all $x, y \in \mathcal{U}$ the quantity $R_x^{-1}(y)$ is well defined and satisfies

$$R_x^{-1}(y) = R_x^{-1}(x) + \text{DR}_x^{-1}(x)[\text{Log}_x(y)] + O(\text{dist}(x, y)^2) = \text{Log}_x(y) + O(\text{dist}(x, y)^2).$$

In particular, using the identity $s_k = R_{x_k}^{-1}(x_{k+1})$, we find that there exists a constant c_2 such that $\|v_k\| \leq c_2 \text{dist}(x_k, x_{k+1})^2$ holds for large enough k . Combining this with (20), we obtain

$$\text{dist}(x_{k+1}, \mathcal{S}) \leq \frac{1}{\mu} \left(\left(\frac{L_H}{2} c_r^2 + \lambda^\sharp c_2 c_r^2 + \varsigma_k \right) \|s_k\|^2 + (\kappa + \beta_H) \|\nabla f(x_k)\| \|s_k\| \right).$$

Finally, using (19) and the upper-bound from Proposition 2.8, we conclude that

$$\begin{aligned} \text{dist}(x_{k+1}, \mathcal{S}) &\leq c_q \text{dist}(x_k, \mathcal{S})^2 \quad \text{where} \quad c_q = \frac{c_1^2}{\mu} \left(\frac{L_H}{2} c_r^2 + \lambda^\sharp c_2 c_r^2 + \varsigma_{\max} \right) \\ &\quad + \frac{\lambda^\sharp}{\mu} (\kappa + \beta_H) c_1. \end{aligned} \quad \square$$

The quadratic convergence rates of the sequences $\{f(x_k)\}$ and $\{\|\nabla f(x_k)\|\}$ follow immediately by QG and EB. Theorem 4.1 is a direct consequence of Corollary 3.7 and Proposition 4.8.

4.2 Trust-region algorithms

In this section we analyze Riemannian trust-region algorithms (TR), which embed Newton iterations in safeguards to ensure global convergence guarantees [3]. They produce sequences $\{(x_k, \Delta_k)\}$, where x_k is the current iterate and Δ_k is the trust-region radius. At iteration k , we define the trust-region model as

$$m_k(s) = f(x_k) + \langle s, \nabla f(x_k) \rangle + \frac{1}{2} \langle s, H_k[s] \rangle, \tag{21}$$

where $H_k: T_{x_k} \mathcal{M} \rightarrow T_{x_k} \mathcal{M}$ is a linear map close to $\nabla^2 f(x_k)$, satisfying A3. The step s_k is chosen by (usually approximately) solving the trust-region subproblem

$$\min_{s_k \in T_{x_k} \mathcal{M}} m_k(s_k) \quad \text{subject to} \quad \|s_k\| \leq \Delta_k. \tag{TRS}$$

The point x_k and radius Δ_k are then updated depending on how good the model is, as measured by the ratio

$$\rho_k = \frac{f(x_k) - f(\mathbf{R}_{x_k}(s_k))}{m_k(0) - m_k(s_k)}.$$

(If the denominator is zero, we let $\rho_k = 1$.) Specifically, given parameters $\rho' \in]0, \frac{1}{4}[$ and $\bar{\Delta} > 0$, the update rules for the state are

$$x_{k+1} = \begin{cases} \mathbf{R}_{x_k}(s_k) & \text{if } \rho_k > \rho', \\ x_k & \text{otherwise,} \end{cases} \quad \Delta_{k+1} = \begin{cases} \frac{1}{4}\Delta_k & \text{if } \rho_k < \frac{1}{4}, \\ \min(2\Delta_k, \bar{\Delta}) & \text{if } \rho_k > \frac{3}{4} \text{ and } \|s_k\| = \Delta_k, \\ \Delta_k & \text{otherwise.} \end{cases}$$

Shortcomings of trust-region with exact subproblem solver. There exist algorithms to efficiently solve the subproblem exactly (up to some accuracy). Can we provide strong guarantees in the presence of non-isolated minima using an exact subproblem solver? Assume for simplicity that $H_k = \nabla^2 f(x_k)$. We recall [77, Thm. 4.1] that a vector $s \in \mathbf{T}_{x_k}\mathcal{M}$ is a global solution of the subproblem (TRS) if and only if $\|s\| \leq \Delta_k$ and there exists a scalar $\lambda \geq 0$ such that

$$(\nabla^2 f(x_k) + \lambda I)s = -\nabla f(x_k), \quad \lambda(\Delta_k - \|s\|) = 0 \quad \text{and} \quad \nabla^2 f(x_k) + \lambda I \geq 0. \tag{22}$$

As mentioned in [66, 94, 95], if f is convex in a neighborhood of \bar{x} then \mathcal{S} is locally convex, hence affine. Assuming that \mathcal{S} is not flat around \bar{x} , it follows that $\nabla^2 f$ must have a negative eigenvalue in any neighborhood of \bar{x} . Consider an iterate x_k close to \bar{x} , and for which $\nabla^2 f(x_k)$ has a negative eigenvalue. Conditions (22) imply $\lambda > 0$ and hence $\|s\| = \Delta_k$, meaning that s is at the border of the trust region. Consequently, even if x_0 is arbitrarily close to \bar{x} , we can arrange for the next iterate to be far away (if the radius Δ_0 is large). This shows that capture results such as Corollary 3.7 fail for this algorithm.

As an example, define $f: \mathbb{R}^2 \rightarrow \mathbb{R}$ as $f(x, y) = (x^2 + y^2 - 1)^2$. The set of minima is the unit circle, on which f satisfies MB. Given $t > 0$, define $x(t) = 0$ and $y(t) = 1 - t$. Consider the subproblem (TRS) at $(x(t), y(t))^\top$ with exact Hessian $\nabla^2 f$ and with parameter Δ . For small enough t there are two solutions:

$$s = \left[\begin{array}{c} \pm \sqrt{\Delta^2 - \frac{t^2(t-2)^2}{4(t-1)^2}} \\ \frac{t(t-2)}{2(t-1)} \end{array} \right].$$

The criterion (22) with $\lambda = 4t(2-t)$ confirms optimality. We find that $s \rightarrow (\pm\Delta, 0)^\top$ as $t \rightarrow 0$, so the tentative step is far even when t is small. We could arrange for that step to be accepted by adjusting the function value at the tentative iterate. That rules out even basic capture-type theorems. This type of behavior does not happen when the Hessian is positive definite at the minimum.

Trust-region with Cauchy steps. As just discussed, TR with an exact subproblem solver can fail in the face of non-isolated minima. However, practical implementations of TR typically solve the subproblem only approximately. We set out to investigate the robustness of such mechanisms to non-isolated minima.

Our investigation is prompted by the empirical observation that TR with a popular approximate subproblem solver known as truncated conjugate gradient (tCG, see [3, 33]) seems to enjoy superlinear convergence under PŁ, even with non-isolated minima. We confirmed this subsequently [85] using significant additional machinery.

As a more direct illustration, we show the following theorem, regarding TR with a crude subproblem solver that computes *Cauchy steps* (see (25) below). It is relevant in particular because tCG generates a sequence of increasingly good tentative steps, the first of which is the Cauchy step.

Theorem 4.9 *Suppose A2, A3 and μ -(PŁ) hold around $\bar{x} \in \mathcal{S}$. Let \mathcal{U} be a neighborhood of \bar{x} . There exists a neighborhood \mathcal{V} of \bar{x} such that if a sequence of iterates generated by TR with Cauchy steps enters \mathcal{V} then the sequence converges linearly to some $x_\infty \in \mathcal{U} \cap \mathcal{S}$ with rate $\sqrt{1 - \frac{\mu}{\lambda_{\max}}}$, where $\lambda_{\max} = \lambda_{\max}(\nabla^2 f(x_\infty))$.*

A local convergence analysis of TR with Cauchy steps is given in [73] for *non-singular* local minima. Here, we prove that the favorable convergence properties also hold if we only assume PŁ.

To prove Theorem 4.9, we first establish a number of intermediate results only assuming the subproblem solver satisfies the properties (23) and (24) defined below. We then secure these properties for Cauchy steps. First, given a local minimum $\bar{x} \in \mathcal{S}$, we assume that the step s_k satisfies the *sufficient decrease* condition

$$m_k(0) - m_k(s_k) \geq c_p \|\nabla f(x_k)\| \min\left(\Delta_k, \frac{\|\nabla f(x_k)\|^3}{|\langle \nabla f(x_k), H_k[\nabla f(x_k)] \rangle|}\right) \tag{23}$$

whenever the iterate x_k is sufficiently close to \bar{x} . (If the denominator is zero, consider the rightmost expression to be infinite.) This condition holds for many practical subproblem solvers and ensures global convergence guarantees in particular (see [4, §7.4] and [24, §6.4]). Second, given a local minimum $\bar{x} \in \mathcal{S}$, we assume that there exists a constant $c_s \geq 0$ such that

$$\|s_k\| \leq c_s \|\nabla f(x_k)\| \tag{24}$$

when x_k is sufficiently close to \bar{x} .

We find that the ratios $\{\rho_k\}$ are large around minima where these two conditions hold. This is because they imply that the trust-region model is an accurate approximation of the local behavior of f . It follows that the steps $\{s_k\}$ decrease f nearly as much as predicted by the model.

Proposition 4.10 *Suppose that A2 and A3 hold around $\bar{x} \in \mathcal{S}$. Also assume that the steps s_k satisfy (23) and (24) around \bar{x} . For all $\varepsilon > 0$ there exists a neighborhood \mathcal{U} of \bar{x} such that if an iterate x_k is in \mathcal{U} then $\rho_k \geq 1 - \varepsilon$.*

Proof We follow and adapt some arguments from [4, Thm. 7.4.11], which is stated there assuming $\nabla^2 f(\bar{x}) \succ 0$. We can dismiss the case where $\nabla f(x_k) = 0$ because it implies $\rho_k = 1$. Using the definitions of m_k and ρ_k we have

$$1 - \rho_k = \frac{f(\mathbf{R}_{x_k}(s_k)) - m_k(s_k)}{m_k(0) - m_k(s_k)}.$$

Assuming x_k is sufficiently close to \bar{x} , we bound the numerator as $f(\mathbf{R}_{x_k}(s_k)) - m_k(s_k) \leq \frac{L'_H}{6} \|s_k\|^3 \leq \frac{L'_H c_s^2}{6} \|s_k\| \|\nabla f(x_k)\|^2$ using A2 and inequality (24). Combining this with the sufficient decrease (23) and A3 gives

$$1 - \rho_k \leq \frac{L'_H c_s^2 \|s_k\| \|\nabla f(x_k)\|}{6c_p \min\left(\Delta_k, \frac{\|\nabla f(x_k)\|}{\|\nabla^2 f(x_k)\| + \beta_H \|\nabla f(x_k)\|}\right)}.$$

If Δ_k is active in the denominator then we obtain $1 - \rho_k \leq \frac{L'_H c_s^2}{6c_p} \|\nabla f(x_k)\|$ because $\|s_k\| \leq \Delta_k$. Otherwise, using (24) we obtain $1 - \rho_k \leq \frac{L'_H c_s^3}{6c_p} \|\nabla f(x_k)\| (\|\nabla^2 f(x_k)\| + \beta_H \|\nabla f(x_k)\|)$. In both cases this yields the result. \square

This result notably implies that the trust-region radius does not decrease in the vicinity of the minimum \bar{x} . It means that the trust region eventually becomes inactive when the iterates converge to \bar{x} . We now employ the particular alignment between the gradient and the top eigenspace of the Hessian induced by PL (see Lemma 2.14) to derive bounds on the inner products $\langle \nabla f(x), \nabla^2 f(x)[\nabla f(x)] \rangle$.

Proposition 4.11 *Suppose that f is μ -(PL) around $\bar{x} \in \mathcal{S}$. Let $\mu^b < \mu$ and $\lambda^\sharp > \lambda_{\max}(\nabla^2 f(\bar{x}))$. Then there exists a neighborhood \mathcal{U} of \bar{x} such that for all $x \in \mathcal{U}$ we have*

$$\mu^b \|\nabla f(x)\|^2 \leq \langle \nabla f(x), \nabla^2 f(x)[\nabla f(x)] \rangle \leq \lambda^\sharp \|\nabla f(x)\|^2.$$

Proof By continuity of eigenvalues, there exists a neighborhood \mathcal{U} of \bar{x} such that for all $x \in \mathcal{U}$ we have $\lambda_{\max}(\nabla^2 f(x)) \leq \lambda^\sharp$. The upper-bound $\langle \nabla f(x), \nabla^2 f(x)[\nabla f(x)] \rangle \leq \lambda^\sharp \|\nabla f(x)\|^2$ follows immediately. We now prove the lower-bound. Let d be the rank of $\nabla^2 f(\bar{x})$. Given x sufficiently close to \bar{x} , we let $P(x) : \mathbf{T}_x \mathcal{M} \rightarrow \mathbf{T}_x \mathcal{M}$ denote the orthogonal projector onto the top d eigenspace of $\nabla^2 f(x)$. From Lemma 2.14 we have $\|(I - P(x))\nabla f(x)\|^2 = o(\|\nabla f(x)\|^2)$ as $x \rightarrow \bar{x}$. If we write $Q(x) = I - P(x)$, we obtain

$$\begin{aligned} & \langle \nabla f(x), \nabla^2 f(x)[\nabla f(x)] \rangle \\ &= \langle P(x)\nabla f(x), \nabla^2 f(x)P(x)\nabla f(x) \rangle + \langle Q(x)\nabla f(x), \nabla^2 f(x)Q(x)\nabla f(x) \rangle \\ &\geq \lambda_d (\nabla^2 f(x)) \|\nabla f(x)\|^2 + o(\|\nabla f(x)\|^2) \end{aligned}$$

as $x \rightarrow \bar{x}$. \square

Combining these results guarantees a linear rate of convergence.

Proposition 4.12 *Suppose that A2, A3 and μ -(PŁ) hold around $\bar{x} \in \mathcal{S}$. Let $\{x_k\}$ be a sequence of iterates produced by TR converging to \bar{x} . Assume that the steps s_k satisfy (23) and (24) around \bar{x} . Then the iterates converge at least linearly with rate $\sqrt{1 - \frac{2c_p\mu}{\lambda_{\max}}}$, where $\lambda_{\max} = \lambda_{\max}(\nabla^2 f(\bar{x}))$.*

Proof We can assume that $\nabla f(x_k)$ is non-zero for all k (otherwise the sequence converges in a finite number of steps). We show that the sequence satisfies the sufficient decrease property (5). Given $\mu^b < \mu$ and $\lambda^\sharp > \lambda_{\max}(\nabla^2 f(\bar{x}))$, Proposition 4.11 and A3 ensure that

$$\frac{1}{\lambda^\sharp} \leq \frac{\|\nabla f(x_k)\|^2}{\langle \nabla f(x_k), H_k[\nabla f(x_k)] \rangle} \leq \frac{1}{\mu^b}$$

for all large enough k . We let $0 < \varepsilon < \frac{3}{4}$ and Proposition 4.10 implies that $\rho_k \geq 1 - \varepsilon$ for all large enough k . In particular, the radii $\{\Delta_k\}$ are bounded away from zero (because the update mechanism does not decrease the radius when $\rho_k \geq \frac{1}{4}$). Combining the definition of ρ_k and the sufficient decrease (23) gives

$$f(x_k) - f(x_{k+1}) = \rho_k(m_k(0) - m_k(s_k)) \geq \frac{(1 - \varepsilon)c_p}{\lambda^\sharp} \|\nabla f(x_k)\|^2$$

for all large enough k . We can now conclude with Proposition 3.9. □

We are now in a position to prove Theorem 4.9. The Cauchy step at iteration x_k is defined as the minimum of (TRS) with the additional constraint that $s_k \in \text{span}(\nabla f(x_k))$. We can find an explicit expression for it: when $\nabla f(x_k) \neq 0$, the Cauchy step is $s_k^c = -t_k^c \nabla f(x_k)$, where

$$t_k^c = \begin{cases} \min\left(\frac{\|\nabla f(x_k)\|^2}{\langle \nabla f(x_k), H_k[\nabla f(x_k)] \rangle}, \frac{\Delta_k}{\|\nabla f(x_k)\|}\right) & \text{if } \langle \nabla f(x_k), H_k[\nabla f(x_k)] \rangle > 0, \\ \frac{\Delta_k}{\|\nabla f(x_k)\|} & \text{otherwise.} \end{cases} \tag{25}$$

Cauchy steps notably satisfy the sufficient decrease property (23) globally with $c_p = \frac{1}{2}$ (see [33, Thm. 6.3.1]). We now prove that they also satisfy (24) around minima where PŁ holds.

Proposition 4.13 *Suppose that A3 and μ -(PŁ) hold around $\bar{x} \in \mathcal{S}$. Given $\mu^b < \mu$, there exists a neighborhood \mathcal{U} of \bar{x} such that if an iterate x_k is in \mathcal{U} then the Cauchy step satisfies $\|s_k^c\| \leq \frac{1}{\mu^b} \|\nabla f(x_k)\|$.*

Proof Given $\mu^b < \mu$, Proposition 4.11 and A3 yield that $\langle \nabla f(x_k), H_k[\nabla f(x_k)] \rangle \geq \mu^b \|\nabla f(x_k)\|^2$ if x_k is sufficiently close to \bar{x} . It implies that the step-sizes defined in (25) are bounded as $t_k^c \leq \frac{1}{\mu^b}$, which gives the result. □

In particular, this proposition shows that TR with Cauchy steps satisfies the (VS) property at \bar{x} with $\eta(x) = \frac{c_r}{\mu^b} \|\nabla f(x)\|$, where c_r is as in (RD). Furthermore, Cauchy steps satisfy the model decrease

$$m_k(0) - m_k(s_k^c) \geq \frac{1}{2} \|\nabla f(x_k)\| \|s_k^c\|,$$

as shown in [2, Lem. 4.3]. It implies that TR with Cauchy steps generates sequences that satisfy the strong decrease property (4) with $\sigma = \frac{\rho'}{2c_r}$, where c_r is as in (RD), and ρ' is defined in the algorithm description in Sect. 4.2. See [2, Thm. 4.4] for details on this. As a consequence, TR with Cauchy steps satisfies the bounded path length property (BPL) at points where a Łojasiewicz inequality holds (Lemma 3.8). Moreover, if the iterates of this algorithm stay in a compact region then they accumulate only at critical points [4, Thm. 7.4.4]. We can finally combine the statements from Corollary 3.7 and Proposition 4.12 to obtain Theorem 4.9.

Remark 4.14 The model decrease (23) is not a sufficient condition for the strong decrease property (4) to hold. As a result, it is not straightforward to determine whether the bounded path length property (BPL) holds for a given subproblem solver: see [2, §4.2] for a discussion of this.

5 Conclusions and perspectives

We showed the (local) equivalence (up to arbitrarily small losses in constants) of MB, PL, EB and QG (Sect. 2). We then revisited classical capture results compatible with non-isolated minima to factor out the roles of vanishing step-sizes and bounded path lengths (Sect. 3). The MB property and the alignment of the gradient with respect to the Hessian eigenspaces (Lemma 2.14) are particularly adapted to analyze second-order algorithms. Accordingly, assuming the above conditions we establish quadratic convergence for ARC with inexact subproblem solvers and linear convergence for TR with Cauchy steps (Sect. 4).

We conclude with a few research directions:

- In Sect. 4.2 we analyze a simple subproblem solver (Cauchy steps) for TR. It is natural to explore more advanced subproblem solvers. For example in [85] we show superlinear convergence for a truncated CG method assuming MB.
- In Sect. 4.2 we argue that TR with an *exact* subproblem solver cannot satisfy a standard capture property in the presence of non-isolated minima. However, establishing capture is only a means to an end. It may still be possible to obtain other satisfactory guarantees.
- More generally, using the tools from Sects. 2 and 3, there is an opportunity to revisit analyses of other algorithms that currently require strong local convexity. For example, Goyens and Royer [47] control the global complexity of hybrid TR algorithms for strict saddle functions, and currently require non-singular minima.
- Likewise, Remark 2.20 illustrates equivalence of MB with the restricted secant inequality (RSI). Showing MB implies RSI is direct. The converse is facilitated

via equivalence of MB with EB. There may be other local properties in the literature that turn out to be equivalent to these.

A Łojasiewicz and function growth

We review here the classical arguments at the basis of Proposition 2.2. Given an initial point $x_0 \in \mathcal{M}$, we let $x : I \rightarrow \mathcal{M}$ denote a solution of the negative gradient flow

$$x'(t) = -\nabla f(x(t)) \quad \text{with} \quad x(0) = x_0 \tag{GF}$$

on the maximum interval I . The following is classical [69], restated succinctly to highlight neighborhood assumptions.

Lemma A.1 *Suppose f satisfies (Ł) with constants $\theta \in [0, 1[$ and $\mu > 0$ in a neighborhood \mathcal{U} of $\bar{x} \in \mathcal{S}$. Also suppose that $f(x) > f_{\mathcal{S}}$ for all $x \in \mathcal{U} \setminus \mathcal{S}$. Let x be a solution to (GF) for some $x_0 \in \mathcal{U} \setminus \mathcal{S}$. Suppose that for all $t \in]0, T[$ we have $x(t) \in \mathcal{U}$ and $\nabla f(x(t)) \neq 0$. Then the path length is bounded as*

$$\int_0^T \|x'(t)\| dt \leq \frac{1}{(1-\theta)\sqrt{2\mu}} |f(x_0) - f_{\mathcal{S}}|^{1-\theta}.$$

Proof Define $h(t) = \frac{1}{(1-\theta)\sqrt{2\mu}} |f(x(t)) - f_{\mathcal{S}}|^{1-\theta}$. For all $t \in [0, T[$, the Łojasiewicz inequality provides:

$$\|x'(t)\| = \|\nabla f(x(t))\| = \frac{\|\nabla f(x(t))\|^2}{\|\nabla f(x(t))\|} \leq \frac{\|\nabla f(x(t))\|^2}{\sqrt{2\mu}(f(x(t)) - f_{\mathcal{S}})^\theta} = -h'(t).$$

It follows that $\int_0^T \|x'(t)\| dt \leq \int_0^T -h'(t) dt = h(0) - h(T) \leq h(0)$. □

As shown below, this bound implies that the trajectories are trapped and have a limit point if x_0 is close enough to \mathcal{S} .

Proposition A.2 (Lyapunov stability) *Suppose that f satisfies (Ł) around $\bar{x} \in \mathcal{S}$ and let \mathcal{U} be a neighborhood of \bar{x} . There exists a neighborhood \mathcal{V} of \bar{x} such that if $x_0 \in \mathcal{V}$ then the solution x to (GF) is defined on $[0, +\infty[$, and $x(t) \in \mathcal{U}$ for all $t \geq 0$.*

Proof The set \mathcal{U} contains a ball \mathcal{B} centered on \bar{x} of radius δ_1 such that (i) f satisfies (Ł) with constants θ and μ in \mathcal{B} , and (ii) $f(y) > f_{\mathcal{S}}$ for all $y \in \mathcal{B} \setminus \mathcal{S}$. By continuity of f there exists an open ball $\mathcal{V} \subseteq \mathcal{B}$ of radius δ_2 around \bar{x} such that for all $y \in \mathcal{V}$ we have

$$\frac{1}{(1-\theta)\sqrt{2\mu}} |f(y) - f_{\mathcal{S}}|^{1-\theta} + \delta_2 < \delta_1.$$

Given $x_0 \in \mathcal{V}$, let $x : I \rightarrow \mathcal{M}$ be the maximal solution to (GF). Suppose that $\{t \in I : t \geq 0 \text{ and } x(t) \notin \mathcal{B}\}$ is non-empty and let T be the infimum of this set. Then $x(t) \in \mathcal{B}$

for all $0 \leq t < T$. Suppose first that there exists $t \in [0, T[$ such that $\nabla f(x(t)) = 0$. Then $x(t') = x(t)$ for all $t' > t$, which is impossible. So $\nabla f(x(t)) \neq 0$ for all $t \in [0, T[$. It follows that the assumptions of Lemma A.1 are satisfied, and the path length is bounded as

$$\int_0^T \|x'(t)\| dt \leq \frac{1}{(1 - \theta)\sqrt{2\mu}} |f(x_0) - f_S|^{1-\theta}.$$

This implies that

$$\text{dist}(x(T), \bar{x}) \leq \int_0^T \|x'(t)\| dt + \text{dist}(x_0, \bar{x}) < \delta_1,$$

and hence that $x(T) \in \mathcal{B}$. This is a contradiction and we deduce that $x(t) \in \mathcal{B}$ for all $t \in I$. Therefore, the total path length of x is bounded, and the escape lemma [58, Lem. A.43] implies that x is defined for all $t \geq 0$. \square

Corollary A.3 *Suppose that f satisfies (\mathbf{L}) around $\bar{x} \in \mathcal{S}$. There exists a neighborhood \mathcal{V} of \bar{x} such that for all $x_0 \in \mathcal{V}$ the solution to (\mathbf{GF}) is defined on $[0, +\infty[$ and has a limit in \mathcal{S} .*

Proof Let \mathcal{U} be a compact neighborhood of \bar{x} such that f satisfies (\mathbf{L}) around all points in $\mathcal{U} \cap \mathcal{S}$, and such that all critical points of f in \mathcal{U} are in \mathcal{S} . Let \mathcal{V} be a neighborhood of \bar{x} associated to \mathcal{U} as in Proposition A.2. Let $x : [0, +\infty[\rightarrow \mathcal{M}$ denote a solution of (\mathbf{GF}) starting from $x_0 \in \mathcal{V}$. Then $x(t) \in \mathcal{U}$ for all $t \geq 0$. The set \mathcal{U} is compact so x has an accumulation point $x_\infty \in \mathcal{U}$. This is a critical point for f so $x_\infty \in \mathcal{S}$. It is also Lyapunov stable because f satisfies (\mathbf{L}) around x_∞ (Proposition A.2). We deduce that $\lim x(t) = x_\infty$. \square

From this we deduce that the Łojasiewicz inequality implies the local growth of f announced in Proposition 2.2.

Proof of Proposition 2.2 Let \mathcal{U} be a neighborhood of \bar{x} where (\mathbf{L}) holds with constants μ and θ . Proposition A.2 and Corollary A.3 give a neighborhood \mathcal{V} of \bar{x} such that for all $x_0 \in \mathcal{V}$ the solution to (\mathbf{GF}) is defined on $[0, \infty[$, stays in \mathcal{U} at all times, and has a limit $x_\infty \in \mathcal{S}$. Then, Lemma A.1 provides the first inequality in

$$\frac{1}{(1 - \theta)\sqrt{2\mu}} |f(x_0) - f_S|^{1-\theta} \geq \int_0^\infty \|x'(t)\| dt \geq \text{dist}(x_0, x_\infty) \geq \text{dist}(x_0, \mathcal{S}),$$

which concludes the proof. \square

B Other Łojasiewicz exponents

We prove here the statements of Remark 2.21. From Proposition 2.2, there exists $c > 0$ such that

$$f(x) - f_S \geq c \text{dist}(x, \mathcal{S})^{\frac{1}{1-\theta}} \tag{26}$$

for all x sufficiently close to \bar{x} .

Assume first that f is C^1 and ∇f is L -Lipschitz continuous around \bar{x} . Then $f(x) - f_S \leq \frac{L}{2} \text{dist}(x, \mathcal{S})^2$ for all x sufficiently close to \bar{x} (see for example [24, Cor. 10.54]). When $\theta < \frac{1}{2}$, this inequality is incompatible with (26) if f is non-constant around \bar{x} .

Now assume that f is C^2 and $\nabla^2 f$ is L -Lipschitz continuous around \bar{x} . Define $h : (y, v) \mapsto \langle v, \nabla^2 f(y)[v] \rangle$. Lipschitz continuity of $\nabla^2 f$ gives that [24, Cor. 10.56]

$$f(\text{Exp}_y(tv)) - f_S - \frac{t^2}{2}h(y, v) \leq \frac{L}{6}t^3$$

for all $y \in \mathcal{S}$ close enough to \bar{x} , all unitary $v \in N_y\mathcal{S}$ and $t > 0$ small enough. Take $t \rightarrow 0$ and invoke Lemma 2.6 to see that $h(y, v)$ must be positive for this inequality to be compatible with the function growth (26) when $\theta \in [\frac{1}{2}, \frac{2}{3}[$. We conclude that there is a compact neighborhood \mathcal{V} of \bar{x} such that $h(y, v) > 0$ for all $y \in \mathcal{V} \cap \mathcal{S}$ and $v \in N_y\mathcal{S}$ unitary. The function h is continuous and the set $\mathcal{D} = \{(y, v) : y \in \mathcal{V} \cap \mathcal{S}, v \in N_y\mathcal{S} \text{ unitary}\}$ is compact so $\mu = \inf_{(y,v) \in \mathcal{D}} h(y, v)$ is positive. Now let \mathcal{W} be a neighborhood of \bar{x} as in Lemma 1.5 such that for all $x \in \mathcal{W}$ the projection $\text{proj}_{\mathcal{S}}(x)$ is included in $\mathcal{V} \cap \mathcal{S}$. Then for all $x \in \mathcal{W}$ we have

$$f(x) - f_S = h(y, v) + o(\|v\|^2) \geq \mu \text{dist}(x, \mathcal{S})^2 + o(\text{dist}(x, \mathcal{S})^2),$$

where $y \in \text{proj}_{\mathcal{S}}(x)$ and $v = \text{Log}_y(x)$. We conclude that QG holds around \bar{x} , which implies PŁ, as shown in Sect. 2.2.

Now assume that f is C^3 . The argument is similar in this case. Let $y \in \mathcal{S}$ and $v \in N_y\mathcal{S}$ unitary and assume that $h(y, v) = 0$. Then a Taylor expansion gives

$$f(\text{Exp}_x(tv)) - f_S = o(t^3)$$

because the third-order term vanishes. This is incompatible with the function growth (26) when y is close to \bar{x} and $\theta \in [\frac{1}{2}, \frac{2}{3}[$. So $h(y, v)$ must be positive and we conclude with the same arguments as above.

C Łojasiewicz and bounded path length

Proof of Lemma 3.8 Let \mathcal{U} be an open neighborhood of \bar{x} such that (i) f satisfies (Ł) with constants $\theta \in [0, 1[$ and $\mu > 0$ in \mathcal{U} , (ii) $f(x) \geq f_S$ for all $x \in \mathcal{U}$ and (iii) condition (4) holds in \mathcal{U} . For all $x \in \mathcal{U}$ we have

$$0 \leq |f(x) - f_S|^{2\theta} \leq \frac{1}{2\mu} \|\nabla f(x)\|^2, \quad \text{and so} \quad \|\nabla f(x)\| \geq \sqrt{2\mu} |f(x) - f_S|^\theta.$$

Let x_L, \dots, x_K be consecutive iterates in \mathcal{U} . For such an iterate x_k , either $\nabla f(x_k) = 0$ and $\text{dist}(x_k, x_{k+1}) = 0$, or $\nabla f(x_k) \neq 0$ and $f(x_k) > f_S$. In this second case, combining the lower-bound on $\|\nabla f(x)\|$ above with the strong decrease condition

in (4), we find that

$$\begin{aligned} \text{dist}(x_k, x_{k+1}) &\leq \frac{1}{\sigma\sqrt{2\mu}} \frac{f(x_k) - f(x_{k+1})}{|f(x_k) - f_S|^\theta} \\ &\leq \frac{1}{\sigma(1-\theta)\sqrt{2\mu}} \left((f(x_k) - f_S)^{1-\theta} - (f(x_{k+1}) - f_S)^{1-\theta} \right), \end{aligned}$$

where the second inequality comes from

$$\begin{aligned} \frac{f(x_k) - f(x_{k+1})}{|f(x_k) - f_S|^\theta} &= \int_{f(x_{k+1})}^{f(x_k)} \frac{1}{|f(x_k) - f_S|^\theta} dt \\ &\leq \int_{f(x_{k+1})}^{f(x_k)} \frac{1}{|t - f_S|^\theta} dt \\ &= \frac{1}{1-\theta} \left((f(x_k) - f_S)^{1-\theta} - (f(x_{k+1}) - f_S)^{1-\theta} \right). \end{aligned}$$

Summing the bound on $\text{dist}(x_k, x_{k+1})$ over $k = L, \dots, K-1$ gives the (BPL) property. \square

D Łojasiewicz and linear convergence rate

Proof of Proposition 3.9 Let \mathcal{U} be a neighborhood of \bar{x} where PL holds. The sufficient decrease (5) and (PL) give

$$f(x_{k+1}) - f(x_k) \leq -\omega \|\nabla f(x_k)\|^2 \leq -2\mu\omega(f(x_k) - f_S)$$

for large enough k . Adding $f(x_k) - f_S$ on both sides yields

$$f(x_{k+1}) - f_S \leq (1 - 2\omega\mu)(f(x_k) - f_S),$$

showing the linear rate for function values. We now prove the rate for $\{\|\nabla f(x_k)\|\}$ and $\{\text{dist}(x_k, \mathcal{S})\}$. The sufficient decrease (5) gives that for large enough k we have

$$\begin{aligned} \omega \|\nabla f(x_k)\|^2 &\leq f(x_k) - f(x_{k+1}) \leq f(x_k) - f_S, \\ \text{and so } \|\nabla f(x_k)\| &\leq \sqrt{\frac{1}{\omega}(f(x_k) - f_S)}. \end{aligned}$$

This shows that $\{\|\nabla f(x_k)\|\}$ converges linearly to zero with rate $\sqrt{1 - 2\omega\mu}$. For $\{\text{dist}(x_k, \mathcal{S})\}$ we use the local quadratic growth (Proposition 2.2). For all x sufficiently close to \bar{x} we have

$$f(x) - f_S \geq \frac{\mu}{2} \text{dist}(x, \mathcal{S})^2, \quad \text{and so} \quad \text{dist}(x_k, \mathcal{S}) \leq \sqrt{\frac{2}{\mu}(f(x_k) - f_S)}.$$

We conclude that $\text{dist}(x_k, \mathcal{S})$ converges linearly to zero with rate $\sqrt{1 - 2\omega\mu}$. \square

Proof of Proposition 3.10 When x_k is sufficiently close to \bar{x} , the Lipschitz-type property (6) implies

$$f(x_k) - f(x_{k+1}) = f(x_k) - f(\mathbf{R}_x(s_k)) \geq -\left(\langle \nabla f(x_k), s_k \rangle + \frac{L}{2} \|s_k\|^2\right). \quad (27)$$

Plugging $s_k = -\gamma \nabla f(x_k)$ yields $f(x_k) - f(x_{k+1}) \geq (1 - \frac{L}{2}\gamma) \|\nabla f(x_k)\| \|s_k\|$. The number $1 - \frac{L}{2}\gamma$ is positive because $\gamma \in]0, \frac{2}{L}[$. So this algorithm satisfies the strong decrease property (4) around \bar{x} with constant $\sigma = \frac{1}{c_r}(1 - \frac{L}{2}\gamma)$, where c_r is as in (RD). It further satisfies (BPL) by Lemma 3.8, so Corollary 3.7 ensures capture of the iterates. Equation (27) additionally gives $f(x_k) - f(x_{k+1}) \geq (\gamma - \frac{L}{2}\gamma^2) \|\nabla f(x_k)\|^2$. It follows that the iterates satisfy the sufficient decrease (5) with $\omega = \gamma - \frac{L}{2}\gamma^2$. We obtain the linear convergence rate with Proposition 3.9. \square

Using performance estimation, it is possible to derive sharper convergence rates for gradient descent with constant step-sizes under the PŁ assumption [1, Thm. 3]. The Lipschitz-type property (6) also implies the sufficient decrease (5) for gradient descent with backtracking line-search.

Proposition D.1 *Suppose that f and \mathbf{R} satisfy (6) around $\bar{x} \in \mathcal{S}$. Let $\{x_k\}$ be a sequence of iterates generated by gradient descent with Armijo backtracking line-search converging to \bar{x} . Then $\{x_k\}$ satisfies (5) with*

$$\omega = \sigma \min\left(\bar{\alpha}, \frac{2\beta(1 - \sigma)}{L}\right).$$

Proof See [24, Lem. 4.12] for the Riemannian case. \square

See also [53] for more results on line-search with Łojasiewicz inequalities.

Acknowledgements We thank Dmitry Drusvyatskiy, Jérôme Bolte and Christopher Criscitiello for insightful discussions and pointers to literature, and Brighton Ancelin for helpful feedback on an earlier version.

Funding Open access funding provided by EPFL Lausanne.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Abbaszadehpavasti, H., de Klerk, E., Zamani, M.: Conditions for linear convergence of the gradient method for non-convex optimization. *Optim. Lett.* **17**(5), 1105–1125 (2023)

2. 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)
3. Absil, P.-A., Baker, C.G., Gallivan, K.A.: Trust-region methods on Riemannian manifolds. *Found. Comput. Math.* **7**(3), 303–330 (2007)
4. Absil, P.-A.: Robert Mahony, and Rodolphe Sepulchre, *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, Princeton (2008)
5. Agarwal, N., Boumal, N., Bullins, B., Cartis, C.: Adaptive regularization with cubics on manifolds. *Math. Program.* **188**(1), 85–134 (2021)
6. Anitescu, M.: Degenerate nonlinear programming with a quadratic growth condition. *SIAM J. Optim.* **10**(4), 1116–1135 (2000)
7. Apidopoulos, V., Ginatta, N., Villa, S.: Convergence rates for the Heavy-Ball continuous dynamics for non-convex optimization, under Polyak–Łojasiewicz conditioning. *J. Global Optim.* **84**, 563–589 (2021)
8. Arbel, M., Mairal, J.: Non-convex bilevel games with critical point selection maps. *Adv. Neural. Inf. Process. Syst.* **35**, 8013–8026 (2022)
9. Attouch, H., Bolte, J., Redont, P., Soubeyran, A.: Proximal alternating minimization and projection methods for nonconvex problems: an approach based on the Kurdyka–Łojasiewicz inequality. *Math. Oper. Res.* **35**(2), 438–457 (2010)
10. 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), 91–129 (2013)
11. Bassily, R., Belkin, M., Ma, S.: On exponential convergence of SGD in non-convex over-parametrized learning. arXiv preprint [arXiv:1811.02564](https://arxiv.org/abs/1811.02564) (2018)
12. Behling, R., Gonçalves, D.S., Santos, S.A.: Local convergence analysis of the Levenberg–Marquardt framework for nonzero-residue nonlinear least-squares problems under an error bound condition. *J. Optim. Theory Appl.* **183**(3), 1099–1122 (2019)
13. Belkin, M.: Fit without fear: remarkable mathematical phenomena of deep learning through the prism of interpolation. *Acta Numer* **30**, 203–248 (2021)
14. Bellavia, S., Morini, B.: Strong local convergence properties of adaptive regularized methods for nonlinear least squares. *IMA J. Numer. Anal.* **35**(2), 947–968 (2015)
15. Bellettini, G., Masala, M., Novaga, M.: A conjecture of De Giorgi on the square distance function. *J. Convex Anal.* **14**(2), 353 (2007)
16. Bertsekas, D.P.: *Nonlinear Programming*. Athena Scientific, Nashua (1995)
17. Birgin, E.G., Gardenghi, J.L., Martínez, J.M., Santos, S.A., Toint, P.L.: Worst-case evaluation complexity for unconstrained nonlinear optimization using high-order regularized models. *Math. Program.* **163**(1), 359–368 (2017)
18. Bolte, J., Pauwels, E.: Curiosities and counterexamples in smooth convex optimization. *Math. Program.* **195**, 553–603 (2021)
19. Bolte, J., Daniilidis, A., Ley, O., Mazet, L.: Characterizations of Łojasiewicz inequalities: subgradient flows, talweg, convexity. *Trans. Am. Math. Soc.* **362**(6), 3319–3363 (2010)
20. Bolte, J., Sabach, S., Teboulle, M.: Proximal alternating linearized minimization for nonconvex and nonsmooth problems. *Math. Program.* **146**(1), 459–494 (2014)
21. Bolte, J., Nguyen, T.P., Peypouquet, J., Suter, B.W.: From error bounds to the complexity of first-order descent methods for convex functions. *Math. Program.* **165**, 471–507 (2017)
22. Bonnans, J.F., Ioffe, A.: Second-order sufficiency and quadratic growth for nonisolated minima. *Math. Oper. Res.* **20**(4), 801–817 (1995)
23. Boos, E., Gonçalves, D.S., Bazán, F.S.V.: Levenberg–Marquardt method with singular scaling and applications. *Appl. Math. Comput.* **474**, 128688 (2024)
24. Boumal, N.: *An Introduction to Optimization on Smooth Manifolds*. Cambridge University Press, London (2023)
25. Burer, S., Monteiro, R.D.C.: A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization. *Math. Program.* **95**(2), 329–357 (2003)
26. Burer, S., Monteiro, R.D.C.: Local minima and convergence in low-rank semidefinite programming. *Math. Program.* **103**(3), 427–444 (2005)
27. Cartis, C., Gould, N.I.M., Toint, P.L.: Adaptive cubic regularisation methods for unconstrained optimization. Part I: motivation, convergence and numerical results. *Math. Program.* **127**(2), 245–295 (2011)

28. Cartis, C., Gould, N.I.M., Toint, P.L.: Adaptive cubic regularisation methods for unconstrained optimization. Part II: worst-case function-and derivative-evaluation complexity. *Math. Program.* **130**(2), 295–319 (2011)
29. Cartis, C., Gould, N.I.M., Toint, P.L.: *Evaluation Complexity of Algorithms for Nonconvex Optimization: Theory, Computation and Perspectives*. SIAM, Philadelphia (2022)
30. Chatterjee, S.: Convergence of gradient descent for deep neural networks. arXiv preprint [arXiv:2203.16462](https://arxiv.org/abs/2203.16462), (2022)
31. Chill, R.: The Łojasiewicz–Simon gradient inequality in Hilbert spaces. In: Jendoubi, M.A. (ed.) *Proceedings of the 5th European-Maghrebian Workshop on Semigroup Theory, Evolution Equations, and Applications*, pp. 25–36. Citeseer (2006)
32. Clarke, F.H., Stern, R.J., Wolenski, P.R.: Proximal smoothness and the lower-C2 property. *J. Convex Anal.* **2**(1–2), 117–144 (1995)
33. Conn, A.R., Gould, N.I., Toint, P.L.: *Trust Region Methods*. SIAM, Philadelphia (2000)
34. Cooper, Y.: Global minima of overparameterized neural networks. *SIAM J. Math. Data Sci.* **3**(2), 676–691 (2021)
35. Corvellec, J.-N., Motreanu, V.V.: Nonlinear error bounds for lower semicontinuous functions on metric spaces. *Math. Program.* **114**, 291–319 (2008)
36. Davis, D., Jiang, L.: A local nearly linearly convergent first-order method for nonsmooth functions with quadratic growth. *Found. Comput. Math.* (2024). <https://doi.org/10.1007/s10208-024-09653-y>
37. Dereich, S., Kassing, S.: Convergence of stochastic gradient descent schemes for Łojasiewicz landscapes. arXiv preprint [arXiv:2102.09385](https://arxiv.org/abs/2102.09385) (2021)
38. Dereich, S., Kassing, S.: Central limit theorems for stochastic gradient descent with averaging for stable manifolds. *Electron. J. Probab.* **28**, 1–48 (2023)
39. Drusvyatskiy, D., Lewis, A.S.: Error bounds, quadratic growth, and linear convergence of proximal methods. *Math. Oper. Res.* **43**(3), 919–948 (2018)
40. Drusvyatskiy, D., Mordukhovich, B.S., Nghia, T.T.A.: Second-order growth, tilt stability, and metric regularity of the subdifferential. arXiv preprint [arXiv:1304.7385](https://arxiv.org/abs/1304.7385) (2013)
41. Drusvyatskiy, D., Ioffe, A.D., Lewis, A.S.: Curves of descent. *SIAM J. Control. Optim.* **53**(1), 114–138 (2015)
42. Drusvyatskiy, D., Ioffe, A.D., Lewis, A.S.: Nonsmooth optimization using Taylor-like models: error bounds, convergence, and termination criteria. *Math. Program.* **185**(1), 357–383 (2021)
43. Fan, J., Yuan, Y.: On the quadratic convergence of the Levenberg–Marquardt method without nonsingularity assumption. *Computing* **74**(1), 23–39 (2005)
44. Feehan, P.: On the Morse–Bott property of analytic functions on Banach spaces with Łojasiewicz exponent one half. *Calc. Var. Partial. Differ. Equ.* **59**(2), 1–50 (2020)
45. Fehrman, B., Gess, B., Jentzen, A.: Convergence rates for the stochastic gradient descent method for non-convex objective functions. *J. Mach. Learn. Res.* **21**, 136 (2020)
46. Frankel, P., Garrigos, G., Peypouquet, J.: Splitting methods with variable metric for Kurdyka–Łojasiewicz functions and general convergence rates. *J. Optim. Theory Appl.* **165**(3), 874–900 (2015)
47. Goyens, F., Royer, C.: Riemannian trust-region methods for strict saddle functions with complexity guarantees. arXiv preprint [arXiv:2402.07614](https://arxiv.org/abs/2402.07614) (2024)
48. Griewank, A.: The modification of Newton’s method for unconstrained optimization by bounding cubic terms. Technical report, Technical report NA/12 (1981)
49. Helmke, U., Moore, J.B.: *Optimization and Dynamical Systems*, 2nd edn. Springer, Berlin (1996)
50. Ioffe, A.D.: Metric regularity and subdifferential calculus. *Russ. Math. Surv.* **55**(3), 501 (2000)
51. Jarre, F., Toint, P.L.: Simple examples for the failure of Newton’s method with line search for strictly convex minimization. *Math. Program.* **158**(1), 23–34 (2016)
52. Karimi, H., Nutini, J., Schmidt, M.: Linear convergence of gradient and proximal-gradient methods under the Polyak–Łojasiewicz condition. In: *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 795–811. Springer (2016)
53. Khanh, P.D., Mordukhovich, B.S., Tran, D.B.: Inexact reduced gradient methods in smooth nonconvex optimization. arXiv preprint [arXiv:2204.01806](https://arxiv.org/abs/2204.01806) (2022)
54. Ko, T., Li, X.: A local convergence theory for the stochastic gradient descent method in non-convex optimization with non-isolated local minima. *J. Mach. Learn.* **2**(2), 138–160 (2023)
55. Kurdyka, K.: On gradients of functions definable in o-minimal structures. *Annales de l’institut Fourier* **48**, 769–783 (1998)

56. Lageman, C.: Convergence of gradient-like dynamical systems and optimization algorithms. Ph.D. thesis, Universität Würzburg (2007)
57. Lageman, C.: Pointwise convergence of gradient-like systems. *Math. Nachr.* **280**(13–14), 1543–1558 (2007)
58. Lee, J.M.: *Introduction to Riemannian Manifolds*. Springer, Berlin (2018)
59. Levin, E., Kileel, J., Boumal, N.: Finding stationary points on bounded-rank matrices: a geometric hurdle and a smooth remedy. *Math. Program.* (2022). <https://doi.org/10.1007/s10107-022-01851-2>
60. Levin, E., Kileel, J., Boumal, N.: The effect of smooth parametrizations on nonconvex optimization landscapes. *Math. Program.* (2024). <https://doi.org/10.1007/s10107-024-02058-3>
61. Lewis, A.S., Tian, T.: Identifiability, the KL property in metric spaces, and subgradient curves. *Found. Comput. Math.* (2024). <https://doi.org/10.1007/s10208-024-09652-z>
62. Li, G., Pong, T.K.: Calculus of the exponent of Kurdyka–Łojasiewicz inequality and its applications to linear convergence of first-order methods. *Found. Comput. Math.* **18**(5), 1199–1232 (2018)
63. Li, S., Li, Q.: Local and global convergence of general Burer–Monteiro tensor optimizations. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36 (2022)
64. Li, Z., Wang, T., Arora, S.: What happens after SGD reaches zero loss?—A mathematical framework. In: *International Conference on Learning Representations* (2022)
65. Liao, F.-Y., Ding, L., Zheng, Y.: Error bounds, PL condition, and quadratic growth for weakly convex functions, and linear convergences of proximal point methods. In: *6th Annual Learning for Dynamics & Control Conference*, pp. 993–1005. PMLR (2024)
66. Liu, C., Zhu, L., Belkin, M.: Loss landscapes and optimization in over-parameterized non-linear systems and neural networks. *Appl. Comput. Harmonic Anal.* **59**, 85–116 (2022)
67. Liu, J., Wright, S.J.: Asynchronous stochastic coordinate descent: parallelism and convergence properties. *SIAM J. Optim.* **25**(1), 351–376 (2015)
68. Łojasiewicz, S.: Une propriété topologique des sous-ensembles analytiques réels. *Les équations aux dérivées partielles* **117**, 87–89 (1963)
69. Łojasiewicz, S.: Sur les trajectoires du gradient d’une fonction analytique. *Seminari di Geometria* **115–117**, 1982 (1983)
70. Luo, Z.-Q., Tseng, P.: Error bounds and convergence analysis of feasible descent methods: a general approach. *Ann. Oper. Res.* **46**(1), 157–178 (1993)
71. Marteau-Ferey, U., Bach, F., Rudi, A.: Second order conditions to decompose smooth functions as sums of squares. *SIAM J. Optim.* **34**(1), 616–641 (2024)
72. Masiha, S., Salehkaleybar, S., He, N., Kiyavash, N., Thiran, P.: Stochastic second-order methods provably beat SGD for gradient-dominated functions. *Adv. Neural. Inf. Process. Syst.* (2022). <https://doi.org/10.48550/arXiv.2205.12856>
73. Muoi, B.N., Yen, N.D.: Local stability and local convergence of the basic trust-region method. *J. Optim. Theory Appl.* **172**(2), 578–593 (2017)
74. Necoara, I., Lupu, D.: General higher-order majorization–minimization algorithms for (non) convex optimization. *arXiv preprint arXiv:2010.13893* (2020)
75. Necoara, I., Nesterov, Yu., Glineur, F.: Linear convergence of first order methods for non-strongly convex optimization. *Math. Program.* **175**, 69–107 (2019)
76. Nesterov, Y., Polyak, B.T.: Cubic regularization of Newton method and its global performance. *Math. Program.* **108**(1), 177–205 (2006)
77. Nocedal, J., Wright, S.J.: *Numerical Optimization*. Springer, New York (2006)
78. Noll, D., Rondepierre, A.: Convergence of line search and trust-region methods using the Kurdyka–Łojasiewicz inequality. In: *Computational and analytical mathematics*, pp. 593–611. Springer (2013)
79. Otto, F., Villani, C.: Generalization of an inequality by Talagrand and links with the logarithmic Sobolev inequality. *J. Funct. Anal.* **173**(2), 361–400 (2000)
80. Oymak, S., Soltanolkotabi, M.: Overparameterized nonlinear learning: gradient descent takes the shortest path? In: *International Conference on Machine Learning*, pp. 4951–4960. PMLR, (2019)
81. Polyak, B.T.: Gradient methods for the minimisation of functionals. *USSR Comput. Math. Math. Phys.* **3**(4), 864–878 (1963)
82. Polyak, B.T.: *Introduction to Optimization*. Optimization Software (1987)
83. Qi, C.: *Numerical optimization methods on Riemannian manifolds*. Ph.D. thesis (2011)
84. Qian, Y., Pan, S.: A superlinear convergence iterative framework for Kurdyka–Łojasiewicz optimization and application. *arXiv preprint arXiv:2210.12449* (2022)

85. Rebjock, Q., Boumal, N.: Fast convergence of trust-regions for non-isolated minima via analysis of CG on indefinite matrices. arXiv preprint [arXiv:2311.07404](https://arxiv.org/abs/2311.07404) (2023)
86. Ring, W., Wirth, B.: Optimization methods on Riemannian manifolds and their application to shape space. *SIAM J. Optim.* **22**(2), 596–627 (2012)
87. Shapiro, A.: Perturbation theory of nonlinear programs when the set of optimal solutions is not a singleton. *Appl. Math. Optim.* **18**(1), 215–229 (1988)
88. Stonyakin, F., Kuruzov, I., Polyak, B.: Stopping rules for gradient methods for non-convex problems with additive noise in gradient. *J. Optim. Theory Appl.* **198**(2), 531–551 (2023)
89. Sun, Y., Flammarion, N., Fazel, M.: Escaping from saddle points on Riemannian manifolds. *Adv. Neural. Inf. Process. Syst.* (2019). <https://doi.org/10.48550/arXiv.1906.07355>
90. Terjék, D., González-Sánchez, D.: A framework for overparameterized learning. arXiv preprint [arXiv:2205.13507](https://arxiv.org/abs/2205.13507) (2022)
91. Tseng, P.: Error bounds and superlinear convergence analysis of some Newton-type methods in optimization. In: *Nonlinear Optimization and Related Topics*, pp. 445–462. Springer (2000)
92. Usevich, K., Li, J., Comon, P.: Approximate matrix and tensor diagonalization by unitary transformations: convergence of Jacobi-type algorithms. *SIAM J. Optim.* **30**(4), 2998–3028 (2020)
93. Vannieuwenhoven, N.: The condition number of singular subspaces, revisited. *Linear Algebra Appl.* **687**, 157–180 (2024)
94. Wojtowytsch, S.: Stochastic gradient descent with noise of machine learning type. Part I: discrete time analysis. *J. Nonlinear Sci.* **33**(3), 45 (2023)
95. Wojtowytsch, S.: Stochastic gradient descent with noise of machine learning type. Part II: continuous time analysis. *J. Nonlinear Sci.* **34**(1), 16 (2024)
96. Yamashita, N., Fukushima, M.: On the rate of convergence of the Levenberg–Marquardt method. In: *Topics in Numerical Analysis*, pp. 239–249. Springer (2001)
97. Yue, M.-C., Zhou, Z., So, A.M.-C.: On the quadratic convergence of the cubic regularization method under a local error bound condition. *SIAM J. Optim.* **29**(1), 904–932 (2019)
98. Yue, P., Fang, C., Lin, Z.: On the lower bound of minimizing Polyak–Łojasiewicz functions. In: *The Thirty Sixth Annual Conference on Learning Theory*, pp. 2948–2968. PMLR (2023)
99. Zeng, Z.: A Newton– ϵ^{TM} s iteration converges quadratically to nonisolated solutions too. *Math. Comput.* **92**(344), 2795–2824 (2023)
100. Zhang, H.: The restricted strong convexity revisited: analysis of equivalence to error bound and quadratic growth. *Optim. Lett.* **11**, 817–833 (2017)
101. Zhang, H., Yin, W.: Gradient methods for convex minimization: better rates under weaker conditions. arXiv preprint [arXiv:1303.4645](https://arxiv.org/abs/1303.4645) (2013)
102. Zhang, J., Zhang, S.: A cubic regularized Newton’s method over Riemannian manifolds. arXiv preprint [arXiv:1805.05565](https://arxiv.org/abs/1805.05565) (2018)
103. Zhang, R.Y.: Improved global guarantees for the nonconvex Burer–Monteiro factorization via rank overparameterization. arXiv preprint [arXiv:2207.01789](https://arxiv.org/abs/2207.01789) (2022)
104. Zhang, Y., Qu, Q., Wright, J.: From symmetry to geometry: tractable nonconvex problems. *ArXiv [arxiv:abs/2007.06753](https://arxiv.org/abs/2007.06753)* (2020)
105. Zhou, Y., Wang, Z., Liang, Y.: Convergence of cubic regularization for nonconvex optimization under KL property. *Adv. Neural Inf. Process. Syst.* **31** (2018)