



The effect of smooth parametrizations on nonconvex optimization landscapes

Eitan Levin¹ · Joe Kileel² · Nicolas Boumal³

Received: 18 August 2022 / Accepted: 12 January 2024
© The Author(s) 2024

Abstract

We develop new tools to study landscapes in nonconvex optimization. Given one optimization problem, we pair it with another by smoothly parametrizing the domain. This is either for practical purposes (e.g., to use smooth optimization algorithms with good guarantees) or for theoretical purposes (e.g., to reveal that the landscape satisfies a strict saddle property). In both cases, the central question is: how do the landscapes of the two problems relate? More precisely: how do desirable points such as local minima and critical points in one problem relate to those in the other problem? A key finding in this paper is that these relations are often determined by the parametrization itself, and are almost entirely independent of the cost function. Accordingly, we introduce a general framework to study parametrizations by their effect on landscapes. The framework enables us to obtain new guarantees for an array of problems, some of which were previously treated on a case-by-case basis in the literature. Applications include: optimizing low-rank matrices and tensors through factorizations; solving semidefinite programs via the Burer–Monteiro approach; training neural networks by optimizing their weights and biases; and quotienting out symmetries.

Keywords Overparametrization · Symmetries in landscapes · Benign nonconvexity · Strict saddles · Low rank optimization · Hadamard parametrization of simplex · Optimization on manifolds

✉ Nicolas Boumal
nicolas.boumal@epfl.ch

Eitan Levin
eitanl@caltech.edu

Joe Kileel
jkileel@math.utexas.edu

¹ Department of Computing and Mathematical Sciences, California Institute of Technology, Pasadena, USA

² Department of Mathematics and Oden Institute for Computational Engineering and Sciences, University of Texas at Austin, Austin, USA

³ Institute of Mathematics, École polytechnique fédérale de Lausanne (EPFL), Lausanne, Switzerland

Contents

- 1 Introduction
- 2 Lifts and their properties
 - 2.1 The sphere-to-simplex Hadamard lift
 - 2.2 Smooth semidefinite programs via Burer–Monteiro
 - 2.3 Low-rank matrices
 - 2.4 Low-rank tensors
 - 2.5 Neural networks
 - 2.6 Submersions and higher order stationary points
- 3 Characterizations of lifts
 - 3.1 Local minima
 - 3.2 Stationary points
 - 3.2.1 “1 ⇒ 1”: lifts preserving 1-critical points
 - 3.2.2 “2 ⇒ 1”: lifts mapping 2-critical points to 1-critical points
 - 3.3 Composition of lifts
 - 3.4 Computing L_y and Q_y
 - 3.5 Low-rank PSD matrices, and computing tangent cones
- 4 Constructing lifts via fiber products
- 5 Analysis of low rank lifts
 - 5.1 Proof of Proposition 2.8 (LR^\top lift)
 - 5.2 Proof of Proposition 2.9 (desingularization lift)
 - 5.3 Multilinear lifts, and tensors
- 6 Conclusions and future work
- A Lifts preserving local minima
- References

1 Introduction

We consider pairs of optimization problems (P) and (Q) as defined below, where \mathcal{E} is a linear space, \mathcal{M} is a smooth manifold,¹ and $\varphi: \mathcal{M} \rightarrow \mathcal{E}$ is a smooth (over)parametrization of the search space $\mathcal{X} = \varphi(\mathcal{M})$ of (P).² Their optimal values are equal:

$$\begin{array}{ccc}
 \mathcal{M} & & \min_{y \in \mathcal{M}} g(y) & \text{(Q)} \\
 \varphi \downarrow & \searrow^{g=f \circ \varphi} & & \\
 \mathcal{X} \subseteq \mathcal{E} & \xrightarrow{f} & \mathbb{R} & \text{(P)}
 \end{array}$$

We usually assume $f: \mathcal{E} \rightarrow \mathbb{R}$ is smooth (C^∞), hence so is $g = f \circ \varphi: \mathcal{M} \rightarrow \mathbb{R}$ by composition.

¹ All linear spaces and manifolds are assumed to be finite-dimensional.

² The two-headed arrow \twoheadrightarrow in the diagram denotes a surjection.

Such pairs of problems (P) and (Q) arise in **two scenarios** (concrete examples follow):

- (a) Our task is to minimize f on \mathcal{X} as in (P), but we lack good algorithms to do so, e.g., because \mathcal{X} lacks regularity. In this case, we choose a smooth parametrization φ of \mathcal{X} and run algorithms on the smooth problem (Q) instead.
- (b) Our task is to minimize g on \mathcal{M} as in (Q), but its landscape is complex (e.g., due to symmetries). In this case, we factor g through a smooth map φ in the hope of revealing a problem (P) whose landscape is simpler and can be leveraged to analyze that of (Q).

In both cases, we run an optimization algorithm on the smooth problem (Q). This algorithm may find desirable points y on \mathcal{M} for (Q) (global or local minima, stationary points). For example, certain trust-region algorithms are guaranteed to accumulate at second-order stationary points—see [18] and an extension to manifolds [40, §3]—and many first- and even zeroth-order methods converge to second-order stationary points from almost all initializations [4, 36, 57]. However, in general such points need not map to desirable points $\varphi(y)$ on \mathcal{X} for (P). Indeed, nonlinear parametrizations may severely distort landscapes, and notably may introduce spurious critical points. Algorithms running on (Q) are liable to terminate at an approximately stationary point near such a spurious point, and return a point whose image through φ is nowhere near any stationary point for (P).

In this paper, we characterize the properties that the parametrization φ needs to satisfy for desirable points of (Q) to map to desirable points of (P), that is, we develop a general framework to relate the landscapes of pairs of problems of the above form. Importantly, we observe that these properties are often entirely independent of the cost function f in (P), since many parametrizations map desirable points for (Q) to those for (P) for any cost function. Our framework enables us to unify and strengthen the analysis of a wealth of parametrizations arising in applications, hitherto studied case-by-case and often only for specific costs.

Parametrizations are ubiquitous. They arise in semidefinite programming [10, 16], low-rank optimization [25, 38, 44], computer vision [19], inverse kinematics and trajectory planning [52, Chaps. 1,4], algebraic geometry [26, Chap. 17], training neural networks [41, 46], and risk minimization [6, 7, 54]. The following are two concrete examples that illustrate the above two scenarios.

For an example of **scenario (a)**, consider minimizing a cost f over the set $\mathcal{X} = \mathbb{R}_{\leq r}^{m \times n}$ of all $m \times n$ matrices of rank at most r . Unfortunately, standard algorithms running on (P) may converge to a non-stationary point because of the nonsmooth geometry of \mathcal{X} [40, 45]. Instead of trying to solve (P) directly, it is common to parametrize \mathcal{X} by the linear space $\mathcal{M} = \mathbb{R}^{m \times r} \times \mathbb{R}^{n \times r}$ using the rank factorization $\varphi(L, R) = LR^\top$, and to solve (Q) instead. The resulting problem (Q) requires minimizing a smooth cost function over a linear space; there are several algorithms that converge to a second-order stationary point for such problems. Furthermore, any second-order stationary point for (Q) maps under φ to a stationary point for (P) by [25, Thm. 1]. Thus, parametrization of \mathcal{X} by φ gives us an algorithm converging to a stationary point for (P) by running standard algorithms on (Q), even though similarly

Table 1 Sets used frequently in the paper

Bounded-rank matrices	$\mathbb{R}_{\leq r}^{m \times n} = \{X \in \mathbb{R}^{m \times n} : \text{rank}(X) \leq r\}$
Fixed-rank matrices	$\mathbb{R}_{=r}^{m \times n} = \{X \in \mathbb{R}^{m \times n} : \text{rank}(X) = r\}$
Symmetric matrices	$\mathbb{S}^m = \{X \in \mathbb{R}^{m \times m} : X = X^\top\}$
Nonnegative orthant	$\mathbb{R}_{\geq 0}^m = \{x \in \mathbb{R}^m : x_i \geq 0 \text{ for all } i\}$
Positive orthant	$\mathbb{R}_{> 0}^m = \{x \in \mathbb{R}^m : x_i > 0 \text{ for all } i\}$
Positive-semidefinite (PSD) matrices	$\mathbb{S}_{\geq 0}^m = \{X \in \mathbb{S}^m : X \geq 0\}$
Positive-definite (PD) matrices	$\mathbb{S}_{> 0}^m = \{X \in \mathbb{S}^m : X > 0\}$
Standard simplex	$\Delta^{n-1} = \left\{x \in \mathbb{R}_{\geq 0}^n : \sum_{i=1}^n x_i = 1\right\}$
Stiefel matrices	$\text{St}(m, r) = \{X \in \mathbb{R}^{m \times r} : X^\top X = I_r\}$
Unit sphere	$\mathbb{S}^{m-1} = \text{St}(m, 1)$
Orthogonal matrices	$\text{O}(m) = \text{St}(m, m)$
Special orthogonal matrices	$\text{SO}(m) = \{X \in \text{O}(m) : \det(X) = 1\}$
Invertible matrices	$\text{GL}(m) = \{X \in \mathbb{R}^{m \times m} : \det(X) \neq 0\}$
Grassmannian	$\text{Gr}(m, r) = \{\mathcal{S} \subseteq \mathbb{R}^m : \mathcal{S} \text{ linear subspace, } \dim \mathcal{S} = r\}$

Other sets include low rank tensors and functions representable by fixed neural network architectures

reasonable algorithms may fail to produce a stationary point when applied directly to (P).

For an illustration of **scenario (b)**, consider finding the smallest eigenvalue of a $d \times d$ symmetric matrix A , which can be written in the form (Q) with \mathcal{M} the unit sphere in \mathbb{R}^d and $g(y) = y^\top A y$. This problem is not convex, hence it could have bad local minima. Here is one way to reason that it does not (as is well known). If $\lambda \in \mathbb{R}^d$ denotes the vector of eigenvalues of A and $U \in \text{O}(d)$ is an orthogonal matrix of eigenvectors satisfying $A = U \text{diag}(\lambda) U^\top$ (both of which are unknown), define $\varphi(y) = \text{diag}(U^\top y y^\top U) \in \mathbb{R}^d$ and $f(x) = \lambda^\top x$. It is easy to check that $g = f \circ \varphi$, and that $\mathcal{X} = \varphi(\mathcal{M})$ is the standard simplex in \mathbb{R}^d . The resulting problem (P) is convex in this case, hence each of its stationary points is a global minimum. A corollary of the theory we develop in this paper is that any second-order stationary point for (Q) with φ as above maps to a stationary point for (P), for any cost function f —see Example 4.13. Thus, we recover the well-known fact that any second-order stationary point for the eigenvalue problem (Q) is globally optimal. Even though the problem (P) cannot be solved directly in this case because f and φ are unknown, their mere existence can be used to show that the nonconvexity of (Q) is “benign”. From this perspective, problem (P) reveals hidden convexity in problem (Q). This hidden convexity is present more generally in lifts arising from Kostant’s convexity theorem, extending this example to optimization of certain linear functions over certain Lie group orbits [35].

We state our main definitions and results relating the landscapes of (P) and (Q) in general, and instantiate these results on a number of specific lifts arising in the literature, in Sect. 2. Table 1 collects the notations and definitions for several sets used throughout the paper.

2 Lifts and their properties

We call the parametrization in (Q) a (smooth) *lift* of \mathcal{X} :

Definition 2.1 A *smooth lift* of $\mathcal{X} \subseteq \mathcal{E}$ is a smooth manifold \mathcal{M} together with a smooth map $\varphi: \mathcal{M} \rightarrow \mathcal{E}$ such that $\varphi(\mathcal{M}) = \mathcal{X}$.

As the two scenarios in Sect. 1 illustrate, understanding when lifts map desirable points for (Q) to desirable points for (P) yields guarantees for algorithms running on (Q). Here desirable points might be minimizers (global or local) and stationary points (of first, second, or higher order). The relation between these two sets of desirable points has been studied for various specific lifts and cost functions. In this paper, we study this relation in general and answer the following question:

*Which lifts have the property that desirable points of (Q) map to desirable points of (P), for **all** cost functions f ?*

Surprisingly, we find that many lifts arising in practice satisfy such properties, yielding guarantees for algorithms running on (Q) that are independent of the particular cost function involved, and only depend on the geometry of the lift. We further show that whenever a lift does not preserve desirable points for all cost functions, then it fails to do so already for quite simple costs. In this case our results identify obstructions to proofs of guarantees for algorithms, which must then exploit the structure of the particular cost function at hand.

To begin answering the above question, we note that global minima of (Q) always map under φ to global minima of (P), for all cost functions f . This holds simply because $\varphi(\mathcal{M}) = \mathcal{X}$, see Proposition 3.5. Global minima are hard to find in general, so we study other types of desirable points such as local minima and stationary points. In contrast to global minima, these types of desirable points are not guaranteed to map³ to each other under smooth lifts. In fact, it is possible for a local minimum of (Q) to map under φ to a non-stationary point for (P), see Example 3.7. Thus, we define the following properties of smooth lifts that, when satisfied, yield a connection between desirable points for (Q) and those for (P).

Definition 2.2 (*Desirable properties of lifts*) Suppose $\varphi: \mathcal{M} \rightarrow \mathcal{X}$ is a smooth lift.

- The lift φ satisfies the “local \Rightarrow local” property at $y \in \mathcal{M}$ if, for all continuous $f: \mathcal{X} \rightarrow \mathbb{R}$, if y is a local minimum for (Q) then $x = \varphi(y)$ is a local minimum for (P). We say φ satisfies the “local \Rightarrow local” property if it does so at all $y \in \mathcal{M}$.
- The lift φ satisfies the “ $k \Rightarrow 1$ ” property at y for $k = 1, 2$ if for all k -times differentiable $f: \mathcal{X} \rightarrow \mathbb{R}$, if y is a k th-order stationary point (“ k -critical” for short) for (Q) then $x = \varphi(y)$ is stationary for (P). We say φ satisfies the “ $k \Rightarrow 1$ ” property if it does so at all $y \in \mathcal{M}$.

The precise definitions of each type of desirable points above is given in Sect. 3. We fully characterize these properties and explain how to check them on specific examples.

³ The converse question is simple: preimages of local minima are local minima by continuity of φ (Proposition 3.6), and preimages of stationary points are stationary by differentiability of φ (Proposition 3.13).

We then apply our results to study lifts arising in applications ranging from low-rank matrices and tensors to neural networks.

Note that “ $2 \Rightarrow 1$ ” at y implies “ $1 \Rightarrow 1$ ” at y since 2-critical points are 1-critical, but no other implication between the different properties holds in general—see Remark 2.13. We also mention that C^∞ smoothness is not necessary for the above properties or for their characterizations. For example, it suffices for the manifold \mathcal{M} and the lift φ to be of class C^k for the definition of “ $k \Rightarrow 1$ ” and its characterization to apply. For “local \Rightarrow local” it suffices for \mathcal{M} to be a topological space satisfying certain properties (see “Appendix 1”) and for φ to be continuous.

Our characterizations of “local \Rightarrow local” and “ $1 \Rightarrow 1$ ” are easy to state as follows.

Theorem 2.3 *The lift $\varphi: \mathcal{M} \rightarrow \mathcal{X}$ satisfies “local \Rightarrow local” at $y \in \mathcal{M}$ if and only if it is open at y . If φ does not satisfy “local \Rightarrow local” at y , there is a smooth cost f such that y is a local minimum for (Q) but $\varphi(y)$ is not a local minimum for (P).*

By definition, the map φ is open at $y \in \mathcal{M}$ if it maps neighborhoods of y (that is, sets containing y in their interior) to neighborhoods of $\varphi(y)$ (in the subspace topology on \mathcal{X} from \mathcal{E})—a purely topological property. Proving that openness is sufficient for “local \Rightarrow local” is easy. Proof of its necessity requires substantial work, deferred to “Appendix 1”. Our proof in the appendix provides the result in a more general, topological setting without using smoothness. It also provides (possibly new) conditions which are equivalent to openness and may be easier to check for some lifts.

Our characterization for “ $1 \Rightarrow 1$ ” involves the image of the differential of the lift map φ , and is proved in Sect. 3.2.1.

Theorem 2.4 *The lift $\varphi: \mathcal{M} \rightarrow \mathcal{X}$ satisfies “ $1 \Rightarrow 1$ ” at $y \in \mathcal{M}$ if and only if $\text{im } D\varphi(y) = T_x\mathcal{X}$, where $x = \varphi(y)$. If φ does not satisfy “ $1 \Rightarrow 1$ ”, there is a linear cost f such that y is 1-critical for (Q) but $\varphi(y)$ is not stationary for (P).*

Here φ is viewed as a smooth map between smooth manifolds $\mathcal{M} \rightarrow \mathcal{E}$, and its differential $D\varphi(y)$ maps the tangent space $T_y\mathcal{M}$ to (in general, a subset of) the tangent cone $T_x\mathcal{X}$, see Definition 3.1. Since $D\varphi(y)$ is a linear map and $T_y\mathcal{M}$ is a linear space, “ $1 \Rightarrow 1$ ” is rarely satisfied: unless all tangent cones of \mathcal{X} are linear subspaces, for every smooth lift φ , there exists a (linear) f such that some stationary point for (Q) does not map to a stationary point for (P).

Our characterization for “ $2 \Rightarrow 1$ ” is more complicated, involving the second derivative of φ as well. We state an equivalent condition for “ $2 \Rightarrow 1$ ”, as well as sufficient and necessary conditions for it that are easier to check in some applications, in Theorem 3.23. If “ $2 \Rightarrow 1$ ” fails at y , we show in Corollary 3.18 that there exists a convex quadratic cost f such that y is 2-critical for (Q) but $\varphi(y)$ is not stationary for (P). Note also that if “ $1 \Rightarrow 1$ ” holds at y then so does “ $2 \Rightarrow 1$ ” by definition.

Understanding stationarity on \mathcal{X} requires knowledge of its tangent cones. These can be hard to characterize. We show that it is sometimes possible to obtain an explicit expression for the tangent cones simultaneously with proving “ $1 \Rightarrow 1$ ” and “ $2 \Rightarrow 1$ ” for some lift of \mathcal{X} , see Sects. 3.5 and 4. This is somewhat surprising since the tangent cones to \mathcal{X} are defined independently of any lift.

Given a set \mathcal{X} , it is also natural to seek *constructions* of a smooth lift $\varphi: \mathcal{M} \rightarrow \mathcal{X}$ satisfying desirable properties. We give a systematic construction of a map $\varphi: \mathcal{M} \rightarrow$

\mathcal{X} in Sect. 4 which maps a set \mathcal{M} surjectively onto \mathcal{X} . When the set \mathcal{M} constructed in this way is a smooth manifold, we obtain a smooth lift and give conditions under which φ satisfies each of the above properties.

We now proceed to apply our results to study various lifts arising in the literature.

2.1 The sphere-to-simplex Hadamard lift

There is growing interest in optimizing over the probability simplex $\mathcal{X} = \Delta^{n-1} \subseteq \mathcal{E} = \mathbb{R}^n$ by lifting it to the sphere via the Hadamard lift

$$\mathcal{M} = \mathbb{S}^{n-1}, \quad \varphi(y) = y \odot y, \tag{Had}$$

where \odot denotes the entrywise (Hadamard) product. Using this lift leads to fast algorithms for high-dimensional problems (Q), see [14, 42, 56]. This is also essentially the lift that appears in the eigenvalue example (scenario (b)) in Sect. 1, see Example 4.13. This lift is particularly natural in applications involving probabilities since the push-forward under φ of the standard metric on the sphere is the Fisher-Rao metric on the simplex [5, Prop. 2.1].

We can characterize precisely where each of our desirable properties holds.

Proposition 2.5 *The lift (Had) satisfies “local \Rightarrow local” everywhere, “1 \Rightarrow 1” at y if and only if $y_i \neq 0$ for all i (i.e., at preimages of the relative interior of the simplex), and “2 \Rightarrow 1” everywhere.*

We prove this proposition in Corollary 4.12 by showing that the lift (Had) is a special case of our construction of lifts in Sect. 4 and can be analyzed using the general results we prove there.

The relation between desirable points for (Q) and for (P) have been previously studied in [42], where the authors show that 2-critical points for (Q) map to 2nd-order KKT points for (P), viewed as a nonlinear program, for any twice-differentiable cost. This is a strengthening of “2 \Rightarrow 1”. The authors of [21] prove similar relations between first- and second-order optimality conditions for problems (P) over $\mathbb{R}_{\geq 0}^n$, and for their lifts (Q) to \mathbb{R}^n via the entrywise-squaring lift in (Had).

The Hadamard lift also induces a lift of the set of column-stochastic matrices $\mathcal{X} = \{X \in \mathbb{R}_{\geq 0}^{n \times m} : X^\top \mathbb{1}_n = \mathbb{1}_m\}$ to the product of spheres (called the oblique manifold [9, §7.2]) via

$$\mathcal{M} = \left(\mathbb{S}^{n-1}\right)^m, \quad \varphi(y_1, \dots, y_m) = [y_1 \odot y_1, \dots, y_m \odot y_m]. \tag{HadProd}$$

Here $\mathbb{1}_n$ is the all-1’s vector of length n and $[x_1, \dots, x_m]$ denotes horizontal concatenation of m vectors x_i of length n to form an $n \times m$ matrix. By studying products of lifts in Proposition 4.11, we characterize the properties of (HadProd) in Example 4.14 and obtain the following result.

Proposition 2.6 *The lift (HadProd) satisfies “local \Rightarrow local” everywhere, “1 \Rightarrow 1” at (y_1, \dots, y_k) if and only if $(y_i)_j \neq 0$ for all i, j , and “2 \Rightarrow 1” everywhere.*

Optimization over stochastic matrices has been applied to nonnegative matrix factorization [22]. Such optimization also arises in information theory [8], where stochastic matrices represent transition probabilities of channels.

2.2 Smooth semidefinite programs via Burer–Monteiro

Consider the domain \mathcal{X} of a rank-constrained semidefinite program (SDP),

$$\mathcal{X} = \{X \in \mathbb{S}_{\geq 0}^n : \text{rank}(X) \leq r, \langle A_i, X \rangle = b_i \text{ for } i = 1, \dots, m\} \subseteq \mathcal{E} = \mathbb{S}^n,$$

where $\langle U, V \rangle = \text{Tr}(U^\top V)$ is the (Frobenius) inner product on \mathcal{E} . The Burer–Monteiro approach [11] to optimizing over \mathcal{X} consists of optimizing over the following parametrization instead:

$$\mathcal{M} = \{R \in \mathbb{R}^{n \times r} : h_i(R) := \langle A_i R, R \rangle - b_i = 0 \text{ for } i = 1, \dots, m\}, \quad \varphi(R) = RR^\top. \quad (\text{BM})$$

Burer and Monteiro prove in [12, Prop. 2.3] that local minima for (Q) map under φ to local minima for (P) for linear costs f .

Under some conditions on the A_i, b_i (which are satisfied generically [16, Prop. 1] as well as for several applications of interest [10]), the constraints $h_i(R) = 0$ constitute (constant-rank) local defining functions (in the sense of [37, Thm. 5.12]) for \mathcal{M} , which is then an embedded submanifold of $\mathbb{R}^{n \times r}$. In that case, \mathcal{M} and φ constitute a smooth lift of \mathcal{X} . In [10], assuming f is linear (as is typical for SDPs), the authors use the assumption that the h_i are local defining functions to prove that (in our terminology) rank-deficient 2-critical points for (Q) map under φ to stationary points for (P). This was also shown for nonlinear f in [30], though under more restrictive conditions on A_i, b_i (e.g., $A_i A_j = 0$ for $i \neq j$). In all cases, these results allow to capture benign non-convexity when f is convex, as then stationary points for (P) are global minima.

Using our framework, we can generalize these results to any (twice-differentiable) costs and remove the restrictions on the rank of the 2-critical points.

Proposition 2.7 *The Burer–Monteiro lift (BM) satisfies “local \Rightarrow local” everywhere. If \mathcal{M} in (BM) is a smooth manifold with (constant-rank) local defining functions $\{h_i\}_{i=1}^m$, then this lift satisfies “1 \Rightarrow 1” at R if and only if $\text{rank}(R) = r$, and “2 \Rightarrow 1” everywhere.*

We prove this result too in Corollary 4.12 using general properties of our lift construction in Sect. 4. Our theory also yields explicit expressions for the tangent cones to \mathcal{X} in (4.4), which (to our knowledge) have not previously appeared in the literature.

2.3 Low-rank matrices

Consider the set $\mathcal{X} = \mathbb{R}_{\leq r}^{m \times n}$ of matrices in $\mathcal{E} = \mathbb{R}^{m \times n}$ with rank at most r . We study several natural lifts of this real algebraic variety. The first one we study is based on the rank factorization of a matrix:

$$\mathcal{M} = \mathbb{R}^{m \times r} \times \mathbb{R}^{n \times r}, \quad \varphi(L, R) = LR^\top. \tag{LR}$$

The authors of [25] showed (in our terminology) that “ $1 \Rightarrow 1$ ” does not hold everywhere, but “ $2 \Rightarrow 1$ ” does. We further proved in [38, Prop. 2.30] that this lift does not satisfy “local \Rightarrow local” everywhere either. We strengthen these results here using our unified framework, by characterizing precisely where each of these properties hold. The proof of the following proposition is given in Sect. 5.1.

Proposition 2.8 *The lift of $\mathcal{X} = \mathbb{R}_{\leq r}^{m \times n}$ given by (LR) satisfies:*

- “local \Rightarrow local” at (L, R) if and only if $\text{rank}(L) = \text{rank}(R) = \text{rank}(LR^\top)$,
- “ $1 \Rightarrow 1$ ” at (L, R) if and only if $\text{rank}(L) = \text{rank}(R) = r$,
- and “ $2 \Rightarrow 1$ ” everywhere on \mathcal{M} [25].

The second lift we study for $\mathbb{R}_{\leq r}^{m \times n}$ is the desingularization lift introduced in [31]. It is given by

$$\mathcal{M} = \{(X, \mathcal{S}) \in \mathbb{R}^{m \times n} \times \text{Gr}(n, n - r) : \mathcal{S} \subseteq \ker X\}, \quad \varphi(X, \mathcal{S}) = X. \tag{Desing}$$

Here $\text{Gr}(n, n - r)$ is the Grassmann manifold of $(n - r)$ -dimensional subspaces of \mathbb{R}^n [9, §9]. We proved in [38, Prop. 2.37] that this lift too does not satisfy “local \Rightarrow local”. The following proposition parallels the one above and is proved in Sect. 5.2.

Proposition 2.9 *The lift of $\mathcal{X} = \mathbb{R}_{\leq r}^{m \times n}$ given by (Desing) satisfies:*

- “local \Rightarrow local” at (X, \mathcal{S}) if and only if $\text{rank}(X) = r$; the same is true for “ $1 \Rightarrow 1$ ”.
- “ $2 \Rightarrow 1$ ” everywhere on \mathcal{M} .

A potential advantage of the desingularization lift over the matrix factorization lift is that the preimage of a matrix, $\varphi^{-1}(X)$, is compact for the former but not for the latter.

Note that both lifts (LR) and (Desing) satisfy “ $1 \Rightarrow 1$ ” and “local \Rightarrow local” at preimages of rank- r matrices, but the lift (LR) further satisfies “local \Rightarrow local” at “balanced” preimages of lower-rank matrices. We also mention that no smooth lift of $\mathbb{R}_{\leq r}^{m \times n}$ can satisfy “ $1 \Rightarrow 1$ ” at preimages of lower-rank matrices by Theorem 2.4, since the tangent cones to such matrices are not linear spaces.

In [44], the authors experiment with various SVD-type lifts for optimization over matrices of rank exactly r . The following proposition, proved in the arxiv version of this paper [39, App. D], gives some of the properties of these lifts, extended to parametrize all of $\mathbb{R}_{\leq r}^{m \times n}$.

Proposition 2.10 *The SVD lift and its modification from [44] satisfy the following.*

- The SVD lift of $\mathbb{R}_{\leq r}^{m \times n}$ given by

$$\mathcal{M} = \text{St}(m, r) \times \mathbb{R}^r \times \text{St}(n, r), \quad \varphi(U, \sigma, V) = U \text{diag}(\sigma) V^\top, \tag{SVD}$$

satisfies “local \Rightarrow local” at (U, σ, V) if and only if $|\sigma_1|, \dots, |\sigma_r|$ are nonzero and distinct; the same holds for “ $1 \Rightarrow 1$ ”.

- *The modified SVD lift*

$$\mathcal{M} = \text{St}(m, r) \times \mathbb{S}^r \times \text{St}(n, r), \quad \varphi(U, M, V) = U M V^\top, \quad (\text{MSVD})$$

satisfies “local \Rightarrow local” at (U, M, V) if and only if the eigenvalues of M satisfy $\lambda_i(M) + \lambda_j(M) \neq 0$ for all i, j ; the same holds for “ $1 \Rightarrow 1$ ”.

In [44, §6.3], the authors observed that Riemannian gradient descent running on (Q) gets stuck in a suboptimal point for a certain matrix completion problem using (SVD) but not using (MSVD). We can use Proposition 2.10 to understand their observation. Their algorithm only generates iterates with strictly positive diagonals σ in (SVD) and strictly positive-definite middle factors M in (MSVD), and can only converge to such points. Proposition 2.10 shows that (MSVD) satisfies “local \Rightarrow local” and “ $1 \Rightarrow 1$ ” in that region, while (SVD) does not.

2.4 Low-rank tensors

Tensor factorization formats correspond to lifts mapping factors to low-rank tensors, for various notions of tensor rank. For example, the canonical polyadic (CP) decomposition of rank at most 1 corresponds to the lift of $\mathcal{X} = \{X \in \mathbb{R}^{n_1 \times \dots \times n_d} : \text{CP-rank}(X) \leq 1\}$ [34] via tensor product \otimes as:

$$\mathcal{M} = \mathbb{R}^{n_1} \times \dots \times \mathbb{R}^{n_d}, \quad \varphi(v_1, \dots, v_d) = v_1 \otimes \dots \otimes v_d.$$

Other examples include CP decompositions of higher rank, Tucker and Tensor Train (TT) decompositions, and more generally tensor networks [15, 34]. Surprisingly, none of these lifts satisfy “ $2 \Rightarrow 1$ ”: we derive this from more general obstructions to “ $2 \Rightarrow 1$ ” for multilinear lifts in Sect. 5.3. Here is one take-away: any stationarity guarantees for algorithms targeting second-order critical points over the factors in a tensor decomposition must exploit the structure of the cost function.

2.5 Neural networks

Training neural networks is done via lifts. Indeed, here \mathcal{M} is the manifold of weights and biases of a fixed neural network architecture (typically a linear space; sometimes a product of spheres if normalization constraints are present). The lift φ maps a choice of weights and biases to *the function* given by the corresponding neural network. The image $\mathcal{X} = \varphi(\mathcal{M})$ of this lift is the set of functions that can be represented by the architecture, viewed as a subset of some linear space \mathcal{E} of functions (e.g., an L^p space⁴).

The authors of [46] show that such φ is not open for any choice of (nonconstant) Lipschitz continuous activation functions. Our Theorem 2.3 then implies that “local \Rightarrow local” fails for all neural network lifts used in practice. Consequently, training such a neural network by optimizing over its weights and biases might yield a spurious local

⁴ Even though these spaces are not finite-dimensional, our results still apply, see “Appendix 1”.

minimum that does not parametrize a local minimum in function space. In that case, a different parametrization of the same function might not be a local minimum for (Q).

When the neural network architecture is linear with three or more layers, the corresponding lift is multilinear, hence does not satisfy “ $2 \Rightarrow 1$ ” by the same general obstructions from Sect. 5.3 we use for tensor decompositions. Similarly to the tensor case, this implies that proofs of guarantees for training algorithms must exploit the structure in specific cost functions (the loss). Additional study of lifts defined by linear neural networks was done in [32, 54], where the authors characterize (in our terminology) “ $1 \Rightarrow 1$ ” for lifts defined by linear and linear convolutional architectures.

2.6 Submersions and higher order stationary points

All the sets \mathcal{X} we consider in this paper contain dense smooth submanifolds. Moreover, even though lifts of such sets \mathcal{X} do not satisfy “ $1 \Rightarrow 1$ ” everywhere on the lift, they do so at preimages of points on this submanifold, allowing us to prove much stronger guarantees. More precisely, we define the following subset of \mathcal{X} .

Definition 2.11 A point $x \in \mathcal{X}$ is *smooth* if there is an open neighborhood $U \subseteq \mathcal{E}$ containing x such that $U \cap \mathcal{X}$ is a smooth embedded submanifold of \mathcal{E} . It is called *nonsmooth* or *singular* otherwise.

The *smooth locus* of \mathcal{X} , denoted $\mathcal{X}^{\text{smth}}$, is the set of smooth points of \mathcal{X} .

For all constraint sets in practical optimization problems we are aware of, $\mathcal{X}^{\text{smth}}$ is itself a smooth embedded submanifold of \mathcal{E} . In general, it is a union of smooth embedded submanifolds, though possibly of different dimensions. For example, if $\mathcal{X} = \mathbb{R}_{\leq r}^{m \times n}$ then $\mathcal{X}^{\text{smth}} = \mathbb{R}_{=r}^{m \times n}$ and if $\mathcal{X} = \Delta^{n-1}$ then $\mathcal{X}^{\text{smth}} = \Delta_{>0}^{n-1}$ consisting of strictly positive simplex vectors. All the lifts we consider for these sets in Sects. 2.1 and 2.3 indeed satisfy “ $1 \Rightarrow 1$ ” on the preimages of $\mathcal{X}^{\text{smth}}$ (though that is not always the case).

If the lift satisfies “ $1 \Rightarrow 1$ ” at preimages of smooth points, then it is a submersion there and hence preserves not only local minima, but also stationary points of all orders. The following proposition, proved in Sect. 3.2.1, formalizes this.

Proposition 2.12 *Let $y \in \varphi^{-1}(\mathcal{X}^{\text{smth}}) \subseteq \mathcal{M}$. If φ satisfies “ $1 \Rightarrow 1$ ” at y , then it also satisfies “local \Rightarrow local” and “ $k \Rightarrow k$ ” for all $k \geq 1$ at y .*

Here “ $k \Rightarrow k$ ” is defined analogously to Definition 2.2, where k th-order stationarity (or “ k -criticality” for short) of $x \in \mathcal{X}^{\text{smth}}$ is defined using curves similarly to 1- and 2-criticality [13, §3.1.1]. This property can be used in proofs of benign nonconvexity.

Remark 2.13 (Relations between lift properties) Aside from Proposition 2.12, the only relation between the three properties in Definition 2.2 is that “ $1 \Rightarrow 1$ ” at y implies “ $2 \Rightarrow 1$ ” at y (since 2-critical points are 1-critical). None of the other possible implications hold in general: The desingularization lift (Desing) shows that “ $2 \Rightarrow 1$ ” at y implies neither “ $1 \Rightarrow 1$ ” nor “local \Rightarrow local” at y in general. The example $\varphi(x) = x^3$ viewed as a lift from $\mathcal{M} = \mathbb{R}$ to $\mathcal{X} = \mathbb{R}$ satisfies “local \Rightarrow local” at the origin but neither “ $2 \Rightarrow 1$ ” nor “ $1 \Rightarrow 1$ ”, hence “local \Rightarrow local” does not imply the other two properties.

Finally, the standard parametrization of the cochleoid curve [58] satisfies “ $1 \Rightarrow 1$ ” but not “local \Rightarrow local” at all preimages of the origin, hence “ $1 \Rightarrow 1$ ” does not imply “local \Rightarrow local”.

Submersions between smooth manifolds, including quotients by group actions [9, §9.2] and several lifts arising in practice (Example 3.14), satisfy “local \Rightarrow local” and “ $k \Rightarrow k$ ” for all $k \geq 1$ [9, Prop. 9.6]. Therefore, a lift φ composed with a submersion ψ as $\varphi \circ \psi$ inherits the properties of φ . We study such compositions of lifts in Sect. 3.3, and apply our results to a composition used in the robotics and computer vision literature [19] in Example 3.29.

3 Characterizations of lifts

In this section, we relate the landscapes of (P) and (Q) and prove the characterizations of our lift properties stated in Sect. 2. To this end, we formally define the different types of desirable points we consider. We first define the (*contingent* or *Bouligand*) tangent cone⁵ [17, §2.7].

Definition 3.1 The *tangent cone* to \mathcal{X} at $x \in \mathcal{X}$ is the set

$$T_x \mathcal{X} = \left\{ v = \lim_{i \rightarrow \infty} \frac{x_i - x}{\tau_i} : x_i \in \mathcal{X}, \tau_i > 0 \text{ for all } i, \tau_i \rightarrow 0 \right\} \subseteq \mathcal{E}.$$

This is a closed (not necessarily convex) cone [50, Lem. 3.12].

In particular, if γ is a differentiable curve in \mathcal{X} with $\gamma(0) = x$, then $\gamma'(0) \in T_x \mathcal{X}$. If x is a smooth point of \mathcal{X} (Definition 2.11), then $T_x \mathcal{X}$ is the usual tangent space to \mathcal{X} at x [49, Ex. 6.8].

Definition 3.2 (*Desirable points for (P)*) A point $x \in \mathcal{X}$ is a

- (a) *global minimum* for (P) if $f(x) = \min_{x' \in \mathcal{X}} f(x')$.
- (b) *local minimum* for (P) if there is a neighborhood $U \subseteq \mathcal{X}$ of x such that $f(x) = \min_{x' \in U} f(x')$.
- (c) (*first-order*) *stationary point* for (P) if $Df(x)[v] \geq 0$ for all $v \in T_x \mathcal{X}$, or equivalently, if $\nabla f(x)$ is in the dual $(T_x \mathcal{X})^*$ of the tangent cone.

In words, x is stationary if the cost function is non-decreasing to first order along all tangent directions at x . Local minima of (P) are stationary [50, Thm. 3.24]. The dual of a cone $K \subseteq \mathcal{E}$ contained in a Euclidean space \mathcal{E} with inner product $\langle \cdot, \cdot \rangle$ is defined by

$$K^* = \{x \in \mathcal{E} : \langle x, x' \rangle \geq 0 \text{ for all } x' \in K\}.$$

The equivalence in part (c) then follows since $Df(x)[v] = \langle \nabla f(x), v \rangle$ by definition of the (Euclidean) gradient $\nabla f(x)$. We use the following properties of dual cones throughout (see [20, Prop. 4.5] for proofs):

- The dual cone is always a closed convex cone.
- If $K_1 \subseteq K_2$, then $K_2^* \subseteq K_1^*$.

⁵ In this paper, a cone is a set K such that $x \in K \implies \alpha x \in K$ for all $\alpha > 0$.

- The bidual cone $K^{**} = (K^*)^*$ of K is equal to the closure of its convex hull: $K^{**} = \overline{\text{conv}}(K)$. In particular, $K^{**} \supseteq K$.
- If K is a linear space, then its dual K^* is equal to its orthogonal complement K^\perp .

Next, we define desirable points for (Q).

Definition 3.3 (*Desirable points for (Q)*)

- (a)+(b) Global and local minima for (Q) are defined exactly as for (P).
- (c) A point $y \in \mathcal{M}$ is *first-order stationary* (or “1-critical”) for (Q) if for each smooth curve $c: \mathbb{R} \rightarrow \mathcal{M}$ satisfying $c(0) = y$, we have $(g \circ c)'(0) \geq 0$, or equivalently,⁶ $(g \circ c)'(0) = 0$.
- (d) A point $y \in \mathcal{M}$ is *second-order stationary* (or “2-critical”) for (Q) if it is 1-critical and $(g \circ c)''(0) \geq 0$ for all smooth curves $c: \mathbb{R} \rightarrow \mathcal{M}$ satisfying $c(0) = y$.

If \mathcal{M} is embedded in a linear space, first-order stationarity in Definition 3.3(c) coincides with Definition 3.2(c) by [49, Ex. 6.8]. Definition 3.3 can be rephrased in terms of the Riemannian gradient and Hessian of g , as follows.

Proposition 3.4 [9, §4.2, §6.1] *A point $y \in \mathcal{M}$ is 1-critical for (Q) if and only if $\nabla g(y) = 0$. It is 2-critical if and only if $\nabla g(y) = 0$ and $\nabla^2 g(y) \geq 0$.*

We proceed to study the connections between desirable points for (Q) and (P). As mentioned in Sect. 2, the connection between global minima of (Q) and (P) is straightforward.

Proposition 3.5 *A point $y \in \mathcal{M}$ is a global minimum of (Q) if and only if $x = \varphi(y)$ is a global minimum of (P).*

Proof Because $\varphi(\mathcal{M}) = \mathcal{X}$, we have $\inf_{y \in \mathcal{M}} g(y) = \inf_{y \in \mathcal{M}} f(\varphi(y)) = \inf_{x \in \mathcal{X}} f(x) =: p^*$. Therefore, y is a global minimum for (Q) iff $g(y) = f(x) = p^*$ which happens iff x is a global minimum for (P). □

Since computing global minima is hard, the remainder of this section is devoted to characterizing the properties in Definition 2.2 that yield connections between the other types of points.

3.1 Local minima

In this section, we investigate the relationship between the local minima of (P) and those of (Q). Preimages of local minima on \mathcal{X} are always local minima on \mathcal{M} merely because φ is continuous.

Proposition 3.6 *Let x be a local minimum for (P). Any $y \in \varphi^{-1}(x)$ is a local minimum for (Q).*

Proof There exists a neighborhood U of x in \mathcal{X} such that $f(x) \leq f(x')$ for all $x' \in U$. Since $\varphi: \mathcal{M} \rightarrow \mathcal{X}$ is continuous, the set $\mathcal{U} = \varphi^{-1}(U)$ is a neighborhood of y in \mathcal{M} . Pick an arbitrary $y' \in \mathcal{U}$: it satisfies $\varphi(y') = x'$ for some $x' \in U$. Hence, $g(y) = f(x) \leq f(x') = g(y')$, i.e., y is a local minimum of (Q). □

⁶ If $(g \circ c)'(0) > 0$, let $\tilde{c}(t) = c(-t)$ and note that $(g \circ \tilde{c})'(0) < 0$.




Fig. 1 Nodal cubic in \mathbb{R}^2 as the shadow of its lift in \mathbb{R}^3 , shaded by the value of the function $f(x) = -x_1 - x_2$. The highlighted points are $x = (0, 0)$ (not stationary for **(P)**), and its two preimages on the lift, including $y = (0, 0, 1)$ (a spurious local minimum for **(Q)**)

Unfortunately, lifting can introduce *spurious* local minima, that is, local minima for **(Q)** that exist only because of the lift and not because they were present in **(P)** to begin with.

Example 3.7 (Nodal cubic) Consider the nodal cubic

$$\mathcal{X} = \{x \in \mathbb{R}^2 : x_2^2 = x_1^2(x_1 + 1)\}, \quad (3.1)$$

and the following lift,⁷ as depicted in Fig. 1:

$$\mathcal{M} = \{y \in \mathbb{R}^3 : y_1 = y_3^2 - 1, y_2 = y_1 y_3\}, \quad \varphi(y_1, y_2, y_3) = (y_1, y_2). \quad (3.2)$$

Let $f(x) = -x_1 - x_2$. Then the point $y = (0, 0, 1)$ is a local minimum for $g = f \circ \varphi$ but $\varphi(y) = (0, 0)$ is not even stationary for f . Indeed, we have $(1, 1) \in T_{(0,0)}\mathcal{X}$ and $Df(0, 0)[(1, 1)] = -2 < 0$.

To ensure that a lift does not introduce spurious local minima, we need to verify that it satisfies the “local \Rightarrow local” property (Definition 2.2(a)). We proceed to prove the easy direction of Theorem 2.3 stating that openness implies “local \Rightarrow local”. The converse is more involved and is deferred to “Appendix 1”.

Proof of Theorem 2.3 Assume φ is open at y , and that y is a local minimum for **(Q)**. Then, there exists a neighborhood \mathcal{U} of y on \mathcal{M} such that $g(y) \leq g(y')$ for all $y' \in \mathcal{U}$. The set $U = \varphi(\mathcal{U})$ is a neighborhood of $x = \varphi(y)$ in \mathcal{X} by openness of φ at y . Moreover, each $x' \in U$ is of the form $x' = \varphi(y')$ for some $y' \in \mathcal{U}$. Therefore, $f(x) = g(y) \leq g(y') = f(x')$ for all $x' \in U$, that is, x is a local minimum of **(P)**. For the converse, see Theorem A.2. \square

⁷ The curve \mathcal{M} is obtained by blowing up \mathcal{X} at the origin in the sense of algebraic geometry [26, Ch. 17].

Not all lifts of interest are open. In particular, all lifts of low-rank matrices in Sect. 2.3 as well as the neural network lifts in Sect. 2.5 fail to be open.

Remark 3.8 In “Appendix 1”, we introduce an equivalent condition for openness of φ at y that we call the Subsequence Lifting Property (SLP), see Definition A.1(3); we find that it is sometimes easier to check. For example, Burer and Monteiro prove that the lift (BM) satisfies “local \Rightarrow local” in [12, Prop. 2.3] by (in our terminology) proving SLP holds.

We note in passing that all continuous, surjective, open maps are quotient maps, hence if φ is a smooth lift of \mathcal{X} satisfying “local \Rightarrow local” then it is a quotient map from \mathcal{M} to \mathcal{X} .

3.2 Stationary points

In this section, we investigate the relationship between the first- and second-order stationary points for (Q) and (first-order) stationary points for (P). To that end, we begin by relating the (Riemannian) gradient and Hessian of $g = f \circ \varphi$ to the (Euclidean) counterparts of f . This relation depends on the first and second derivatives of the lift φ .

Definition 3.9 Let $\varphi: \mathcal{M} \rightarrow \mathcal{X}$ be a smooth lift and fix $y \in \mathcal{M}$. For each $v \in T_y\mathcal{M}$, choose a curve c_v on \mathcal{M} satisfying $c(0) = y$ and $c'(0) = v$. Define maps $\mathbf{L}_y, \mathbf{Q}_y: T_y\mathcal{M} \rightarrow \mathcal{E}$ by

$$\mathbf{L}_y(v) = (\varphi \circ c_v)'(0), \quad \mathbf{Q}_y(v) = (\varphi \circ c_v)''(0).$$

We write \mathbf{L}_y^φ and \mathbf{Q}_y^φ when we wish to emphasize the lift.

As a point of notation: $\varphi \circ c_v$ is a curve in \mathcal{E} hence $(\varphi \circ c_v)''$ denotes its Euclidean acceleration. In contrast, c_v is a curve on \mathcal{M} hence c_v'' denotes its Riemannian acceleration, see [9, §5.8, §8.12].

Of course, \mathbf{L}_y is simply the differential $D\varphi(y)$, and is therefore linear and independent of the choice of curves c_v . The map \mathbf{Q}_y will play an important role in characterizing “2 \Rightarrow 1” in Sect. 3.2.2, where we also clarify its inconsequential dependence on the choice of curve c_v . We explain how to compute \mathbf{L}_y and \mathbf{Q}_y without explicitly choosing curves c_v in Sect. 3.4.

The gradients and Hessians of f and $g = f \circ \varphi$ are neatly related as follows in terms of \mathbf{L}_y .

Definition 3.10 For any $w \in \mathcal{E}$, define $\varphi_w: \mathcal{M} \rightarrow \mathbb{R}$ by $\varphi_w(y) = \langle w, \varphi(y) \rangle$.

Lemma 3.11 For any twice differentiable cost $f: \mathcal{E} \rightarrow \mathbb{R}$, any $y \in \mathcal{M}$, and $x = \varphi(y)$, we have

$$\nabla g(y) = \mathbf{L}_y^*(\nabla f(x)), \quad \nabla^2 g(y) = \mathbf{L}_y^* \circ \nabla^2 f(x) \circ \mathbf{L}_y + \nabla^2 \varphi_{\nabla f(x)}(y),$$

where $\mathbf{L}_y^*: \mathcal{E} \rightarrow T_y\mathcal{M}$ is the adjoint of \mathbf{L}_y .

Proof For any $v \in T_y\mathcal{M}$, let c_v be a smooth curve on \mathcal{M} satisfying $c_v(0) = y$, $c'_v(0) = v$ and $c''_v(0) = 0$ (e.g., let c_v be a geodesic). Let $\gamma_v = \varphi \circ c_v$: it satisfies $\gamma_v(0) = x$ and $\gamma'_v(0) = \mathbf{L}_y(v)$. Then

$$\langle \nabla g(y), v \rangle = (g \circ c_v)'(0) = (f \circ \gamma_v)'(0) = \langle \nabla f(x), \mathbf{L}_y(v) \rangle = \langle \mathbf{L}_y^*(\nabla f(x)), v \rangle.$$

Since this holds for all $v \in T_y\mathcal{M}$, we conclude that $\nabla g(y) = \mathbf{L}_y^*(\nabla f(x))$. Next,

$$\begin{aligned} \langle \nabla^2 g(y)[v], v \rangle &= (g \circ c_v)''(0) = (f \circ \gamma_v)''(0) \\ &= \langle \nabla^2 f(x)[\mathbf{L}_y(v)], \mathbf{L}_y(v) \rangle + \langle \nabla f(x), \gamma_v''(0) \rangle, \end{aligned}$$

where the first equality uses $c''_v(0) = 0$, see [9, §5.9]. On the other hand, with Definition 3.10,

$$\begin{aligned} \langle \nabla^2 \varphi_{\nabla f(x)}(y)[v], v \rangle &= (\varphi_{\nabla f(x)} \circ c_v)''(0) \\ &= \left. \frac{d^2}{dt^2} \langle \nabla f(x), \gamma_v(t) \rangle \right|_{t=0} = \langle \nabla f(x), \gamma_v''(0) \rangle, \end{aligned}$$

hence

$$\begin{aligned} \langle \nabla^2 g(y)[v], v \rangle &= \langle \nabla^2 f(x)[\mathbf{L}_y(v)], \mathbf{L}_y(v) \rangle + \langle \nabla^2 \varphi_{\nabla f(x)}(y)[v], v \rangle \\ &= \left\langle \left(\mathbf{L}_y^* \circ \nabla^2 f(x) \circ \mathbf{L}_y + \nabla^2 \varphi_{\nabla f(x)}(y) \right) [v], v \right\rangle. \end{aligned}$$

Since this holds for all $v \in T_y\mathcal{M}$ and both $\nabla^2 g(y)$ and $\mathbf{L}_y^* \circ \nabla^2 f(x) \circ \mathbf{L}_y + \nabla^2 \varphi_{\nabla f(x)}(y)$ are self-adjoint linear maps on $T_y\mathcal{M}$, we conclude that they are equal. \square

We turn to proving our characterizations of “ $k \Rightarrow 1$ ” for $k = 1, 2$ announced in Sect. 2.

3.2.1 “1 \Rightarrow 1’’: lifts preserving 1-critical points

Preimages of stationary points on \mathcal{X} are always 1-critical on \mathcal{M} . We show this after a helpful lemma.

Lemma 3.12 Fix $y \in \mathcal{M}$ and let $x = \varphi(y)$. Then $\text{im } \mathbf{L}_y \subseteq T_x\mathcal{X}$. Moreover, y is 1-critical for (Q) if and only if $\nabla f(x) \in (\text{im } \mathbf{L}_y)^\perp = (\text{im } \mathbf{L}_y)^*$.

Proof The first claim follows from Definition 3.1 for the tangent cone $T_x\mathcal{X}$ and the fact that $\mathbf{L}_y(v) = (\varphi \circ c_v)'(0)$ for a curve c_v as in Definition 3.9. For the second claim, y is 1-critical for (Q) iff $\nabla g(y) = \mathbf{L}_y^*(\nabla f(x)) = 0$, or equivalently, $\nabla f(x) \in \ker(\mathbf{L}_y^*) = (\text{im } \mathbf{L}_y)^\perp$. \square

Proposition 3.13 If $x \in \mathcal{X}$ is stationary for (P), then any $y \in \varphi^{-1}(x)$ is 1-critical for (Q).

Proof If $x \in \mathcal{X}$ is stationary for (P), then $\nabla f(x) \in (T_x \mathcal{X})^*$. Since $T_x \mathcal{X} \supseteq \text{im } \mathbf{L}_y$, taking duals on both sides we get that $\nabla f(x) \in (T_x \mathcal{X})^* \subseteq (\text{im } \mathbf{L}_y)^\perp$, hence y is 1-critical for (Q) by Lemma 3.12. \square

The converse to Proposition 3.13 is false in general. In fact, Example 3.7 shows that a lift need not even map local minima to stationary points on \mathcal{X} . We therefore proceed to prove Theorem 2.4 characterizing the “ $1 \Rightarrow 1$ ” property.

Proof of Theorem 2.4 Suppose $\text{im } \mathbf{L}_y = T_x \mathcal{X}$, so $(\text{im } \mathbf{L}_y)^\perp = (\text{im } \mathbf{L}_y)^* = (T_x \mathcal{X})^*$. If y is 1-critical for (Q), then $\nabla f(x) \in (\text{im } \mathbf{L}_y)^\perp$ by Lemma 3.12. Therefore, $\nabla f(x) \in (T_x \mathcal{X})^*$, which is the definition of x being stationary for (P). Thus, “ $1 \Rightarrow 1$ ” holds.

Now suppose $\text{im } \mathbf{L}_y \neq T_x \mathcal{X}$. This implies $(T_x \mathcal{X})^* \neq (\text{im } \mathbf{L}_y)^\perp$. Indeed, otherwise we would have

$$\text{im } \mathbf{L}_y = (\text{im } \mathbf{L}_y)^{\perp\perp} = \left((\text{im } \mathbf{L}_y)^\perp \right)^* = (T_x \mathcal{X})^{**} \supseteq T_x \mathcal{X} \supseteq \text{im } \mathbf{L}_y,$$

which would imply $\text{im } \mathbf{L}_y = T_x \mathcal{X}$. (The right-most inclusion above is by Lemma 3.12.) Using $\text{im } \mathbf{L}_y \subseteq T_x \mathcal{X}$ again, we see that $(T_x \mathcal{X})^* \subseteq (\text{im } \mathbf{L}_y)^\perp$. Therefore, the above observations imply that $(T_x \mathcal{X})^* \subsetneq (\text{im } \mathbf{L}_y)^\perp$. Pick $w \in (\text{im } \mathbf{L}_y)^\perp \setminus (T_x \mathcal{X})^*$ and define $f(x') = \langle w, x' \rangle$ for $x' \in \mathcal{E}$. Then $\nabla f(x) = w \in (\text{im } \mathbf{L}_y)^\perp$ so y is 1-critical for (Q) by Lemma 3.12, but $\nabla f(x) \notin (T_x \mathcal{X})^*$, so x is not stationary for (P). Hence “ $1 \Rightarrow 1$ ” is not satisfied at y .

This argument also shows that if φ does not satisfy “ $1 \Rightarrow 1$ ” at y , then y is 1-critical for (Q) but x is not stationary for (P) if and only if $\nabla f(x) \in (\text{im } \mathbf{L}_y)^\perp \setminus (T_x \mathcal{X})^*$, showing that if “ $1 \Rightarrow 1$ ” fails at y then this is witnessed by a linear cost f . \square

As discussed in Sect. 2, Theorem 2.4 implies that “ $1 \Rightarrow 1$ ” rarely holds on all of \mathcal{M} . Nevertheless, “ $1 \Rightarrow 1$ ” does usually hold at preimages of smooth points, that is, points around which \mathcal{X} is a smooth embedded submanifold of \mathcal{E} as in Definition 2.11. We now prove Proposition 2.12, stating that if “ $1 \Rightarrow 1$ ” holds at such points then “local \Rightarrow local” and “ $k \Rightarrow k$ ” hold there as well.

Proof of Proposition 2.12 Let $U \subseteq \mathcal{E}$ be an open neighborhood of $\varphi(y)$ in \mathcal{E} such that $U \cap \mathcal{X}$ is a smooth embedded submanifold of \mathcal{E} . Since $\varphi(\mathcal{M}) = \mathcal{X}$, we have $\varphi^{-1}(U \cap \mathcal{X}) = \varphi^{-1}(U) =: V$, which is open in \mathcal{M} by continuity of φ . Therefore, V is also a smooth manifold, since it is an open subset of \mathcal{M} , and $\varphi|_V : V \rightarrow U \cap \mathcal{X}$ is a smooth map between smooth manifolds. By Theorem 2.4, φ satisfies “ $1 \Rightarrow 1$ ” at y iff $T_x \mathcal{X} = \text{im } D\varphi(y) = \text{im } D(\varphi|_V)(y)$, where $\varphi|_V$ is viewed as a map $V \rightarrow \mathcal{E}$. Since $U \cap \mathcal{X}$ is an embedded submanifold of \mathcal{E} , the differential of $\varphi|_V$ viewed as a map $V \rightarrow \mathcal{E}$ coincides with its differential viewed as a map $V \rightarrow U \cap \mathcal{X}$, hence the latter is a submersion near y [37, Prop. 4.1]. By [37, Prop. 4.28], this implies φ is open at y , hence it satisfies “local \Rightarrow local” at y by Theorem 2.3. To see that φ further satisfies “ $k \Rightarrow k$ ” for all $k \geq 1$, note that any curve passing through $\varphi(y)$ is the image under φ of a curve passing through y [37, Thm. 4.26], and apply Definition 3.3 for $k = 1, 2$ and [13, Eq. (3.11)] for $k > 2$. \square

The converse of Proposition 2.12 fails. For example, $\varphi(y) = y^3$ viewed as a map $\mathbb{R} \rightarrow \mathbb{R}$ satisfies “local \Rightarrow local” at $y = 0$ but not “ $1 \Rightarrow 1$ ” since $\mathbf{L}_y = 0$ for this

y . That example also shows that “ $1 \Rightarrow 1$ ” can fail at the preimage of a smooth point. Likewise, “ $1 \Rightarrow 1$ ” can hold at the preimage of a nonsmooth point, as the standard parametrization of the cochleoid curve [58] shows. The only examples of lifts we know of that satisfy “ $1 \Rightarrow 1$ ” everywhere are smooth maps between smooth manifolds that are submersions.

Example 3.14 (Submersions) Examples of submersions in optimization, that is, lifts of the form $\varphi: \mathcal{M} \rightarrow \mathcal{X}$ where \mathcal{X} is an embedded submanifold of \mathcal{E} and $\text{im } D\varphi(y) = T_{\varphi(y)}\mathcal{X}$ for all $y \in \mathcal{M}$, include:

- Quotient maps by smooth, free, and proper Lie group actions [9, §9], [2, §3.4], used in particular to optimize over Grassmannians by lifting to Stiefel manifolds [23].
- The map $\text{SO}(p) \rightarrow \text{St}(p, d)$ taking the first d columns of a rotation matrix, which is used in the rotation averaging algorithm of the robotics paper [19], see Example 3.29 below.
- Deep orthogonal linear networks, mapping $\text{O}(p)^n \rightarrow \text{O}(p)$ by $\varphi(Q_1, \dots, Q_n) = Q_1 \cdots Q_n$, whose properties are studied in [1].

Theorem 2.4 and Proposition 2.12 show that these lifts satisfy “ $1 \Rightarrow 1$ ” and “local \Rightarrow local”, and hence also “ $k \Rightarrow k$ ” for all $k \geq 1$.

Failure of a lift to satisfy “ $1 \Rightarrow 1$ ” means that it may introduce spurious critical points. In the next section, we characterize the “ $2 \Rightarrow 1$ ” property, which allows algorithms to avoid these spurious points by using second-order information.

3.2.2 “ $2 \Rightarrow 1$ ”: lifts mapping 2-critical points to 1-critical points

Since “ $1 \Rightarrow 1$ ” fails on many sets of interest, we proceed to study “ $2 \Rightarrow 1$ ”. As Sect. 2 demonstrates, this property is satisfied for many interesting lifts. We begin by stating an equivalent characterization for “ $2 \Rightarrow 1$ ” involving the following set.

Definition 3.15 For $y \in \mathcal{M}$ and $x = \varphi(y) \in \mathcal{X}$, define

$$W_y = \left\{ w \in \mathcal{E} : \text{there exists a twice differentiable function } f: \mathcal{E} \rightarrow \mathbb{R} \text{ such that } \nabla f(x) = w \text{ and } y \text{ is } 2\text{-critical for } (Q) \right\}.$$

We write W_y^φ when we wish to emphasize the lift.

Theorem 3.16 *The lift $\varphi: \mathcal{M} \rightarrow \mathcal{X}$ satisfies “ $2 \Rightarrow 1$ ” at y if and only if $W_y \subseteq (T_x \mathcal{X})^*$ where $x = \varphi(y)$.*

Proof Say φ satisfies “ $2 \Rightarrow 1$ ” at y and let $w \in W_y$. Pick f such that y is 2-critical for (Q) and $\nabla f(x) = w$. By “ $2 \Rightarrow 1$ ”, we know x is stationary for f , hence $w = \nabla f(x) \in (T_x \mathcal{X})^*$. Conversely, say $W_y \subseteq (T_x \mathcal{X})^*$ and let y be 2-critical for (Q) with some cost f . Then $\nabla f(x) \in W_y \subseteq (T_x \mathcal{X})^*$, hence x is stationary for (P) . This shows “ $2 \Rightarrow 1$ ”. \square

Since the 2-criticality of y for (Q) only depends on the first two derivatives of f , we can restrict the functions f in Definition 3.15 to be of class C^∞ or even quadratic polynomials whose Hessians are a multiple of the identity, as the following proposition shows.

Proposition 3.17 *For $y \in \mathcal{M}$ and $x = \varphi(y)$, the set W_y in Definition 3.15 satisfies:*

$$W_y = \left\{ w \in \mathcal{E} : \exists \alpha > 0 \text{ s.t. } y \text{ is 2-critical for (Q) with } f(x') = \langle x', w \rangle + \frac{\alpha}{2} \|x' - x\|^2 \right\}.$$

In particular, W_y is a convex cone.

Proof The inclusion \supseteq is clear from the definition of W_y . Conversely, if w is in W_y then y is 2-critical for (Q) for some f with $\nabla f(x) = w$. Let $g = f \circ \varphi$ and $\alpha = \lambda_{\max}(\nabla^2 f(x))$, and define

$$f_\alpha(x') = \langle w, x' \rangle + \frac{\alpha}{2} \|x - x'\|^2, \quad g_\alpha = f_\alpha \circ \varphi.$$

Note that $\nabla f_\alpha(x) = w$ and, by Lemma 3.11, we have $\nabla g_\alpha(y) = \mathbf{L}_y^*(w) = \nabla g(y) = 0$ and

$$\begin{aligned} \nabla^2 g_\alpha(y) &= \mathbf{L}_y^* \circ \nabla^2 f_\alpha(x) \circ \mathbf{L}_y + \nabla^2 \varphi_{\nabla f_\alpha(x)}(y) = \mathbf{L}_y^* \circ (\alpha I) \circ \mathbf{L}_y + \nabla^2 \varphi_w(y) \\ &\geq \mathbf{L}_y^* \circ \nabla^2 f(x) \circ \mathbf{L}_y + \nabla^2 \varphi_w(y) = \nabla^2 g(y) \geq 0. \end{aligned}$$

Thus, y is 2-critical for g_α . This shows the reverse inclusion.

W_y is a convex cone since if w_1, w_2 are in W_y as witnessed by functions f_1, f_2 , then any $w = \lambda_1 w_1 + \lambda_2 w_2$ with $\lambda_1, \lambda_2 \geq 0$ is in W_y as witnessed by $f = \lambda_1 f_1 + \lambda_2 f_2$. □

Proposition 3.17 shows that if “2 \Rightarrow 1” is not satisfied at y , then there exists a simple strongly convex quadratic cost f for which y is 2-critical for (Q) but $x = \varphi(y)$ is not stationary for (P).

Corollary 3.18 *Suppose φ does not satisfy “2 \Rightarrow 1” at $y \in \mathcal{M}$ and denote $x = \varphi(y)$. Then $W_y \setminus (T_x \mathcal{X})^* \neq \emptyset$ and for any w in that set there exists $\alpha > 0$ such that if $f(x') = \langle w, x' \rangle + \frac{\alpha}{2} \|x' - x\|^2$, then y is 2-critical for (Q) but x is not stationary for (P).*

We conjecture that the reverse inclusion in Theorem 3.16 always holds (it does for all the lifts in Sect. 2). If this is indeed true, then φ satisfies the “2 \Rightarrow 1” property at y if and only if $(T_x \mathcal{X})^* = W_y$, neatly echoing the condition for “1 \Rightarrow 1”, namely, $(T_x \mathcal{X})^* = (\text{im } \mathbf{L}_y)^\perp$.

Conjecture 3.19 *It always holds that $(T_x \mathcal{X})^* \subseteq W_y$.*

The description of W_y can be complicated. It is therefore worthwhile to derive sufficient conditions for “2 \Rightarrow 1” that are easier to check. We do so by identifying two (admittedly technical) sets whose duals contain $\nabla f(x)$ if $x = \varphi(y)$ and y is 2-critical for (Q). The sufficient conditions then require the duals of these two subsets to be contained in $(T_x \mathcal{X})^*$.

Definition 3.20 For $y \in \mathcal{M}$, define

$$A_y = \{w \in \mathcal{E} : \exists c: \mathbb{R} \rightarrow \mathcal{M} \text{ smooth s.t. } c(0) = y, (\varphi \circ c)'(0) = 0, (\varphi \circ c)''(0) = w\},$$

$$B_y = \{w \in \mathcal{E} : \exists c_i: \mathbb{R} \rightarrow \mathcal{M} \text{ smooth s.t. } c_i(0) = y, \lim_{i \rightarrow \infty} (\varphi \circ c_i)'(0) = 0, \lim_{i \rightarrow \infty} (\varphi \circ c_i)''(0) = w\}.$$

We write A_y^φ, B_y^φ when we wish to emphasize the lift.

The following are the basic properties these two sets satisfy. We give further expressions for A_y, B_y and W_y in Proposition 3.26 below.

Proposition 3.21 Fix $y \in \mathcal{M}$ and denote $x = \varphi(y)$.

- (a) A_y and B_y are cones, and B_y is closed.
- (b) $A_y \subseteq T_x \mathcal{X}$ and $\text{im } \mathbf{L}_y \subseteq A_y \subseteq B_y$. Moreover, $\text{im } \mathbf{L}_y + A_y = A_y$ and $\text{im } \mathbf{L}_y + B_y = B_y$.
- (c) If y is 2-critical for $g = f \circ \varphi$, then $\nabla f(x) \in B_y^*$.
- (d) $W_y \subseteq B_y^* \subseteq A_y^* \subseteq (\text{im } \mathbf{L}_y)^\perp$.

Proof Part (a) and the second half of part (b) are straightforward, see [39, App. B].

- (b) If $c: \mathbb{R} \rightarrow \mathcal{M}$ satisfies $c(0) = y$ and $(\varphi \circ c)'(0) = 0$, then $(\varphi \circ c)(t) \in \mathcal{X}$ for all t and $(\varphi \circ c)(t) = x + (t^2/2)(\varphi \circ c)''(0) + \mathcal{O}(t^3)$, hence by Definition 3.1 we have

$$(\varphi \circ c)''(0) = \lim_{t \rightarrow 0} \frac{(\varphi \circ c)(t) - x}{t^2/2} \in T_x \mathcal{X}.$$

This shows $A_y \subseteq T_x \mathcal{X}$. If $w \in \text{im } \mathbf{L}_y$ so $w = \mathbf{L}_y(v)$ for some $v \in T_y \mathcal{M}$, let $c: \mathbb{R} \rightarrow \mathcal{M}$ be a curve satisfying $c(0) = y$ and $c'(0) = v$. Define $\tilde{c}: \mathbb{R} \rightarrow \mathcal{M}$ by $\tilde{c}(t) = c(t^2/2)$, and note that $\tilde{c}(0) = y$, $(\varphi \circ \tilde{c})'(0) = 0$, and $(\varphi \circ \tilde{c})''(0) = (\varphi \circ c)'(0) = w$. Hence w is in A_y . This shows $\text{im } \mathbf{L}_y \subseteq A_y$.

It is clear that $A_y \subseteq B_y$ from Definition 3.20.

- (c) Suppose y is 2-critical for $g = f \circ \varphi$ and $w \in B_y$. Let $c_i: \mathbb{R} \rightarrow \mathcal{M}$ witness $w \in B_y$. Because y is 1-critical, $(g \circ c_i)'(0) = 0$ for all i . Because y is 2-critical, for all i we have

$$(g \circ c_i)''(0) = \langle \nabla f(x), (\varphi \circ c_i)''(0) \rangle + \langle \nabla^2 f(x)[(\varphi \circ c_i)'(0)], (\varphi \circ c_i)'(0) \rangle \geq 0.$$

Taking $i \rightarrow \infty$, we conclude that $\langle \nabla f(x), w \rangle \geq 0$ and hence $\nabla f(x) \in B_y^*$ as claimed.

- (d) If $w \in W_y$ then there exists f such that $\nabla f(x) = w$ and y is 2-critical for (Q), hence $w \in B_y^*$ by part (c). The other inclusions follow by taking duals in part (b).

□

We remark that neither the inclusion $B_y \subseteq T_x \mathcal{X}$ nor $T_x \mathcal{X} \subseteq B_y$ hold in general, see [39, App. B]. Combining part (d) above with Theorem 3.16 yields the following sufficient conditions for “2 \Rightarrow 1”.

Corollary 3.22 Fix $y \in \mathcal{M}$ and denote $x = \varphi(y)$.

- (a) If $A_y^* \subseteq (T_x \mathcal{X})^*$, and in particular if $A_y = T_x \mathcal{X}$, then “ $2 \Rightarrow 1$ ” holds at y .
- (b) If $B_y^* \subseteq (T_x \mathcal{X})^*$, then “ $2 \Rightarrow 1$ ” holds at y .

The conditions in Corollary 3.22 yield simpler proofs of “ $2 \Rightarrow 1$ ” for many lifts. For example, the condition in Corollary 3.22(a) holds for the Burer–Monteiro lift of Sect. 2.2. While it does not hold for the lifts of low-rank matrices in Sect. 2.3, they do satisfy the stronger condition in Corollary 3.22(b). In fact, the condition in Corollary 3.22(b) holds in all the examples satisfying “ $2 \Rightarrow 1$ ” that we consider. It would be interesting to determine whether it is necessary as well.

We now state and prove the chain of implications we find the most useful for verifying or refuting “ $2 \Rightarrow 1$ ”, as well as for computing tangent cones (see Sect. 3.5).

Theorem 3.23 *Let $\varphi: \mathcal{M} \rightarrow \mathcal{X}$ be a smooth lift and fix $y \in \mathcal{M}$. We have the following chain of implications for “ $2 \Rightarrow 1$ ”:*

$$\begin{aligned}
 & T_x \mathcal{X} \subseteq A_y \\
 \iff & T_x \mathcal{X} = A_y \\
 \implies & B_y^* \subseteq (T_x \mathcal{X})^* \\
 \implies & W_y \subseteq (T_x \mathcal{X})^* \\
 \iff & \varphi \text{ satisfies “} 2 \Rightarrow 1 \text{” at } y \\
 \implies & (\text{im } \mathbf{L}_y)^\perp \cap (\mathbf{Q}_y(T_y \mathcal{M}))^* \subseteq (T_x \mathcal{X})^*.
 \end{aligned}$$

Proof The equivalence of the first two conditions follows by Proposition 3.21(b). The second condition implies the third by Proposition 3.21(b) as well. The third condition implies the fourth by Proposition 3.21(d), which itself is equivalent to “ $2 \Rightarrow 1$ ” at y by Theorem 3.16.

The last implication gives a necessary condition for “ $2 \Rightarrow 1$ ” to hold. Suppose there exists $w \in (\text{im } \mathbf{L}_y)^\perp \cap (\mathbf{Q}_y(T_y \mathcal{M}))^* \setminus (T_x \mathcal{X})^*$. Define $f(x') = \langle w, x' \rangle$, whose gradient and Hessian at x are $\nabla f(x) = w$ and $\nabla^2 f(x) = 0$. For any curve $c: I \rightarrow \mathcal{M}$ satisfying $c(0) = y$, denote $v = c'(0) \in T_y \mathcal{M}$. Let $g = f \circ \varphi$. Note that $(g \circ c)'(0) = \langle w, \mathbf{L}_y(v) \rangle = 0$ since $w \in (\text{im } \mathbf{L}_y)^\perp$ and

$$(g \circ c)''(0) = (f \circ \varphi \circ c)''(0) = \langle w, (\varphi \circ c)''(0) \rangle = \langle w, \mathbf{Q}_y(v) \rangle \geq 0,$$

where the second equality follows from the chain rule, the third equality follows from Lemma 3.24(a) below, and the inequality follows from $w \in (\mathbf{Q}_y(T_y \mathcal{M}))^*$. Thus, y is 2-critical for (\mathbf{Q}) . However, $\nabla f(x) = w \notin (T_x \mathcal{X})^*$ so x is not stationary for (\mathbf{P}) , hence “ $2 \Rightarrow 1$ ” does not hold at y . □

Our goal now is to derive more explicit expressions for the sets A_y, B_y, W_y in terms of the maps \mathbf{L}_y and \mathbf{Q}_y from Definition 3.9. Such expressions allow us to compute these sets in specific examples. To do so, we first recall that the value of $\mathbf{Q}_y(v)$ depends on the choice of curve c_v in Definition 3.9. Before proceeding, we characterize the ambiguity in $\mathbf{Q}_y(v)$ arising from different such choices, verifying that it causes no issues.

Lemma 3.24 For each $y \in \mathcal{M}$ and $v \in T_y\mathcal{M}$, let $c_v: I \rightarrow \mathcal{M}$ be a curve satisfying $c_v(0) = y$ and $c'_v(0) = v$, so we can set $\mathbf{Q}_y(v) = (\varphi \circ c_v)''(0)$ according to Definition 3.9.

- (a) For any other curve c satisfying $c(0) = y$ and $c'(0) = v$, we have $(\varphi \circ c)''(0) - (\varphi \circ c_v)''(0) = \mathbf{L}_y(c''(0) - c''_v(0)) \in \text{im } \mathbf{L}_y$.
- (b) For any $u \in T_y\mathcal{M}$, there exists a curve c as in part (a) satisfying $c''(0) - c''_v(0) = u$, hence $(\varphi \circ c)''(0) - (\varphi \circ c_v)''(0) = \mathbf{L}_y(u)$.

In particular, $\{(\varphi \circ c)''(0) : c(0) = y \text{ and } c'(0) = v\} = \mathbf{Q}_y(v) + \text{im } \mathbf{L}_y$.

Proof (a) For any $w \in \mathcal{E}$, recall the function $\varphi_w(y) = \langle w, \varphi(y) \rangle$ from Definition 3.10.

Let $c: I \rightarrow \mathcal{M}$ be a curve satisfying $c(0) = y$ and $c'(0) = v$. Then, on the one hand,

$$\left. \frac{d^2}{dt^2} \varphi_w(c(t)) \right|_{t=0} = \left. \frac{d^2}{dt^2} \langle w, (\varphi \circ c)(t) \rangle \right|_{t=0} = \langle w, (\varphi \circ c)''(0) \rangle.$$

On the other hand, using the Riemannian structure on \mathcal{M} ,

$$\left. \frac{d^2}{dt^2} \varphi_w(c(t)) \right|_{t=0} = \langle \nabla^2 \varphi_w(y)[c'(0)], c'(0) \rangle + \langle \nabla \varphi_w(y), c''(0) \rangle.$$

By Lemma 3.11, we have $\nabla \varphi_w(y) = \mathbf{L}_y^*(w)$, so $\langle \nabla \varphi_w(y), c''(0) \rangle = \langle w, \mathbf{L}_y(c''(0)) \rangle$.

We conclude that

$$\langle w, (\varphi \circ c)''(0) \rangle = \langle \nabla^2 \varphi_w(y)[v], v \rangle + \langle w, \mathbf{L}_y(c''(0)) \rangle. \tag{3.3}$$

The first term on the right-hand side is independent of c . Thus, for any $w \in \mathcal{E}$ we have

$$\langle w, (\varphi \circ c)''(0) - (\varphi \circ c_v)''(0) \rangle = \langle w, \mathbf{L}_y(c''(0) - c''_v(0)) \rangle,$$

which proves the claim.

- (b) For the first claim, set $c(t) = \exp_y(tv + t^2(c''_v(0) - u)/2)$ where \exp is the exponential map of \mathcal{M} [9, Exer. 5.46]. The second claim follows from part (a). □

Lemma 3.24 shows that the possible values of $\mathbf{Q}_y(v)$ (depending on the choice of curve c_v in Definition 3.9) differ by an element of $\text{im } \mathbf{L}_y$, and conversely, every element of $\mathbf{Q}_y(v) + \text{im } \mathbf{L}_y$ can be obtained by an appropriate choice of c_v . Consequently, if $w \in (\text{im } \mathbf{L}_y)^\perp$, then the inner product $\langle w, \mathbf{Q}_y(v) \rangle$ is independent of the choice of c_v in Definition 3.9. In fact, (3.3) shows that it is a quadratic form in $v \in T_y\mathcal{M}$ given by:

$$\langle w, \mathbf{Q}_y(v) \rangle = \langle \nabla^2 \varphi_w(y)[v], v \rangle \quad \forall v \in T_y\mathcal{M}. \tag{3.4}$$

We stress that this identity requires $w \in (\text{im } \mathbf{L}_y)^\perp$ in general. It allows us to view $\langle w, \mathbf{Q}_y(v) \rangle$ interchangeably as either a quadratic form in v on $T_y\mathcal{M}$ or a linear form in w on $(\text{im } \mathbf{L}_y)^\perp$.

Remark 3.25 (Disambiguation of \mathbf{Q}_y) Lemma 3.24 implies that it is natural to define \mathbf{Q}_y on the quotient $\mathcal{E}/\text{im } \mathbf{L}_y$ to avoid the above ambiguities, but it is less convenient to use in practice—see [39, Rmk. 3.25]. Accordingly, in the remainder of the paper we often refer to \mathbf{Q}_y modulo $\text{im } \mathbf{L}_y$.

We now express the sets A_y, B_y, W_y appearing in our conditions for “2 \Rightarrow 1” in terms of the maps \mathbf{L}_y and \mathbf{Q}_y . We explain how to compute \mathbf{L}_y and \mathbf{Q}_y in various settings in Sect. 3.4 below.

Proposition 3.26 For any $y \in \mathcal{M}$,

- (a) $A_y = \mathbf{Q}_y(\ker \mathbf{L}_y) + \text{im } \mathbf{L}_y$.
- (b) $B_y = \bigcup_{(v_i)_{i \geq 1}: \mathbf{L}_y(v_i) \rightarrow 0} \lim_{i \rightarrow \infty} (\mathbf{Q}_y(v_i) + \text{im } \mathbf{L}_y)$.⁸
- (c) $W_y = \left\{ w \in A_y^* : \forall v \in \ker \mathbf{L}_y, \langle \nabla^2 \varphi_w(y)[v], v \rangle = 0 \implies \nabla^2 \varphi_w(y)[v] = 0 \right\}$.

Proof (a) If $w \in A_y$, then $w = (\varphi \circ c)''(0)$ for some smooth curve c on \mathcal{M} such that $c(0) = y$ and $0 = (\varphi \circ c)'(0) = \mathbf{L}_y(c'(0))$, so $c'(0) \in \ker \mathbf{L}_y$. By Lemma 3.24(a), we have $w \in \mathbf{Q}_y(c'(0)) + \text{im } \mathbf{L}_y$, showing $A_y \subseteq \mathbf{Q}_y(\ker \mathbf{L}_y) + \text{im } \mathbf{L}_y$.

Conversely, suppose $w = \mathbf{Q}_y(v) + \mathbf{L}_y(u)$ for some $v \in \ker \mathbf{L}_y$ and $u \in T_y \mathcal{M}$. By Lemma 3.24(b), there is a smooth curve c on \mathcal{M} satisfying $c(0) = y, c'(0) = v$ and $(\varphi \circ c)''(0) = w$. Since $(\varphi \circ c)'(0) = \mathbf{L}_y(v) = 0$, this shows $\mathbf{Q}_y(\ker \mathbf{L}_y) + \text{im } \mathbf{L}_y \subseteq A_y$.

- (b) If $w \in B_y$, then there are smooth curves c_i such that $c_i(0) = y, \mathbf{L}_y(c'_i(0)) \rightarrow 0$ and $(\varphi \circ c_i)''(0) \rightarrow w$. By Lemma 3.24(a), we have $(\varphi \circ c_i)''(0) \in \mathbf{Q}_y(c'_i(0)) + \text{im } \mathbf{L}_y$. Because $\lim_i (\varphi \circ c_i)''(0) = w$ exists, we conclude that $\lim_i (\mathbf{Q}_y(c'_i(0)) + \text{im } \mathbf{L}_y) = w + \text{im } \mathbf{L}_y$ exists as well, and w is contained in this limit. This shows the inclusion \subseteq in the claim.

Conversely, suppose $w \in \lim_i (\mathbf{Q}_y(v_i) + \text{im } \mathbf{L}_y)$ for some sequence $(v_i)_{i \geq 1} \subseteq T_y \mathcal{M}$ such that $\mathbf{L}_y(v_i) \rightarrow 0$. Then there exist $u_i \in T_y \mathcal{M}$ such that $w = \lim_i (\mathbf{Q}_y(v_i) + \mathbf{L}_y(u_i))$. By Lemma 3.24(b), there exist curves c_i satisfying $c_i(0) = y, c'_i(0) = v_i$ and $(\varphi \circ c_i)''(0) = \mathbf{Q}_y(v_i) + \mathbf{L}_y(u_i)$. Then $(\varphi \circ c_i)'(0) = \mathbf{L}_y(v_i) \rightarrow 0$ and $(\varphi \circ c_i)''(0) \rightarrow w$, so $w \in B_y$ and hence the reverse inclusion in the claim also holds.

- (c) Let $x = \varphi(y)$. By Proposition 3.17, a vector $w \in \mathcal{E}$ is contained in W_y iff there exists $\alpha > 0$ such that y is 2-critical for (\mathbf{Q}) with cost $g_\alpha = f_\alpha \circ \varphi$ where $f_\alpha(x') = \langle w, x' \rangle + \frac{\alpha}{2} \|x' - x\|^2$. By Lemma 3.11, this is equivalent to

$$\nabla g_\alpha(y) = \mathbf{L}_y^*(w) = 0, \quad \nabla^2 g_\alpha(y) = \alpha \mathbf{L}_y^* \circ \mathbf{L}_y + \nabla^2 \varphi_w(y) \geq 0.$$

In other words, $w \in W_y$ iff $w \in (\text{im } \mathbf{L}_y)^\perp$ and $\nabla^2 \varphi_w(y) + \alpha \mathbf{L}_y^* \circ \mathbf{L}_y \geq 0$ for some $\alpha > 0$. To understand when the second condition holds, we decompose $T_y \mathcal{M} = \ker \mathbf{L}_y \oplus (\ker \mathbf{L}_y)^\perp$ and express the relevant self-adjoint operators on $T_y \mathcal{M}$ in block matrix form with respect to a basis compatible with this decomposition.

⁸ A sequence $(v_i + W)_{i \geq 1}$ of translates of a subspace W of a (topological) vector space V converges (necessarily to another translate of W) iff there exist $w_i \in W$ such that $(v_i + w_i)_{i \geq 1} \subseteq V$ converges in V .

More explicitly, choose a basis as described above and denote $n = \dim \ker \mathbf{L}_y$ and $m = \dim(\ker \mathbf{L}_y)^\perp$. Assume first that $m > 0$. We represent $\nabla^2\varphi_w(y)$ and $\alpha \mathbf{L}_y^* \circ \mathbf{L}_y$ in that basis as

$$\begin{aligned} \nabla^2\varphi_w(y) &= \begin{bmatrix} \Phi_1 & \Phi_2 \\ \Phi_2^\top & \Phi_3 \end{bmatrix}, & \text{with } \Phi_1 \in \mathbb{S}^n, \Phi_3 \in \mathbb{S}^m. \\ \alpha \mathbf{L}_y^* \circ \mathbf{L}_y &= \begin{bmatrix} 0 & 0 \\ 0 & \alpha\Psi \end{bmatrix}, & \text{with } \Psi \in \mathbb{S}_{>0}^m. \end{aligned}$$

Thus,

$$w \in W_y \iff w \in (\operatorname{im} \mathbf{L}_y)^\perp \text{ and } \exists \alpha > 0 \text{ such that } \begin{bmatrix} \Phi_1 & \Phi_2 \\ \Phi_2^\top & \Phi_3 + \alpha\Psi \end{bmatrix} \succeq 0.$$

By the generalized Schur complement theorem [59, Thm. 1.20], the block-matrix on the right-hand side is positive semidefinite exactly if

$$\Phi_1 \succeq 0, \quad \operatorname{im} \Phi_2 \subseteq \operatorname{im} \Phi_1 \quad \text{and} \quad \Phi_3 + \alpha\Psi \succeq \Phi_2^\top \Phi_1^\dagger \Phi_2,$$

where Φ_1^\dagger is the Moore–Penrose pseudo-inverse of Φ_1 . The last condition holds upon choosing $\alpha \geq \lambda_{\max}(\Phi_2^\top \Phi_1^\dagger \Phi_2 - \Phi_3) / \lambda_{\min}(\Psi)$. Thus, we deduce the following expression for W_y :

$$W_y = \{w \in (\operatorname{im} \mathbf{L}_y)^\perp : \Phi_1 \succeq 0 \text{ and } \operatorname{im} \Phi_2 \subseteq \operatorname{im} \Phi_1\},$$

with Φ_1 and Φ_2 as defined above. We now work out basis-free characterizations of the properties $\Phi_1 \succeq 0$ and $\operatorname{im} \Phi_2 \subseteq \operatorname{im} \Phi_1$.

First, notice that $\Phi_1 \succeq 0$ iff

$$\begin{bmatrix} v_1^\top & \mathbf{0}_m^\top \end{bmatrix} \begin{bmatrix} \Phi_1 & \Phi_2 \\ \Phi_2^\top & \Phi_3 \end{bmatrix} \begin{bmatrix} v_1 \\ \mathbf{0}_m \end{bmatrix} \geq 0 \quad \text{for all } v_1 \in \mathbb{R}^n,$$

or in basis-free terms, $\langle \nabla^2\varphi_w(y)[v], v \rangle \geq 0$ for all $v \in \ker \mathbf{L}_y$. If $w \in (\operatorname{im} \mathbf{L}_y)^\perp$ then (3.4) shows that this is also equivalent to $\langle w, \mathbf{Q}_y(v) \rangle \geq 0$ for all $v \in \ker \mathbf{L}_y$, which is in turn equivalent to $\langle w, \mathbf{Q}_y(v) + \mathbf{L}_y(u) \rangle \geq 0$ for all $v \in \ker \mathbf{L}_y$ and $u \in T_y\mathcal{M}$. This last condition is just $w \in A_y^*$ by part (a).

Second, we must understand for which vectors w it holds that $\operatorname{im} \Phi_2 \subseteq \operatorname{im} \Phi_1$, or equivalently, $\ker \Phi_1 \subseteq \ker \Phi_2^\top$ (recall that $\Phi_1^\top = \Phi_1$). If $\Phi_1 \succeq 0$, then $v_1 \in \ker \Phi_1$ iff $v_1^\top \Phi_1 v_1 = 0$. Moreover, if $v_1 \in \ker \Phi_1$ then

$$\begin{bmatrix} \Phi_1 & \Phi_2 \\ \Phi_2^\top & \Phi_3 \end{bmatrix} \begin{bmatrix} v_1 \\ \mathbf{0}_m \end{bmatrix} = \begin{bmatrix} \mathbf{0}_n \\ \Phi_2^\top v_1 \end{bmatrix},$$

which vanishes iff $v_1 \in \ker \Phi_2^\top$. Thus, assuming $\Phi_1 \geq 0$, the inclusion $\text{im } \Phi_2 \subseteq \text{im } \Phi_1$ is equivalent to the implication

$$\begin{bmatrix} v_1^\top & \mathbf{0}_m^\top \end{bmatrix} \begin{bmatrix} \Phi_1 & \Phi_2 \\ \Phi_2^\top & \Phi_3 \end{bmatrix} \begin{bmatrix} v_1 \\ \mathbf{0}_m \end{bmatrix} = 0 \implies \begin{bmatrix} \Phi_1 & \Phi_2 \\ \Phi_2^\top & \Phi_3 \end{bmatrix} \begin{bmatrix} v_1 \\ \mathbf{0}_m \end{bmatrix} = 0, \quad \text{for all } v_1 \in \mathbb{R}^n.$$

In basis-free terms, we have shown that, if $\Phi_1 \geq 0$, then $\text{im } \Phi_2 \subseteq \text{im } \Phi_1$ is equivalent to the implication $\langle \nabla^2 \varphi_w(y)[v], v \rangle = 0 \implies \nabla^2 \varphi_w(y)[v] = 0$ holding for all $v \in \ker \mathbf{L}_y$. Putting everything together,

$$\begin{aligned} W_y &= \{w \in (\text{im } \mathbf{L}_y)^\perp : \Phi_1 \geq 0 \text{ and } \text{im } \Phi_2 \subseteq \text{im } \Phi_1\} \\ &= \{w \in (\text{im } \mathbf{L}_y)^\perp : w \in A_y^* \text{ and } \forall v \in \ker \mathbf{L}_y, \\ &\quad \langle \nabla^2 \varphi_w(y)[v], v \rangle = 0 \implies \nabla^2 \varphi_w(y)[v] = 0\} \\ &= \{w \in A_y^* : \forall v \in \ker \mathbf{L}_y, \langle \nabla^2 \varphi_w(y)[v], v \rangle = 0 \implies \nabla^2 \varphi_w(y)[v] = 0\}, \end{aligned}$$

where the last equality holds because $A_y^* \subseteq (\text{im } \mathbf{L}_y)^\perp$ by Proposition 3.21(b). This is the claimed expression for W_y .

If $m = 0$, or equivalently, if $\mathbf{L}_y = 0$, then $w \in W_y$ iff $w \in (\text{im } \mathbf{L}_y)^\perp = \mathcal{E}$ and $\nabla^2 \varphi_w(y) \geq 0$. This in turn is equivalent to $w \in A_y^* = (\mathbf{Q}_y(\ker \mathbf{L}_y))^* \cap (\text{im } \mathbf{L}_y)^\perp$ by (3.4), so $W_y = A_y^*$ in this case. Conversely, if $w \in A_y^*$ and $\mathbf{L}_y = 0$ then $\nabla^2 \varphi_w(y) \geq 0$ so the condition in the claimed expression for W_y is satisfied automatically: it also evaluates to A_y^* . This verifies that the claimed expression for W_y holds for $m = 0$ as well. \square

3.3 Composition of lifts

In this section, we ask: when are lift properties preserved under composition? We use the following proposition both to compute \mathbf{L}_y and \mathbf{Q}_y in various settings, and to study some of the lifts appearing in the literature in Sects. 4 and 5.

Proposition 3.27 *Let $\varphi : \mathcal{M} \rightarrow \mathcal{X}$ be a smooth lift, and let $\psi : \mathcal{N} \rightarrow \mathcal{M}$ be a smooth map between smooth manifolds such that $\varphi \circ \psi : \mathcal{N} \rightarrow \mathcal{X}$ is surjective. Both φ and $\varphi \circ \psi$ are smooth lifts for \mathcal{X} . For $z \in \mathcal{N}$ and $y = \psi(z) \in \mathcal{M}$, the following hold.*

- (a) *If $\varphi \circ \psi$ satisfies “local \implies local” at z , then φ satisfies “local \implies local” at y . If ψ is open (in particular, if ψ is a submersion) at z , and if φ satisfies “local \implies local” at y , then $\varphi \circ \psi$ satisfies “local \implies local” at z .*
- (b) *If $\varphi \circ \psi$ satisfies “1 \implies 1” or “2 \implies 1” at z , then φ satisfies the corresponding property at y . If ψ is a submersion at z and φ satisfies “1 \implies 1” or “2 \implies 1” at y , then $\varphi \circ \psi$ satisfies the corresponding property at z .*
- (c) *If ψ is a submersion at z , then*

$$\mathbf{L}_z^{\varphi \circ \psi} = \mathbf{L}_y^\varphi \circ \mathbf{L}_z^\psi \quad \text{and} \quad \mathbf{Q}_z^{\varphi \circ \psi} \equiv \mathbf{Q}_y^\varphi \circ \mathbf{L}_z^\psi \pmod{\text{im } \mathbf{L}_z^{\varphi \circ \psi}}.$$

Moreover, $\text{im } \mathbf{L}_z^{\varphi \circ \psi} = \text{im } \mathbf{L}_y^\varphi$, $A_z^{\varphi \circ \psi} = A_y^\varphi$, $B_z^{\varphi \circ \psi} = B_y^\varphi$, and $W_z^{\varphi \circ \psi} = W_y^\varphi$.

The proof is straightforward, see [39, App. C.1]. Here we denote $v \equiv w \pmod{\text{im } \mathbf{L}_y}$ to mean $v - w \in \text{im } \mathbf{L}_y$. By Lemma 3.24, equality of $\mathbf{Q}_z^{\varphi \circ \psi}$ and \mathbf{Q}_y^φ modulo $\text{im } \mathbf{L}_y^\varphi = \text{im } \mathbf{L}_z^{\varphi \circ \psi}$ means that either one can be used to verify “ $2 \Rightarrow 1$ ”.

Proposition 3.27 shows that, given a smooth lift $\varphi: \mathcal{M} \rightarrow \mathcal{X}$, there is no benefit to further lifting \mathcal{M} to another smooth manifold through $\psi: \mathcal{N} \rightarrow \mathcal{M}$ in terms of our properties. Indeed, if φ does not satisfy one of our properties, then neither does $\varphi \circ \psi$ for any smooth ψ (we cannot ‘fix’ a bad lift by lifting it further). On the other hand, this proposition also tells us that our properties, as well as the sets involved in their characterization, are preserved under submersions. This notably means lift properties can be checked through charts of \mathcal{M} . Moreover, for lifts to a manifold \mathcal{M} which is a quotient of another manifold $\overline{\mathcal{M}}$ (these arise naturally when quotienting by group actions, see [9, §9]), Proposition 3.27 allows us to verify our properties on the total space $\overline{\mathcal{M}}$, which is often easier.

Remark 3.28 If $\psi: \mathcal{N} \rightarrow \mathcal{M}$ is a submersion, for each $z \in \mathcal{N}$, let $V_z = \ker D\psi(z)$ and $H_z = (\ker D\psi(z))^\perp$ be the so-called vertical and horizontal spaces at z , which satisfy $T_z\mathcal{N} = V_z \oplus H_z$ and $H_z \cong T_{\psi(z)}\mathcal{M}$. Proposition 3.27 implies that $\mathbf{L}_z^{\varphi \circ \psi} = \mathbf{L}_z^{\varphi \circ \psi} \circ \text{Proj}_{H_z}$ and $\mathbf{Q}_z^{\varphi \circ \psi} \equiv \mathbf{Q}_z^{\varphi \circ \psi} \circ \text{Proj}_{H_z}$ where Proj_{H_z} denotes orthogonal projection onto H_z , so it suffices to consider the restrictions of $\mathbf{L}_z^{\varphi \circ \psi}$ and $\mathbf{Q}_z^{\varphi \circ \psi}$ to the horizontal space at z . The latter is often simpler than $T_z\mathcal{N}$, see [9, §9.4].

We end this section with an implicit application of Proposition 3.27 seen in the robotics and computer vision literature.

Example 3.29 (*Shohan rotation averaging*) In [19], Dellaert et al. estimate a set of n rotations of \mathbb{R}^d from noisy measurements of pairs of relative rotations. Their algorithm involves the Burer–Monteiro lift (BM) with $\mathcal{M} = \text{St}(p, d)^n$ for appropriate $p \geq d$, composed with the submersion $\psi: \text{SO}(p)^n \rightarrow \text{St}(p, d)^n$ extracting the first d columns of each matrix. Using our framework, we can analyze this composition and thereby strengthen the guarantees proved in [19]. Indeed, since the Burer–Monteiro lift satisfies “ $2 \Rightarrow 1$ ” and “local \Rightarrow local” by Proposition 2.7 and ψ is a submersion, Proposition 3.27 shows that the composed lift satisfies “local \Rightarrow local” and “ $2 \Rightarrow 1$ ” as well. Furthermore, if every stationary point for the original low-rank SDP of [19] is globally optimal (e.g., if the conditions of [10] hold), then any 2-critical point for their lifted problem is globally optimal and therefore the lifted problem enjoys benign nonconvexity.

3.4 Computing \mathbf{L}_y and \mathbf{Q}_y

Theorem 3.23 gives several conditions for “ $2 \Rightarrow 1$ ” that (together with Proposition 3.26) can be checked using \mathbf{L}_y and \mathbf{Q}_y from Definition 3.9. We therefore consider various strategies for computing \mathbf{L}_y and \mathbf{Q}_y depending on how we can access \mathcal{M} . Since \mathbf{Q}_y is only defined modulo $\text{im } \mathbf{L}_y$, different methods may yield different expressions, any of which can be used to verify “ $2 \Rightarrow 1$ ”.

\mathcal{M} through charts: Suppose we are given a chart $\psi: U \rightarrow \mathcal{M}$ on \mathcal{M} , which is a diffeomorphism from an open subset $U \subseteq \mathcal{E}'$ of some linear space \mathcal{E}' onto its image,

and let $y \in \psi(U)$. Then we can compose φ with ψ to obtain a lift to a linear space $\tilde{\varphi} = \varphi \circ \psi$. By Proposition 3.27, the lift φ satisfies “1 \Rightarrow 1” or “2 \Rightarrow 1” at $y \in \mathcal{M}$ if and only if $\tilde{\varphi}$ satisfies the corresponding property at $z = \psi^{-1}(y)$. Thus, it suffices to compute $\mathbf{L}_z^{\tilde{\varphi}}$ and $\mathbf{Q}_z^{\tilde{\varphi}}$ and use them to check “2 \Rightarrow 1” at z . Since U is an open subset of a linear space \mathcal{E}' , it is natural to compute $\mathbf{L}_z^{\tilde{\varphi}}$ and $\mathbf{Q}_z^{\tilde{\varphi}}$ directly from Definition 3.9 using curves $\tilde{c}_{\tilde{v}}(t) = z + t\tilde{v}$ which are straight lines through z in direction $\tilde{v} \in \mathcal{E}' = T_zU$. This choice yields the expressions

$$\mathbf{L}_z^{\tilde{\varphi}}(\tilde{v}) = (\tilde{\varphi} \circ \tilde{c}_{\tilde{v}})'(0) = D\tilde{\varphi}(z)[\tilde{v}], \quad \text{and} \quad \mathbf{Q}_z^{\tilde{\varphi}}(\tilde{v}) = (\tilde{\varphi} \circ \tilde{c}_{\tilde{v}})''(0) = D^2\tilde{\varphi}(z)[\tilde{v}, \tilde{v}], \tag{3.5}$$

where $D\tilde{\varphi}(z)$ and $D^2\tilde{\varphi}(z)$ are the ordinary first- and second-order derivative maps of $\tilde{\varphi}$ viewed as a map between linear spaces $\mathcal{E}' \rightarrow \mathcal{E}$. In particular, if \mathcal{M} is itself a linear space (e.g., for the (LR) lift), we may take $U = \mathcal{E}' = \mathcal{M}$ and $\psi = \text{id}$ and use (3.5) with $\tilde{\varphi} = \varphi$.

\mathcal{M} embedded in a linear space: Suppose now that \mathcal{M} is an embedded submanifold of another linear space \mathcal{E}' . By [9, Prop. 3.31], the lift φ can be extended to a smooth map on a neighborhood V of \mathcal{M} in \mathcal{E}' , denoted by $\bar{\varphi}: V \rightarrow \mathcal{E}$. This means $\bar{\varphi}$ is a smooth map defined on an open subset $V \subseteq \mathcal{E}'$ containing \mathcal{M} and it satisfies $\bar{\varphi}|_{\mathcal{M}} = \varphi$. If c_v is a curve on \mathcal{M} passing through $y \in \mathcal{M} \subseteq V$ with velocity $v \in T_y\mathcal{M} \subseteq T_yV = \mathcal{E}'$, then $\varphi \circ c_v = \bar{\varphi} \circ c_v$ because the curve is contained in \mathcal{M} where $\bar{\varphi}$ agrees with φ . Denote by $u_v = \ddot{c}_v(0)$ the ordinary (extrinsic) acceleration of c_v at $t = 0$, viewed as a curve in \mathcal{E}' . Then Definition 3.9 and the chain rule give

$$\begin{aligned} \mathbf{L}_y(v) &= (\bar{\varphi} \circ c_v)'(0) = D\bar{\varphi}(y)[v], \quad \mathbf{Q}_y(v) \\ &= (\bar{\varphi} \circ c_v)''(0) = D^2\bar{\varphi}(y)[v, v] + D\bar{\varphi}(y)[u_v]. \end{aligned}$$

To better understand u_v , let $h: \mathcal{E}' \rightarrow \mathbb{R}^k$ be a local defining function for \mathcal{M} around y , that is, h is smooth, \mathcal{M} is locally its zero-set, and $\text{rank } Dh(y') = k$ for all y' around y . For a curve c_v as above, we have $h(c_v(t)) \equiv 0$ around $t = 0$, so in particular $(h \circ c_v)'(0) = (h \circ c_v)''(0) = 0$. By the chain rule, the latter equations can be written as

$$Dh(y)[v] = 0 \quad \text{and} \quad D^2h(y)[v, v] + Dh(y)[u_v] = 0. \tag{3.6}$$

Conversely, for any $v, u_v \in \mathcal{E}'$ satisfying (3.6), there exists a curve on \mathcal{M} passing through y with velocity v and extrinsic acceleration u_v by [49, Prop. 13.13]. Thus, the expressions (3.6) describe all possible velocities and extrinsic accelerations of curve as they pass through y . This set of all possible such accelerations of curves on \mathcal{M} passing through $y \in \mathcal{M}$ with velocity $v \in T_y\mathcal{M}$ is the *second-order tangent set* to \mathcal{M} at y for v , and is denoted by $T_{y,v}^2\mathcal{M}$ [49, Def. 13.11]. The above discussion shows that

$$T_y\mathcal{M} = \ker Dh(y) \quad \text{and} \quad T_{y,v}^2\mathcal{M} = \left\{ u \in \mathcal{E}' : Dh(y)[u] = -D^2h(y)[v, v] \right\}. \tag{3.7}$$

As a result, for any extension $\bar{\varphi}$ of φ and all $v \in T_y\mathcal{M}$, we have

$$\mathbf{L}_y(v) = D\bar{\varphi}(y)[v] \quad \text{and} \quad \mathbf{Q}_y(v) = D^2\bar{\varphi}(y)[v, v] + D\bar{\varphi}(y)[u_v] \quad \text{for some } u_v \in T_{y,v}^2\mathcal{M}. \tag{3.8}$$

Note that $T_{y,v}^2\mathcal{M}$ is an affine subspace of \mathcal{E}' which is a translate of the subspace $T_y\mathcal{M}$, as can be seen from (3.7). Therefore, while different choices of u_v lead to different expressions for \mathbf{Q}_y , they are all equal modulo $\text{im } \mathbf{L}_y$.

\mathcal{M} as a quotient manifold: Suppose next that \mathcal{M} is a quotient manifold of $\overline{\mathcal{M}}$ with quotient map $\pi: \overline{\mathcal{M}} \rightarrow \mathcal{M}$ [9, §9]. Then $\bar{\varphi} = \varphi \circ \pi$ gives a smooth lift of \mathcal{X} to $\overline{\mathcal{M}}$. Since π is a submersion, Proposition 3.27 and Remark 3.28 imply that to check “ $2 \Rightarrow 1$ ”, we need only compute \mathbf{L}_z and \mathbf{Q}_z for $\bar{\varphi}$ restricted to the horizontal spaces $(\ker D\pi(z))^\perp$ using the preceding two methods.

Computing $\nabla^2\varphi_w$: To check the equivalent condition in Theorem 3.23 for any presentation of \mathcal{M} , we need to compute W_y . If we use Proposition 3.26 to do so, we need an expression for the Riemannian Hessian $\nabla^2\varphi_w(y)$ where $\varphi_w(y) = \langle w, \varphi(y) \rangle$ for $w \in (\text{im } \mathbf{L}_y)^\perp$. Given \mathbf{Q}_y , we can obtain $\nabla^2\varphi_w(y)$ as the unique self-adjoint operator on $T_y\mathcal{M}$ that defines the quadratic form (3.4). Conversely, if we compute $\nabla^2\varphi_w(y)$ for all $w \in (\text{im } \mathbf{L}_y)^\perp$, e.g., using the techniques from [9, §5.5], we can set $\mathbf{Q}_y(v)$ to be the unique element of $(\text{im } \mathbf{L}_y)^\perp$ satisfying (3.4), providing another way to compute \mathbf{Q}_y .

We now illustrate the above techniques for computing \mathbf{L}_y and \mathbf{Q}_y . The following example uses charts.

Example 3.30 (Desingularization of $\mathbb{R}_{\leq r}^{m \times n}$) Consider the desingularization lift (Desing) of bounded rank matrices $\mathcal{X} = \mathbb{R}_{\leq r}^{m \times n}$. We compute \mathbf{L} and \mathbf{Q} using charts. For $(X_0, S_0) \in \mathcal{M}$, let $Y_0 \in \mathbb{R}^{n \times (n-r)}$ be a matrix satisfying $\text{col}(Y_0) = S_0$, so $X_0Y_0 = 0$. Since $\text{rank}(Y_0) = n - r$, we can find $n - r$ linearly independent rows in Y_0 . Let $J \in \mathbb{R}^{(n-r) \times (n-r)}$ be the invertible submatrix of Y_0 obtained by extracting these rows, and let $\Pi \in \mathbb{R}^{n \times n}$ be a permutation matrix sending those $n - r$ rows to the first rows. Then there exist

$$Z_0 \in \mathbb{R}^{m \times r} \text{ and } W_0 \in \mathbb{R}^{r \times (n-r)} \text{ satisfying } \Pi Y_0 J^{-1} = \begin{bmatrix} I_{n-r} \\ W_0 \end{bmatrix} \quad \text{and} \quad X_0 = [-Z_0 W_0, Z_0] \Pi,$$

where the second identity is implied by $0 = X_0Y_0J^{-1} = X_0\Pi^\top(\Pi Y_0J^{-1})$. Accordingly, a chart $\psi: \mathbb{R}^{m \times r} \times \mathbb{R}^{r \times (n-r)} \rightarrow \mathcal{M}$ containing (X_0, S_0) is given by

$$\psi(Z, W) = \left([-ZW, Z] \Pi, \text{col} \left(\Pi^\top \begin{bmatrix} I_{n-r} \\ W \end{bmatrix} \right) \right).$$

Composing with φ , we obtain the lift $\tilde{\varphi}(Z, W) = \varphi(\psi(Z, W)) = [-ZW, Z] \Pi$, and by (3.5),

$$\begin{aligned} \mathbf{L}_{\tilde{\varphi}(Z,W)}(\dot{Z}, \dot{W}) &= D\tilde{\varphi}(Z, W)[\dot{Z}, \dot{W}] = [-\dot{Z}W - Z\dot{W}, \dot{Z}] \Pi, \\ \mathbf{Q}_{\tilde{\varphi}(Z,W)}(\dot{Z}, \dot{W}) &= D^2\tilde{\varphi}(Z, W)[(\dot{Z}, \dot{W}), (\dot{Z}, \dot{W})] = [-2\dot{Z}\dot{W}, 0] \Pi. \end{aligned} \tag{3.9}$$

For $V = [V_1, V_2] \Pi \in \mathbb{R}^{m \times n}$ where V_2 is $m \times r$, the Hessian $\nabla^2\tilde{\varphi}_V(Z, W)$ is the ordinary Euclidean Hessian of $\tilde{\varphi}_V(Z, W) = \langle V, \tilde{\varphi}(Z, W) \rangle$, given by

$$\nabla^2\tilde{\varphi}_V(Z, W)[\dot{Z}, \dot{W}] = [-V_1\dot{W}^\top, \dot{Z}^\top V_1]. \tag{3.10}$$

We use the above expressions to show that this lift satisfies “2 \Rightarrow 1” everywhere on \mathcal{M} in Sect. 5.2.

Next, we illustrate the embedded submanifold case on the following low-dimensional example, which shows that the necessary condition for “2 \Rightarrow 1” in the last implication of Theorem 3.23 is not sufficient.

Example 3.31 Consider the lift of the unit disk \mathcal{X} in \mathbb{R}^2 to

$$\mathcal{M} = \{y \in \mathbb{R}^3 : y_1^2 + y_2^2 + y_3^4 = 1\}, \quad \varphi(y) = (y_1, y_2).$$

Note that \mathcal{M} is an embedded submanifold of \mathbb{R}^3 with defining function $h(y) = y_1^2 + y_2^2 + y_3^4 - 1$, since $\nabla h(y) \neq 0$ for all $y \in \mathcal{M}$. The first two derivatives of h are

$$Dh(y)[\dot{y}] = 2y_1\dot{y}_1 + 2y_2\dot{y}_2 + 4y_3^3\dot{y}_3, \quad D^2h(y)[\dot{y}, \dot{y}] = 2\dot{y}_1^2 + 2\dot{y}_2^2 + 12y_3^2\dot{y}_3^2.$$

Let $y = (1, 0, 0)$ and $x = \varphi(y) = (1, 0)$. We get from (3.7) that

$$\mathbf{T}_y\mathcal{M} = \{\dot{y} \in \mathbb{R}^3 : \dot{y}_1 = 0\}, \quad \mathbf{T}_{y,\dot{y}}^2\mathcal{M} = (-\dot{y}_2^2, 0, 0) + \mathbf{T}_y\mathcal{M}.$$

Because φ extends to a linear map $\bar{\varphi}(y) = (y_1, y_2)$ on all of \mathbb{R}^3 , whose first two derivatives are $D\bar{\varphi}(y)[\dot{y}] = (\dot{y}_1, \dot{y}_2)$ and $D^2\bar{\varphi}(y)[\dot{y}, \dot{y}] = 0$, we have from (3.8) that

$$\mathbf{L}_y(\dot{y}) = (0, \dot{y}_2) \quad \text{and} \quad \mathbf{Q}_y(\dot{y}) = (-\dot{y}_2^2, 0).$$

On the other hand,

$$\mathbf{T}_x\mathcal{X} = \{\dot{x} \in \mathbb{R}^2 : \dot{x}_1 \leq 0\} \supseteq \text{im } \mathbf{L}_y = \{\dot{x} \in \mathbb{R}^2 : \dot{x}_1 = 0\}.$$

Since $\text{im } \mathbf{L}_y \neq \mathbf{T}_x\mathcal{X}$, “1 \Rightarrow 1” does not hold at y . All the sufficient conditions for “2 \Rightarrow 1” in Theorem 3.23 fail as well, see [39, Ex. 3.31].

For the equivalent condition in Theorem 3.23, if $w \in A_y^*$, or equivalently, $w = (w_1, 0)$, then $\nabla^2\varphi_w(y)$ is the unique self-adjoint operator on $\mathbf{T}_y\mathcal{M}$ satisfying

$$\langle \nabla^2\varphi_w(y)[\dot{y}], \dot{y} \rangle = \langle w, \mathbf{Q}_y(\dot{y}) \rangle = -w_1\dot{y}_2^2, \quad \text{hence} \quad \nabla^2\varphi_w(y)[\dot{y}] = (0, -w_1\dot{y}_2, 0).$$

For $u \in \ker \mathbf{L}_y$, or equivalently, $u = (0, 0, u_3)$, we get $\langle \nabla^2 \varphi_w(y)[u], u \rangle = 0$ and $\nabla^2 \varphi_w(y)[u] = 0$. Proposition 3.26 then shows that $W_y = A_y^* \not\subseteq (\mathbf{T}_x \mathcal{X})^*$, hence “ $2 \Rightarrow 1$ ” does not hold.

Nevertheless, the necessary condition in Theorem 3.23 does hold. Indeed,

$$\mathbf{Q}_y(\mathbf{T}_y \mathcal{M}) + \text{im } \mathbf{L}_y = \{(-\dot{y}_2^2, 0) : \dot{y}_2 \in \mathbb{R}\} + \{(0, \dot{y}_2) : \dot{y}_2 \in \mathbb{R}\} = \mathbf{T}_x \mathcal{X},$$

so taking duals on both sides yields the desired condition.

3.5 Low-rank PSD matrices, and computing tangent cones

As an application of the theory we developed so far, we consider the set of bounded-rank PSD matrices

$$\mathcal{X} = \{X \in \mathbb{R}^{n \times n} : X^\top = X, X \geq 0, \text{rank}(X) \leq r\} = \mathbb{S}_{\geq 0}^n \cap \mathbb{R}_{\leq r}^{n \times n},$$

with $r < n$ together with its lift to $\mathcal{M} = \mathbb{R}^{n \times r}$ via the factorization map $\varphi(R) = RR^\top$. This is a special case of the Burer–Monteiro lift for SDPs (BM) without constraints ($m = 0$). Constructions in the next section enable us to deduce the properties of the general lift from this special case.

Proposition 3.32 *The lift $\varphi(R) = RR^\top$ from $\mathcal{M} = \mathbb{R}^{n \times r}$ to $\mathcal{X} = \mathbb{R}_{\leq r}^{n \times n} \cap \mathbb{S}_{\geq 0}^n$ satisfies*

- “local \Rightarrow local” and “ $2 \Rightarrow 1$ ” everywhere, and
- “ $1 \Rightarrow 1$ ” at $R \in \mathcal{M}$ if and only if $\text{rank}(R) = r$.

Moreover, with $X = RR^\top$, the sufficient condition $A_R = \mathbf{T}_X \mathcal{X}$ for “ $2 \Rightarrow 1$ ” holds everywhere, and

$$\begin{aligned} \mathbf{T}_X \mathcal{X} &= \mathbf{T}_X \mathbb{R}_{\leq r}^{n \times n} \cap \mathbf{T}_X \mathbb{S}_{\geq 0}^n \\ &= \{V \in \mathbb{S}^n : V_\perp \geq 0, \text{rank}(V_\perp) \leq r - \text{rank}(X) \text{ where } V_\perp = \text{Proj}_{\text{col}(X)^\perp} V \text{Proj}_{\text{col}(X)}^\perp\} \\ &= \left\{ U \begin{pmatrix} V_1 & V_2 \\ V_2^\top & V_3 \end{pmatrix} U^\top : V_1 \in \mathbb{S}^{\text{rank}(X)}, V_3 \geq 0, \text{rank}(V_3) \leq r - \text{rank}(X) \right\}, \end{aligned}$$

where in the last line $U \in O(n)$ is an eigenmatrix for X satisfying $X = U \Sigma U^\top$ with $\Sigma \in \mathbb{R}^{n \times n}$ diagonal containing the eigenvalues of X in descending order.

Proof The “local \Rightarrow local” property for this lift was proved in [12, Prop. 2.3] by (in our terminology) establishing SLP (the condition from Remark 3.8).

For $V \in \mathbb{R}^{n \times n}$ and $X = RR^\top \in \mathcal{X}$, define

$$V_\perp := \text{Proj}_{\ker(X)} V \text{Proj}_{\ker(X)} = \text{Proj}_{\text{col}(R)^\perp} V \text{Proj}_{\text{col}(R)^\perp},$$

where we used the fact that $\text{col}(R) = \text{col}(X) = \ker(X)^\perp$. The tangent cone at X to $\mathbb{S}_{\geq 0}^n$ is given by [27, Eq. (9)]

$$\begin{aligned} T_X \mathbb{S}_{\geq 0}^n &= \{V \in \mathbb{S}^n : \langle Vu, u \rangle \geq 0, \text{ for all } u \in \ker(X)\} = \{V \in \mathbb{S}^n : V_\perp \geq 0\} \\ &= \left\{ U \begin{pmatrix} V_1 & V_2 \\ V_2^\top & V_3 \end{pmatrix} U^\top : V_1 \in \mathbb{R}^{\text{rank}(X) \times \text{rank}(X)}, V_3 \geq 0 \right\}, \end{aligned} \tag{3.11}$$

and the tangent cone at X to $\mathbb{R}_{\leq r}^{n \times n}$ is given by [51, Thm. 3.2]

$$\begin{aligned} T_X \mathbb{R}_{\leq r}^{n \times n} &= \{V \in \mathbb{R}^{n \times n} : \text{rank}(V_\perp) \leq r - \text{rank}(X)\} \\ &= \left\{ U \begin{pmatrix} V_1 & V_2 \\ \tilde{V}_2 & V_3 \end{pmatrix} U^\top : V_1 \in \mathbb{R}^{\text{rank}(X) \times \text{rank}(X)}, \text{rank}(V_3) \leq r - \text{rank}(X) \right\}. \end{aligned}$$

Hence the intersection $T_X \mathbb{R}_{\leq r}^{n \times n} \cap T_X \mathbb{S}_{\geq 0}^n$ is given by the claimed expression. Furthermore, the tangent cone to an intersection is always included in the intersection of the tangent cones, which follows easily from Definition 3.1. Hence $T_X \mathcal{X} \subseteq T_X \mathbb{R}_{\leq r}^{n \times n} \cap T_X \mathbb{S}_{\geq 0}^n$ and it suffices to show the reverse inclusion. We do so simultaneously with proving “ $2 \Rightarrow 1$ ”.

Since \mathcal{M} is a linear space, the expressions (3.5) with the identity chart give

$$\begin{aligned} \mathbf{L}_R(\dot{R}) &= D\varphi(R)[\dot{R}] = R\dot{R}^\top + \dot{R}R^\top, \\ \mathbf{Q}_R(\dot{R}) &= D^2\varphi(R)[\dot{R}, \dot{R}] = 2\dot{R}\dot{R}^\top. \end{aligned}$$

Therefore,

$$\text{im } \mathbf{L}_R = \{V \in \mathbb{S}^n : V_\perp = 0\}. \tag{3.12}$$

Indeed, if $V \in \text{im } \mathbf{L}_R$ then $V = R\dot{R}^\top + \dot{R}R^\top$ for some $\dot{R} \in \mathbb{R}^{n \times r}$ and hence $V_\perp = 0$ since $\text{Proj}_{\text{col}(R)^\perp} R = 0$, while if $V \in \mathbb{S}^n$ satisfies $V_\perp = 0$ then

$$\begin{aligned} V &= \text{Proj}_{\text{col}(R)} V + V \text{Proj}_{\text{col}(R)} - \text{Proj}_{\text{col}(R)} V \text{Proj}_{\text{col}(R)} \\ &= \mathbf{L}_R \left(V R^{\dagger\top} - \frac{1}{2} R R^\dagger V R^{\dagger\top} \right) \in \text{im } \mathbf{L}_R, \end{aligned}$$

where $R R^\dagger = \text{Proj}_{\text{col}(R)}$.

Furthermore, we have

$$\mathbf{Q}_R(\ker \mathbf{L}_R) \supseteq \{V \in \mathbb{S}_{\geq 0}^n : \text{rank}(V) \leq r - \text{rank}(X)\}.$$

Indeed, if $V \in \mathbb{S}_{\geq 0}^n$ with $\text{rank}(V) \leq r - \text{rank}(X)$, let $V = 2\dot{R}_0 \dot{R}_0^\top$ be a (rescaled) Cholesky decomposition with $\dot{R}_0 \in \mathbb{R}^{n \times r}$, so $\text{rank}(\dot{R}_0) = \text{rank}(V)$. Since

$$\dim \text{col}(R^\top) = \text{rank}(R^\top) = \text{rank}(R) = \text{rank}(X) \leq r - \text{rank}(\dot{R}_0) = \dim \ker(\dot{R}_0),$$

there is an orthogonal matrix $Q \in O(r)$ satisfying $Q\text{col}(R^\top) \subseteq \ker(\dot{R}_0)$. Let $\dot{R} = \dot{R}_0 Q$ and note that $V = 2\dot{R}\dot{R}^\top$ so $\mathbf{Q}_R(\dot{R}) = V$, and $\dot{R}R^\top = 0$ so $\dot{R} \in \ker \mathbf{L}_R$. Thus, with Proposition 3.26(a),

$$A_R = \mathbf{Q}_R(\ker \mathbf{L}_R) + \text{im } \mathbf{L}_R \supseteq T_X \mathbb{R}_{\leq r}^{n \times n} \cap T_X \mathbb{S}_{\geq 0}^n.$$

On the other hand, by Proposition 3.21(b), we have $A_R \subseteq T_X \mathcal{X}$. Thus, we have the chain of inclusions

$$T_X \mathbb{R}_{\leq r}^{n \times n} \cap T_X \mathbb{S}_{\geq 0}^n \subseteq \mathbf{Q}_R(\ker \mathbf{L}_R) + \text{im } \mathbf{L}_R \subseteq T_X \mathcal{X} \subseteq T_X \mathbb{R}_{\leq r}^{n \times n} \cap T_X \mathbb{S}_{\geq 0}^n,$$

so all the above inclusions are equalities. In particular, we obtain the claimed expression for $T_X \mathcal{X}$ and “ $2 \Rightarrow 1$ ” everywhere on \mathcal{M} by Theorem 3.23. Our claims about “ $1 \Rightarrow 1$ ” follow from (3.12) and Theorem 2.4. \square

Finding an explicit expression for tangent cones can be difficult in general. In Proposition 3.32, the set \mathcal{X} was an intersection of two sets whose tangent cones are known, namely $\mathbb{R}_{\leq r}^{n \times n}$ and $\mathbb{S}_{\geq 0}^n$, which gave us an inclusion $T_X \mathcal{X} \subseteq T_X \mathbb{R}_{\leq r}^{n \times n} \cap T_X \mathbb{S}_{\geq 0}^n$. However, the tangent cone to an intersection can be strictly contained in the intersection of the tangent cones.⁹ The proof of “ $2 \Rightarrow 1$ ” in Proposition 3.32 proceeds by showing $A_R = T_X \mathbb{R}_{\leq r}^{n \times n} \cap T_X \mathbb{S}_{\geq 0}^n$, which gives $T_X \mathcal{X} = T_X \mathbb{R}_{\leq r}^{n \times n} \cap T_X \mathbb{S}_{\geq 0}^n = A_R$ because $A_R \subseteq T_X \mathcal{X}$ by Proposition 3.21(b). This simultaneously gives us “ $2 \Rightarrow 1$ ” and an expression for the tangent cone.

This illustrates a more general and, as far as we know, novel technique of getting expressions for the tangent cones using lifts. Generalizing the above discussion, if we have an inclusion $T_X \mathcal{X} \subseteq S$ for some set S and we are able to prove $A_y \supseteq S$ for some $y \in \varphi^{-1}(x)$, then we must have $T_X \mathcal{X} = S$ by Proposition 3.21(b). In this case, we also conclude that “ $2 \Rightarrow 1$ ” holds at y by Theorem 3.23. In Sect. 4, we shall see another setting in which we naturally have a superset for $T_X \mathcal{X}$ (see Lemma 4.8), and which allows us to derive expressions for $T_X \mathcal{X}$ from lifts satisfying “ $1 \Rightarrow 1$ ” and “ $2 \Rightarrow 1$ ”.

A general condition implying that the tangent cone to an intersection is the intersection of the tangent cones is given in [49, Thm. 6.42]. That condition does not apply to $\mathcal{X} = \mathbb{R}_{\leq r}^{n \times n} \cap \mathbb{S}_{\geq 0}^n$ because $\mathbb{R}_{\leq r}^{n \times n}$ is not Clarke-regular in the sense of [49, Def. 6.4]. Our approach circumvents Clarke regularity, exploiting the existence of an appropriate lift instead.

4 Constructing lifts via fiber products

In this section, we give a systematic construction of lifts for a large class of sets \mathcal{X} . If the resulting lifted space is a smooth manifold, we also give conditions under which the lift satisfies our desirable properties. Moreover, under these conditions we can obtain expressions for the tangent cones to \mathcal{X} . We shall see that several natural lifts, including the Hadamard and Burer–Monteiro lifts from Sect. 2, are special cases of this construction.

⁹ For example, consider intersecting the circle in the plane with one of its tangent lines.

Suppose the set \mathcal{X} is presented in the form

$$\mathcal{X} = \{x \in \mathcal{E} : F(x) \in \mathcal{Z}\} = F^{-1}(\mathcal{Z}),$$

where $\mathcal{Z} \subseteq \mathcal{E}'$ is some subset of a linear space and $F : \mathcal{E} \rightarrow \mathcal{E}'$ is smooth. This form is general—any set \mathcal{X} can be written in this form by letting F be the identity and $\mathcal{Z} = \mathcal{X}$. However, we shall see that our framework is most useful when \mathcal{Z} is a product of simple sets for which we have smooth lifts satisfying desirable properties. For example, any set defined by k smooth equalities $g_i(x) = 0$ and ℓ smooth inequalities $h_j(x) \geq 0$ can be written in this form by letting $F(x) = (g_1(x), \dots, g_k(x), h_1(x), \dots, h_\ell(x))$ and $\mathcal{Z} = \{0\}^k \times \mathbb{R}_{\geq 0}^\ell$. We can also incorporate semidefiniteness and rank constraints of smooth functions of x by taking Cartesian products of \mathcal{Z} with $\mathbb{R}_{\leq r}^{m \times n}$ or $\mathbb{S}_{\geq 0}^n$.

Suppose now that we have a smooth lift $\psi : \mathcal{N} \rightarrow \mathcal{Z}$. We can use this lift of \mathcal{Z} to construct a lift of \mathcal{X} by taking the *fiber product* of F and ψ .

Definition 4.1 Let \mathcal{X} be a subset of \mathcal{E} defined by a smooth map $F : \mathcal{E} \rightarrow \mathcal{E}'$ and a set $\mathcal{Z} \subseteq \mathcal{E}'$ as $\mathcal{X} = F^{-1}(\mathcal{Z})$. Suppose $\psi : \mathcal{N} \rightarrow \mathcal{Z}$ is a smooth lift of \mathcal{Z} to the smooth manifold \mathcal{N} . Then the *fiber product lift* of \mathcal{X} with respect to F and ψ is $\varphi : \mathcal{M}_{F,\psi} \rightarrow \mathcal{X}$ where

$$\mathcal{M}_{F,\psi} = \{(x, y) \in \mathcal{E} \times \mathcal{N} : F(x) = \psi(y)\} \quad \text{and} \quad \varphi(x, y) = x.$$

Here $\mathcal{M}_{F,\psi}$ is the (set-theoretic) *fiber product* of the maps $F : \mathcal{E} \rightarrow \mathcal{E}'$ and $\psi : \mathcal{N} \rightarrow \mathcal{E}'$.

The following commutative diagram illustrates Definition 4.1. Its top horizontal arrow is the coordinate projection $\pi(x, y) = y$.

$$\begin{array}{ccc} \mathcal{M}_{F,\psi} & \xrightarrow{\pi} & \mathcal{N} \\ \varphi \downarrow & & \downarrow \psi \\ \mathcal{E} \supseteq \mathcal{X} & \xrightarrow{F} & \mathcal{Z} \subseteq \mathcal{E}' \end{array}$$

The fiber product $\mathcal{M}_{F,\psi}$ need not be a smooth manifold even when both F and ψ are smooth maps between smooth manifolds. Accordingly, we make the following assumption:

Assumption 4.2 The differential of $(x, y) \mapsto F(x) - \psi(y)$ has constant rank in a neighborhood of $\mathcal{M}_{F,\psi}$ in $\mathcal{E} \times \mathcal{N}$.

Assumption 4.2 not only implies $\mathcal{M}_{F,\psi}$ is a smooth embedded submanifold of $\mathcal{E} \times \mathcal{N}$, but also that $F(x) - \psi(y) = 0$ is a (constant-rank) defining function for it (in the sense of [37, Thm. 5.12]). Under this assumption, the tangent space to $\mathcal{M}_{F,\psi}$ is given by

$$\mathbf{T}_{(x,y)}\mathcal{M}_{F,\psi} = \{(\dot{x}, \dot{y}) \in \mathcal{E} \times \mathbf{T}_y\mathcal{N} : \mathbf{D}F(x)[\dot{x}] = \mathbf{D}\psi(y)[\dot{y}]\}. \quad (4.1)$$

We proceed to give some examples of the above construction. We then study fiber product lifts in general and instantiate our results on these examples.

Example 4.3 (*Sphere to ball*) Let $\mathcal{X} = \{x \in \mathbb{R}^n : \|x\|_2 \leq 1\}$ be the unit Euclidean ball. Let $\mathcal{Z} = \mathbb{R}_{\geq 0}$ and $F(x) = 1 - x^\top x$ so $\mathcal{X} = F^{-1}(\mathcal{Z})$. Let $\psi: \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ be the smooth lift $\psi(y) = y^2$. Then $\mathcal{M}_{F,\psi} = \{(x, y) \in \mathbb{R}^n \times \mathbb{R} : 1 - x^\top x = y^2\}$, which is just the unit sphere in \mathbb{R}^{n+1} , and $\varphi(x, y) = x$ projects onto the first n coordinates. This lift is used in [47, §2.7] to apply a solver for quadratic programming over the sphere (Q) to quadratic programs over the ball (P).

Example 4.4 (*Sphere to simplex*) The Hadamard lift (Had) from Sect. 2.1 can be obtained as a special case of Definition 4.1. Indeed, let $\mathcal{X} = \Delta^{n-1} = \{x \in \mathbb{R}^n : x \geq 0, \sum_{i=1}^n x_i = 1\}$ be the standard simplex, let $\mathcal{Z} = \mathbb{R}_{\geq 0} \times \{0\}$ and $F(x) = (x, \sum_{i=1}^n x_i - 1)$ so $\mathcal{X} = F^{-1}(\mathcal{Z})$. Let $\psi: \mathbb{R}^n \rightarrow \mathcal{Z}$ be $\psi(y) = (y^{\odot 2}, 0)$ where superscript $\odot 2$ denotes entrywise squaring. Then

$$\begin{aligned} \mathcal{M}_{F,\psi} &= \left\{ (x, y) \in \mathbb{R}^n \times \mathbb{R}^n : x = y^{\odot 2}, \sum_{i=1}^n x_i = 1 \right\} \\ &= \left\{ (x, y) \in \mathbb{R}^n \times \mathbb{R}^n : x = y^{\odot 2}, \|y\|_2 = 1 \right\}, \end{aligned}$$

and $\varphi(x, y) = x$. It is easy to check that the coordinate projection π defines a diffeomorphism of $\mathcal{M}_{F,\psi}$ with S^{n-1} , and that the composition $\varphi \circ \pi^{-1}: S^{n-1} \rightarrow \mathcal{X}$ yields the Hadamard lift (Had) from Sect. 2.1. By Proposition 3.27, the fiber product lift of the simplex is equivalent (for the purposes of checking our desirable properties) to the lift (Had).

Example 4.5 (*Torus to annulus*) Let $\mathcal{X} = \{x \in \mathbb{R}^n : r_1 \leq \|x\|_2 \leq r_2\}$, where we assume $0 < r_1 < r_2$. Let $\mathcal{Z} = \mathbb{R}_{\geq 0}^2$ and $F(x) = (x^\top x - r_1^2, r_2^2 - x^\top x)$. Let $\psi: \mathbb{R}^2 \rightarrow \mathcal{Z}$ be $\psi(y) = y^{\odot 2}$ so

$$\begin{aligned} \mathcal{M}_{F,\psi} &= \{(x, y) \in \mathbb{R}^n \times \mathbb{R}^2 : x^\top x - r_1^2 = y_1^2, r_2^2 - x^\top x = y_2^2\}, \\ &= \left\{ (x, y) \in \mathbb{R}^n \times \mathbb{R}^2 : \|x\|_2 = \sqrt{r_1^2 + y_1^2}, \|y\|_2 = \sqrt{r_2^2 - r_1^2} \right\}. \end{aligned}$$

This is an n -dimensional manifold diffeomorphic to $S^{n-1} \times S^1$, with diffeomorphism

$$\Phi(x, y) = [(r_1^2 + y_1^2)^{-1/2}x, (r_2^2 - r_1^2)^{-1/2}y].$$

Viewed differently, the equivalent (by Proposition 3.27) lift $\varphi \circ \Phi^{-1}$ is the composition

$$S^{n-1} \times S^1 \rightarrow S^{n-1} \times \Delta^1 \rightarrow \mathcal{X},$$

where the first map is the Hadamard lift from the sphere to the simplex from the preceding example, and the second map is $(y, \theta) \mapsto \sqrt{\theta_1 r_1^2 + \theta_2 r_2^2} y$. If $n = 2$, then \mathcal{X} is an annulus and \mathcal{M} is a torus.

Example 4.6 (*Smooth SDPs*) The Burer–Monteiro lift (BM) from Sect. 2.2 is also a special case of Definition 4.1. To see this, in the notation of Sect. 2.2, let

$$\mathcal{Z} = (\mathbb{S}_{\geq 0}^n \cap \mathbb{R}_{\leq r}^{n \times n}) \times \{0\}^m, \quad F(X) = (X, \langle A_1, X \rangle - b_1, \dots, \langle A_m, X \rangle - b_m),$$

and $\psi(R) = (RR^\top, 0, \dots, 0)$ defined on $\mathcal{N} = \mathbb{R}^{n \times r}$. In this case,

$$\mathcal{M}_{F,\psi} = \{(X, R) \in \mathbb{S}^n \times \mathbb{R}^{n \times r} : X = RR^\top, \langle A_i R, R \rangle = b_i\},$$

which the projection π maps diffeomorphically onto the set \mathcal{M} in (BM). Furthermore, Assumption 4.2 is equivalent in this case to the assumption in Proposition 2.7 that $h_i(R) = \langle A_i R, R \rangle - b_i$ are local defining functions. Thus, proving Proposition 2.7 is equivalent to proving the corresponding properties for the fiber product lift above under Assumption 4.2.

Now that we have seen several examples of fiber product lifts, we ask: when do desirable properties of the lift $\psi : \mathcal{N} \rightarrow \mathcal{Z}$ imply the corresponding properties for the fiber product lift $\varphi : \mathcal{M}_{F,\psi} \rightarrow \mathcal{X}$? This is answered by the next few propositions.

Proposition 4.7 *Under Assumption 4.2, if $\psi : \mathcal{N} \rightarrow \mathcal{Z}$ satisfies “local \Rightarrow local” at $y \in \mathcal{N}$, then $\varphi : \mathcal{M}_{F,\psi} \rightarrow \mathcal{X}$ satisfies “local \Rightarrow local” at $(x, y) \in \mathcal{M}_{F,\psi}$ for any $x \in \pi^{-1}(y)$.*

Proof By Theorem 2.3, it is equivalent to show that openness of ψ at y implies openness of φ at (x, y) . Assumption 4.2 implies that $\mathcal{M}_{F,\psi}$ is an embedded submanifold of $\mathcal{E} \times \mathcal{N}$, hence its manifold topology coincides with the subspace topology induced from $\mathcal{E} \times \mathcal{N}$. Thus, to show φ is open at (x, y) , it suffices to show that $\varphi((U \times V) \cap \mathcal{M}_{F,\psi})$ is open for any open $U \subseteq \mathcal{E}$ containing x and open $V \subseteq \mathcal{N}$ containing y , since such sets form a basis for the subspace topology on $\mathcal{M}_{F,\psi}$. Since ψ is open at y , we get that $\psi(V) \subseteq \mathcal{Z}$ is open. Since F is continuous, $F^{-1}(\psi(V)) \subseteq \mathcal{X}$ is open. Since $\varphi(x, y) = x$, we have

$$\begin{aligned} \varphi((U \times V) \cap \mathcal{M}_{F,\psi}) &= \{x \in U \cap \mathcal{X} : \exists y \in V \text{ s.t. } F(x) = \psi(y)\} \\ &= \{x \in U \cap \mathcal{X} : F(x) \in \psi(V)\} = (U \cap \mathcal{X}) \cap F^{-1}(\psi(V)), \end{aligned}$$

which is open in \mathcal{X} as the intersection of two open sets. Thus, φ is open at (x, y) . \square

Note that the above proof, and hence the conclusion of Proposition 4.7, apply more generally whenever $\mathcal{M}_{F,\psi}$ is endowed with the subspace topology induced from $\mathcal{E} \times \mathcal{N}$ (but is not necessarily a smooth manifold) and when all maps involved are continuous (but not necessarily smooth).

We now turn to studying “1 \Rightarrow 1” and “2 \Rightarrow 1”. Along the way, we give another instance of the technique for finding tangent cones via lifts outlined in Sect. 3.5. To do so, we begin by giving a superset of the tangent cone, obtained from the fact that \mathcal{X} is an inverse image [49, Thm. 6.31]. As usual, $DF(x)^{-1}$ denotes the preimage under the differential (which may not be invertible).

Lemma 4.8 *The following inclusion holds for all $x \in \mathcal{X}$:*

$$T_x \mathcal{X} \subseteq \{\dot{x} \in \mathcal{E} : DF(x)[\dot{x}] \in T_{F(x)} \mathcal{Z}\} = DF(x)^{-1}(T_{F(x)} \mathcal{Z}).$$

Proof If $v \in T_x \mathcal{X}$ then by Definition 3.1 there exist sequences $(x_i)_{i \geq 1} \subseteq \mathcal{X}$ converging to x and $(\tau_i)_{i \geq 1} \subseteq \mathbb{R}_{>0}$ converging to zero satisfying $v = \lim_{i \rightarrow \infty} \frac{x_i - x}{\tau_i}$. Because F is differentiable at x , we have $F(x_i) = F(x) + DF(x)[x_i - x] + o(\|x_i - x\|)$, so

$$DF(x)[v] = \lim_{i \rightarrow \infty} \frac{F(x_i) - F(x)}{\tau_i}.$$

Since $F(x_i) \in \mathcal{Z}$ for all i , we conclude that $DF(x)[v] \in T_{F(x)} \mathcal{Z}$ by Definition 3.1. \square

Proposition 4.9 *Let $(x, y) \in \mathcal{M}_{F, \psi}$. Under Assumption 4.2, if ψ satisfies “ $1 \Rightarrow 1$ ” at $y \in \mathcal{N}$, then φ satisfies “ $1 \Rightarrow 1$ ” at (x, y) , and equality holds in Lemma 4.8.*

Proof Since ψ satisfies “ $1 \Rightarrow 1$ ” at y , Theorem 2.4 yields $\text{im } \mathbf{L}_y^\psi = T_{\psi(y)} \mathcal{Z} = T_{F(x)} \mathcal{Z}$. Assumption 4.2 implies that $\mathcal{M}_{F, \psi}$ is an embedded submanifold of $\mathcal{E} \times \mathcal{N}$. Since φ extends to $\bar{\varphi}(x, y) = x$ defined on all of $\mathcal{E} \times \mathcal{N}$, we get from (3.8) that $\mathbf{L}_{(x, y)}^\varphi(\dot{x}, \dot{y}) = \dot{x}$ for all $(\dot{x}, \dot{y}) \in T_{(x, y)} \mathcal{M}_{F, \psi}$. By (4.1),

$$\text{im } \mathbf{L}_{(x, y)}^\varphi = DF(x)^{-1}(\text{im } \mathbf{L}_y^\psi) = DF(x)^{-1}(T_{F(x)} \mathcal{Z}).$$

Using Lemma 4.8 and Proposition 3.21(b), we get the chain of inclusions

$$T_x \mathcal{X} \subseteq DF(x)^{-1}(T_{F(x)} \mathcal{Z}) = \text{im } \mathbf{L}_{(x, y)}^\varphi \subseteq T_x \mathcal{X}.$$

We conclude that all these sets are equal and hence “ $1 \Rightarrow 1$ ” holds for φ at (x, y) . \square

Proposition 4.10 *Under Assumption 4.2, if ψ satisfies the sufficient condition $A_y^\psi = T_{\psi(y)} \mathcal{Z}$ for “ $2 \Rightarrow 1$ ” at $y \in \mathcal{N}$ (recall Theorem 3.23), then φ satisfies the sufficient condition $A_{(x, y)}^\varphi = T_x \mathcal{X}$ for “ $2 \Rightarrow 1$ ” at $(x, y) \in \mathcal{M}_{F, \psi}$, and equality holds in Lemma 4.8.*

Proof By Lemma 4.8, we always have $T_x \mathcal{X} \subseteq DF(x)^{-1}(T_{F(x)} \mathcal{Z})$. For the reverse inclusion and the desired sufficient condition for “ $2 \Rightarrow 1$ ”, it suffices to prove that $DF(x)^{-1}(T_{F(x)} \mathcal{Z}) \subseteq A_{(x, y)}^\varphi$ since $A_{(x, y)}^\varphi \subseteq T_x \mathcal{X}$ by Proposition 3.21(b).

Suppose $DF(x)[\dot{x}] \in T_{F(x)} \mathcal{Z}$. By hypothesis, $T_{F(x)} \mathcal{Z} = T_{\psi(y)} \mathcal{Z} = A_y^\psi$, so by Proposition 3.26(a):

$$DF(x)[\dot{x}] = \mathbf{Q}_y^\psi(v) + \mathbf{L}_y^\psi(u), \quad \text{for some } v \in \ker \mathbf{L}_y^\psi \text{ and } u \in T_y \mathcal{N}. \quad (4.2)$$

Because $v \in \ker \mathbf{L}_y^\psi$, we have $(0, v) \in T_{(x, y)} \mathcal{M}_{F, \psi}$ by (4.1). Let $t \mapsto c(t) = (c_x(t), c_y(t))$ be a smooth curve on $\mathcal{M}_{F, \psi}$ passing through (x, y) with velocity $(0, v)$

at $t = 0$. Because $F(c_x(t)) = \psi(c_y(t))$ for all t near 0, differentiating this expression twice we get

$$D^2F(x)[c'_x(0), c'_x(0)] + DF(x)[c''_x(0)] = (\psi \circ c_y)''(0).$$

The first term vanishes since $c'_x(0) = 0$. Using Definition 3.9 for \mathbf{Q}_y^ψ with Lemma 3.24, we obtain

$$DF(x)[c''_x(0)] = \mathbf{Q}_y^\psi(v) + \mathbf{L}_y^\psi(u'), \quad \text{for some } u' \in T_y\mathcal{N}. \tag{4.3}$$

Subtracting (4.3) from (4.2) yields (using Definition 3.9 for \mathbf{L}_y^ψ and (4.1) for $T_{(x,y)}\mathcal{M}_{F,\psi}$)

$$\begin{aligned} DF(x)[\dot{x} - c''_x(0)] &= \mathbf{L}_y^\psi(u - u') \\ &= D\psi(y)[u - u'], \quad \text{hence } (\dot{x} - c''_x(0), u - u') \in T_{(x,y)}\mathcal{M}_{F,\psi}. \end{aligned}$$

Since $D\varphi(x, y)[\dot{x}, \dot{y}] = \dot{x}$, it follows that $\dot{x} - c''_x(0) \in \text{im } \mathbf{L}_{(x,y)}^\varphi$. By Definition 3.9 and Lemma 3.24 again, there exists $w \in T_{(x,y)}\mathcal{M}_{F,\psi}$ satisfying

$$c''_x(0) + \mathbf{L}_{(x,y)}^\varphi(w) = \mathbf{Q}_{(x,y)}^\varphi(0, v) \in \mathbf{Q}_{(x,y)}^\varphi(\ker \mathbf{L}_{(x,y)}^\varphi),$$

from which we see $\dot{x} \in \mathbf{Q}_{(x,y)}^\varphi(\ker \mathbf{L}_{(x,y)}^\varphi) + \text{im } \mathbf{L}_{(x,y)}^\varphi = A_{(x,y)}^\varphi$ (again with Proposition 3.26(a)). □

We remark that other sufficient conditions for equality in Lemma 4.8 to be achieved are given in [49, Exer. 6.7, Thm. 6.31]. However, they do not apply to Example 4.6 (\mathcal{Z} is not Clarke-regular and $DF(X)$ may not be surjective). In contrast, our approach via lifts does apply to this example, and gives “ $2 \Rightarrow 1$ ” and an expression for the tangent cones simultaneously, see Corollary 4.12 below.

As the examples in the beginning of this section illustrate, \mathcal{Z} is often a product of sets. It is therefore useful to note that a product of lifts satisfying desirable properties also satisfies those properties:

Proposition 4.11 *Suppose $\mathcal{Z}_i \subseteq \mathcal{E}_i$ for $i = 1, \dots, k$ are subsets admitting smooth lifts $\psi_i: \mathcal{N}_i \rightarrow \mathcal{Z}_i$. Let $\mathcal{Z} = \mathcal{Z}_1 \times \dots \times \mathcal{Z}_k$ and $\psi = \psi_1 \times \dots \times \psi_k: \mathcal{N}_1 \times \dots \times \mathcal{N}_k \rightarrow \mathcal{Z}$, which is a smooth lift of \mathcal{Z} . Then the following hold.*

- (a) $T_{\mathcal{Z}}\mathcal{Z} \subseteq T_{z_1}\mathcal{Z}_1 \times \dots \times T_{z_k}\mathcal{Z}_k$ (the inclusion may be strict, see [49, Prop. 6.41]).
- (b) ψ satisfies “local \Rightarrow local” at y if and only if ψ_i satisfies “local \Rightarrow local” at y_i for all i .
- (c) We have $\text{im } \mathbf{L}_y^\psi = \text{im } \mathbf{L}_{y_1}^{\psi_1} \times \dots \times \text{im } \mathbf{L}_{y_k}^{\psi_k}$. In particular, ψ satisfies “ $1 \Rightarrow 1$ ” at y if ψ_i satisfies “ $1 \Rightarrow 1$ ” at y_i for all i , in which case equality in (a) holds.
- (d) We have $\mathbf{Q}_y^\psi \equiv \mathbf{Q}_{y_1}^{\psi_1} \times \dots \times \mathbf{Q}_{y_k}^{\psi_k} \pmod{\text{im } \mathbf{L}_y^\psi}$. Moreover, $A_y^\psi = A_{y_1}^{\psi_1} \times \dots \times A_{y_k}^{\psi_k}$ and likewise for B_y^ψ and W_y^ψ . In particular, ψ satisfies “ $2 \Rightarrow 1$ ” at y if ψ_i satisfies “ $2 \Rightarrow 1$ ” at y_i for all i .

(e) ψ satisfies the sufficient condition $A_y^\psi = T_{\psi(y)}\mathcal{Z}$ for “ $2 \Rightarrow 1$ ” if ψ_i satisfies the corresponding conditions $A_{y_i}^{\psi_i} = T_{\psi_i(y_i)}\mathcal{Z}_i$ for all i .

The proof is straightforward, see [39, App. C.2]. By Remark 3.25, equality of \mathbf{Q}_y^ψ and $\mathbf{Q}_{y_1}^{\psi_1} \times \cdots \times \mathbf{Q}_{y_k}^{\psi_k}$ modulo $\text{im } \mathbf{L}_y^\psi$ means that either one can be used to verify “ $2 \Rightarrow 1$ ”.

We can now revisit the examples from the beginning of this section.

Corollary 4.12 *The lifts in Examples 4.3, 4.4, 4.5 and 4.6 satisfy the following.*

- The sphere to ball lift in Example 4.3 satisfies “local \Rightarrow local” everywhere, “ $1 \Rightarrow 1$ ” at y if and only if $y \neq 0$ (i.e., at preimages of the interior of the ball), and “ $2 \Rightarrow 1$ ” everywhere.
- The sphere to simplex lift (Had) satisfies “local \Rightarrow local” everywhere, “ $1 \Rightarrow 1$ ” at y if and only if $y_i \neq 0$ for all i (i.e., at preimages of the relative interior of the simplex), and “ $2 \Rightarrow 1$ ” everywhere.
- The lift of the annulus in Example 4.5 satisfies “local \Rightarrow local” everywhere, “ $1 \Rightarrow 1$ ” at y if and only if $y_1, y_2 \neq 0$ (i.e., at preimages of the interior of the annulus), and “ $2 \Rightarrow 1$ ” everywhere.
- The Burer–Monteiro lift (BM) under the smoothness assumption satisfies “local \Rightarrow local” everywhere, “ $1 \Rightarrow 1$ ” at R if and only if $\text{rank}(R) = r$ (i.e., at preimages of points of rank r), and “ $2 \Rightarrow 1$ ” everywhere. Moreover, we get an expression for the tangent cones to \mathcal{X} :

$$T_X \mathcal{X} = \{V \in \mathbb{S}^n : V \in T_X \mathbb{S}_{\geq 0}^n \cap T_X \mathbb{R}_{\leq r}^{n \times n} \text{ and } \langle A_i, V \rangle = 0 \text{ for all } i\}. \quad (4.4)$$

In particular, this proves Propositions 2.5 and 2.7. Note that an expression for $T_X \mathbb{S}_{\geq 0}^n \cap T_X \mathbb{R}_{\leq r}^{n \times n}$ is derived in Proposition 2.7 (incidentally, also as a consequence of the sufficient condition for “ $2 \Rightarrow 1$ ” used in Proposition 4.10).

Proof For the first three bullet points, consider the lift $\psi(y) = y^2$ from $\mathcal{N} = \mathbb{R}$ to $\mathcal{Z} = \mathbb{R}_{\geq 0}$. Observe that it satisfies “local \Rightarrow local” everywhere, “ $1 \Rightarrow 1$ ” at $y \neq 0$, and the sufficient condition $A_y = T_{\psi(y)}\mathcal{Z}$ for “ $2 \Rightarrow 1$ ” at $y = 0$. Indeed, at $y \neq 0$ we have $\mathbf{L}_y(\dot{y}) = 2y\dot{y}$ which is an isomorphism of $T_y \mathcal{N} = \mathbb{R}$ and $T_{y^2} \mathcal{Z} = \mathbb{R}$; and at $y = 0$ we have $\mathbf{L}_y = 0$ and $\mathbf{Q}_y(\dot{y}) = 2\dot{y}^2$ by (3.5) so $A_y = \mathbf{Q}_y(\ker \mathbf{L}_y) + \text{im } \mathbf{L}_y = \mathbb{R}_{\geq 0} = T_0 \mathcal{Z}$. Propositions 4.7 and 4.9–4.11 imply that the first three lifts satisfy “local \Rightarrow local” and “ $2 \Rightarrow 1$ ” everywhere, and give the claimed “if” directions for “ $1 \Rightarrow 1$ ”. The “only if” directions follow from Theorem 2.4.

For the Burer–Monteiro lift, consider the lift $\psi(R) = RR^\top$ from $\mathcal{N} = \mathbb{R}^{n \times r}$ to $\mathcal{Z} = \mathbb{S}_{\geq 0}^n \cap \mathbb{R}_{\leq r}^{n \times n}$. Proposition 3.32 shows that ψ satisfies “local \Rightarrow local” and the sufficient condition $A_R = T_{RR^\top} \mathcal{Z}$ for “ $2 \Rightarrow 1$ ” everywhere, as well as “ $1 \Rightarrow 1$ ” at points R of rank r . Therefore, Propositions 4.7 and 4.9–4.11 imply that the Burer–Monteiro lift satisfies “local \Rightarrow local” and “ $2 \Rightarrow 1$ ” everywhere and “ $1 \Rightarrow 1$ ” at R if $\text{rank}(R) = r$. The “ $1 \Rightarrow 1$ ” property does not hold at other points by Theorem 2.4. Proposition 4.10 gives the claimed expression for the tangent cones to \mathcal{X} . \square

Example 4.13 We can now revisit the example from Sect. 1 about computing the smallest eigenvalue of a symmetric matrix $A = U \text{diag}(\lambda) U^\top$ with U orthogonal. There,

$$\mathcal{X} = \Delta^{d-1}, \quad \mathcal{M} = \mathbb{S}^{d-1} \quad \text{and} \quad \varphi(y) = \text{diag}(U^\top y y^\top U).$$

Observe that $\varphi(y) = (U^\top y)^{\odot 2}$, which is the composition of the diffeomorphism $y \mapsto U^\top y$ from the sphere to itself and the Hadamard lift from Example 4.4. We conclude that this lift satisfies “ $2 \Rightarrow 1$ ” everywhere on \mathcal{M} by Proposition 3.27 and Corollary 4.12. Therefore, any 2-critical point for (Q) maps to a stationary point for (P), for any cost f . If f is convex, then since \mathcal{X} is also convex any stationary point for (P) is globally optimal. Thus, in this case any 2-critical point for (Q) is globally optimal and its nonconvexity is benign. This is well-known for the eigenvalue problem, which corresponds to the case of linear f .

Example 4.14 Proposition 4.11 together with Corollary 4.12 implies the properties stated in Proposition 2.6 for the lift (HadProd) of stochastic matrices. Indeed, the set of stochastic matrices $\mathcal{X} = \{X \in \mathbb{R}_{\geq 0}^{n \times m} : X^\top \mathbb{1}_n = \mathbb{1}_m\}$ is just the product of simplices $\mathcal{X} = (\Delta^{n-1})^m$, and the lift (HadProd) is precisely the m -fold power lift of (Had). Thus, Proposition 4.11(b)-(d) yields “local \Rightarrow local” and “ $2 \Rightarrow 1$ ” everywhere and “ $1 \Rightarrow 1$ ” at tuples (y_i) with no zero entries. Furthermore, since simplices are closed convex sets and hence Clarke-regular [49, Thm. 6.9], the tangent cone to their product is equal to the product of tangent cones (i.e., equality in Proposition 4.11(a) holds) by [49, Prop. 6.41], and hence “ $1 \Rightarrow 1$ ” does not hold elsewhere.

5 Analysis of low rank lifts

In this section, we use our theory to prove the remaining results from Sect. 2 concerning lifts of low rank matrices and tensors. Some of the straightforward but technical arguments are omitted here and are given in the arxiv version of this paper [39].

5.1 Proof of Proposition 2.8 (LR^\top lift)

The “local \Rightarrow local” property at “balanced” factorizations (L, R) satisfying $\text{rank}(L) = \text{rank}(R) = \text{rank}(LR^\top)$ was proved in [38, Prop. 2.34] by showing that (in our terminology) SLP holds there. We show “local \Rightarrow local” does not hold at other pairs (L, R) anywhere else by disproving SLP there via an explicit construction, see the arxiv version for details [39, Sec. 5.1].

For “ $1 \Rightarrow 1$ ” and “ $2 \Rightarrow 1$ ”, note that \mathcal{M} is a linear space, hence (3.5) gives

$$\mathbf{L}_{(L,R)}(\dot{L}, \dot{R}) = \dot{L}R^\top + L\dot{R}^\top, \quad \mathbf{Q}_{(L,R)}(\dot{L}, \dot{R}) = 2\dot{L}\dot{R}^\top.$$

If $\text{rank}(LR^\top) = r$, then [38, Prop. 2.15] (which is a slight generalization of the proof of “ $1 \Rightarrow 1$ ” in Proposition 2.7) shows $\text{im } \mathbf{L}_{(L,R)} = \mathbf{T}_{LR^\top} \mathcal{X}$ hence “ $1 \Rightarrow 1$ ” holds. If $\text{rank}(LR^\top) < r$ then $\mathbf{T}_{LR^\top} \mathcal{X}$ is not a linear space [28, Thm. 2.2], hence “ $1 \Rightarrow 1$ ” does not hold at (L, R) by Theorem 2.4.

To show “ $2 \Rightarrow 1$ ” holds everywhere on \mathcal{M} , it suffices to show $B_{(L,R)}^* \subseteq (T_X \mathcal{X})^*$ whenever $\text{rank}(LR^\top) < r$ by Theorem 3.23. Since $\text{rank}(LR^\top) < r$ we must have either $\text{rank}(L) < r$ or $\text{rank}(R) < r$, assume the former (the case $\text{rank}(R) < r$ is similar). Then there exists $w \in \mathbb{R}^r$ such that $Lw = 0$ and $\|w\|^2 = 1$. For any $u \in \mathbb{R}^m$, $v \in \mathbb{R}^n$, and $i \in \mathbb{N}$, let $\dot{L}_i = i^{-1}uw^\top$ and $\dot{R}_i = (i/2)vw^\top$. Then

$$\mathbf{L}_{(L,R)}(\dot{L}_i, \dot{R}_i) = i^{-1}uw^\top R^\top \xrightarrow{i \rightarrow \infty} 0, \quad \mathbf{Q}_{(L,R)}(\dot{L}_i, \dot{R}_i) = uv^\top,$$

showing that $uv^\top \in B_{(L,R)}$. Thus, $B_{(L,R)}$ contains all rank-1 matrices, showing that $B_y^* = \{0\} = (T_X \mathcal{X})^*$.

5.2 Proof of Proposition 2.9 (desingularization lift)

We show that “local \Rightarrow local” does not hold at $(X, \mathcal{S}) \in \mathcal{M}$ if $\text{rank}(X) < r$ by disproving SLP via an explicit construction, see the arxiv version [39, App. 5.2]. We show “local \Rightarrow local” does hold at $(X, \mathcal{S}) \in \mathcal{M}$ if $\text{rank}(X) = r$ by showing that “ $1 \Rightarrow 1$ ” holds there, which suffices by Proposition 2.12.

For “ $1 \Rightarrow 1$ ” and “ $2 \Rightarrow 1$ ”, we use the results of Example 3.30. Using the notation of that example, recall that every $(X, \mathcal{S}) \in \mathcal{M}$ is in the image of a chart $\psi(Z, W)$, and that $\mathbf{L}_{(Z,W)}$ and $\mathbf{Q}_{(Z,W)}$ in this chart are given by (3.9).

Suppose $\text{rank}(X) = r$ and $(X, \mathcal{S}) = \psi(Z, W)$. Note that $\text{col}(X) = \text{col}(Z)$, so $\text{rank}(Z) = r$. If $\mathbf{L}_{(Z,W)}(\dot{Z}, \dot{W}) = 0$, then $\dot{Z} = 0$ and $Z\dot{W} = 0$. This implies $\dot{W} = 0$ since Z has full column rank. Thus, $\mathbf{L}_{(Z,W)}$ is injective, but since its domain has dimension $(m + n - r)r = \dim \mathbb{R}_{\leq r}^{m \times n}$, we conclude that it is an isomorphism. Thus, “ $1 \Rightarrow 1$ ” holds at (Z, W) by Theorem 2.4. If $\text{rank}(X) < r$ then $T_X \mathcal{X}$ is not a linear space [28, Thm. 2.2], hence “ $1 \Rightarrow 1$ ” cannot hold for any lift by Theorem 2.4.

Suppose $\text{rank}(X) < r$. We show “ $2 \Rightarrow 1$ ” holds at (Z, W) by showing that $B_{(Z,W)}^* \subseteq (T_X \mathcal{X})^*$. To that end, note that $\text{rank}(Z) = \text{rank}(X) < r$, so there is a unit vector $w \in \mathbb{R}^r$ satisfying $Zw = 0$. Let $u \in \mathbb{R}^m$ and $v \in \mathbb{R}^{n-r}$ be arbitrary. For any $i \in \mathbb{N}$, let $\dot{Z}_i = i^{-1}uw^\top$ and $\dot{W}_i = i w v^\top$. Then

$$\begin{aligned} \mathbf{L}_{(Z,W)}(\dot{Z}_i, \dot{W}_i) &= [-i^{-1}uw^\top W - i(Zw)v^\top, i^{-1}uw^\top] \Pi \xrightarrow{i \rightarrow \infty} 0, \\ \mathbf{Q}_{(Z,W)}(\dot{Z}_i, \dot{W}_i) &\equiv [-2uv^\top, 0] \Pi. \end{aligned}$$

We conclude that

$$\begin{aligned} B_{(Z,W)} &\supseteq (\mathbb{R}_{\leq 1}^{m \times (n-r)} \times \{0\}) \Pi + \text{im } \mathbf{L}_{(Z,W)} \\ &\implies B_{(Z,W)}^* \subseteq (\{0\} \times \mathbb{R}^{m \times r}) \Pi \cap (\text{im } \mathbf{L}_{(Z,W)})^\perp. \end{aligned}$$

To characterize $(\text{im } \mathbf{L}_{(Z,W)})^\perp$, observe that $V = [V_1, V_2] \Pi \in \mathbb{R}^{m \times n}$ with $V_1 \in \mathbb{R}^{m \times (n-r)}$ satisfies $V \in (\text{im } \mathbf{L}_{(Z,W)})^\perp$ iff the following holds for all $(\dot{Z}, \dot{W}) \in \mathbb{R}^{m \times r} \times$

$\mathbb{R}^{r \times (n-r)}$:

$$\begin{aligned} \langle V, \mathbf{L}_{(Z,W)}(\dot{Z}, \dot{W}) \rangle &= \langle V_1, -\dot{Z}W - Z\dot{W} \rangle + \langle V_2, \dot{Z} \rangle \\ &= \langle \dot{Z}, V_2 - V_1W^\top \rangle - \langle \dot{W}, Z^\top V_1 \rangle = 0. \end{aligned}$$

This is equivalent to $V_2 = V_1W^\top$ and $Z^\top V_1 = 0$. Thus, if $V_1 = 0$ then $V = 0$, hence $B_{(Z,W)}^* = \{0\} = (\mathbb{T}_X \mathcal{X})^*$. This shows “2 \Rightarrow 1” holds at (Z, W) , and hence also at (X, \mathcal{S}) by Proposition 3.27.

5.3 Multilinear lifts, and tensors

In this section, we prove two obstructions to “2 \Rightarrow 1” for multilinear lifts, which apply in particular to lifts defined by tensor factorizations and linear neural networks as discussed in Sects. 2.4–2.5.

Proposition 5.1 *Suppose $\varphi: \mathcal{M} \rightarrow \mathcal{X} \subseteq \mathcal{E}$ is a smooth lift where $\mathcal{M} \subseteq \mathcal{E}' = \mathcal{E}_1 \times \dots \times \mathcal{E}_d$ is a smooth embedded submanifold of a product of Euclidean spaces \mathcal{E}_i , and φ is defined on all of \mathcal{E}' and is multilinear in its d arguments. If \mathcal{M} contains a point (y_1, \dots, y_d) such that $y_i = 0$ for three indices i , and $0 = \varphi(y_1, \dots, y_d)$ is not an isolated point of \mathcal{X} , then φ does not satisfy “2 \Rightarrow 1” at (y_1, \dots, y_d) .*

Proof Let $(y_1, \dots, y_d) \in \mathcal{E}'$. Note that if $y_i = 0$ for some i , then $\varphi(y_1, \dots, y_d) = 0$ by the multilinearity of φ . Similarly, multilinearity gives $D\varphi(y_1, \dots, y_d) = D^2\varphi(y_1, \dots, y_d) = 0$ if $y_i = 0$ for at least three indices i . Hence (3.8) gives $\mathbf{L}_{(y_1, \dots, y_d)} = 0$ and $\mathbf{Q}_{(y_1, \dots, y_d)} = 0$. This implies

$$(\text{im } \mathbf{L}_{(y_1, \dots, y_d)})^\perp \cap (\mathbf{Q}_{(y_1, \dots, y_d)}(\mathbb{T}_{(y_1, \dots, y_d)} \mathcal{M}))^* = \mathcal{E}.$$

The necessary condition for “2 \Rightarrow 1” given by the last implication in Theorem 3.23 is satisfied iff $(\mathbb{T}_0 \mathcal{X})^* = \mathcal{E}$, or equivalently $\mathbb{T}_0 \mathcal{X} = \{0\}$. This holds if and only if 0 is an isolated point of \mathcal{X} , since if $(x_i) \subseteq \mathcal{X} \setminus \{0\}$ is a sequence converging to 0, then after passing to a subsequence $(x_i / \|x_i\|)$ converges and gives a nonzero element of $\mathbb{T}_0 \mathcal{X}$. \square

Proposition 5.1 implies that the lifts corresponding to linear neural networks, as well as standard tensor decompositions such as CPD, Tensor Train (TT), and Tuckett, all do not satisfy “2 \Rightarrow 1” as points with at least three zero factors.

Proposition 5.1 might suggest that failure of “2 \Rightarrow 1” can be avoided by normalizing the arguments of the lift to have unit norm. Specifically, by multilinearity of φ we have

$$\varphi(y_1, \dots, y_d) = \left(\prod_{i=1}^d \|y_i\| \right) \varphi\left(\frac{y_1}{\|y_1\|}, \dots, \frac{y_d}{\|y_d\|} \right), \quad \text{whenever } y_i \neq 0 \text{ for all } i.$$

Using this observation, one could replace a lift $\varphi: \mathbb{R}^{n_1} \times \dots \times \mathbb{R}^{n_d} \rightarrow \mathcal{X}$ to a product of Euclidean spaces by a lift $\psi: \mathbb{R} \times S^{n_1-1} \times \dots \times S^{n_d-1}$ to a product of \mathbb{R} and several spheres, satisfying $\psi(\lambda, x_1, \dots, x_d) = \lambda\varphi(x_1, \dots, x_d)$. Only one factor can be zero

in this new lift, so Proposition 5.1 does not apply and we might hope that “ $2 \Rightarrow 1$ ” holds. Unfortunately, this may not resolve the problem as there is another obstruction to “ $2 \Rightarrow 1$ ” for the following specific form of a lift.

Proposition 5.2 *Suppose $\varphi: \mathcal{M} \rightarrow \mathcal{X}$ is a smooth lift of the form*

$$\varphi(\lambda, Y_1, \dots, Y_d) = \sum_{i=1}^r \lambda_i \cdot (Y_1)_{:,i} \otimes \cdots \otimes (Y_d)_{:,i},$$

where $\mathcal{M} \subseteq \mathbb{R}^r \times \mathbb{R}^{n_1 \times r} \times \cdots \times \mathbb{R}^{n_d \times r}$. Denote $X = \varphi(\lambda, Y_1, \dots, Y_d)$. If $d \geq 3$ and

$$\text{col}(Y_1)^\perp \otimes \cdots \otimes \text{col}(Y_d)^\perp \not\subseteq (\text{T}_X \mathcal{X})^*, \quad (5.1)$$

then φ does not satisfy “ $2 \Rightarrow 1$ ” at $(\lambda, Y_1, \dots, Y_d)$ for any $\lambda \in \mathbb{R}^r$. If $d = 2$ and (5.1) holds, then φ does not satisfy “ $2 \Rightarrow 1$ ” at $(0, Y_1, Y_2)$.

Proof of Proposition 5.2 For any $W \in \text{col}(Y_1)^\perp \otimes \cdots \otimes \text{col}(Y_d)^\perp$, we have

$$\begin{aligned} & \langle W, \text{D}\varphi(\lambda, Y_1, \dots, Y_d)[\dot{\lambda}, \dot{Y}_1, \dots, \dot{Y}_d] \rangle \\ &= \langle W, \text{D}^2\varphi(\lambda, Y_1, \dots, Y_d)[(\dot{\lambda}, \dot{Y}_1, \dots, \dot{Y}_d), (\dot{\lambda}, \dot{Y}_1, \dots, \dot{Y}_d)] \rangle = 0, \end{aligned}$$

for all $(\dot{Y}_1, \dots, \dot{Y}_d)$ if $d \geq 3$ or $d = 2$ and $\lambda = 0$, by multilinearity. Since $\mathbf{L}_{(\lambda, Y_1, \dots, Y_d)}$ is the restriction of $\text{D}\varphi(\lambda, Y_1, \dots, Y_d)$ to $\text{T}_{(\lambda, Y_1, \dots, Y_d)}\mathcal{M}$ and $\mathbf{Q}_{(\lambda, Y_1, \dots, Y_d)}$ is given by (3.8), we get

$$\text{col}(Y_1)^\perp \otimes \cdots \otimes \text{col}(Y_d)^\perp \subseteq (\text{im } \mathbf{L}_{(\lambda, Y_1, \dots, Y_d)})^\perp \cap (\mathbf{Q}_{(\lambda, Y_1, \dots, Y_d)}(\text{T}_{(\lambda, Y_1, \dots, Y_d)}\mathcal{M}))^*,$$

if either $d \geq 3$ or $d = 2$ and $\lambda = 0$. Thus, if $\text{col}(Y_1)^\perp \otimes \cdots \otimes \text{col}(Y_d)^\perp \not\subseteq (\text{T}_X \mathcal{X})^*$ then the necessary condition for “ $2 \Rightarrow 1$ ” from Theorem 3.23 does not hold. \square

Proposition 5.2 applies in particular to lifts corresponding to symmetric and normalized CP decompositions and ODECO tensors [48], as well as the SVD lift (SVD). As discussed in Sect. 2, these obstructions to “ $2 \Rightarrow 1$ ” imply that guarantees for second-order optimization algorithms running on (Q) must use the structure in the particular cost function involved. This is particularly significant since our obstructions apply to a broad class of lifts arising naturally in several applications.

6 Conclusions and future work

For the pair of problems (Q) and (P), we characterized the properties the lift $\varphi: \mathcal{M} \rightarrow \mathcal{X}$ needs to satisfy in order to map desirable points of (Q) to desirable points of (P). We noted that global minima for (Q) always map to global minima for (P) (Proposition 3.5), and showed that local minima for (Q) map to local minima for (P) if and only if φ is open (Theorem 2.3). We also showed that 1-critical points for (Q) map to stationary points for (P) if and only if the differential of φ , viewed as a map from tangent spaces of

\mathcal{M} to tangent cones of \mathcal{X} , is surjective (Theorem 2.4). This requires the tangent cones of \mathcal{X} to be linear spaces. We then characterized when 2-critical points for (Q) map to stationary points for (P), and gave two sufficient conditions and a necessary condition that may be easier to check for some examples (Theorem 3.23). We explained several techniques to compute all quantities involved in these conditions in Sect. 3.4.

Using our theory, we studied the above properties for a variety of lifts, including several lifts of low-rank matrices and tensors (Sect. 5) and the Burer–Monteiro lift for smooth SDPs (Corollary 4.12). We also proposed a systematic construction of lifts using fiber products that applies when \mathcal{X} is the preimage of a smooth function (Sect. 4). We gave conditions under which it satisfies our desirable properties. In some cases, we can also obtain an expression for the tangent cones of \mathcal{X} simultaneously with “ $2 \Rightarrow 1$ ”, as explained in Sect. 3.5.

We end by listing several future directions suggested by this work.

- (a) “ $k \Rightarrow 1$ ” for general k : Several lifts of interest, notably tensor factorizations with more than two factors, do not satisfy “ $2 \Rightarrow 1$ ”. It would therefore be interesting to characterize “ $k \Rightarrow 1$ ” for general k , i.e., when do k -critical points for (Q) map to stationary points for (P) for any k times differentiable cost f ?

Do lifts that are multilinear in k arguments, such as order- k tensor lifts, satisfy “ $k \Rightarrow 1$ ”?

What can be said about “ $k \Rightarrow \ell$ ” for $\ell > 1$? Already for $\ell = 2$, the second-order optimality conditions on \mathcal{X} can be involved [50, Thm. 3.45]. On the positive side, if “ $1 \Rightarrow 1$ ” holds at a preimage of a smooth point, then “ $k \Rightarrow k$ ” holds there for all $k \geq 1$ by Proposition 2.12.

- (b) **Robust “ $k \Rightarrow 1$ ”**: Algorithms run for finitely many iterations in practice, hence can only find approximate k -critical points for (Q). It is therefore important to characterize “robust” versions of “ $k \Rightarrow 1$ ”, guaranteeing that approximate k -critical points for (Q) map to approximate stationary points for (P).

Note that if \mathcal{X} lacks regularity, care is needed when defining approximate stationarity for (P), see [40].

- (c) **Obstructions to “local \Rightarrow local” and “ $k \Rightarrow 1$ ”**:

For some sets \mathcal{X} , we are not aware of any lifts which satisfy, say, “local \Rightarrow local” or “ $2 \Rightarrow 1$ ”. Are there fundamental obstructions which preclude existence of such lifts for those sets and others? For example, is there a lift for low-rank tensors satisfying “ $2 \Rightarrow 1$ ”? Is there a lift for $\mathbb{R}_{\leq r}^{m \times n}$ satisfying “local \Rightarrow local”?

- (d) **Regularization on the lift**: It is common to modify (Q) by adding a regularizer to $g = f \circ \varphi$, see [33, 53, 55]. For example, with the lift $(L, R) \mapsto LR^\top$, we may regularize (Q) by adding $\frac{1}{2} (\|L\|_F^2 + \|R\|_F^2)$, motivated by the fact that its minimum over a fiber $\{(L, R) : LR^\top = X\}$ is the nuclear norm $\|X\|_*$ [53]. Our framework does not directly apply in this case (because the regularizer may not be constant over fibers, hence may not factor through φ). Can it be extended to relate the landscape of the regularized (Q) to that of (P)?

- (e) **Bypassing tangent cones via lifts**:

To verify “ $1 \Rightarrow 1$ ” and “ $2 \Rightarrow 1$ ” on concrete examples of \mathcal{X} using the theory in this paper, we need to understand the tangent cones to \mathcal{X} , which is often challenging. Many sets \mathcal{X} encountered in applications are only defined implicitly via

a lift $\varphi: \mathcal{M} \rightarrow \mathcal{E}$. Examples include the set of tensors admitting a certain type of factorization, the set of functions parametrized by a given neural network architecture, and the set of positions and orientations attainable by a robotic arm with a given joint configuration. Are there sufficient conditions for “ $k \Rightarrow 1$ ” that can be checked using φ and \mathcal{M} alone, without an explicit expression for the tangent cones to \mathcal{X} ?

- (f) **Dynamical systems on \mathcal{M} and their image on \mathcal{X} :** This paper is focused on comparing properties of points on \mathcal{M} and their images on \mathcal{X} . In contrast, several applications are concerned with properties of entire trajectories of dynamical systems on \mathcal{M} , and it may be interesting to compare these properties with their counterparts for the images of the trajectories on \mathcal{X} . Examples of such comparisons include relating gradient flow on the weights of a neural network to gradient flow in function or measure spaces [6, 7, 29], and the “algorithmic equivalence” technique used in [3, 24, 43] to study mirror descent by showing that its continuous-time analogue is equivalent to gradient flow on a reparametrized problem.

Acknowledgements We thank Christopher Criscitiello and Quentin Rejzock for numerous conversations and comments on drafts of this paper. NB was supported by the Swiss State Secretariat for Education, Research and Innovation (SERI) under contract number MB22.00027. JK was supported in part by NSF DMS 2309782 and NSF CISE-IIS 2312746.

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/>.

A Lifts preserving local minima

We characterize the lifts that map local minima of (\mathbf{Q}) to local minima of (\mathbf{P}) . To this end, we introduce a number of properties related to preservation of local minima and then prove that they are all equivalent. Recall that \bar{S} is our notation for the closure of a set S .

Definition A.1 Let $\varphi: \mathcal{M} \rightarrow \mathcal{X}$ be a continuous, surjective map from a topological space \mathcal{M} to a metric space \mathcal{X} with distance dist , and let $x = \varphi(y)$.

1. φ is *open* at y if $\varphi(U)$ is a neighborhood of x in \mathcal{X} for all neighborhoods U of y in \mathcal{M} .
2. φ is *approximately open* at y if $\overline{\varphi(U)}$ is a neighborhood of x in \mathcal{X} for all neighborhoods U of y in \mathcal{M} .
3. φ satisfies the *Subsequence Lifting Property (SLP)* at y if for every sequence $(x_i)_{i \geq 1} \subseteq \mathcal{X}$ converging to x there exists a subsequence indexed by $(i_j)_{j \geq 1}$ and a sequence $(y_{i_j})_{j \geq 1} \subseteq \mathcal{M}$ converging to y such that $\varphi(y_{i_j}) = x_{i_j}$ for all $j \geq 1$.

4. φ satisfies the *Approximate Subsequence Lifting Property (ASLP)* at y if for every sequence $(x_i)_{i \geq 1} \subseteq \mathcal{X}$ converging to x and every sequence $(\epsilon_i)_{i \geq 1} \subseteq \mathbb{R}_{>0}$ converging to 0 there exists a subsequence indexed by $(i_j)_{j \geq 1}$ and a sequence $(y_{i_j})_{j \geq 1} \subseteq \mathcal{M}$ converging to y such that $\text{dist}(\varphi(y_{i_j}), x_{i_j}) \leq \epsilon_{i_j}$ for all $j \geq 1$.

Theorem A.2 *If \mathcal{M} is Hausdorff, second-countable and locally compact (all of which hold if \mathcal{M} is a topological manifold), then the four properties of φ at $y \in \mathcal{M}$ in Definition A.1 are equivalent to each other and to the “local \Rightarrow local” property at y (Definition 2.2(a)).*

Proof We show that $\text{ASLP} \implies \text{approximate openness} \implies \text{openness} \implies \text{SLP} \implies \text{“local} \Rightarrow \text{local”} \implies \text{ASLP}$.

Suppose φ satisfies ASLP at y . Suppose there exists a neighborhood U of y such that $\varphi(U)$ is not a neighborhood of $x = \varphi(y)$. Then we can find a sequence $(x_i)_{i \geq 1} \subseteq \mathcal{X}$ such that $x_i \rightarrow x$ but $x_i \notin \overline{\varphi(U)}$ for all i . Set $\epsilon_i = \frac{1}{2} \text{dist}(x_i, \overline{\varphi(U)}) > 0$ and apply ASLP to find a sequence $(y_i)_{i \geq 1} \subseteq \mathcal{M}$ such that $y_i \rightarrow y$ and $\text{dist}(\varphi(y_i), x_i) \leq \epsilon_i$. Because $\text{dist}(\varphi(y_i), x_i) < \text{dist}(x_i, \overline{\varphi(U)})$, we have $\varphi(y_i) \notin \overline{\varphi(U)}$ for all i . However, because U is a neighborhood of y and $y_i \rightarrow y$, we must have $y_i \in U$ for all large i , a contradiction. Thus, $\overline{\varphi(U)}$ is a neighborhood of x , so φ is approximately open at y .

Suppose φ is approximately open at y , and let U be a neighborhood of y in \mathcal{M} . Because \mathcal{M} is locally compact, we can find a compact neighborhood $V \subseteq U$ of x . Since φ is continuous and V is compact, we have that $\varphi(V)$ is compact; since \mathcal{X} is Hausdorff (it is a metric space), it follows that $\varphi(V)$ is closed. Combining with the fact that φ is approximately open at y , we deduce that $\varphi(V)$ is a neighborhood of x . Since $\varphi(U) \supseteq \varphi(V)$, we conclude that $\varphi(U)$ is a neighborhood of x as well. Thus, φ is open at y .

Suppose φ is open at y , and $(x_j)_{j \geq 1} \subseteq \mathcal{X}$ converges to $x = \varphi(y)$. Owing to the topological properties of \mathcal{M} , there is a sequence of open neighborhoods U_i of y with compact closures such that $U_i \supseteq \overline{U_{i+1}}$ and $\bigcap_{i=1}^{\infty} U_i = \{y\}$, see Lemma A.3 following this proof. Because φ is open, each $\varphi(U_i)$ is an open neighborhood of x such that $\varphi(U_i) \supseteq \varphi(U_{i+1})$ and $x \in \bigcap_{i=1}^{\infty} \varphi(U_i)$. Moreover, because $\varphi(U_i)$ is a neighborhood of x and $x_j \rightarrow x$, there exists index $J(i)$ such that $x_j \in \varphi(U_i)$ for all $j \geq J(i)$. After passing to a subsequence of (x_j) , we may assume $x_j \in \varphi(U_j)$ and pick $y_j \in U_j$ satisfying $x_j = \varphi(y_j)$. Because (y_j) is an infinite sequence contained in the compact set $\overline{U_1}$, after passing to a subsequence again we may assume that $\lim_j y_j$ exists. With i arbitrary, we have for all $j > i$ that $y_j \in U_j \subseteq U_{i+1}$, hence that $\lim_j y_j \in \overline{U_{i+1}} \subseteq U_i$. This holds for all i , hence $\lim_j y_j \in \bigcap_i U_i = \{y\}$. Thus, $y = \lim_i y_i$ and $\varphi(y_i) = x_i$, so φ satisfies SLP.

Suppose φ satisfies SLP at y . Let $f: \mathcal{X} \rightarrow \mathbb{R}$ be a cost function on \mathcal{X} and $g = f \circ \varphi$. Suppose $x = \varphi(y)$ is not a local minimum for f on \mathcal{X} , that is, there exists a sequence $(x_i)_{i \geq 1} \subseteq \mathcal{X}$ converging to x such that $f(x_i) < f(x)$ for all i . Applying SLP, after passing to a subsequence we can find a sequence $(y_i)_{i \geq 1} \subseteq \mathcal{M}$ converging to y such that $\varphi(y_i) = x_i$. Since $g(y_i) = f(x_i) < f(x) = g(y)$ and $y_i \rightarrow y$, we conclude that y is not a local minimum for g . By contrapositive, this shows that φ satisfies the “local \Rightarrow local” property at y .

For the last implication, we proceed by contrapositive once again. Suppose φ does not satisfy ASLP at y . Then, we can find sequences $(x_i)_{i \geq 1} \subseteq \mathcal{X}$ converging to x and

$(\epsilon_i)_{i \geq 1} \subseteq \mathbb{R}_{>0}$ converging to 0 such that no subsequence of (x_i) can be approximately lifted to \mathcal{M} in the sense of ASLP. Let $\bar{B}(x, \epsilon) = \{x' \in \mathcal{X} : \text{dist}(x, x') \leq \epsilon\}$. Notice that $x = \varphi(y) \notin \bar{B}(x_i, \epsilon_i)$ for all but finitely many indices i , as otherwise the constant sequence $y_i \equiv y$ would give an approximate lift of a subsequence. Since $x_i \rightarrow x$ and $\epsilon_i \rightarrow 0$, after passing to a subsequence we may assume that the closed balls $\bar{B}(x_i, \epsilon_i)$ are pairwise disjoint and none contain x . Define the following sum of smooth bump functions centered at the x_i

$$f(x') = \begin{cases} -\exp\left(1 - \frac{1}{1 - (\text{dist}(x_i, x')/\epsilon_i)^2}\right) & \text{if } x' \in \bar{B}(x_i, \epsilon_i) \text{ for some } i, \\ 0 & \text{otherwise.} \end{cases}$$

This is well defined because the balls $\bar{B}(x_i, \epsilon_i)$ are disjoint. (As a side note, we remark that if \mathcal{X} is a metric subspace of a Euclidean space \mathcal{E} as in our general treatment, then f extends to a smooth function on \mathcal{E} .) Note that x is not a local minimum for f since $x_i \rightarrow x$ and $f(x_i) = -1 < 0 = f(x)$. However, y is a local minimum for $g = f \circ \varphi$. Indeed, if there was a sequence (y_i) converging to y such that $g(y_i) < g(y) = 0$, then we would have $\varphi(y_i) \in \bar{B}(x_{n_i}, \epsilon_{n_i})$ for an infinite subsequence (n_i) with $n_i \rightarrow \infty$ (since we must have $\varphi(y_i) \rightarrow x$ by continuity of φ), showing that (y_i) is an approximate lift of the subsequence (x_{n_i}) : a contradiction to our assumptions about (x_i) , (ϵ_i) . Thus, φ does not satisfy the “local \Rightarrow local” property at y . \square

Lemma A.3 *Suppose \mathcal{M} is Hausdorff, second-countable, and locally compact. Then for any $y \in \mathcal{M}$ there is a sequence of open neighborhoods U_i of y with compact closures such that $U_i \supseteq \overline{U_{i+1}}$ and $\bigcap_{i=1}^{\infty} U_i = \{y\}$.*

Proof Because \mathcal{M} is second-countable and locally compact, we can find a countable basis of open neighborhoods with compact closures $\{V_j\}_{j \geq 1}$ for y . Since \mathcal{M} is Hausdorff and $\{V_j\}$ is a basis for y , we have $\bigcap_{j=1}^{\infty} V_j = \{y\}$. Indeed, if $y' \neq y$, then there exists a neighborhood of y not containing y' , and this neighborhood contains V_i for some i by definition of a local basis. By replacing V_i by $\bigcap_{j=1}^i V_j$ (which preserves their intersection), we may assume $V_i \supseteq V_{i+1}$. We construct $\{U_i\}_{i \geq 1}$ inductively. Set $U_1 = V_1$, which is an open neighborhood of y with compact closure by assumption. Having constructed U_1, \dots, U_i , use local compactness to find a compact neighborhood $K_{i+1} \subseteq V_{i+1} \cap U_i$ of y and let U_{i+1} be the interior of K_{i+1} . Then U_{i+1} is an open neighborhood of y by construction, and $\overline{U_{i+1}} \subseteq K_{i+1} \subseteq U_i$ which also shows $\overline{U_{i+1}}$ is compact as a closed subset of the compact set K_{i+1} . Finally, we have $\{y\} \subseteq \bigcap_{i=1}^{\infty} U_i \subseteq \bigcap_{i=1}^{\infty} V_i = \{y\}$ hence $\bigcap_{i=1}^{\infty} U_i = \{y\}$. \square

References

1. Ablin, P.: Deep orthogonal linear networks are shallow. arXiv preprint [arXiv:2011.13831](https://arxiv.org/abs/2011.13831) (2020)
2. Absil, P.-A., Mahony, R., Sepulchre, R.: Optimization Algorithms on Matrix Manifolds. Princeton University Press, Princeton (2008)
3. Amid, E., Warmuth, M.K.K.: Reparameterizing mirror descent as gradient descent. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., Lin, H. (eds.) Advances in Neural Information Processing Systems, vol. 33, pp. 8430–8439. Curran Associates Inc (2020)

4. Antonakopoulos, K., Mertikopoulos, P., Piliouras, G., Wang, X.: AdaGrad avoids saddle points. In: International Conference on Machine Learning, pp. 731–771. PMLR (2022)
5. Ay, N., Jost, J., Van Lê, H., Schwachhöfer, L.: Information Geometry, vol. 64. Springer, New York (2017)
6. Bach, F., Chizat, L.: Gradient descent on infinitely wide neural networks: global convergence and generalization. arXiv preprint [arXiv:2110.08084](https://arxiv.org/abs/2110.08084) (2021)
7. Bah, B., Rauhut, H., Terstiege, U., Westdickenberg, M.: Learning deep linear neural networks: Riemannian gradient flows and convergence to global minimizers. *Inf. Inference A J IMA* **11**, 307–353 (2021)
8. Ben-Tal, A., Teboulle, M., Charnes, A.: The role of duality in optimization problems involving entropy functionals with applications to information theory. *J. Optim. Theory Appl.* **58**(2), 209–223 (1988)
9. Boumal, N.: An Introduction to Optimization on Smooth Manifolds. Cambridge University Press, Cambridge (2023)
10. Boumal, N., Voroninski, V., Bandeira, A.S.: Deterministic guarantees for Burer–Monteiro factorizations of smooth semidefinite programs. *Commun. Pure Appl. Math.* **73**(3), 581–608 (2020)
11. Burer, S., Monteiro, R.D.: A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization. *Math. Program.* **95**(2), 329–357 (2003)
12. Burer, S., Monteiro, R.D.: Local minima and convergence in low-rank semidefinite programming. *Math. Program.* **103**(3), 427–444 (2005)
13. Cartis, C., Gould, N.I., Toint, P.L.: Second-order optimality and beyond: characterization and evaluation complexity in convexly constrained nonlinear optimization. *Found. Comput. Math.* **18**, 1073–1107 (2018)
14. Chok, J., Vasil, G.M.: Convex optimization over a probability simplex. arXiv preprint [arXiv:2305.09046](https://arxiv.org/abs/2305.09046) (2023)
15. Cichocki, A.: Tensor networks for big data analytics and large-scale optimization problems. arXiv preprint [arXiv:1407.3124](https://arxiv.org/abs/1407.3124) (2014)
16. Cifuentes, D.: On the Burer–Monteiro method for general semidefinite programs. *Optim. Lett.* **15**, 1–11 (2021)
17. Clarke, F.H., Ledyae, Y.S., Stern, R.J., Wolenski, P.R.: Nonsmooth Analysis and Control Theory, vol. 178. Springer, New York (2008)
18. Curtis, F., Lubberts, Z., Robinson, D.: Concise complexity analyses for trust region methods. *Optim. Lett.* **12**(8), 1713–1724 (2018)
19. Dellaert, F., Rosen, D.M., Wu, J., Mahony, R., Carlone, L.: Shonan rotation averaging: global optimality by surfing $SO(p)^n$. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.-M. (eds.) Computer vision—ECCV 2020, pp. 292–308. Springer, Cham (2020)
20. Deutsch, F.R.: Best Approximation in Inner Product Spaces. Springer (2012)
21. Ding, L., Wright, S.J.: On squared-variable formulations. arXiv preprint [arXiv:2310.01784](https://arxiv.org/abs/2310.01784) (2023)
22. Douik, A., Hassibi, B.: Non-negative matrix factorization via low-rank stochastic manifold optimization. In: 2020 Information Theory and Applications Workshop (ITA), pp. 1–5 (2020)
23. Edelman, A., Arias, T., Smith, S.: The geometry of algorithms with orthogonality constraints. *SIAM J. Matrix Anal. Appl.* **20**(2), 303–353 (1998)
24. Ghai, U., Lu, Z., Hazan, E.: Non-convex online learning via algorithmic equivalence. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, Advances in Neural Information Processing Systems, volume 35, pages 22161–22172. Curran Associates, Inc., (2022)
25. Ha, W., Liu, H., Foygel Barber, R.: An equivalence between critical points for rank constraints versus low-rank factorizations. *SIAM J. Optim.* **30**(4), 2927–2955 (2020)
26. Harris, J.: Algebraic Geometry: A First Course. Springer (1992)
27. Hiriart-Urruty, J.-B., Malick, J.: A fresh variational-analysis look at the positive semidefinite matrices world. *J. Optim. Theory Appl.* **153**(3), 551–577 (2012)
28. Hosseini, S., Luke, D.R., Uschmajew, A.: Tangent and Normal Cones for Low-Rank Matrices, pp. 45–53. Springer, Cham (2019)
29. Jacot, A., Gabriel, F., Hongler, C.: Neural tangent kernel: convergence and generalization in neural networks. In: Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R. (eds.) Advances in Neural Information Processing Systems, vol. 31. Curran Associates Inc (2018)
30. Journée, M., Bach, F., Absil, P.-A., Sepulchre, R.: Low-rank optimization on the cone of positive semidefinite matrices. *SIAM J. Optim.* **20**(5), 2327–2351 (2010)

31. Khrulkov, V., Oseledets, I.: Desingularization of bounded-rank matrix sets. *SIAM J. Matrix Anal. Appl.* **39**(1), 451–471 (2018)
32. Kohn, K., Merkh, T., Montúfar, G., Trager, M.: Geometry of linear convolutional networks. arXiv preprint [arXiv:2108.01538](https://arxiv.org/abs/2108.01538) (2021)
33. Kolb, C., Müller, C.L., Bischl, B., Rügamer, D.: Smoothing the edges: a general framework for smooth optimization in sparse regularization using Hadamard overparametrization. arXiv preprint [arXiv:2307.03571](https://arxiv.org/abs/2307.03571) (2023)
34. Kolda, T.G., Bader, B.W.: Tensor decompositions and applications. *SIAM Rev.* **51**(3), 455–500 (2009)
35. Leake, J., Vishnoi, N.K.: Optimization and sampling under continuous symmetry: examples and Lie theory. arXiv preprint [arXiv:2109.01080](https://arxiv.org/abs/2109.01080) (2021)
36. Lee, J.D., Panageas, I., Piliouras, G., Simchowitz, M., Jordan, M.I., Recht, B.: First-order methods almost always avoid strict saddle points. *Math. Program.* **176**(1), 311–337 (2019)
37. Lee, J.M.: *Introduction to Smooth Manifolds*. Springer (2012)
38. Levin, E.: *Towards Optimization on Varieties*. Undergraduate Senior Thesis. Princeton University, Princeton (2020)
39. Levin, E., Kileel, J., Boumal, N.: The effect of smooth parametrizations on nonconvex optimization landscapes. arXiv preprint [arXiv:2207.03512](https://arxiv.org/abs/2207.03512) (2022)
40. Levin, E., Kileel, J., Boumal, N.: Finding stationary points on bounded-rank matrices: a geometric hurdle and a smooth remedy. *Math. Program.* **199**, 831–864 (2022)
41. Lezcano Casado, M.: *Geometric optimisation on manifolds with applications to deep learning*. Ph.D. thesis, University of Oxford (2021)
42. Li, Q., McKenzie, D., Yin, W.: From the simplex to the sphere: faster constrained optimization using the Hadamard parametrization. *Inf. Inference A J. IMA* **12**(3), iaad017 (2023)
43. Li, Z., Wang, T., Lee, J.D., Arora, S.: Implicit bias of gradient descent on reparametrized models: On equivalence to mirror descent. arXiv preprint [arXiv:2207.04036](https://arxiv.org/abs/2207.04036) (2022)
44. Mishra, B., Meyer, G., Bonnabel, S., Sepulchre, R.: Fixed-rank matrix factorizations and Riemannian low-rank optimization. *Comput. Stat.* **29**(3–4), 591–621 (2014)
45. Ollikier, G., Gallivan, K.A., Absil, P.-A.: An apocalypse-free first-order low-rank optimization algorithm. arXiv preprint [arXiv:2201.03962](https://arxiv.org/abs/2201.03962) (2022)
46. Petersen, P., Raslan, M., Voigtlaender, F.: Topological properties of the set of functions generated by neural networks of fixed size. *Found. Comput. Math.* **21**(2), 375–444 (2021)
47. Phan, A.-H., Yamagishi, M., Mandic, D., Cichocki, A.: Quadratic programming over ellipsoids with applications to constrained linear regression and tensor decomposition. *Neural Comput. Appl.* **32**(11), 7097–7120 (2019)
48. Robeva, E.: Orthogonal decomposition of symmetric tensors. *SIAM J. Matrix Anal. Appl.* **37**(1), 86–102 (2016)
49. Rockafellar, R.T., Wets, R.J.-B.: *Variational Analysis*, vol. 317. Springer (2009)
50. Ruszczyński, A.: *Nonlinear Optimization*. Princeton University Press (2006)
51. Schneider, R., Uschmajew, A.: Convergence results for projected line-search methods on varieties of low-rank matrices via Łojasiewicz inequality. *SIAM J. Optim.* **25**(1), 622–646 (2015)
52. Siciliano, B., Sciacivco, L., Villani, L., Oriolo, G.: *Robotics: Modelling, Planning and Control*. Springer (2009)
53. Srebro, N., Rennie, J., Jaakkola, T.: Maximum-margin matrix factorization. In: Saul, L., Weiss, Y., Bottou, L. (eds.) *Advances in Neural Information Processing Systems*, vol. 17. MIT Press, New York (2004)
54. Trager, M., Kohn, K., Bruna, J.: Pure and spurious critical points: a geometric study of linear networks. In: *International Conference on Learning Representations* (2020)
55. Umenberger, J., Simchowitz, M., Perdomo, J., Zhang, K., Tedrake, R.: Globally convergent policy search for output estimation. In: Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., Oh, A. (eds.) *Advances in Neural Information Processing Systems*, vol. 35, pp. 22778–22790. Curran Associates Inc, New York (2022)
56. Vaskevicius, T., Kanade, V., Rebeschini, P.: Implicit regularization for optimal sparse recovery. In: Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*, vol. 32. Curran Associates Inc (2019)
57. Vlatakis-Gkaragkounis, E.-V., Flokas, L., Piliouras, G.: Efficiently avoiding saddle points with zero order methods: no gradients required. In: Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc,

- F., Fox, E., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*, vol. 32. Curran Associates Inc (2019)
58. Woods, R.: The cochlioid. *Am. Math. Mon.* **31**(5), 222–227 (1924)
59. Zhang, F.: *The Schur Complement and Its Applications*. Springer (2005)

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