THE CONDITION NUMBER OF RIEMANNIAN APPROXIMATION PROBLEMS*

PAUL BREIDING[†] AND NICK VANNIEUWENHOVEN[‡]

Abstract. We consider the local sensitivity of least-squares formulations of inverse problems. The sets of inputs and outputs of these problems are assumed to have the structures of Riemannian manifolds. The problems we consider include the approximation problem of finding the nearest point on a Riemannian embedded submanifold from a given point in the ambient space. We characterize the first-order sensitivity, i.e., condition number, of local minimizers and critical points to arbitrary perturbations of the input of the least-squares problem. This condition number involves the Weingarten map of the input manifold, which measures the amount by which the input manifold curves in its ambient space. We validate our main results through experiments with the n -camera triangulation problem in computer vision.

Key words. Riemannian least-squares problem, sensitivity, condition number, local minimizers, Weingarten map, second fundamental form

AMS subject classifications. 90C31, 53A55, 65H10, 65J05, 65D18, 65D19

1. Introduction. A prototypical instance of the Riemannian optimization problems we consider here arises from the following inverse problem (IP): Given an input $x \in \mathcal{I}$, compute $y \in \mathcal{O}$ such that $F(y) = x$. Herein, we are given a *system* modeled as a smooth map F between a Riemannian manifold $\mathcal{O}$, the *output manifold*, and a Riemannian embedded submanifold $\mathcal{I} \subset \mathbb{R}^n$, the *input manifold*. In practice, however, we are often given a point $a \in \mathbb{R}^n$ in the *ambient space* rather than $x \in \mathcal{I}$, for example because a was obtained from noisy or inaccurate physical measurements. This leads to the Riemannian least squares optimization problem associated to the foregoing IP:

$$(\text{PIP}) \quad \operatorname{argmin}_{y \in \mathcal{O}} \frac{1}{2} \|F(y) - a\|^2.$$

In [16] we called this type of problem a *parameter identification problem* (PIP). Application areas where PIPs originate are data completion with low-rank matrix and tensor decompositions [15, 24, 39, 45, 62, 66], geometric modeling [19, 49, 63], computer vision [33, 38, 51], and phase retrieval problems in signal processing [9, 12].

The manifold $\mathcal{O}$ is called the output manifold because its elements are the outputs of the optimization problem. The outputs are the quantities of interest; for instance, they could be the unknown system parameters that one wants to identify based on a given observation a and the forward description of the system F . In order to distinguish between the inputs $x \in \mathcal{I}$ of the IP and the inputs $a \in \mathbb{R}^n$ of the optimization formulation (PIP), we refer to the latter as *ambient inputs*. In this paper, we study the sensitivity of the output of the optimization problem (PIP) (and an implicit generalization thereof) when the ambient input $a \in \mathbb{R}^n$ is perturbed by considering its *condition number* κ at global and local minimizers and critical points.

1.1. The condition number. Condition numbers are one of the cornerstones in numerical analysis [7, 18, 27, 37, 41, 71, 72]. Its definition was introduced for Riemannian manifolds already in the early days of numerical analysis. In 1966, Rice [60] presented the definition in the case where input and output space are general nonlinear *metric spaces*; the special case of Riemannian manifolds is [60, Theorem 3]. Let $\mathcal{I}$ and $\mathcal{O}$ be

*Submitted to the editors.

Funding: P. Breiding has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 787840). N. Vannieuwenhoven was supported by a Postdoctoral Fellowship of the Research Foundation—Flanders (FWO) with project 12E8119N.

[†]Institute of Mathematics, TU Berlin, Berlin, Germany. (breiding@math.tu-berlin.de)

[‡]KU Leuven, Department of Computer Science, Leuven, Belgium. (nick.vannieuwenhoven@kuleuven.be)

Riemannian manifolds and assume that the map $f : \mathcal{I} \rightarrow \mathcal{O}$ models a computational problem. Then, Rice’s definition of the condition number of f at a point $x \in \mathcal{I}$ is

$$(1.1) \quad \kappa[f](x) := \lim_{\epsilon \rightarrow 0} \sup_{\substack{x' \in \mathcal{I}, \\ \text{dist}_{\mathcal{I}}(x, x') \leq \epsilon}} \frac{\text{dist}_{\mathcal{O}}(f(x), f(x'))}{\text{dist}_{\mathcal{I}}(x, x')},$$

where $\text{dist}_{\mathcal{I}}(\cdot, \cdot)$ denotes the distance given by the Riemannian metric on $\mathcal{I}$ and likewise for $\mathcal{O}$. It follows immediately that the condition number yields an asymptotically sharp upper bound on the perturbation of the output of a computational problem when subjected to a perturbation of the input x . Indeed, we have

$$\text{dist}_{\mathcal{O}}(f(x), f(x')) \leq \kappa[f](x) \cdot \text{dist}_{\mathcal{I}}(x, x') + o(\text{dist}_{\mathcal{I}}(x, x')).$$

The number $\kappa[f](x)$ is an *absolute* condition number. On the other hand, if $\mathcal{I}$ is a submanifold of a vector space equipped with the norm $\|\cdot\|_{\mathcal{I}}$ and analogously for $\mathcal{O}$ and $\|\cdot\|_{\mathcal{O}}$, then the *relative* condition number is $\kappa_{\text{rel}}[f](x) := \kappa[f](x) \frac{\|x\|_{\mathcal{I}}}{\|f(x)\|_{\mathcal{O}}}$. We focus on computing the absolute condition number, as the relative condition number can be inferred from the absolute one.

1.2. Condition number of approximation problems. Our starting point is Rice’s classic definition (1.1), which is often easy to compute for differentiable maps, rather than Demmel and Renegar’s inverse distance to ill-posedness, which is generally more difficult to compute. Nevertheless, Rice’s definition cannot be applied unreservedly to measure how sensitive a solution of (PIP) is to perturbations of the ambient input $a \in \mathbb{R}^n$. At first glance, it seems that we could apply (1.1) to

$$\phi : \mathbb{R}^n \rightarrow \mathcal{O}, \quad a \mapsto \underset{y \in \mathcal{O}}{\text{argmin}} \|F(y) - a\|^2.$$

However, this is an ill-posed expression, as ϕ is not a map in general! Indeed, there can be several global minimizers for an ambient input $a \in \mathbb{R}^n$ and their number can depend on a . Even if we could restrict the domain so that ϕ becomes a map, the foregoing formulation has two additional limitations. First, when (PIP) is a nonconvex Riemannian optimization problem, it is unrealistic to expect that we can solve it globally. Instead we generally find local minimizers whose sensitivity to perturbations of the input is also of interest. For our theory to be practical, we need to consider the condition number of every local minimizer separately. Exceeding this goal, we will present the condition number of each critical point in this paper. Second, requiring that the systems in (PIP) can be modeled by a smooth map $F : \mathcal{O} \rightarrow \mathcal{I}$ excludes some interesting computational problems. For example, computing the eigenvalues of a symmetric matrix cannot be modeled as a map from outputs to inputs: there are infinitely many symmetric matrices with a prescribed set of eigenvalues.

The solution to foregoing complications consists of allowing *multivalued* or *set-valued* maps $g : \mathcal{I} \rightrightarrows \mathcal{O}$ in which one input $x \in \mathcal{I}$ can produce a set of outputs $Y \subset \mathcal{O}$ [5, 29]. Such maps are defined implicitly by the graph

$$G = \{(x, y) \in \mathcal{I} \times \mathcal{O} \mid y \in g(x)\}.$$

In numerical analysis, Wilkinson [71] and Woźniakowski [73] realized that the condition number of *single-valued localizations* or local maps of G could still be studied using the standard techniques. Blum, Cucker, Shub, and Smale [10] developed a geometric framework for studying the condition number of single-valued localizations when the graph G is an algebraic subvariety. More recently, Bürgisser and Cucker [18] considered condition numbers when the graph G is a smooth manifold. Because these graphs are topologically severely restricted, the condition number at an input–output pair $(x, y) \in G$ can be computed efficiently using only linear algebra.

The heart of this paper consists of characterizing the condition number of critical points of PIPs (and an implicit generalization thereof) by employing the geometric approach to conditioning from [10, 18]. Therefore, we summarize its main points next.

1.3. Implicit smooth computational problems. Following [10, 18], we model an inverse problem with input manifold $\mathcal{I}$ and output manifold $\mathcal{O}$ implicitly by a *solution manifold*:

$$\mathcal{S} \subset \mathcal{I} \times \mathcal{O}.$$

Given an input $x \in \mathcal{I}$, the problem is finding a corresponding output $y \in \mathcal{O}$ such that $(x, y) \in \mathcal{S}$. As $\mathcal{S}$ models a computational problem, we can consider its condition number. Now it is defined at $(x, y) \in \mathcal{S}$ and not only at the input x , because the sensitivity to perturbations in the output depends on *which* output is considered.

The condition number of (the computational problem modeled by) $\mathcal{S}$ is derived in [10, 18]. Let $\pi_{\mathcal{I}} : \mathcal{S} \rightarrow \mathcal{I}$ and $\pi_{\mathcal{O}} : \mathcal{S} \rightarrow \mathcal{O}$ be the projections on the input and output manifolds, respectively. Recall that [10, 18] assume that $\dim \mathcal{S} = \dim \mathcal{I}$, because otherwise almost all inputs x have infinitely many outputs ($\dim \mathcal{S} > \dim \mathcal{I}$) or no outputs ($\dim \mathcal{S} < \dim \mathcal{I}$). Furthermore, we will assume that $\pi_{\mathcal{I}}$ is surjective, which means that every input $x \in \mathcal{I}$ has at least one output. When the derivative¹ $D_{(x,y)}\pi_{\mathcal{I}}$ is invertible, the inverse function theorem [48, Theorem 4.5] implies that $\pi_{\mathcal{I}}$ is locally invertible. Hence, there is a local smooth *solution map* $\pi_{\mathcal{O}} \circ \pi_{\mathcal{I}}^{-1}$, which locally around (x, y) makes the computational problem explicit. This local map has condition number $\kappa[\pi_{\mathcal{O}} \circ \pi_{\mathcal{I}}^{-1}](x)$ from (1.1). Moreover, because the map is differentiable, Rice gives an expression for (1.1) in terms of the spectral norm of the differential in [60, Theorem 4]. In summary, [10, 18] conclude that the condition number of $\mathcal{S}$ at $(x, y) \in \mathcal{S}$ is

$$\kappa[\mathcal{S}](x, y) = \begin{cases} \kappa[\pi_{\mathcal{O}} \circ \pi_{\mathcal{I}}^{-1}](x) & \text{if } D_{(x,y)}\pi_{\mathcal{I}} \text{ is invertible,} \\ \infty & \text{otherwise,} \end{cases}$$

$$(1.2) \quad = \|(D_{(x,y)}\pi_{\mathcal{O}})(D_{(x,y)}\pi_{\mathcal{I}})^{-1}\|_{\mathcal{I} \rightarrow \mathcal{O}},$$

where the latter *spectral norm* is defined below in (2.2). Note that if $\mathcal{S}$ is the graph of a smooth map $f : \mathcal{I} \rightarrow \mathcal{O}$, then $\kappa[\mathcal{S}](x, f(x)) = \kappa[f](x)$.

REMARK 1.1. By definition, (1.2) evaluates to ∞ if $D_{(x,y)}\pi_{\mathcal{I}}$ is not invertible.

The points in $\mathcal{W} := \{(x, y) \in \mathcal{S} \mid \kappa[\mathcal{S}](x, y) < \infty\}$ are called *well-posed tuples*. Points in $\mathcal{S} \setminus \mathcal{W}$ are called *ill-posed tuples*. An infinitesimal perturbation in the input x of an ill-posed tuple (x, y) can cause an arbitrary change in the output y . It is therefore natural to assume that $\mathcal{W} \neq \emptyset$, because otherwise the computational problem is always ill-posed. If there exists $(x, y) \in \mathcal{W}$, it follows from the inverse function theorem that $\mathcal{W}$ is an open submanifold of $\mathcal{S}$, so $\dim \mathcal{S} = \dim \mathcal{W}$. We also assume that $\mathcal{W}$ is dense. Otherwise, since $\mathcal{S}$ is Hausdorff, around any data point outside of $\overline{\mathcal{W}}$ there would be an *open* neighborhood on which the computational problem is ill-posed; we believe these input-output pairs should be removed from the definition of the problem.

In summary, we make the following assumption throughout this paper.

ASSUMPTION 1. The m -dimensional input manifold $\mathcal{I}$ is an embedded smooth submanifold of $\mathbb{R}^n$ equipped with the Euclidean inner product inherited from $\mathbb{R}^n$ as Riemannian metric. The output manifold $\mathcal{O}$ is a smooth Riemannian manifold. The solution manifold $\mathcal{S} \subset \mathcal{I} \times \mathcal{O}$ is a smoothly embedded m -dimensional submanifold. The projection $\pi_{\mathcal{I}} : \mathcal{S} \rightarrow \mathcal{I}$ is surjective. The well-posed tuples $\mathcal{W}$ form an open dense embedded submanifold of $\mathcal{S}$.

This assumption is typically not too stringent. It is satisfied for example if $\mathcal{S}$ is the graph of a smooth map $f : \mathcal{I} \rightarrow \mathcal{O}$ or $F : \mathcal{O} \rightarrow \mathcal{I}$ (as for IPs) by [48, Proposition 5.7]. Another example is when $\mathcal{S}$ is the smooth locus of an algebraic subvariety of $\mathcal{I} \times \mathcal{O}$, as most problems in [18]. Moreover, the condition of an input-output pair (x, y) is a local property. Since every immersed submanifold is locally embedded [48, Proposition 5.22], taking a restriction of the computational problem $\mathcal{S}$ to a neighborhood of (x, y) , in its topology as immersed submanifold, yields an embedded submanifold. The assumption and proposed theory always apply in this local sense.

¹See Section 2 for a definition.

1.4. Related work. The condition number (1.1) was computed for several computational problems. Already in the 1960s, the condition numbers of solving linear systems, eigenproblems, linear least-squares problems, and polynomial rootfinding, among others were determined [18, 27, 37, 71, 72]. An astute observation by Demmel [25, 26] in 1987 paved the way for the introduction of condition numbers in optimization. He realized that many condition numbers from linear algebra can be interpreted as an inverse distance to the nearest ill-posed problem. This characterization generalizes to contexts where an explicit map as required in (1.1) is not available. Renegar [58, 59] exploited this observation to determine the condition number of linear programming instances and he connected it to the complexity of interior-point methods for solving them. This ignited further research into the sensitivity of optimization problems, leading to (Renegar’s) condition numbers for linear programming [20, 58, 59], conic optimization [8, 21, 32, 54, 68], and convex optimization [4, 22, 35] among others. In many cases, the computational complexity of methods for solving these problems can be linked to the condition number; see [23, 31, 36, 58, 59, 67]. In our earlier work [16], we connected the complexity of Riemannian Gauss–Newton methods for solving (PIP) to the condition number of the exact inverse problem $F(y) = x$. We anticipate that the analysis from [16] could be refined to reveal the dependence on the condition number of the optimization problem (PIP).

Beside Renegar’s condition number, other measures of sensitivity are studied in optimization. The sensitivity of a set-valued map $g : \mathcal{I} \rightrightarrows \mathcal{O}$ at (x, y) with $y \in g(x)$ is often measured by the *Lipschitz modulus* [29, Chapter 3E]. This paper only studies the sensitivity of a single-valued localization $\hat{g}$ of the set-valued map g , in which case the Lipschitz modulus K of g at (x, y) is a Lipschitz constant of $\hat{g}$ [29, Proposition 3E.2]. Consequently, Rice’s condition number $\kappa[\hat{g}](x)$, which we characterize in this paper for Riemannian approximation problems, is always upper bounded by K ; see, e.g., [44, Proposition 6.3.10].

1.5. Main contributions. We characterize the condition number of computing global minimizers, local minimizers, and critical points of the following implicit generalization of (PIP):

$$(IPIP) \quad \pi_{\mathcal{O}} \circ \operatorname{argmin}_{(x,y) \in \mathcal{S}} \frac{1}{2} \|a - \pi_{\mathcal{I}}(x, y)\|^2,$$

where $a \in \mathbb{R}^n$ is a given ambient input and the norm is the Euclidean norm. We call this implicit formulation of (PIP) a Riemannian implicit PIP.

Our main contributions are Theorem 4.3, 5.2 and 6.2. The condition number of critical points of the Riemannian approximation problem $\min_{x \in \mathcal{O}} \|a - x\|^2$, is treated in Theorem 4.3. Theorem 5.2 contains the condition number of (IPIP) at global minimizers when a is sufficiently close to the input manifold. Finally, the most general case is Theorem 6.2: it gives the condition number of (IPIP) at all critical points including local and global minimizers.

Aforementioned theorems will show that the way in which the input manifold $\mathcal{I}$ curves in its ambient space $\mathbb{R}^n$, as measured by classic differential invariants, enters the picture. To the best of our knowledge, the role of curvature in the theory of condition of smooth computational problems involving approximation was not previously identified. We see this as one of the main contributions of this article.

We apply our theory in Section 9 to triangulation problems in computer vision and in [17] to low-rank matrix recovery in compressed sensing.

1.6. Approximation versus idealized problems. The results of this paper raise a pertinent practical question: If one has a computational problem modeled by a solution manifold $\mathcal{S}$ that is solved via an optimization formulation as in (IPIP), which condition number is appropriate? The one of the idealized problem $\mathcal{S}$, where the inputs are constrained to $\mathcal{I}$? Or the more complicated condition number of the optimization problem, where we allow ambient inputs from $\mathcal{I}$ ’s ambient space $\mathbb{R}^n$? We argue that the choice is only an illusion.

The first approach should be taken if the problem is *defined* implicitly by a solution manifold $\mathcal{S} \subset \mathcal{I} \times \mathcal{O}$ and the ambient inputs are perturbations $a \in \mathbb{R}^n$ of inputs $x \in \mathcal{I}$. The idealized framework will still apply because of [Corollary 5.5](#). Consider for example the n -camera triangulation problem, discussed in [Section 9](#), where the problem $\mathcal{S}_{\text{MV}}$ is defined by computing the inverse of the camera projection process. If the measurements a were taken with a perfect pinhole camera for which the theoretical model is exact, then even when $a \notin \mathcal{I}_{\text{MV}}$ (e.g., errors caused by pixelization), as long as it is a sufficiently small perturbation of some $x \in \mathcal{I}_{\text{MV}}$, the first approach is suited.

On the other hand, if it is only *postulated* that the computational problem can be well approximated by the *idealized problem* $\mathcal{S}$, then the newly proposed framework for Riemannian approximation problems is the more appropriate choice. In this case, the “perturbation” to the input is a modeling error. An example of this is applying n -camera triangulation to images obtained by cameras with lenses, i.e., most real-world cameras. In this case, the computational problem is not $\mathcal{S}_{\text{MV}}$, but rather a more complicated problem that takes into account lens distortions. When using $\mathcal{S}_{\text{MV}}$ as a proxy, as is usual [\[33, 38, 51\]](#), the computational problem becomes an optimization problem. Curvature should be taken into account then, as developed here.

1.7. Outline. Before developing the main theory, we briefly recall the necessary concepts from differential geometry and fix the notation in the next section. A key role will be played by the Riemannian Hessian of the distance from an ambient input in $\mathbb{R}^n$ to the input manifold $\mathcal{I}$. Its relation to classic differential-geometric objects such as the Weingarten map is stated in [Section 3](#). The condition number of critical points of [\(IPIP\)](#) is first characterized in [Section 4](#) for the special case where the solution manifold represents the identity map. The general case is studied in [Sections 5](#) and [6](#); the former deriving the condition number of global minimizers, while the latter treats critical points. The expressions of the condition numbers in [Section 4](#) to [6](#) are obtained from implicit formulations of the Riemannian IPIPs. For clarity, in the case of local minimizers, we also express them as condition numbers of global minimizers of localized Riemannian optimization problems in [Section 7](#). [Section 8](#) explains a few standard techniques for computing the condition number in practice along with their computational complexity. Finally, an application of our theory to the triangulation problem in multiview geometry is presented in [Section 9](#).

2. Preliminaries. We briefly recall the main concepts from Riemannian geometry that we will use. The notation introduced here will be used throughout the paper. Proofs of the fundamental results appearing below can be found in [\[28, 47, 48, 52, 53, 55\]](#).

2.1. Manifolds. By *smooth m -dimensional manifold* $\mathcal{M}$ we mean a C^∞ topological manifold that is second-countable, Hausdorff, and locally Euclidean of dimension m . The tangent space $T_p\mathcal{M}$ is the m -dimensional linear subspace of differential operators at p . A differential operator at $p \in \mathcal{M}$ is a linear map v_p from the vector space $C^\infty(\mathcal{M})$ over $\mathbb{R}$ of smooth functions $f : \mathcal{M} \rightarrow \mathbb{R}$ to $\mathbb{R}$ that satisfies the product rule: $v_p(f \cdot g) = (v_p f) \cdot g(p) + f(p) \cdot (v_p g)$ for all $f, g \in C^\infty(\mathcal{M})$. If $\mathcal{M} \subset \mathbb{R}^n$ is *embedded*, the tangent space can be identified with the linear span of all tangent vectors $\frac{d}{dt}|_{t=0}\gamma(t)$ where $\gamma \subset \mathcal{M}$ is a smooth curve passing through p at 0.

A differentiable map $F : \mathcal{M} \rightarrow \mathcal{N}$ between manifolds $\mathcal{M}$ and $\mathcal{N}$ induces a linear map between tangent spaces $D_p F : T_p\mathcal{M} \rightarrow T_{F(p)}\mathcal{N}$, called the *derivative* of F at p . If $v_p \in T_p\mathcal{M}$, then $w_{F(p)} := (D_p F)(v_p)$ is the differential operator $w_{F(p)}(f) := v_p(f \circ F)$ for all $f \in C^\infty(\mathcal{N})$.

The identity map $\mathbf{1}_{\mathcal{M}} : \mathcal{M} \rightarrow \mathcal{M}$ will be denoted by $\mathbf{1}$, the space on which it acts being clear from the context.

2.2. Tubular neighborhoods. Let $\mathcal{M}^m \subset \mathbb{R}^n$ be an embedded submanifold. The disjoint union of tangent spaces to $\mathcal{M}^m$ is the *tangent bundle* of $\mathcal{M}$:

$$T\mathcal{M} := \coprod_{p \in \mathcal{M}} T_p\mathcal{M} = \{(p, v) \mid p \in \mathcal{M} \text{ and } v \in T_p\mathcal{M}\}.$$

The tangent bundle is a smooth manifold of dimension $2m$. The *normal bundle* is constructed similarly. The *normal space* of $\mathcal{M}$ in $\mathbb{R}^n$ at p is the orthogonal complement in $\mathbb{R}^n$ of the tangent space $T_p\mathcal{M}$, namely $N_p\mathcal{M} := (T_p\mathbb{R}^n)^\perp$. The normal bundle is a smooth embedded submanifold of $\mathbb{R}^n \times \mathbb{R}^n$ of dimension n . Formally, we write

$$N\mathcal{M} := \coprod_{p \in \mathcal{M}} N_p\mathcal{M} = \{(p, \eta) \mid p \in \mathcal{M} \text{ and } \eta \in N_p\mathcal{M}\}.$$

Let $\delta : \mathcal{M} \rightarrow \mathbb{R}_+$ be a positive, continuous function. Consider the following open neighborhood of the normal bundle:

$$(2.1) \quad \mathcal{V}_\delta = \{(p, \eta_p) \in N\mathcal{M} \mid \|\eta_p\| < \delta(p)\},$$

where the norm is the one induced from the Riemannian metric; in our setting it is the standard norm on $\mathbb{R}^n$. The map $E : N\mathcal{M} \rightarrow \mathbb{R}^n$, $(p, \eta) \mapsto p + \eta$ is smooth and there exists a δ such that the restriction $E|_{\mathcal{V}_\delta}$ becomes a diffeomorphism onto its image [48, Theorem 6.24]. Consequently, $\mathcal{T} = E(\mathcal{V}_\delta)$ is an n -dimensional, open, smooth, embedded submanifold of $\mathbb{R}^n$ that forms a neighborhood of $\mathcal{M}$. The submanifold $\mathcal{T}$ is called a *tubular neighborhood* of $\mathcal{M}$ and δ is called its *height* (function).

2.3. Vector fields. A *smooth vector field* is a smooth map $X : \mathcal{M} \rightarrow T\mathcal{M}$ from the manifold to its tangent bundle taking a point p to a tangent vector v_p . By the notation $X|_p$ we mean $X(p)$. A smooth vector field X on a properly embedded submanifold $\mathcal{M} \subset \mathcal{N}$ can be extended to a smooth vector field $\hat{X}$ on $\mathcal{N}$ such that it agrees on $\mathcal{M}$: $\hat{X}|_{\mathcal{M}} = X$.

A *smooth frame* of $\mathcal{M}^m$ is a tuple of m smooth vector fields $(X_1, \dots, X_m)$ that are linearly independent: $(X_1|_p, \dots, X_m|_p)$ is linearly independent for all $p \in \mathcal{M}$. If these tangent vectors are orthonormal for all p , then the frame is called orthonormal.

Let X be a smooth vector field on $\mathcal{M}$. An *integral curve* of X is a smooth curve $\gamma \subset \mathcal{M}$ such that there is an $\epsilon > 0$ such that $D_t\gamma = X|_{\gamma(t)}$ for all $t \in (-\epsilon, \epsilon)$. For every $p \in \mathcal{M}$ the vector field X generates an integral curve γ with starting point $p = \gamma(0)$. There is always a smooth curve *realizing* a tangent vector $v_p \in T_p\mathcal{M}$, i.e., a curve γ with $p = \gamma(0)$ and $v = D_0\gamma$ by [48, Lemma 8.6 and Proposition 9.2].

2.4. Riemannian manifolds. A *Riemannian manifold* is a smooth manifold $\mathcal{M}$ equipped with a *Riemannian metric* g : a positive definite symmetric bilinear form $g_p : T_p\mathcal{M} \times T_p\mathcal{M} \rightarrow \mathbb{R}$ that varies smoothly with p . The metric g_p induces a norm on $T_p\mathcal{M}$ denoted $\|t_p\|_{\mathcal{M}} := \sqrt{g_p(t_p, t_p)}$. The only Riemannian metric we will explicitly use in computations is the standard Riemannian metric of $\mathbb{R}^n$, i.e., the Euclidean inner product $g_p(x, y) = \langle x, y \rangle = x^T y$.

If $\mathcal{M}$ and $\mathcal{N}$ are Riemannian manifolds with induced norms $\|\cdot\|_{\mathcal{M}}$ and $\|\cdot\|_{\mathcal{N}}$, and if $F : \mathcal{M} \rightarrow \mathcal{N}$ is a differentiable map, then the *spectral norm* of $D_p F$ is

$$(2.2) \quad \|D_p F\|_{\mathcal{M} \rightarrow \mathcal{N}} := \sup_{t_p \in T_p\mathcal{M} \setminus \{0\}} \frac{\|(D_p F)(t_p)\|_{\mathcal{N}}}{\|t_p\|_{\mathcal{M}}}.$$

If the manifolds are clear from the context we can drop the subscript and write $\|D_p F\|$.

2.5. Riemannian gradient and Hessian. Let $\mathcal{M}$ be a Riemannian manifold with metric g , and let $f : \mathcal{M} \rightarrow \mathbb{R}$ be a smooth function. The *Riemannian gradient* of f at $p \in \mathcal{M}$ is the tangent vector $\nabla f|_p \in T_p\mathcal{M}$ that is dual to the derivative $D_p f : T_p\mathcal{M} \rightarrow \mathbb{R}$ under g , i.e., $g_p(\nabla f|_p, t) = (D_p f)(t)$ for all $t \in T_p\mathcal{M}$. It is known that ∇f is a vector field, hence explaining the notation $\nabla f|_p$ for its value at $p \in \mathcal{M}$.

When $\mathcal{M} \subset \mathbb{R}^n$ is a Riemannian embedded submanifold with the inherited Euclidean metric, the *Riemannian Hessian* [2, Definition 5.5.1] can be defined without explicit reference to connections, as in [13, Chapter 5]:

$$\text{Hess}_p(f) : T_p\mathcal{M} \rightarrow T_p\mathcal{M}, \quad \eta_p \mapsto P_{T_p\mathcal{M}}((D_p \nabla f)(\eta_p)),$$

where ∇f now has to be viewed as a vector field on $\mathbb{R}^n$ and $P_{T_p\mathcal{M}}$ projects orthogonally onto the tangent space $T_p\mathcal{M}$. In other words, the Hessian takes η_p to $P_{T_p\mathcal{M}}\left(\frac{d}{dt}\big|_{t=0}\nabla f|_{\gamma(t)}\right)$ where $\gamma(t) \subset \mathcal{M}$ is a curve realizing the tangent vector η_p .

The Hessian $\text{Hess}_p(f)$ is a symmetric bilinear form on the tangent space $T_p\mathcal{M}$ whose (real) eigenvalues contain essential information about the objective function f at p . If the Hessian is positive definite, negative definite, or indefinite at a *critical point* p of f (i.e., $\nabla f|_p = 0$ or, equivalently, $D_p f = 0$), then p is respectively a local minimum, local maximum, or saddle point of f .

3. Measuring curvature with the Weingarten map. A central role in this paper is played by the Riemannian Hessian of the distance function

$$(3.1) \quad d : \mathcal{I} \rightarrow \mathbb{R}, \quad x \mapsto \frac{1}{2}\|x - a\|^2,$$

which measures the Euclidean distance between a fixed point $a \in \mathbb{R}^n$ and points on the embedded submanifold $\mathcal{I} \subset \mathbb{R}^n$. It is the inverse of this Riemannian Hessian matrix that will appear as an additional factor to be multiplied with $D_{(x,y)}\pi_{\mathcal{O}}(D_{(x,y)}\pi_{\mathcal{I}})^{-1}$ from (1.2) in the condition numbers of the problems studied in Section 4 to 6.

The Riemannian Hessian of d contains classic differential-geometric information about the way $\mathcal{I}$ curves inside of $\mathbb{R}^n$ [55]. Indeed, from [3, Equation (10)] we conclude that it can be expressed at $x \in \mathcal{I}$ as

$$\text{Hess}_x(d) = \mathbf{1}_{T_x\mathcal{I}} + II_x(P_{N_x\mathcal{M}}(x - a)),$$

where $P_{N_x\mathcal{M}}$ is the projection onto the normal space, and $II_x : N_x\mathcal{M} \rightarrow S^2(T_x\mathcal{M})$ is a classical differential invariant called the *second fundamental form* [28, 47, 52, 53, 55], which takes a normal vector in $N_x\mathcal{M}$ and sends it to a symmetric bilinear form on the tangent space $T_x\mathcal{M}$; note that $S^2(A)$ is the space of all bilinear forms on the linear space A . If we let $\eta_x := a - x$, then at a critical point of d we have $\eta_x \in N_x\mathcal{I}$, so that the Hessian of interest satisfies

$$(H) \quad H_\eta := \text{Hess}_x(d) = \mathbf{1} - S_\eta, \quad \text{where } S_\eta := II_x(\eta_x).$$

Note that we dropped the base point $x \in \mathcal{I}$ from the notation in H_η and S_η . In the above, S_η is interpreted as a symmetric endomorphism on $T_x\mathcal{I}$ and it is classically called the *shape operator* or *Weingarten map* [28, 47, 52, 53, 55].

The Weingarten map measures how much the embedding into the ambient space $\mathbb{R}^n$ bends the manifold $\mathcal{I}^m$ in the following concrete way [47, Chapter 8]. For $\eta \neq 0$, let $w := \frac{\eta}{\|\eta\|}$ and let $c_1, \dots, c_m$ be the eigenvalues of $S_w = \frac{1}{\|\eta\|}S_\eta$. As S_w is self-adjoint, all c_i 's are all real. The eigenvalues of S_η are then $c_1\|\eta\|, \dots, c_m\|\eta\|$. The c_i 's are called the *principal curvatures* of $\mathcal{I}$ at x in direction of η . They have the following classic geometric meaning [61, Chapter 1]: If $u_i \in T_x\mathcal{I}$ is a unit length eigenvector of the eigenvalue c_i , then locally around x and in the direction of u_i the manifold $\mathcal{I}$ contains an infinitesimal arc of a circle with center $x + c_i^{-1}w = x + (c_i\|\eta\|)^{-1}\eta$ through x . This circle is called an *osculating circle*. The radii $r_i := |c_i|^{-1}$ of those circles are called the *critical radii* of $\mathcal{I}$ at x in direction η . See also Figure 4.1.

For several submanifolds $\mathcal{I} \subset \mathbb{R}^n$ relevant in optimization applications, the Weingarten map S_η was computed. Absil, Mahoney, and Trumpf [3] provide expressions for the sphere, Stiefel manifold, orthogonal group, Grassmann manifold, and low-rank matrix manifold. Feppon and Lermusiaux [34] further extended these results by computing the Weingarten maps and explicit expressions for the principal curvatures of the isospectral manifold and “bi-Grassmann” manifold, among others. Heidel and Schultz [39] computed the Weingarten map of the manifold of higher-order tensors of fixed multilinear rank [65]. The Weingarten map of the manifold of tensors of fixed tensor trains rank [43] can be obtained from the work of Psenka and Boumal [57].

4. The condition number of critical point problems. Prior to treating (IPIP) in general, we focus on an important and simpler special case. This case arises

when the solution manifold $\mathcal{S} \subset \mathcal{I} \times \mathcal{O}$ is the graph of the identity map $\mathbf{1} : \mathcal{I} \rightarrow \mathcal{I}$, so $\mathcal{O} = \mathcal{I}$. Given an ambient input $a \in \mathbb{R}^n$, the problem (IPIP) now comprises finding an output $x \in \mathcal{I}$ so that the distance function d from (3.1) is minimized. This can be formulated as the standard Riemannian least-squares problem

$$(4.1) \quad \min_{x \in \mathcal{I}} \frac{1}{2} \|x - a\|^2.$$

Computing *global* minimizers of this optimization problem is usually too ambitious a task when $\mathcal{I}$ is a complicated nonlinear manifold. Most Riemannian optimization methods will converge to local minimizers by seeking to satisfy the first-order necessary optimality conditions of (4.1), namely $a - x \in N_x \mathcal{I}$. We say that such methods attempt to solve the *critical point problem* (CPP) whose implicit formulation is

$$(CPP) \quad \mathcal{S}_{CPP} := \{(a, x) \in \mathbb{R}^n \times \mathcal{I} \mid a - x \in N_x \mathcal{I}\}.$$

We call $\mathcal{S}_{CPP}$ the *critical point locus*. It is an exercise to show that $\mathcal{S}_{CPP}$ is linearly diffeomorphic to the normal bundle $N\mathcal{I}$ so that we have the following result.

LEMMA 4.1. $\mathcal{S}_{CPP}$ is an embedded submanifold of $\mathbb{R}^n \times \mathcal{I}$ of dimension n .

Consequently, the CPP falls into the realm of the theory of condition from (1.2). We denote the coordinate projections of $\mathcal{S}_{CPP}$ by $\Pi_{\mathbb{R}^n} : \mathcal{S}_{CPP} \rightarrow \mathbb{R}^n$ and $\Pi_{\mathcal{I}} : \mathcal{S}_{CPP} \rightarrow \mathcal{I}$ for distinguishing them from the projections associated to $\mathcal{S}$. The corresponding submanifold of well-posed tuples is

$$\mathcal{W}_{CPP} := \{(a, x) \in \mathcal{S}_{CPP} \mid D_{(a,x)} \Pi_{\mathbb{R}^n} \text{ is invertible}\}.$$

The fact that it is a manifold is the following result, proved in Appendix A.1.

LEMMA 4.2. $\mathcal{W}_{CPP}$ is an open dense embedded submanifold of $\mathcal{S}_{CPP}$.

We show in Appendix A.1 that the condition number of this problem can be characterized succinctly as follows.²

THEOREM 4.3 (Critical points). Let $(a, x) \in \mathcal{S}_{CPP}$ and $\eta := a - x$. Then,

$$\kappa[\mathcal{S}_{CPP}](a, x) = \|H_\eta^{-1}\|_{\mathcal{I} \rightarrow \mathcal{I}},$$

where H_η is given by (H). Moreover, if $(a, x) \in \mathcal{S}_{CPP} \setminus \mathcal{W}_{CPP}$, then H_η is not invertible.

COROLLARY 4.4. $\kappa[\mathcal{S}_{CPP}] : \mathcal{S}_{CPP} \rightarrow \mathbb{R} \cup \{+\infty\}$ is continuous and $\kappa[\mathcal{S}_{CPP}](x, y)$ is finite if and only if $(x, y) \in \mathcal{W}_{CPP}$.

Curvature has entered the picture in Theorem 4.3 in the form of the Riemannian Hessian H_η from (H). It is an *extrinsic* property of $\mathcal{I}$, however, so that the condition number depends on the Riemannian embedding of $\mathcal{I}$ into $\mathbb{R}^n$. Intuitively this is coherent, as different embeddings give rise to different computational problems each with their own sensitivity to perturbations in the ambient space.

Let us investigate the meaning of the foregoing theorem a little further. Recall, from Section 3, the definition of the principal curvatures c_i as the eigenvalues of the Weingarten map S_w where $w = \frac{\eta}{\|\eta\|}$ and $\eta = a - x$. Then, Theorem 4.3 implies that

$$(4.2) \quad \kappa[\mathcal{S}_{CPP}](a, x) = \max_{1 \leq i \leq m} \frac{1}{|1 - c_i \|\eta\||},$$

so its condition is completely determined by the amount by which $\mathcal{I}$ curves inside of $\mathbb{R}^n$ and the distance of the ambient input a to the critical point $x \in \mathcal{I}$. We see that $\kappa[\mathcal{S}_{CPP}](a, x) = \infty$ if and only if $\mathcal{I}$ contains an infinitesimal arc of the circle with center $x + \eta = a$ (so that there is an i with $c_i \|\eta\| = 1$). Consequently, if $\|\eta\|$ lies close to any of the critical radii $r_i = |c_i|^{-1}$ then the CPP is ill-conditioned at (a, x) .

²We take the convention that whenever an expression $\|\cdot\|_{\mathcal{I} \rightarrow \mathcal{O}}$ is ill-posed it evaluates to ∞ .

FIGURE 4.1. The graphs shows ambient inputs a for the CPP of the parabola. The osculating circle with critical radius lies above the parabola, and its center (the gray point) is a focal point of the parabola, i.e., it is a point in $\Sigma_{\text{CPP}} = \mathcal{S}_{\text{CPP}} \setminus \mathcal{W}_{\text{CPP}}$. In graph (a), the ambient input a also lies above the parabola, so $\eta = a - x$ points towards the focal point. This means that the principal curvature at x in direction of η is $c_1 > 0$. Moreover, $c_1 \|\eta\| < 1$, and thus, by (4.2), we have $\kappa[\mathcal{S}_{\text{CPP}}](a, x) > 1$. In other words, the curvature of the parabola amplifies the perturbation $\|\Delta a\|$. In graph (b), the ambient input a lies below the parabola, so $\eta = a - x$ points away from the focal point. This means that the principal curvature at x in direction of η is negative, so that by (4.2) we have $\kappa[\mathcal{S}_{\text{CPP}}](a, x) < 1$. The curvature of the parabola shrinks the perturbation $\|\Delta a\|$. In graph (c), the ambient input a lies above the parabola, so that the corresponding critical curvature is positive. However, now $c_1 \|\eta\| > 1$, and so by (4.2) we have $\kappa[\mathcal{S}_{\text{CPP}}](a, x) < 1$. Again, the curvature of the parabola shrinks the perturbation $\|\Delta a\|$.

The formula from (4.2) makes it apparent that for some inputs we may have that $\kappa[\mathcal{S}_{\text{CPP}}](a, x) < 1$. In other words, for some ambient inputs $a \in \mathbb{R}^n$ the infinitesimal error $\|\Delta a\|$ can shrink! Figure 4.1 gives a geometric explanation of when the error is amplified (panel (a)) and when it is shrunk (panels (b) and (c)).

We also have an interpretation of $\kappa[\mathcal{S}_{\text{CPP}}](a, x)$ as a normalized inverse distance to ill-posedness: for $(a, x) \in \mathcal{S}_{\text{CPP}}$, let $\eta = a - x$ be the normal vector that connects x to a . The locus of ill-posed inputs for the CPP is $\Sigma_{\text{CPP}} = \mathcal{S}_{\text{CPP}} \setminus \mathcal{W}_{\text{CPP}}$. The projection $\Pi_{\mathbb{R}^n}(\Sigma_{\text{CPP}})$ is well studied in the literature. Thom [64] calls it the *target envelope* of $\mathcal{S}_{\text{CPP}}$, while Porteous [56] calls it the *focal set* of $\mathcal{I}$, and for curves in the plane it is called the *evolute*; see Chapter 10 of [52]. The points in the intersection of $\{x\} \times (x + \mathbb{R}\eta)$ with Σ_{CPP} are given by the multiples of the normal vector η whose length equals one of the critical radii $r_1, \dots, r_m$; this is also visible in Figure 9.3. Comparing with (4.2) we find

$$\frac{1}{\kappa[\mathcal{S}_{\text{CPP}}](a, x)} = \min_{\eta_0 \in \mathbb{R}\eta: (x, a + \eta_0) \in \Sigma_{\text{CPP}}} \frac{\|\eta_0\|}{\|\eta_0\| - \|\eta\|}, \text{ where } \eta = a - x.$$

This is an interpretation of Rice's condition number in the spirit of Demmel [25, 26] and Renegar [58, 59].

REMARK 4.5. If $\mathcal{I}$ is a smooth algebraic variety in $\mathbb{R}^n$, then solving $\mathcal{S}_{\text{CPP}}$ reduces to solving a system of polynomial equations. For almost all inputs $a \in \mathbb{R}^n$ this system has a constant number of solutions over the complex numbers, called the Euclidean distance degree [30]. Theorem 4.3 thus characterizes the condition number of real solutions of this special system of polynomial equations.

5. The condition number of approximation problems. Next, we analyse the condition number of global minimizers of (IPIP). This setting will be crucial for a finer interpretation of Theorem 4.3 and 6.2 in Section 7.

By suitably choosing the height function δ [42, Chapter 4, Section 5], we can assume for each ambient input a in a tubular neighborhood $\mathcal{T} \subset \mathbb{R}^n$ of $\mathcal{I}$ that there is a unique point x on $\mathcal{I}$ that minimizes the distance from $\mathcal{I}$ to a . In this case, the computational problem that consists of finding global minimizers of (IPIP) can be modeled implicitly with the graph

$$(\text{AP}) \quad \mathcal{S}_{\text{AP}} := \{(a, y) \in \mathcal{T} \times \mathcal{O} \mid (P_{\mathcal{I}}(a), y) \in \mathcal{S}\},$$

where $P_{\mathcal{I}}: \mathcal{T} \rightarrow \mathcal{I}, a \mapsto \operatorname{argmin}_{x \in \mathcal{I}} \|x - a\|$ is the nonlinear projection that maps $a \in \mathcal{T}$ to the unique point on $\mathcal{I}$ that minimizes the distance to a . We refer to the problem

$\mathcal{S}_{\text{AP}}$ as the *approximation problem* (AP) to distinguish it from the problem of finding critical points of (IPIP), which is treated in Section 6.

If $\mathcal{S}_{\text{AP}}$ were a smooth submanifold of $\mathcal{T} \times \mathcal{O}$, its condition number would be given by (1.2). However, this is not the case in general. We can nevertheless consider the subset of well-posed tuples for (AP):

$$\mathcal{W}_{\text{AP}} := \{(a, y) \in \mathcal{T} \times \mathcal{O} \mid (P_{\mathcal{I}}(a), y) \in \mathcal{W}\}.$$

In Appendix A.2, we prove the following result under Assumption 1.

LEMMA 5.1. $\mathcal{W}_{\text{AP}}$ is a smooth embedded submanifold of $\mathcal{T} \times \mathcal{O}$ of dimension n . Moreover, $\mathcal{W}_{\text{AP}}$ is dense in $\mathcal{S}_{\text{AP}}$.

The condition number of (AP) is now given by (1.2) on $\mathcal{W}_{\text{AP}}$. For the remaining points, i.e., $\mathcal{S}_{\text{AP}} \setminus \mathcal{W}_{\text{AP}}$, the usual definition (1.1) still applies. From this, we obtained the following characterization, which is proved in Appendix A.2.

THEOREM 5.2 (Global minimizers). *Let the ambient input $a \in \mathcal{T} \subset \mathbb{R}^n$ be sufficiently close to $\mathcal{I}$, specifically lying in a tubular neighborhood of $\mathcal{I}$, so that it has a unique closest point $x \in \mathcal{I}$. Let $\eta = a - x$. Then,*

$$\kappa[\mathcal{S}_{\text{AP}}](a, y) = \|(\mathcal{D}_{(x,y)}\pi_{\mathcal{O}})(\mathcal{D}_{(x,y)}\pi_{\mathcal{I}})^{-1}H_{\eta}^{-1}\|_{\mathcal{I} \rightarrow \mathcal{O}},$$

where $\pi_{\mathcal{I}}$ and $\pi_{\mathcal{O}}$ are the projections onto the input and output manifolds respectively, and H_{η} is as in (H). Moreover, if $(a, y) \in \mathcal{S}_{\text{AP}} \setminus \mathcal{W}_{\text{AP}}$ then $\mathcal{D}_{(x,y)}\pi_{\mathcal{I}}$ is not invertible.

COROLLARY 5.3. $\kappa[\mathcal{S}_{\text{AP}}] : \mathcal{T} \rightarrow \mathbb{R} \cup \{+\infty\}$ is continuous and $\kappa[\mathcal{S}_{\text{AP}}](a, y)$ is finite if and only if $(a, y) \in \mathcal{W}_{\text{AP}}$.

REMARK 5.4. If $f : \mathcal{I} \rightarrow \mathcal{O}$ is a smooth map and $\mathcal{S}$ is the graph of f , then the foregoing specializes to

$$(\mathcal{D}_{(x,y)}\pi_{\mathcal{O}})(\mathcal{D}_{(x,y)}\pi_{\mathcal{I}})^{-1} = \mathcal{D}_x f, \quad \text{so} \quad \kappa[\mathcal{S}_{\text{AP}}](x, y) = \|(\mathcal{D}_x f)H_{\eta}^{-1}\|_{\mathcal{I} \rightarrow \mathcal{O}}$$

Similarly, if $F : \mathcal{O} \rightarrow \mathcal{I}$ is a smooth map with graph $\mathcal{S}$, then

$$(\mathcal{D}_{(x,y)}\pi_{\mathcal{O}})(\mathcal{D}_{(x,y)}\pi_{\mathcal{I}})^{-1} = (\mathcal{D}_y F)^{-1}, \quad \text{so} \quad \kappa[\mathcal{S}_{\text{AP}}](x, y) = \|(\mathcal{D}_y F)^{-1}H_{\eta}^{-1}\|_{\mathcal{I} \rightarrow \mathcal{O}}$$

by the inverse function theorem. These observations apply to Theorem 6.2 as well.

The contribution $(\mathcal{D}_{(x,y)}\pi_{\mathcal{O}})(\mathcal{D}_{(x,y)}\pi_{\mathcal{I}})^{-1}$ to the condition number originates from the solution manifold $\mathcal{S}$; it is intrinsic and does not depend on the embedding. On the other hand, the curvature factor H_{η}^{-1} is extrinsic in the sense that it depends on the embedding of $\mathcal{I}$ into $\mathbb{R}^n$, as in Theorem 4.3.

The curvature factor in Theorem 5.2 disappears when $S_{\eta} = 0$. This occurs when the ambient input a lies on the input manifold, so $\eta = 0$. In this case, $H_0^{-1} = \mathbf{1}$, so we have the following result.

COROLLARY 5.5. *Let $(x, y) \in \mathcal{S}$. Then, we have $\kappa[\mathcal{S}](x, y) = \kappa[\mathcal{S}_{\text{AP}}](x, y)$.*

In other words, if the ambient input, given by its coordinates in $\mathbb{R}^n$, lies exactly on $\mathcal{I}$, the condition number of (AP) equals the condition number of the computational problem $\mathcal{S}$, even though there can be many more directions of perturbation in $\mathbb{R}^n$ than in $\mathcal{I}$! By continuity, the effect of curvature on the sensitivity of the computational problem can essentially be ignored for small $\|\eta\|$, i.e., for points very close to $\mathcal{I}$. This is convenient in practice because computing the Weingarten map is often more complicated than obtaining the derivatives of the projection maps $\pi_{\mathcal{O}}$ and $\pi_{\mathcal{I}}$.

6. The condition number of general critical point problems. We now consider the sensitivity of critical points of (IPIP). As mentioned before, our motivation stems from solving (IPIP) when a closed-form solution is lacking. Applying Riemannian optimization methods to them, one can in general only guarantee to find points satisfying the first-order optimality conditions [2]. Consequently, in theory,

they solve a *generalized critical point problem*: Given an ambient input $a \in \mathbb{R}^n$, produce an output $y \in \mathcal{O}$ such that there is a corresponding $x \in \mathcal{I}$ that satisfies the first-order optimality conditions of (IPIP). We can implicitly formulate this computational problem as follows:

$$(\text{GCPP}) \quad \mathcal{S}_{\text{GCPP}} := \{(a, x, y) \in \mathbb{R}^n \times \mathcal{I} \times \mathcal{O} \mid (a, x) \in \mathcal{S}_{\text{CPP}} \text{ and } (x, y) \in \mathcal{S}\}.$$

The graph $\mathcal{S}_{\text{GCPP}}$ is defined using three factors, where the first is the input and the third is the output. The second factor $\mathcal{I}$ encodes how the output is obtained from the input. Triples are ill-posed if either $(x, y) \in \mathcal{S}$ is ill-posed for the original problem or $(a, x) \in \mathcal{S}_{\text{CPP}}$ is ill-posed for the (CPP). Therefore, the well-posed locus is

$$\mathcal{W}_{\text{GCPP}} := \{(a, x, y) \in \mathbb{R}^n \times \mathcal{I} \times \mathcal{O} \mid (a, x) \in \mathcal{W}_{\text{CPP}} \text{ and } (x, y) \in \mathcal{W}\}.$$

We prove in Appendix A.3 that the well-posed tuples form a manifold.

LEMMA 6.1. $\mathcal{W}_{\text{GCPP}}$ is an embedded submanifold of $\mathbb{R}^n \times \mathcal{W}$ of dimension n . Moreover, $\mathcal{W}_{\text{GCPP}}$ is dense in $\mathcal{S}_{\text{GCPP}}$.

Consequently, on $\mathcal{W}_{\text{GCPP}}$ we can use the definition of condition from (1.2). Our final main result generalizes Theorem 5.2; it is proved in Appendix A.3.

THEOREM 6.2 (Generalized critical points). Let $(a, x, y) \in \mathcal{S}_{\text{GCPP}}$ and $\eta = a - x$. Then, we have

$$\kappa[\mathcal{S}_{\text{GCPP}}](a, x, y) = \|(\mathcal{D}_{(x,y)}\pi_{\mathcal{O}})(\mathcal{D}_{(x,y)}\pi_{\mathcal{I}})^{-1}H_{\eta}^{-1}\|_{\mathcal{I} \rightarrow \mathcal{O}},$$

where $\pi_{\mathcal{I}} : \mathcal{S} \rightarrow \mathcal{I}$ and $\pi_{\mathcal{O}} : \mathcal{S} \rightarrow \mathcal{O}$ are coordinate projections, and H_{η} is given by (H). Moreover, if $(a, x, y) \in \mathcal{S}_{\text{GCPP}} \setminus \mathcal{W}_{\text{GCPP}}$, then either $\mathcal{D}_{(x,y)}\pi_{\mathcal{I}}$ or H_{η} is not invertible.

COROLLARY 6.3. $\kappa[\mathcal{S}_{\text{GCPP}}] : \mathcal{S}_{\text{GCPP}} \rightarrow \mathbb{R} \cup \{+\infty\}$ is continuous and the condition number $\kappa[\mathcal{S}_{\text{GCPP}}](a, x, y)$ is finite if and only if $(a, x, y) \in \mathcal{W}_{\text{GCPP}}$.

Since $\mathcal{I}$ inherits the Euclidean metric from $\mathbb{R}^n$, the following result is immediately obtained from Theorem 6.2.

COROLLARY 6.4. Let $(a, x, y) \in \mathcal{S}_{\text{GCPP}}$ and $\eta = a - x$. We have

$$\frac{\kappa[\mathcal{S}](x, y)}{\max_{1 \leq i \leq m} |1 - c_i \|\eta\|}| \leq \kappa[\mathcal{S}_{\text{GCPP}}](a, x, y) \leq \frac{\kappa[\mathcal{S}](x, y)}{\min_{1 \leq i \leq m} |1 - c_i \|\eta\|}},$$

where $c_1, \dots, c_m$ are the principal curvatures of $\mathcal{I}^m$ at x in direction $\eta = a - x$.

Proof. Since $\kappa[\mathcal{S}](x, y) = \|(\mathcal{D}_{(x,y)}\pi_{\mathcal{O}})(\mathcal{D}_{(x,y)}\pi_{\mathcal{I}})^{-1}\|_{\mathcal{I} \rightarrow \mathcal{O}}$, the submultiplicativity of the spectral norm both yields $\kappa[\mathcal{S}](x, y) \leq \kappa[\mathcal{S}_{\text{GCPP}}](a, x, y)\|H_{\eta}\|$ and also $\kappa[\mathcal{S}_{\text{GCPP}}](a, x, y) \leq \kappa[\mathcal{S}](x, y)\|H_{\eta}^{-1}\|$, concluding the proof. $\square$

7. Consequences for Riemannian optimization. In Section 4, we established a connection between the CPP and the Riemannian optimization problem (4.1): $\mathcal{S}_{\text{CPP}}$ is the manifold of *critical tuples* $(a, x) \in \mathbb{R}^n \times \mathcal{I}$, such that x is a critical point of the squared distance function from $\mathcal{I}$ to a . For the general implicit formulation in (GCPP), however, the connection to Riemannian optimization is more complicated. For clarity, we therefore also characterize the condition number of local minimizers of (IPIP) as the condition number of the *unique* global minimizer of a localized Riemannian optimization problem.

The solutions of the GCPP were defined as the generalization of the AP to critical points. However, we did not show that they too can be seen as the critical tuples of a distance function optimized over in a Riemannian optimization problem. This connection is established next and proved in Appendix A.4.

PROPOSITION 7.1. $(a, x, y) \in \mathcal{W}_{\text{GCPP}}$ if and only if $(a, (x, y))$ is a well-posed critical tuple of the function optimized over in the Riemannian optimization problem

$$\min_{(x,y) \in \mathcal{W}} \frac{1}{2} \|a - \pi_{\mathcal{I}}(x, y)\|^2.$$

Our main motivation for studying critical points instead of only the global minimizer as in (AP) stems from our desire to predict the sensitivity of outputs of Riemannian optimization methods for solving approximation problems like (4.1), (PIP), and (IPIP). When applied to a well-posed optimization problem, these methods in general only guarantee convergence to critical points [2]. However, in practice, convergence to critical points that are not local minimizers is extremely unlikely [2]. The reason is that they are *unstable outputs* of these algorithms: Applying a tiny perturbation to such critical points will cause the optimization method to escape their vicinity, and with high probability, converge to a local minimizer instead.

REMARK 7.2. *One should be careful to separate condition from numerical stability. The former is a property of a problem, the latter the property of an algorithm. We do not claim that computing critical points of GCPPs other than local minimizers are ill-conditioned problems. We only state that many Riemannian optimization methods are unstable algorithms for computing them via (IPIP). This is not a bad property, since the goal was to find minimizers, not critical points!*

For critical points of the GCPP that are local minimizers of the problem in Proposition 7.1 we can characterize their condition number as the condition number of the unique global minimizer of a localized problem. This result is proved in Appendix A.4.

THEOREM 7.3 (Condition of local minimizers). *Let the input manifold $\mathcal{I}$, the output manifold $\mathcal{O}$, and well-posed loci $\mathcal{W}$ and $\mathcal{W}_{\text{GCPP}}$ be as in (GCPP). Assume that $(a^*, x^*, y^*) \in \mathcal{W}_{\text{GCPP}}$ and x^* is a local minimizer of $d_{a^*} : \mathcal{I} \rightarrow \mathbb{R}, x \mapsto \frac{1}{2}\|a^* - x\|^2$. Then, there exist open neighborhoods $\mathcal{A}_{a^*} \subset \mathbb{R}^n$ around a^* and $\mathcal{N}_{(x^*, y^*)} \subset \mathcal{W}$ around (x^*, y^*) , so that*

$$\rho_{(a^*, x^*, y^*)} : \mathcal{A}_{a^*} \rightarrow \mathcal{O}, \quad a \mapsto \pi_{\mathcal{O}} \circ \underset{(x, y) \in \mathcal{N}_{(x^*, y^*)}}{\operatorname{argmin}} \frac{1}{2}\|a - \pi_{\mathcal{I}}(x, y)\|^2$$

is an (AP) and the condition number of this smooth map satisfies

$$\kappa[\rho_{(a^*, x^*, y^*)}](a) = \kappa[\mathcal{S}_{\text{GCPP}}](a, P_{\mathcal{I}_{x^*}}(a), \rho_{(a^*, x^*, y^*)}(a)),$$

where $\mathcal{I}_{x^*} = \pi_{\mathcal{I}}(\mathcal{N}_{(x^*, y^*)})$ is the projection onto the first factor, and $P_{\mathcal{I}_{x^*}}(a) = \underset{x \in \mathcal{I}_{x^*}}{\operatorname{argmin}} \|a - x\|^2$ is the nonlinear projection onto $\mathcal{I}_{x^*}$.

Taking $\mathcal{I} = \mathcal{O}$ and $\mathcal{S}$ as the graph of the identity map $\mathbf{1}$, the next result follows.

COROLLARY 7.4. *Let the input manifold $\mathcal{I}$, and well-posed loci $\mathcal{W}$ and $\mathcal{W}_{\text{CPP}}$ be as in (CPP). Assume that $(a^*, x^*) \in \mathcal{W}_{\text{CPP}}$ is well-posed and that x^* is a local minimizer of $d_{a^*} : \mathcal{I} \rightarrow \mathbb{R}, x \mapsto \frac{1}{2}\|a^* - x\|^2$. Then, there exist open neighborhoods $\mathcal{A}_{a^*} \subset \mathbb{R}^n$ around a^* and $\mathcal{I}_{x^*} \subset \mathcal{I}$ around x^* , so that*

$$\rho_{(a^*, x^*)} : \mathcal{A}_{a^*} \rightarrow \mathcal{I}, \quad a \mapsto \underset{x \in \mathcal{I}_{x^*}}{\operatorname{argmin}} \frac{1}{2}\|a - x\|^2$$

is an (AP) and its condition number is $\kappa[\rho_{(a^*, x^*)}](a) = \kappa[\mathcal{S}_{\text{CPP}}](a, \rho_{(a^*, x^*)}(a))$.

An interpretation of these results is as follows. The condition number of a generalized critical point (a^*, x^*, y^*) that happens to correspond to a local minimum of (IPIP) equals the condition number of the map $\rho_{(a^*, x^*, y^*)}$. This smooth map models a Riemannian AP on the localized computational problem $\mathcal{N}_{(x^*, y^*)} \subset \mathcal{W}$ which has a unique global minimum for all ambient inputs from $\mathcal{A}_{a^*}$. Consequently, every local minimum of (IPIP) is *robust* in the sense that it moves *smoothly* according to $\rho_{(a^*, x^*, y^*)}$ in neighborhoods of the input a^* and $(x^*, y^*) \in \mathcal{W}$. In this open neighborhood, the critical point remains a local minimum.

The foregoing entails that if a local minimum $(x^*, y^*) \in \mathcal{W}$ is computed for (IPIP) with ambient input a^* using Riemannian optimization, then computing a local minimum (x, y) for (IPIP) with perturbed input $a^* + \Delta$ and starting point (x^*, y^*) results in $\operatorname{dist}_{\mathcal{O}}(y, y^*) \leq \kappa[\mathcal{S}_{\text{GCPP}}](a^*, x^*, y^*)\|\Delta\| + o(\|\Delta\|)$, where $\operatorname{dist}_{\mathcal{O}}$ is as in (1.1), provided that $\|\Delta\|$ is sufficiently small.

8. Computing the condition number.

Three main ingredients are required for evaluating the condition number of (IPIP) numerically:

- (i) the inverse of the Riemannian Hessian H_η from (H),
- (ii) the derivative of the computational problem $(D_{(x,y)}\pi_{\mathcal{O}})(D_{(x,y)}\pi_{\mathcal{I}})^{-1}$, and
- (iii) the spectral norm (2.2).

The only additional item for evaluating the condition number of (IPIP) relative to the idealized inverse problem “given $x \in \mathcal{I}$, find a $y \in \mathcal{O}$ such that $(x, y) \in \mathcal{S}$ ” is the curvature term H_η^{-1} from item (i). Therefore, we focus on evaluating (i) and only briefly mention the standard techniques for (ii) and (iii).

The derivative in (ii) is often available analytically as in Remark 5.4. If no analytical expression can be derived, it can be numerically approximated when $\mathcal{O} \subset \mathbb{R}^q$ is an embedded submanifold by computing the standard Jacobian of $\mathcal{S}$, viewed as the graph of a map $\mathbb{R}^n \rightarrow \mathbb{R}^q$, and composing it with the projections onto $T_x\mathcal{I}$ and $T_y\mathcal{O}$.

Evaluating the spectral norm in (iii) is a linear algebra problem. Let $\mathcal{I} \subset \mathbb{R}^n$ be embedded with the standard inherited metric and let the Riemannian metric of $\mathcal{O}$ be given in coordinates by a matrix G . If $G = R^T R$ is its Cholesky factorization [37], then $\|A\|_{\mathcal{I} \rightarrow \mathcal{O}} = \|RA\|_2 = \sigma_1(RA)$, where $A : T_x\mathcal{I} \rightarrow T_{A(x)}\mathcal{O}$ is a linear map and $\|\cdot\|_2$ is the usual spectral norm, i.e., the largest singular value $\sigma_1(\cdot)$. The largest singular value can be computed numerically in $\mathcal{O}(\min\{m^2p, p^2m\})$ operations, where $m = \dim \mathcal{I}$ and $p = \dim \mathcal{O}$, using standard algorithms [37]. If the linear map RA can be applied in fewer than $\mathcal{O}(mp)$ operations, a power iteration can be more efficient [6].

The Riemannian Hessian H_η can be computed analytically using standard techniques from differential geometry for computing Weingarten maps [28, 47, 52, 53, 55], or it can be approximated numerically [11]. Numerical approximation of the Hessian with forward differences is implemented in the Matlab package Manopt [14]. Its inverse can in both cases be computed with $\mathcal{O}(m^3)$ operations.

We found two main analytical approaches helpful to compute Weingarten maps of embedded submanifolds $\mathcal{I} \subset \mathbb{R}^n$. The first relies on the following characterization of the Weingarten map S_{η_x} in the direction of a normal vector $\eta_x = a - x$. It is defined in [28, Section 6.2] as

$$(8.1) \quad S_{\eta_x} : T_x\mathcal{I} \rightarrow T_x\mathcal{I}, \quad v_x \mapsto P_{T_x\mathcal{I}} \left(-(\tilde{\nabla}_{v_x} N)|_x \right),$$

where $P_{T_x\mathcal{I}}$ denotes the orthogonal projection onto $T_x\mathcal{I}$, N is a smooth extension of η to a normal vector field on an open neighborhood of x in $\mathcal{I}$, and $\tilde{\nabla}_{v_x}$ is the Euclidean covariant derivative [28, 47, 52, 53, 55]. In practice, the latter can be computed as

$$(\tilde{\nabla}_{v_x} Y)|_x = \frac{d}{dt} Y|_{\gamma(t)},$$

where Y is an arbitrary vector field on $\mathcal{I}$ and $\gamma(t)$ is a smooth curve realizing v_x [47, Lemma 4.9]. In other words, S_{η_x} maps η_x to the tangential part of the usual directional derivative of N in direction of $v_x \in T_x\mathcal{I} \subset \mathbb{R}^n$.

The second approach relies on the *second fundamental form* of $\mathcal{I}$, which is defined as the projection of the Euclidean covariant derivative onto the normal space $N_x\mathcal{I}$:

$$II_x(X, Y) := (II(X, Y))|_x := P_{N_x\mathcal{I}} \left((\tilde{\nabla}_{X|_x} Y)|_x \right),$$

where now both X and Y are vector fields on $\mathcal{I}$. The second fundamental form is symmetric in X and Y , and $II_x(X, Y)$ only depends on the tangent vectors $X|_x$ and $Y|_x$. It can be viewed as a smooth map $II_x : T_x\mathcal{M} \times T_x\mathcal{M} \rightarrow N_x\mathcal{M}$. Contracting the second fundamental form with a normal vector $\eta_x \in N_x\mathcal{I}$ yields an alternative definition of the Weingarten map: For all $v, w \in T_x\mathcal{I}$ we have $\langle S_{\eta_x}(v), w \rangle = \langle II_x(v, w), \eta_x \rangle$, where $\langle \cdot, \cdot \rangle$ is the Euclidean inner product. Computing the second fundamental form is often facilitated by *pushing forward* vector fields through a diffeomorphism $F : \mathcal{M} \rightarrow \mathcal{N}$. In this case, there exists a vector field Y on $\mathcal{N}$ that is *F-related* to a vector field X

FIGURE 9.1. Setup of the 2-camera triangulation problem. The world coordinates of $y \in \mathbb{R}^3$ are to be reconstructed from the projections $x_1, x_2 \in \mathbb{R}^2$ (in the respective image coordinates) of y onto the image planes of the cameras with centers at c_1 and c_2 (in world coordinates) respectively.

on $\mathcal{M}$: $Y|_p = (F_*X)|_p := (D_{F^{-1}(p)}F)(X|_{F^{-1}(p)})$. The integral curves generated by X and $Y = F_*X$ are related by Proposition 9.6 of [48]. This approach will be illustrated in the next section.

As a rule of thumb, we find that computation of the condition number of (IPIP) is approximately as expensive as one iteration of a Riemannian Newton method for solving this problem. Typically, the cost is not worse than $\mathcal{O}(m^3 + nm^2 + \min\{m^2p, p^2m\})$ operations, where the terms correspond one-to-one to items (i)–(iii), and $\mathcal{I}^m \subset \mathbb{R}^n$ and $p = \dim \mathcal{O}$.

9. Triangulation in computer vision. A rich source of approximation problems whose condition can be studied with the proposed framework is multiview geometry in computer vision; for an introduction to this domain see [33, 38, 51]. The tasks consist of recovering information from $r \geq 2$ camera projections of a scene in the world $\mathbb{R}^3$. We consider the *pinhole camera model*. In this model the image formation process is modeled by the following transformation, as visualised in Figure 9.1:

$$\mu_r : y \mapsto \left[\frac{A_\ell y + b_\ell}{c_\ell^T y + d_\ell} \right]_{\ell=1}^r,$$

where $A_\ell \in \mathbb{R}^{2 \times 3}$, $b_\ell \in \mathbb{R}^2$, $c_\ell \in \mathbb{R}^3$, and $d_\ell \in \mathbb{R}$; see [33, 38, 51]. Clearly, μ_r is not defined on all of $\mathbb{R}^3$. We clarify its domain in Lemma 9.1 below.

The vector $x = \mu_r(y) \in \mathbb{R}^{2r}$ obtained by this image formation process is called a *consistent point correspondence*. Information that can often be identified from consistent point correspondences include camera parameters and scene structure [33, 38, 51].

As an example application of our theory we compute the condition number of the *triangulation problem* in computer vision [38, Chapter 12]. In this computational problem we have $r \geq 2$ stationary projective cameras and a consistent point correspondence $x = (x_1, x_2, \dots, x_r) \in \mathbb{R}^{2r}$. The goal is to retrieve the world point $y \in \mathbb{R}^3$ from which they originate. Since the imaging process is subject to noisy measurements and the above idealized projective camera model holds only approximately [38], we expect that instead of x we are only given $a = (a_1, a_2, \dots, a_r)$ close to x . Thus, x is the true input of the problem, y is the output and a is the ambient input.

According to [38, p. 314], the “gold standard algorithm” for triangulation solves the (Riemannian) optimization problem

$$(9.1) \quad \min_{y \in \mathbb{R}^3} \frac{1}{2} \|a - \mu_r(y)\|^2.$$

Consequently, the triangulation problem can be cast as a GCCP provided that two conditions hold: (i) the set of consistent point correspondences is an embedded manifold, and (ii) intersecting the back-projected rays is a smooth map from aforementioned manifold to $\mathbb{R}^3$. We verify both conditions in the following subsection.

9.1. The multiview manifold. The image formation process can be interpreted as a projective transformation from the projective space $\mathbb{P}^3$ to $\mathbb{P}^2$ in terms of the 4×3 camera matrices $P_\ell = \begin{bmatrix} A_\ell & b_\ell \\ c_\ell^T & d_\ell \end{bmatrix}$ [33, 38, 51]. It yields homogeneous coordinates

of $x_\ell = (z_1^\ell/z_3^\ell, z_2^\ell/z_3^\ell) \in \mathbb{R}^2$ where $z^\ell = P_\ell \begin{bmatrix} y \\ 1 \end{bmatrix}$. Note that if $z_3^\ell = 0$ then the point y has no projection³ onto the image plane of the ℓ -th camera, which occurs precisely when y lies on the *principal plane* of the camera [38, p. 160]. This is the plane parallel to the image plane through the camera center c_ℓ ; see Figure 9.1. It is also known that points on the baseline of two cameras, i.e., the line connecting the camera centers visualised by the dashed line in Figure 9.1, all project to the same two points on the two cameras, called the epipoles [38, Section 10.1]. Such points cannot be triangulated from only two images. For simplicity, let $\mathcal{B} \subset \mathbb{R}^3$ be the union of the principal planes of the first two cameras and their baseline, so $\mathcal{B}$ is a 2-dimensional subvariety. The next result shows that a subset of the consistent point correspondences outside of $\mathcal{B}$ forms a smooth embedded submanifold of $\mathbb{R}^{2r}$.

LEMMA 9.1. *Let $\mathcal{O}_{\text{MV}} = \mathbb{R}^3 \setminus \mathcal{B}$ with $\mathcal{B}$ as above. The map $\mu_r : \mathcal{O}_{\text{MV}} \rightarrow \mathbb{R}^{2r}$ is a diffeomorphism onto its image $\mathcal{I}_{\text{MV}} = \mu_r(\mathcal{O}_{\text{MV}})$.*

Proof. Clearly μ_r is a smooth map between manifolds ($c_\ell^T y + d_\ell = z_3^\ell \neq 0$). It only remains to show that it has a smooth inverse. Theorem 4.1 of [40] states that μ_2 's projectivization is a birational map, entailing that μ_2 is a diffeomorphism onto its image. Let $\pi_{1:4} : \mathbb{R}^{2r} \rightarrow \mathbb{R}^4$ denote projection onto the first 4 coordinates. Then, $\mu_2^{-1} \circ \pi_{1:4}$ is a smooth map such that $(\mu_2^{-1} \circ \pi_{1:4}) \circ \mu_r = \mu_2^{-1} \circ \mu_2 = \mathbf{1}_{\mathcal{O}_{\text{MV}}}$, having used that the domain $\mathcal{O}_{\text{MV}}$ is the same for all r . For the right inverse, we see that any element of $\mathcal{I}_{\text{MV}}$ can be written as $\mu_r(y)$ for some $y \in \mathcal{O}_{\text{MV}}$. Hence,

$$(\mu_r \circ (\mu_2^{-1} \circ \pi_{1:4}))(\mu_r(y)) = (\mu_r \circ \mathbf{1}_{\mathcal{O}_{\text{MV}}})(y) = \mu_r(y),$$

so it is the identity on $\mathcal{I}_{\text{MV}}$. This proves that μ_r has $\mu_2^{-1} \circ \pi_{1:4}$ as smooth inverse. $\square$

As μ_r^{-1} is required in the statement of the Weingarten map, we give an algorithm for computing $\mu_r^{-1} = \mu_2^{-1} \circ \pi_{1:4}$. Assume that we are given $x \in \mathcal{I}_{\text{MV}} \subset \mathbb{R}^{2r}$ in the r -camera multiview manifold and let its first four coordinates be (x_1, y_1, x_2, y_2) . Then, by classic results [38, Section 12.2], the unique element in the kernel of

$$\begin{bmatrix} (x_1 e_3^T - e_1^T) P_1 \\ (y_1 e_3^T - e_2^T) P_1 \\ (x_2 e_3^T - e_1^T) P_2 \\ (y_2 e_3^T - e_2^T) P_2 \end{bmatrix} \in \mathbb{R}^{4 \times 4}$$

yields homogeneous coordinates of the back-projected point in $\mathbb{R}^3$.

The solution manifold $\mathcal{S}_{\text{MV}} \subset \mathcal{I}_{\text{MV}} \times \mathcal{O}_{\text{MV}}$ is the graph of μ_r^{-1} . Therefore, it is a properly embedded smooth submanifold; see, e.g., [48, Proposition 5.7].

9.2. The second fundamental form. Vector fields and integral curves on $\mathcal{I}_{\text{MV}}$ can be viewed through the lens of μ_r : We construct a local smooth frame of $\mathcal{I}_{\text{MV}}$ by pushing forward a frame from $\mathcal{O}_{\text{MV}}$ by μ_r . Then, the integral curves of each of the local smooth vector fields are computed, after which we apply the Gauss formula for curves [47, Lemma 8.5] to compute the second fundamental form.

Because of Lemma 9.1 a local smooth frame for the *multiview manifold* $\mathcal{I}_{\text{MV}} \subset \mathbb{R}^{2r}$ is obtained by pushing forward the constant global smooth orthonormal frame (e_1, e_2, e_3) of $\mathbb{R}^3$ by the derivative of μ_r . Its derivative is

$$D_y \mu_r : T_y \mathcal{O}_{\text{MV}} \rightarrow T_{\mu_r(y)} \mathcal{I}_{\text{MV}}, \quad \dot{y} \mapsto \left[\frac{A_\ell \dot{y}}{c_\ell^T y + d_\ell} - (c_\ell^T \dot{y}) \frac{A_\ell y + b_\ell}{(c_\ell^T y + d_\ell)^2} \right]_{\ell=1}^r,$$

and so a local smooth frame of $\mathcal{I}_{\text{MV}}$ is given by

$$E_i : \mathcal{I}_{\text{MV}} \rightarrow T \mathcal{I}_{\text{MV}}, \quad \mu_r(y) \mapsto \left[\frac{A_\ell e_i}{c_\ell^T y + d_\ell} - c_{\ell,i} \frac{A_\ell y + b_\ell}{(c_\ell^T y + d_\ell)^2} \right]_{\ell=1}^r, \quad i = 1, 2, 3,$$

³Actually, it has a projection if we allow points at infinity, i.e., if we consider the triangulation problem in projective space.

where $c_{\ell,i} = c_\ell^T e_i$. It is generally neither orthonormal nor orthogonal.

We compute the second fundamental form of the multiview manifold by differentiation along the integral curves generated by the smooth local frame (E_1, E_2, E_3) . The integral curves through $\mu_r(y)$ generated by this frame are the images of the integral curves passing through y generated by the e_i 's due to [48, Proposition 9.6]. The latter are seen to be $g_i(t) = y + te_i$ by elementary results on linear differential equations. Therefore, the integral curves generated by E_i are

$$\gamma_i(t) = \mu_r(g_i(t)) = \left[\frac{A_\ell(y + te_i) + b_\ell}{c_\ell^T(y + te_i) + d_\ell} \right]_{\ell=1}^r.$$

The components of the second fundamental form at $x = \mu_r(y) \in \mathcal{I}_{\text{MV}}$ are then

$$\begin{aligned} II_x(E_i, E_j) &= \text{P}_{\text{NT}_{\text{MV}}} \left(\frac{d}{dt} E_j |_{\gamma_i(t)} \right) \\ &= \text{P}_{\text{NT}_{\text{MV}}} \left(\frac{d}{dt} \left[\frac{A_\ell e_j}{c_\ell^T(y + te_i) + d_\ell} - c_{\ell,j} \frac{A_\ell(y + te_i) + b_\ell}{(c_\ell^T(y + te_i) + d_\ell)^2} \right]_{\ell=1}^r \right) \\ &= \text{P}_{\text{NT}_{\text{MV}}} \left(\left[-\frac{c_{\ell,i}}{\alpha_\ell^2(y)} A_\ell e_j - \frac{c_{\ell,j}}{\alpha_\ell^2(y)} A_\ell e_i + 2 \frac{c_{\ell,i} c_{\ell,j}}{\alpha_\ell^3(y)} (A_\ell y + b_\ell) \right]_{\ell=1}^r \right), \end{aligned}$$

where $\alpha_\ell(y) = c_\ell^T y + d_\ell$ and $i, j = 1, 2, 3$.

9.3. A practical algorithm. The Weingarten map of $\mathcal{I}_{\text{MV}}$ in the direction of the normal vector $\eta \in \text{N}_x \mathcal{I}_{\text{MV}}$ is obtained by contracting the second fundamental form with η ; that is, $\hat{S}_\eta = \langle II_x(E_i, E_j), \eta \rangle$. This can be computed efficiently using linear algebra operations. Partitioning $\eta = [\eta_\ell]_{\ell=1}^r$ with $\eta_\ell \in \mathbb{R}^2$, the symmetric coefficient matrix of the Weingarten map relative to the frame $\mathcal{E} = (E_1, E_2, E_3)$ becomes

$$\hat{S}_\eta = \sum_{\ell=1}^r \left[2 \frac{c_{\ell,i} c_{\ell,j}}{\alpha_\ell^3(y)} \eta_\ell^T (A_\ell y + b_\ell) - \frac{c_{\ell,i}}{\alpha_\ell^2(y)} \eta_\ell^T A_\ell e_j - \frac{c_{\ell,j}}{\alpha_\ell^2(y)} \eta_\ell^T A_\ell e_i \right]_{i,j=1}^3 \in \mathbb{R}^{3 \times 3},$$

where $y = \mu_r^{-1}(x)$.

To compute the spectrum of the Weingarten map using efficient linear algebra algorithms, we need to express it with respect to an orthonormal local smooth frame by applying Gram–Schmidt orthogonalization to $\mathcal{E}$. This is accomplished by placing the tangent vectors E_1, E_2, E_3 as columns of a $2r \times 3$ matrix J and computing its compact QR decomposition $QR = J$. The coefficient matrix of the Weingarten map expressed with respect to the orthogonalized frame $Q = (Q_1, Q_2, Q_3)$ is then $S_\eta = R^{-T} \hat{S}_\eta R^{-1}$.

Since $\mu_r : \mathcal{O}_{\text{MV}} \rightarrow \mathcal{I}_{\text{MV}}$ is a diffeomorphism, $(D_{(x,y)} \pi_{\mathcal{O}_{\text{MV}}}) (D_{(x,y)} \pi_{\mathcal{I}_{\text{MV}}})^{-1}$ equals the inverse of the derivative of μ_r by Remark 5.4. As R is the matrix of $D_y \mu_r : T_y \mathcal{O}_{\text{MV}} \rightarrow T_x \mathcal{I}_{\text{MV}}$ expressed with respect to the orthogonal basis (Q_1, Q_2, Q_3) of $\mathcal{I}_{\text{MV}}$ and the standard basis of $\mathbb{R}^3$, we get that $\kappa[\mathcal{S}_{\text{GCPP}}](x + \eta, x, y)$ equals

$$\|(D_{(x,y)} \pi_{\mathcal{O}_{\text{MV}}}) (D_{(x,y)} \pi_{\mathcal{I}_{\text{MV}}})^{-1} H_\eta^{-1}\|_{\mathcal{I} \rightarrow \mathcal{O}} = \|R^{-1} (I - S_\eta)^{-1}\|_2 = \frac{1}{\sigma_3((I - S_\eta)R)},$$

where I is the 3×3 identity matrix, and σ_3 is the third largest singular value.

The computational complexity of this algorithm grows linearly with the number of cameras r . Indeed, the Weingarten map $\hat{S}_\eta$ can be constructed in $\mathcal{O}(r)$ operations, the smooth frame $\mathcal{E}$ is orthogonalized in $\mathcal{O}(r)$ operations, the change of basis from $\hat{S}_\eta$ to S_η requires a constant number of operations, and the singular values of a 3×3 matrix adds another constant.

9.4. Numerical experiments. The computations below were implemented in Matlab R2017b [50]. The code we used is provided as supplementary files accompanying the arXiv version of this article. It uses functionality from Matlab's optimization toolbox. The experiments were performed on a computer running Ubuntu 18.04.3 LTS that consisted of an Intel Core i7-5600U CPU with 8GB main memory.

We present a few numerical examples illustrating the behavior of the condition number of triangulation. The basic setup is described next. We take all 10 camera matrices $P_i \in \mathbb{R}^{3 \times 4}$ from the “model house” data set of the Visual Geometry Group of the University of Oxford [69]. These cameras are all pointing roughly in the same direction. In our experiments we reconstruct the point $y := p + 1.5^{10}v \in \mathbb{R}^3$, where $p = (-1.85213, -0.532959, -5.65752) \in \mathbb{R}^3$ is one point of the house and $v = (-0.29292, -0.08800, -0.95208)$ is the unit-norm vector that points from the center of the first camera p . Given a data point $a \in \mathbb{R}^{2r}$, it is triangulated as follows. First a linear triangulation method is applied, finding the right singular vector corresponding to the least singular value of a matrix whose construction is described (for two points) in [38, Section 12.2]. We then use it as a starting point for solving optimization problem (9.1) with Matlab’s nonlinear least squares solver `lsqnonlin`. For this solver, the following settings were used: `TolFun` and `TolX` set to 10^{-28} , `StepTolerance` equal to 10^{-14} , and both `MaxFunEvals` and `MaxIter` set to 10^4 . We provided the Jacobian to the algorithm, which it can evaluate numerically.

Since the multiview manifold $\mathcal{I}_{\text{MV}}$ is only a three-dimensional submanifold of $\mathbb{R}^{2r}$ and the computational complexity is linear in r , we do not focus on the computational efficiency of computing $\kappa[\mathcal{S}_{\text{GCPP}}]$. In practical cases, with $r \ll 1000$, the performance is excellent. For example, the average time to compute the condition number for $r = 1000$ cameras was less than 0.1 seconds in our setup. This includes the time to compute the frame $\mathcal{E}$, the Weingarten map expressed with respect to $\mathcal{E}$, the orthogonalization of the frame, and the computation of the smallest singular value.

Experiment 1. In the first experiment, we provide numerical evidence for the correctness of our theoretical results. We project y to the image planes of the $r = 10$ cameras: $x = \mu_r(y)$. A random normal direction $\eta \in N_x \mathcal{I}_{\text{MV}}$ is sampled by taking a random vector with i.i.d. standard normal entries, projecting it to the normal space, and then scaling it to unit norm. We investigate the sensitivity of the points along the ray $a(t) := x - t\eta$.

The theoretical condition number of the triangulation problem at $a(t)$ can be computed numerically for every t using the algorithm from subsection 9.3. We can also estimate the condition number experimentally, for example by generating a large number of small perturbations $a(t) + E(t)$ with $\|E(t)\| \approx 0$, solving (9.1) numerically using `lsqnonlin` and then checking the distance to the true critical point y . However, from the theory we know that the worst direction of infinitesimal perturbation is given by the left singular vector of $(I - S_{-t\eta})R$ corresponding the smallest (i.e. third) singular value, where R is as in subsection 9.3. Let $u(t) \in \mathbb{R}^3$ denote this vector, which contains the coordinates of the perturbation relative to the orthonormal basis $Q = (Q_1, Q_2, Q_3)$; see subsection 9.3. Then, it suffices to consider only a small (relative) perturbation in this direction; we took $E(t) := 10^{-6}\|a(t)\|Qu(t)$. We solve (9.1) with input $a(t) + E(t)$ and output $y_{\text{est}}(t)$ using `lsqnonlin` with the exact solution of the unperturbed problem y as starting point. The experimental estimate of the condition number is then $\kappa_{\text{est}}(t) = \frac{\|y - y_{\text{est}}(t)\|}{\|E(t)\|}$.

In Figure 9.2 we show the experimental results for the above setup, where we chose $t = 10^i\|x\|$ with 100 values of i uniformly spaced between -3 and 2 . It is visually evident that the experimental results support the theory. Excluding the three last values of t , the arithmetic and geometric means of the condition number divided by the experimental estimate are approximately 1.006664 and 1.006497 respectively. This indicates a near perfect match with the theory.

There appears to be one anomaly in Figure 9.2, however: the three crosses to the right of the singularity at $t^* = 79.64416\|x\|$. What happens is that past the singularity, y is no longer a local minimizer of the optimization problem $\min_{y \in \mathcal{O}_{\text{MV}}} \frac{1}{2}\|a(t) - \mu_{10}(y)\|^2$. This is because $x(t)$ is no longer a local minimizer of $\min_{x \in \mathcal{I}_{\text{MV}}} \frac{1}{2}\|a(t) - x\|^2$. Since we are applying an optimization method, the local minimizer happens to move away from the nearby critical point and instead converges to a different point. Indeed, one can show that the Riemannian Hessians (see [2, Chapter 5.5]) of these problems

FIGURE 9.2. A comparison of theoretical and experimental data of the condition number of the triangulation problem with all cameras for several points along the ray $x - t\eta$. Herein, η is a randomly chosen unit-norm normal direction at $x \in \mathcal{I}_{\text{MV}}$. The numerically computed theoretical condition numbers are indicated by circles, while the experimentally estimated condition numbers, as described in the text, are plotted with crosses.

FIGURE 9.3. The GCPP condition number in function of the relative distance $\frac{t}{\|x\|}$ from a specified $x \in \mathcal{I}_{\text{MV}}$ and a randomly chosen unit-length normal direction η . Each of the four lines corresponds to a different number k of cameras taking the pictures, namely $k = 2, 3, 5, 10$.

are positive definite for small t , but become indefinite past the singularity t^* .

Experiment 2. The next experiment illustrates how the condition number varies with the distance t along a normal direction $t\eta \in \text{NZ}_{\text{MV}}$. We consider the setup from the previous paragraph. The projection of the point y onto the image planes of the first k cameras is $x_k := \mu_k(y) = (P_1 y, \dots, P_k y)$. A random unit-norm normal vector η_k in $\text{N}_{x_k} \mathcal{I}_{\text{MV}}$ is chosen as described before. We investigate the sensitivity of the perturbed inputs $a_k(t) := x_k + t\eta_k$ for $t \in \mathbb{R}$. The condition number is computed for these inputs for $k = 2, 3, 5, 10$ and $\frac{\pm t}{\|x_k\|} = 10^i$ for 10^4 values of i uniformly spaced between -3 and 4 . The results are shown in Figure 9.3. We have 3 peaks for a fixed k because $\dim \mathcal{I}_{\text{MV}} = 3$. The two-camera system has only two singularities in the range $|t| < 10^4 \|x_2\|$. These figures illustrate the continuity of κ_{GCPP} from Corollary 6.3.

Up to about $|t| \leq 10^{-2} \|x_k\|$ we observe that curvature hardly affects the condition number of the GCPP for the triangulation problem. Beyond this value, curvature plays the dominant role: both the singularities and the tapering off for high relative errors are caused by it. Moreover, the condition number decreases as the number of cameras increases. This indicates that the use of *minimal problems* [46] in computer vision could potentially introduce numerical instability.

Acknowledgements. The authors thank Nicolas Boumal for posing a question on unconstrained errors, which eventually started the project that led to this article. We are greatly indebted to Carlos Beltrán who meticulously studied the manuscript and asked pertinent questions about the theory. Furthermore, the authors would like to thank Peter Bürgisser for hosting the second author at TU Berlin in August 2019, and Joeri Van der Veken for fruitful suggestions and pointers to the literature.

We thank two anonymous reviewers whose remarks enticed us to immensely im-

prove the presentation of our results, specifically in [Section 1](#) and [3](#), and to include a discussion of computational aspects in [Section 8](#). We also thank the editor, Anthony So, for adeptly handling the reviewing process.

Appendix A. Proofs of the technical results. We recall two classic results that we could not locate in modern references in differential geometry. A common theme in the three computational problems considered in [Section 4](#) to [6](#) is that they involve the projection $\Pi_{\mathcal{I}} : N\mathcal{I} \rightarrow \mathcal{I}$, $(x, \eta) \mapsto x$ from the normal bundle $N\mathcal{I}$ to the base of the bundle, the embedded submanifold $\mathcal{I} \subset \mathbb{R}^n$. This projection is smooth [[48](#), Corollary 10.36] and its derivative is characterized by the following classical result, which dates back at least to Weyl [[70](#)].

PROPOSITION A.1. *Let $(x, \eta) \in N\mathcal{I}$ and $(\dot{x}, \dot{\eta}) \in T_{(x, \eta)}N\mathcal{I}$. Then, $P_{T_x\mathcal{I}}(\dot{x} + \dot{\eta}) = H_\eta \dot{x}$, where H_η is as in [\(H\)](#).*

A useful consequence is that it allows us to characterize when H_η is singular.

LEMMA A.2. *Let $(a, x) \in \mathcal{S}_{\text{CPP}}$, $\eta = a - x$, and H_η be as in [\(H\)](#). Then, H_η is invertible if and only if $(a, x) \in \mathcal{W}_{\text{CPP}}$.*

Proof. Let $(\dot{a}, \dot{x}) \in T_{(a, x)}\mathcal{S}_{\text{CPP}}$. Then, we have $\dot{x} = (D_{(a, x)}\Pi_{\mathcal{I}})(\dot{a}, \dot{x})$ and $\dot{a} = (D_{(a, x)}\Pi_{\mathbb{R}^n})(\dot{a}, \dot{x})$, so that [Proposition A.1](#) yields the equality of operators

$$(A.1) \quad P_{T_x\mathcal{I}} D_{(a, x)}\Pi_{\mathbb{R}^n} = H_\eta D_{(a, x)}\Pi_{\mathcal{I}}.$$

On the one hand, if $(a, x) \in \mathcal{W}_{\text{CPP}}$ then the map $D_{(a, x)}\Pi_{\mathbb{R}^n}$ is invertible. By multiplying with its inverse we get $P_{T_x\mathcal{I}} = H_\eta (D_{(a, x)}\Pi_{\mathcal{I}})(D_{(a, x)}\Pi_{\mathbb{R}^n})^{-1}$. As H_η is an endomorphism on $T_x\mathcal{I}$, which is the range of the left-hand side, the equality requires H_η to be an automorphism, i.e., it is invertible.

On the other hand, if $(a, x) \notin \mathcal{W}_{\text{CPP}}$, there is some $(\dot{a}, \dot{x}) \neq 0$ in the kernel of $D_{(a, x)}\Pi_{\mathbb{R}^n}$. Since $\dot{a} = (D_{(a, x)}\Pi_{\mathbb{R}^n})(\dot{a}, \dot{x}) = 0$, we must have $\dot{x} \neq 0$. Consequently, applying both sides of [\(A.1\)](#) to $(\dot{a}, \dot{x})$ gives $0 = H_\eta \dot{x}$, which implies that H_η is singular. This finishes the proof. $\square$

The next result follows almost directly from [Proposition A.1](#) and appears several times as an independent statement in the literature. The earliest reference we could locate is Abatzoglou [[1](#), Theorem 4.1], who presented it for C^2 embedded submanifolds of $\mathbb{R}^n$ in a local coordinate chart.

COROLLARY A.3. *Let $\mathcal{I} \subset \mathbb{R}^n$ be a Riemannian embedded submanifold. Then, there exists a tubular neighborhood $\mathcal{T}$ of $\mathcal{I}$ such that the projection map $P_{\mathcal{I}} : \mathcal{T} \rightarrow \mathcal{I}$, $a \mapsto \operatorname{argmin}_{x \in \mathcal{I}} \|a - x\|^2$ is a smooth submersion with derivative*

$$D_a P_{\mathcal{I}} = H_\eta^{-1} P_{T_x\mathcal{I}},$$

where $x = P_{\mathcal{I}}(a)$, $\eta = a - x$, and H_η is given by [\(H\)](#).

In the next subsections, proofs are supplied for the main results.

A.1. Proofs for [Section 4](#).

Proof of [Lemma 4.2](#). Recall that $\Pi_{\mathbb{R}^n}$ is a smooth map, so if its derivative is nonsingular at a point $(a, x) \in \mathcal{S}_{\text{CPP}}$, then there exists an open neighborhood of (a, x) where this property remains valid. Therefore, applying the inverse function theorem [[48](#), Theorem 4.5] at $(a, x) \in \mathcal{W}_{\text{CPP}}$, the fact that $D_{(a, x)}\Pi_{\mathbb{R}^n}$ is invertible implies that there is a neighborhood $\mathcal{N}_{(a, x)}$ of (a, x) in $\mathcal{W}_{\text{CPP}}$ and a neighborhood of a in $\mathbb{R}^n$ such that $\mathcal{N}_{(a, x)}$ is diffeomorphic to (an open ball in) $\mathbb{R}^n$. Hence, $\mathcal{W}_{\text{CPP}}$ is open in $\mathcal{S}_{\text{CPP}}$, and, hence, it is an open submanifold of dimension n .

To show that $\mathcal{W}_{\text{CPP}}$ is dense in $\mathcal{S}_{\text{CPP}}$, take any $(a, x) \in \mathcal{S}_{\text{CPP}} \setminus \mathcal{W}_{\text{CPP}}$. Let $\eta = a - x$. Then, $(x, x) + (\alpha\eta, 0) \in \mathcal{S}_{\text{CPP}}$ because $\eta \in N_x\mathcal{I}$. By [Lemma A.2](#), $(x, x) + (\alpha\eta, 0) \in \mathcal{W}_{\text{CPP}}$ with $\alpha \in (0, \infty)$ if and only if $H_{\alpha\eta} = \mathbf{1} - S_{\alpha\eta} = \mathbf{1} - \alpha S_\eta$ is singular. This occurs only if $\frac{1}{\alpha}$ equals one of the eigenvalues of the Weingarten map

S_η . Consequently, (a, x) can be reached as the limit of a sequence $(x, x) + (\alpha_n \eta, 0)$, which is wholly contained in $\mathcal{W}_{\text{CPP}}$ by taking α_n outside of the discrete set.

Applying [48, Proposition 5.1], we can conclude that $\mathcal{W}_{\text{CPP}}$ is even an embedded submanifold. This concludes the argument. $\square$

Proof of Theorem 4.3. Let $(a, x) \in \mathcal{W}_{\text{CPP}}$ be arbitrary. On $\mathcal{W}_{\text{CPP}}$, the coordinate projection $\Pi_{\mathbb{R}^n} : \mathcal{S}_{\text{CPP}} \rightarrow \mathbb{R}^n$ has an invertible derivative. Consequently, by the inverse function theorem [48, Theorem 4.5] there exist open neighborhoods $\mathcal{X} \subset \mathcal{S}_{\text{CPP}}$ of (a, x) and $\mathcal{Y}$ of $a \in \mathbb{R}^n$ such that $\Pi_{\mathbb{R}^n}|_{\mathcal{X}} : \mathcal{X} \rightarrow \mathcal{Y}$ has a smooth inverse function that we call $\phi_{(a,x)}$. Its derivative is $D_a \phi_{(a,x)} = (D_{(a,x)} \Pi_{\mathbb{R}^n})^{-1}$. Consider the smooth map $\Phi_{(a,x)} := \Pi_{\mathcal{I}} \circ \phi_{(a,x)}$, where $\Pi_{\mathcal{I}} : \mathcal{S}_{\text{CPP}} \rightarrow \mathcal{I}$ is the coordinate projection. The CPP condition number at $(a, x) \in \mathcal{W}_{\text{CPP}}$ is $\kappa[\Phi_{(a,x)}](a) = \kappa[\mathcal{W}_{\text{CPP}}](a, x) = \|D_a \Phi_{(a,x)}\|_{\mathbb{R}^n \rightarrow \mathcal{I}}$. Since $D_a \Phi_{(a,x)} = (D_{(a,x)} \Pi_{\mathcal{I}})(D_{(a,x)} \Pi_{\mathbb{R}^n})^{-1}$, we have that $\kappa[\mathcal{W}_{\text{CPP}}](a, x) = \|(D_{(a,x)} \Pi_{\mathcal{I}})(D_{(a,x)} \Pi_{\mathbb{R}^n})^{-1}\|_{\mathbb{R}^n \rightarrow \mathcal{I}}$. Now it follows from Lemma A.2 and (A.1) that $(D_{(a,x)} \Pi_{\mathcal{I}})(D_{(a,x)} \Pi_{\mathbb{R}^n})^{-1} = H_\eta^{-1} P_{\mathcal{I}x}$. As the metrics on $\mathcal{I}$ and $\mathbb{R}^n$ are identical, $\|H_\eta^{-1} P_{\mathcal{I}x}\|_{\mathbb{R}^n \rightarrow \mathcal{I}} = \|H_\eta^{-1}\|_{\mathcal{I} \rightarrow \mathcal{I}}$. This concludes the first part.

The second part is Lemma A.2. $\square$

A.2. Proofs for Section 5.

Proof of Lemma 5.1. By Corollary A.3 the projection $P_{\mathcal{I}} : \mathcal{T} \rightarrow \mathcal{I}$ is a smooth map on $\mathcal{T}$. Let us denote the graph of this map by $\mathcal{V} := \{(a, P_{\mathcal{I}}(a)) \in \mathcal{T} \times \mathcal{I} \mid a \in \mathcal{T}\}$. It is a smooth embedded submanifold of dimension $\dim \mathcal{T} = \dim \mathcal{N}\mathcal{I} = n$; see, e.g., [48, Proposition 5.4]. Let $x = P_{\mathcal{I}}(a)$, $\eta = a - x$ and $(\dot{a}, \dot{x}) \in T_{(a,x)}\mathcal{V}$ be a nonzero tangent vector. Because of Corollary A.3 it satisfies $\dot{x} = H_\eta^{-1} P_{\mathcal{I}x} \dot{a}$. Hence, $(D_{(a,x)} \Pi_{\mathbb{R}^n})(\dot{a}, H_\eta^{-1} P_{\mathcal{I}x} \dot{a}) = \dot{a} \neq 0$, so that its kernel is trivial. It follows that $\mathcal{V} \subset \mathcal{W}_{\text{CPP}}$. Since their dimensions match by Lemma 4.1, $\mathcal{V}$ is an open submanifold. As $\mathcal{W}_{\text{CPP}}$ is embedded by Lemma 4.2, $\mathcal{V}$ is embedded as well by [48, Proposition 5.1]. Moreover, by construction, we have

$$\mathcal{W}_{\text{AP}} = (\mathcal{V} \times \mathcal{O}) \cap (\mathbb{R}^n \times \mathcal{W}) \subset \mathbb{R}^n \times \mathcal{I} \times \mathcal{O}.$$

The first part of the proof is concluded by observing that the proof of Lemma 6.1 applies by replacing $\mathcal{W}_{\text{CPP}}$ with $\mathcal{V}$ and $\mathcal{W}_{\text{GCPP}}$ with $\mathcal{W}_{\text{AP}}$.

Finally, we show that $\mathcal{W}_{\text{AP}}$ is dense in $\mathcal{S}_{\text{AP}}$. Let $(x + \eta, y) \in \mathcal{S}_{\text{AP}}$ with $x \in \mathcal{I}$ and $\eta \in \mathcal{N}_x \mathcal{I}$ be arbitrary. Then $P_{\mathcal{I}}(x + \eta) = x$ and $(x, y) \in \mathcal{S}$. As $\mathcal{W}$ is an open dense submanifold of $\mathcal{S}$, there exists a sequence $(x_i, y_i) \in \mathcal{W}$ such that $\lim_{i \rightarrow \infty} (x_i, y_i) = (x, y)$. Consider $(x_i + \eta_i, y_i) \in \mathcal{W}_{\text{AP}}$ with $\eta_i \in \mathcal{N}_{x_i} \mathcal{I}$ and $\eta_i \rightarrow \eta$ chosen in such a way that the first component lies in the tubular neighborhood; this is possible as $\mathcal{T}$ is an open submanifold of $\mathbb{R}^n$. Then, the limit of this sequence is $(x + \eta, y)$. Q.E.D. $\square$

Proof of Theorem 5.2. For $(a, y) \in \mathcal{W}_{\text{AP}}$, let $x = P_{\mathcal{I}}(a)$ and $\eta = a - x$. By (1.2), $\kappa[\mathcal{W}_{\text{AP}}](a, y) = \|D_a(\pi_{\mathcal{O}} \circ \pi_{\mathcal{I}}^{-1} \circ P_{\mathcal{I}})\|_{\mathbb{R}^n \rightarrow \mathcal{O}}$. Since $(x, y) \in \mathcal{W}$, we know that $D_{(x,y)} \pi_{\mathcal{I}}$ is invertible. Moreover, by Corollary A.3 the derivative of the projection $P_{\mathcal{I}} : \mathcal{T} \rightarrow \mathcal{I}$ at a is $D_a P_{\mathcal{I}} = H_\eta^{-1} P_{\mathcal{I}x}$. Therefore,

$$\kappa[\mathcal{W}_{\text{AP}}](a, y) = \|(D_{(x,y)} \pi_{\mathcal{O}})(D_{(x,y)} \pi_{\mathcal{I}})^{-1} H_\eta^{-1}\|_{\mathcal{I} \rightarrow \mathcal{O}}.$$

For $(a, y) \in \mathcal{S}_{\text{AP}} \setminus \mathcal{W}_{\text{AP}}$, we have $(x, y) \notin \mathcal{W}$ with $x = P_{\mathcal{I}}(a)$. By definition of $\mathcal{W}$, this entails that $D_{(x,y)} \pi_{\mathcal{I}}$ is not invertible, concluding the proof. $\square$

A.3. Proofs for Section 6.

Proof of Lemma 6.1. Combining Assumption 1 and Lemma 4.2, it is seen that $\mathcal{W}_{\text{GCPP}}$ is realized as the intersection of two smoothly embedded product manifolds:

$$\mathcal{W}_{\text{GCPP}} = (\mathcal{W}_{\text{CPP}} \times \mathcal{O}) \cap (\mathbb{R}^n \times \mathcal{W}) \subset \mathbb{R}^n \times \mathcal{I} \times \mathcal{O}.$$

Note that $\text{codim}(\mathcal{W}_{\text{CPP}} \times \mathcal{O}) = \dim \mathcal{I}$, by Lemma 4.1, and $\text{codim}(\mathbb{R}^n \times \mathcal{W}) = \dim \mathcal{O}$, because $\dim \mathcal{W} = \dim \mathcal{S}$ and Assumption 1 stating that $\dim \mathcal{S} = \dim \mathcal{I}$. Thus, it

suffices to prove that the intersection is *transversal* [48, Theorem 6.30], because then $\mathcal{W}_{\text{GCPP}}$ would be an embedded submanifold of codimension $\dim \mathcal{I} + \dim \mathcal{O}$.

Consider a point $(a, x, y) \in \mathcal{W}_{\text{GCPP}}$. Transversality means that we need to show

$$T_{(a,x,y)}(\mathcal{W}_{\text{GCPP}} \times \mathcal{O}) + T_{(a,x,y)}(\mathbb{R}^n \times \mathcal{W}) = T_{(a,x,y)}(\mathbb{R}^n \times \mathcal{I} \times \mathcal{O}).$$

Fix an arbitrary $(\dot{a}, \dot{x}, \dot{y}) \in T_a \mathbb{R}^n \times T_x \mathcal{I} \times T_y \mathcal{O}$. Then, it suffices to show there exist $((\dot{a}_1, \dot{x}_1), \dot{y}_1) \in T_{(a,x)} \mathcal{W}_{\text{GCPP}} \times T_y \mathcal{O}$ and $(\dot{a}_2, (\dot{x}_2, \dot{y}_2)) \in T_a \mathbb{R}^n \times T_{(x,y)} \mathcal{W}$ such that $(\dot{a}, \dot{x}, \dot{y}) = (\dot{a}_1 + \dot{a}_2, \dot{x}_1 + \dot{x}_2, \dot{y}_1 + \dot{y}_2)$. It follows from [Proposition A.1](#) and [Lemma A.2](#) that $\dot{a}_1 = H_\eta \dot{x}_1 + \dot{\eta}$, where $\eta = a - x$ and where $\dot{\eta} \in N_x \mathcal{I}$. The tangent vectors $\dot{y}_1 \in T_y \mathcal{O}$, $\dot{a}_2 \in T_a \mathbb{R}^n$ and $\dot{x}_1 \in T_x \mathcal{I}$ can be chosen freely. Choosing $\dot{x}_1 = \dot{x}$ and $\dot{x}_2 = 0$; $\dot{y}_1 = \dot{y}$ and $\dot{y}_2 = 0$ and $\dot{a}_1 = H_\eta \dot{x}_1$ and $\dot{a}_2 = \dot{a} - \dot{a}_1$ yields $(\dot{a}, \dot{x}, \dot{y})$. This concludes the first part of the proof.

It remains to prove $\mathcal{W}_{\text{GCPP}}$ is dense in $\mathcal{S}_{\text{GCPP}}$. For this, we recall that

$$\mathcal{S}_{\text{GCPP}} = (\mathcal{S}_{\text{CPP}} \times \mathcal{O}) \cap (\mathbb{R}^n \times \mathcal{S}).$$

Let $(a, x, y) \in \mathcal{S}_{\text{GCPP}}$ and let $\eta = a - x \in N_x \mathcal{I}$. Then, $(a, x) \in \mathcal{S}_{\text{CPP}}$ and $(x, y) \in \mathcal{S}$. As $\mathcal{W}$ is an open dense submanifold of $\mathcal{S}$ by [Assumption 1](#), there exists a sequence $(x_i, y_i) \in \mathcal{W}$ such that $\lim_{i \rightarrow \infty} (x_i, y_i) = (x, y)$. Moreover, there exists a sequence $\eta_i \rightarrow \eta$ with $\eta_i \in T_{x_i} \mathcal{I}$ so that $(x_i + \eta_i, x_i) \in \mathcal{S}_{\text{CPP}}$, because $\Pi_{\mathcal{I}}(\mathcal{S}_{\text{CPP}}) = \mathcal{I}$. Since $\mathcal{W}_{\text{CPP}}$ is open dense in $\mathcal{S}_{\text{CPP}}$ by [Lemma 4.2](#), only finitely many elements of the sequence are not in $\mathcal{W}_{\text{CPP}}$. We can pass to a subsequence $(x_j + \eta_j, x_j) \in \mathcal{W}_{\text{CPP}}$. Now $(x_j + \eta_j, x_j, y_j) \in \mathcal{W}_{\text{GCPP}}$ and moreover its limit is $(x + \eta, x, y) \in \mathcal{S}_{\text{GCPP}}$. Q.E.D. $\square$

Proof of Theorem 6.2. For a well-posed triple $(a, x, y) \in \mathcal{W}_{\text{GCPP}}$ we obtain from (1.2) that $\kappa[\mathcal{W}_{\text{GCPP}}](a, x, y) = \|(\mathcal{D}_{(a,x,y)} \Pi_{\mathcal{O}})(\mathcal{D}_{(a,x,y)} \Pi_{\mathbb{R}^n})^{-1}\|_{\mathbb{R}^n \rightarrow \mathcal{O}}$. We compute the right-hand side of this equation. The situation appears as follows:

$$\begin{array}{ccccc} \mathcal{W}_{\text{CPP}} & \xrightarrow{\Pi_{\mathbb{R}^n}} & \mathbb{R}^n & \xleftarrow{\Pi_{\mathbb{R}^n}} & \mathcal{W}_{\text{GCPP}} \\ \Pi_{\mathcal{I}} \downarrow & & & & \downarrow \Pi_{\mathcal{O}} \\ \mathcal{I} & \xleftarrow{\pi_{\mathcal{I}}} & \mathcal{W} & \xrightarrow{\pi_{\mathcal{O}}} & \mathcal{O} \end{array}$$

We have in addition the projections $(\mathbf{1} \times \Pi_{\mathcal{I}}) : \mathcal{W}_{\text{GCPP}} \rightarrow \mathcal{W}_{\text{CPP}}$, $(a, x, y) \mapsto (a, x)$ and $(\Pi_{\mathcal{I}} \times \mathbf{1}) : \mathcal{W}_{\text{GCPP}} \rightarrow \mathcal{W}$, $(a, x, y) \mapsto (x, y)$. With this we have

$$(\Pi_{\mathbb{R}^n} \circ (\mathbf{1} \times \pi_{\mathcal{I}}))(a, x, y) = \Pi_{\mathbb{R}^n}(a, x, y) \quad \text{and} \quad (\pi_{\mathcal{O}} \circ (\Pi_{\mathcal{I}} \times \mathbf{1}))(a, x, y) = \Pi_{\mathcal{O}}(a, x, y).$$

Taking derivatives we get

$$(\mathcal{D}_{(a,x)} \Pi_{\mathbb{R}^n}) \mathcal{D}_{(a,x,y)}(\mathbf{1} \times \pi_{\mathcal{I}}) = \mathcal{D}_{(a,x,y)} \Pi_{\mathbb{R}^n} \quad \text{and} \quad (\mathcal{D}_{(a,x)} \pi_{\mathcal{O}}) \mathcal{D}_{(a,x,y)}(\Pi_{\mathcal{I}} \times \mathbf{1}) = \mathcal{D}_{(a,x,y)} \Pi_{\mathcal{O}}.$$

Since $(x, y) \in \mathcal{W}$, the derivative $\mathcal{D}_{(x,y)} \pi_{\mathcal{I}}$ is invertible, and so $\mathcal{D}_{(a,x,y)}(\mathbf{1} \times \pi_{\mathcal{I}})$ is also invertible. On other hand, since $(a, x) \in \mathcal{W}_{\text{CPP}}$, the derivative $\mathcal{D}_{(a,x)} \Pi_{\mathbb{R}^n}$ is invertible. Altogether we see that $\mathcal{D}_{(a,x,y)} \Pi_{\mathbb{R}^n}$ is invertible, and that

$$\begin{aligned} & (\mathcal{D}_{(a,x,y)} \Pi_{\mathcal{O}})(\mathcal{D}_{(a,x,y)} \Pi_{\mathbb{R}^n})^{-1} \\ &= (\mathcal{D}_{(x,y)} \pi_{\mathcal{O}}) (\mathcal{D}_{(a,x,y)} (\Pi_{\mathcal{I}} \times \mathbf{1})) (\mathcal{D}_{(a,x,y)} (\mathbf{1} \times \pi_{\mathcal{I}}))^{-1} (\mathcal{D}_{(a,x)} \Pi_{\mathbb{R}^n})^{-1}. \end{aligned}$$

The derivatives in the middle satisfy

$$(\dot{x}, \dot{y}) = \mathcal{D}_{(a,x,y)} (\Pi_{\mathcal{I}} \times \mathbf{1}) (\mathcal{D}_{(a,x,y)} (\mathbf{1} \times \pi_{\mathcal{I}}))^{-1} (\dot{a}, \dot{x}),$$

while on the other hand $(\dot{x}, \dot{y}) = (\mathcal{D}_{(x,y)} \pi_{\mathcal{I}})^{-1} (\mathcal{D}_{(a,x)} \Pi_{\mathcal{I}}) (\dot{a}, \dot{x})$. This implies that the two linear maps on the right hand side in the previous equations are equal. Hence,

$$\begin{aligned} & (\mathcal{D}_{(a,x,y)} \Pi_{\mathcal{O}})(\mathcal{D}_{(a,x,y)} \Pi_{\mathbb{R}^n})^{-1} = (\mathcal{D}_{(x,y)} \pi_{\mathcal{O}})(\mathcal{D}_{(x,y)} \pi_{\mathcal{I}})^{-1} (\mathcal{D}_{(a,x)} \Pi_{\mathcal{I}})(\mathcal{D}_{(a,x)} \Pi_{\mathbb{R}^n})^{-1} \\ & \text{(A.2)} \quad \quad \quad = (\mathcal{D}_{(x,y)} \pi_{\mathcal{O}})(\mathcal{D}_{(x,y)} \pi_{\mathcal{I}})^{-1} H_\eta^{-1} P_{T_x \mathcal{I}}, \end{aligned}$$

where we used (A.1) and [Lemma A.2](#). Taking spectral norms on both sides completes the proof for points in $\mathcal{W}_{\text{GCPP}}$. Finally, for points outside $\mathcal{W}_{\text{GCPP}}$ the proof follows from the definition of $\mathcal{W}_{\text{GCPP}}$ combined with [Theorem 4.3](#) and [5.2](#). $\square$

FIGURE A.1. A sketch of the geometric construction in the proof of Theorem 7.3.

A.4. Proofs for Section 7.

Proof of Proposition 7.1. Taking the derivative of the objective function at a point $(x, y) \in \mathcal{W}$ we see that the critical points satisfy $\langle a - x, D_{(x,y)}\pi_{\mathcal{I}}(\dot{x}, \dot{y}) \rangle = 0$ for all $(\dot{x}, \dot{y}) \in T_{(x,y)}\mathcal{W}$. By the definition of $\mathcal{W}$, the derivative $D_{(x,y)}\pi_{\mathcal{I}}$ is surjective, which implies that $(D_{(x,y)}\pi_{\mathcal{I}})(T_{(x,y)}\mathcal{W}) = T_x\mathcal{I}$. Therefore, the condition of being a critical point is equivalent to $a - x \in N_x\mathcal{I}$. $\square$

Proof of Theorem 7.3. By assumption $(a^*, x^*, y^*) \in \mathcal{W}_{\text{GCPP}}$. In particular, this implies $(a^*, x^*) \in \mathcal{W}_{\text{GCPP}}$, so that $D_{(a^*, x^*)}\Pi_{\mathbb{R}^n} : T_{(a^*, x^*)}\mathcal{W}_{\text{GCPP}} \rightarrow \mathbb{R}^n$ is invertible. By the inverse function theorem [48, Theorem 4.5], there is an open neighborhood $\mathcal{U}_{(a^*, x^*)} \subset \mathcal{W}_{\text{GCPP}}$ of (a^*, x^*) such that $\Pi_{\mathbb{R}^n}$ restricts to a diffeomorphism on

$$\mathcal{A}_{a^*} = \Pi_{\mathbb{R}^n}(\mathcal{U}_{(a^*, x^*)}).$$

Let $\Pi_{\mathbb{R}^n}^{-1}$ denote the inverse of $\Pi_{\mathbb{R}^n}$ on $\mathcal{A}_{a^*}$. Combining (A.1) and Lemma A.2, the derivative of $\Pi_{\mathcal{I}} \circ \Pi_{\mathbb{R}^n}^{-1}$ is given by $(D_{(a,x)}\Pi_{\mathcal{I}})(D_{(a,x)}\Pi_{\mathbb{R}^n})^{-1} = H_{\eta}^{-1}P_{T_x\mathcal{I}}$ and thus has constant rank on $\mathcal{A}_{a^*}$ equal to $\dim \mathcal{I}$. This implies that

$$(1 \times (D_{(x,y)}\pi_{\mathcal{O}})(D_{(x,y)}\pi_{\mathcal{I}})^{-1})(D_{(a,x)}\Pi_{\mathcal{I}})(D_{(a,x)}\Pi_{\mathbb{R}^n})^{-1} : T_a\mathcal{A}_{a^*} \rightarrow T_{(x,y)}\mathcal{W}$$

has constant rank on $\mathcal{A}_{a^*}$ equal to $\dim \mathcal{I} = \dim \mathcal{W}$, by Assumption 1. Consequently, there is a smooth function $\Phi : \mathcal{A}_{a^*} \rightarrow \mathcal{W}$ with $\Phi(a^*) = (x^*, y^*)$. The above showed that its derivative has constant rank equal to $\dim \mathcal{W}$, so Φ is a smooth submersion. By Proposition 4.28 of [48], $\Phi(\mathcal{A}_{a^*}) \subset \mathcal{W}$ is an open subset, so it is a submanifold of dimension $\dim \mathcal{I} = \dim \mathcal{W}$. Similarly, the derivative $D_{(x,y)}\pi_{\mathcal{I}}$ at $(x, y) \in \mathcal{W}$ has maximal rank equal to $\dim \mathcal{I}$, by definition of $\mathcal{W}$. Hence, $\pi_{\mathcal{I}} \circ \Phi$ is also a smooth submersion. After possibly passing to an open subset of $\mathcal{A}_{a^*}$, we can thus assume

$$\mathcal{N}_{(x^*, y^*)} := \Phi(\mathcal{A}_{a^*}) \subset \mathcal{W} \quad \text{and} \quad \mathcal{I}_{x^*} = \pi_{\mathcal{I}}(\mathcal{N}_{(x^*, y^*)}) \subset \mathcal{I}$$

are open submanifolds each of dimension $\dim \mathcal{I}$ and that

- (i) x^* is the global minimizer of the distance to a^* on $\mathcal{I}_{x^*}$, and hence,
- (ii) all $x \in \mathcal{I}_{x^*}$ lie on the same side of $T_{x^*}\mathcal{I}_{x^*}$, considering the tangent space $T_{x^*}\mathcal{I}_{x^*}$ as an affine linear subspace of $\mathbb{R}^n$ with base point x^* .

Now let $a \in \mathcal{A}_{a^*}$ be a point in the neighborhood of a^* . The restricted squared distance function $d_a : \mathcal{N}_{(x^*, y^*)} \rightarrow \mathbb{R}$, $(x, y) \mapsto \frac{1}{2}\|a - \pi_{\mathcal{I}}(x, y)\|^2$ that is minimized on $\mathcal{N}_{(x^*, y^*)}$ is smooth. Its critical points satisfy

$$\langle a - \pi_{\mathcal{I}}(x, y), D_{(x,y)}\pi_{\mathcal{I}}(\dot{x}, \dot{y}) \rangle = 0 \quad \text{for all} \quad (\dot{x}, \dot{y}) \in T_{(x,y)}\mathcal{N}_{(x^*, y^*)}.$$

Since $(x^*, y^*) \in \mathcal{N}_{(x^*, y^*)} \subset \mathcal{W}$, the derivative $D_{(x^*, y^*)}\pi_{\mathcal{I}}$ of $\pi_{\mathcal{I}} : \mathcal{W} \rightarrow \mathcal{I}$ is invertible. Then, for all $(x, y) \in \mathcal{N}_{(x^*, y^*)}$ the image of $D_{(x,y)}\pi_{\mathcal{I}}$ is the whole tangent space $T_x\mathcal{I}$. Consequently, the critical points must satisfy $a - x \perp T_x\mathcal{I}$.

By construction, for every $a \in \mathcal{A}_{a^*}$ there is a unique $(x, y) \in \mathcal{N}_{(x^*, y^*)}$ so that $(a, x, y) \in \mathcal{W}_{\text{GCPP}}$, namely $(x, y) = \Phi(a)$. This implies that (x, y) is the unique critical point of the squared distance function restricted to $\mathcal{N}_{(x^*, y^*)}$.

It only remains to show that for all $a \in \mathcal{A}_{a^*}$, the minimizer of d_a is attained in the interior of $\mathcal{N}_{(x^*, y^*)}$. This entails that the unique critical point is also the unique

global minimizer on $\mathcal{N}_{(x^*, y^*)}$. We show that there exists a $\delta > 0$ such that restricting $\mathcal{A}_{a^*}$ to the open ball $B_\delta(a^*)$ of radius δ centered at a^* in $\mathbb{R}^n$ yields the desired result.

Here is the geometric setup for the rest of the proof, depicted in Figure A.1: let $a \in B_\delta(a^*)$ and let $(x, y) \in \overline{\mathcal{N}_{(x^*, y^*)}}$ be the minimizer for d_a in the closure of $\mathcal{N}_{(x^*, y^*)}$. Then, x is in the closure of $\mathcal{I}_{x^*}$ by continuity of $\pi_{\mathcal{I}}$. Furthermore, let $x' = P(x)$, where $P := P_{T_{x^*}\mathcal{I}_{x^*}}$ is the orthogonal projection onto the tangent space $T_{x^*}\mathcal{I}_{x^*}$ considered as an affine linear subspace of $\mathbb{R}^n$ with base point x^* .

The image of $B_\delta(a^*)$ under P is an open ball $B_\delta(x^*) \cap T_{x^*}\mathcal{I}_{x^*} \subset T_{x^*}\mathcal{I}_{x^*}$. The distance function $h : \mathcal{I}_{x^*} \rightarrow \mathbb{R}$, $x \mapsto \|x - P(x)\|$ measures the “height” of x from the tangent space $T_{x^*}\mathcal{I}_{x^*}$. Now, we can bound the distance from a to x as follows.

$$\begin{aligned} \|a - x\| &= \|a - a^* + a^* - x^* + x^* - x' + x' - x\| \\ &\leq \|a - a^*\| + \|a^* - x^*\| + \|x^* - x'\| + \|x' - x\| \\ &\leq \|a^* - x^*\| + 2\delta + h(x) \\ (A.3) \quad &\leq \|a^* - x^*\| + 2\delta + \max_{z \in \mathcal{P}_\delta(x^*)} h(z), \end{aligned}$$

where the first bound is the triangle inequality, and where $\mathcal{P}_\delta(x^*) \subset \mathcal{I}_{x^*}$ is the preimage of $B_\delta(x^*) \cap T_{x^*}\mathcal{I}_{x^*}$ under the projection P restricted to $\mathcal{I}_{x^*}$.

Now let $x \in \partial\mathcal{I}_{x^*}$ be a point in the boundary of $\mathcal{I}_{x^*}$ and, again, $x' = P(x)$. We want to show that the distance $\|a - x\|$ is strictly larger than (A.3). This would imply that local minimizer for d_a is indeed contained in the interior of $\mathcal{N}_{(x^*, y^*)}$. We have

$$\begin{aligned} \|a - x\| &= \|a - a^* + a^* - x^* + x^* - x' + x' - x\| \\ &\geq \|a^* - x^* + x^* - x' + x' - x\| - \|a - a^*\|. \end{aligned}$$

Note that $x^* = P(a^*)$. Hence we have the orthogonal decomposition

$$a^* - x^* + x^* - x' + x' - x = (\mathbf{1} - P)(a^* - x) + P(a^* - x) = a^* - x.$$

Plugging this into the above, we find

$$\begin{aligned} \|a - x\| &\geq \|(\mathbf{1} - P)(a^* - x) + P(a^* - x)\| - \delta \\ &= \sqrt{\|(\mathbf{1} - P)(a^* - x)\|^2 + \|P(a^* - x)\|^2} - \delta \\ &\geq \sqrt{\|(\mathbf{1} - P)(a^* - x)\|^2 + \epsilon^2} - \delta, \end{aligned}$$

where $\epsilon = \min_{z \in P(\partial\mathcal{I}_{x^*})} \|x^* - z\| > 0$.

Observe that $a^* - x^*$ points from x^* to a^* and likewise $x' - x$ points from x to x' . Therefore, $\langle a^* - x^*, x' - x \rangle \geq 0$ because both point away from the manifold $\mathcal{I}_{x^*}$, which lies entirely on the side of x by assumption (ii) above; see also Figure A.1. Consequently, $\|(\mathbf{1} - P)(a^* - x)\| = \|a^* - x^* + x' - x\| \geq \|a^* - x^*\|$, which implies

$$\|a - x\| \geq \sqrt{\|a^* - x^*\|^2 + \epsilon^2} - \delta = \|a^* - x^*\| + \frac{\epsilon^2}{2\|a^* - x^*\|} - \delta + \mathcal{O}(\epsilon^4).$$

Note that h is a continuous function with $\lim_{x \rightarrow x^*} h(x) = 0$. Hence, as $\delta \rightarrow 0$ the maximum height tends to zero as well. Moreover, ϵ and $\|a^* - x^*\|$ are independent of δ . All this implies that we can choose $\delta > 0$ sufficiently small so that

$$\frac{\epsilon^2}{2\|a^* - x^*\|} - \delta + \mathcal{O}(\epsilon^4) \geq 2\delta + \max_{z \in \mathcal{P}_\delta(x^*)} h(z).$$

It follows that $\|a - x\|$ is lower bounded by (A.3), so that the minimizer for d_a must be contained in the interior of $\mathcal{N}_{(x^*, y^*)}$.

The foregoing shows that on $B_\delta(a^*) \cap \mathcal{A}_{a^*}$, the map $\Pi_{\mathcal{I}} \circ \Pi_{\mathbb{R}^n}^{-1}$ equals the nonlinear projection $P_{\mathcal{I}_{x^*}}$ onto the manifold $\mathcal{I}_{x^*}$. Putting everything together, we obtain

$$\rho_{(a^*, x^*, y^*)}(a) = (\Pi_{\mathcal{O}} \circ \Pi_{\mathbb{R}^n}^{-1})(a) = (\pi_{\mathcal{O}} \circ \pi_{\mathcal{I}}^{-1} \circ \Pi_{\mathcal{I}} \circ \Pi_{\mathbb{R}^n}^{-1})(a) = (\pi_{\mathcal{O}} \circ \pi_{\mathcal{I}}^{-1} \circ P_{\mathcal{I}_{x^*}})(a);$$

the second equality following (A.2). The condition number of this smooth map is given by (1.2). Comparing with Theorem 5.2 and 6.2 concludes the proof. $\square$

- [1] T. J. ABATZOGLOU, *The minimum norm projection on C^2 manifolds in R^n* , Trans. Amer. Math. Soc., 243 (1978), pp. 115–122.
- [2] P.-A. ABSIL, R. MAHONY, AND R. SEPULCHRE, *Optimization Algorithms on Matrix Manifolds*, Princeton University Press, 2008.
- [3] P.-A. ABSIL, R. MAHONY, AND J. TRUMPF, *An extrinsic look at the Riemannian Hessian*, in Geometric Science of Information, Lecture Notes in Computer Science, Springer, Berlin, Heidelberg, 2013.
- [4] D. AMELUNXEN AND P. BÜRGISSEER, *A coordinate-free condition number for convex programming*, SIAM J. Optim., 22 (2012), pp. 1029–1041.
- [5] J.-P. AUBIN AND H. FRANKOWSKA, *Set-Valued Analysis*, Birkhäuser Boston, reprint of the 1990 ed., 2009.
- [6] Z. BAI, J. DEMMEL, J. DONGARRA, A. RUHE, AND H. VAN DER VORST, eds., *Templates for the Solution of Algebraic Eigenvalue Problems: A Practical Guide*, SIAM, Philadelphia, PA, 2000.
- [7] D. BAU AND L. N. TREFETHEN, *Numerical Linear Algebra*, SIAM, 1997.
- [8] A. BELLONI AND R. M. FREUND, *A geometric analysis of Renegar’s condition number, and its interplay with conic curvature*, Math. Program., Ser. A, 119 (2009), pp. 95–107.
- [9] T. BENDORY, Y. C. ELDAR, AND N. BOUMAL, *Non-convex phase retrieval from STFT measurements*, IEEE Trans. Inform. Theory, 64 (2018), pp. 467–484.
- [10] L. BLUM, F. CUCKER, M. SHUB, AND S. SMALE, *Complexity and Real Computation*, Springer-Verlag, New York, 1998.
- [11] N. BOUMAL, *Riemannian trust regions with finite-difference Hessian approximations are globally convergent*, in Geometric Science of Information, F. Nielsen and F. Barbaresco, eds., vol. 2, 2015, pp. 467–475.
- [12] ———, *Nonconvex phase synchronization*, SIAM J. Optim., 26 (2016), pp. 2355–2377.
- [13] ———, *An introduction to optimization on smooth manifolds*. Available online, 2020.
- [14] N. BOUMAL, B. MISHRA, P.-A. ABSIL, AND R. SEPULCHRE, *Manopt, a Matlab toolbox for optimization on manifolds*, J. Mach. Learn. Res., 15 (2014), pp. 1455–1459.
- [15] P. BREIDING AND N. VANNIEUWENHOVEN, *The condition number of join decompositions*, SIAM J. Matrix Anal. Appl., 39 (2018), pp. 287–309.
- [16] ———, *Convergence analysis of Riemannian Gauss–Newton methods and its connection with the geometric condition number*, Appl. Math. Letters, 78 (2018), pp. 42–50.
- [17] ———, *Sensitivity of low-rank matrix recovery*, arXiv, (2021). in preparation.
- [18] P. BÜRGISSEER AND F. CUCKER, *Condition: The Geometry of Numerical Algorithms*, Springer, Heidelberg, 2013.
- [19] N. CHERNOV AND H. MA, *Least squares fitting of quadratic curves and surfaces*, in Computer Vision, S. Yoshida, ed., Nova Science Publishers, 2011.
- [20] D. CHEUNG AND F. CUCKER, *A new condition number for linear programming*, Math. Program., 91 (2001), pp. 163–174.
- [21] D. CHEUNG, F. CUCKER, AND J. PEÑA, *A condition number for multifold conic systems*, SIAM J. Optim., 19 (2008), pp. 261–280.
- [22] A. COULIBALY AND J.-P. CROUZEIX, *Condition numbers and error bounds in convex programming*, Math. Program., 116 (2009), pp. 79–113.
- [23] F. CUCKER AND J. PEÑA, *A primal-dual algorithm for solving polyhedral conic systems with a finite-precision machine*, SIAM J. Optim., 12 (2001), pp. 522–554.
- [24] C. DA SILVA AND F. J. HERRMANN, *Optimization on the hierarchical Tucker manifold – applications to tensor completion*, Linear Algebra Appl., 481 (2015), pp. 131–173.
- [25] J. W. DEMMEL, *The geometry of ill-conditioning*, J. Complexity, 3 (1987), pp. 201–229.
- [26] ———, *On condition numbers and the distance to the nearest ill-posed problem*, Numer. Math., 51 (1987), pp. 251–289.
- [27] ———, *Applied Numerical Linear Algebra*, SIAM, 1997.
- [28] M. DO CARMO, *Riemannian Geometry*, Birkhäuser, 1993.
- [29] A. L. DONTCHEV AND R. T. ROCKAFELLER, *Implicit Functions and Solution Mappings: A View From Variational Analysis*, Springer Monographs in Mathematics, Springer, 2009.
- [30] J. DRAISMA, E. HOROBET, G. OTTAVIANI, B. STURMFELS, AND R. THOMAS, *The Euclidean distance degree of an algebraic variety*, Found. Comput. Math., 16 (2016), pp. 99–149.
- [31] M. EPELMAN AND R. M. FREUND, *Condition number complexity of an elementary algorithm for computing a reliable solution of a conic linear system*, Math. Program., 88 (2000), pp. 451–485.
- [32] ———, *A new condition measure, preconditioners, and relations between different measures of conditioning for conic linear systems*, SIAM J. Optim., 12 (2002), pp. 627–655.
- [33] O. FAUGERAS AND Q. LUONG, *The Geometry of Multiple Images*, MIT Press, Cambridge, MA, 2001.
- [34] F. FEPPON AND P. F. J. LERMUSIAUX, *The extrinsic geometry of dynamical systems tracking nonlinear matrix projections*, SIAM J. Matrix Anal. Appl., 40 (2019), pp. 814–844.
- [35] R. M. FREUND AND F. ORDÓÑEZ, *On an extension of condition number theory to nonconvex convex optimization*, Math. Oper. Res., 30 (2005), pp. 173–194.
- [36] R. M. FREUND AND J. R. VERA, *Condition-based complexity of convex optimization in conic*

- linear form via the ellipsoid algorithm, SIAM J. Optim., 10 (1999), pp. 155–176.
- [37] G. H. GOLUB AND C. VAN LOAN, *Matrix Computations*, The John Hopkins University Press, 4 ed., 2013.
- [38] R. HARTLEY AND A. ZISSERMAN, *Multiple View Geometry in Computer Vision*, Cambridge University Press, Cambridge, second ed., 2003.
- [39] G. HEIDEL AND V. SCHULZ, *A Riemannian trust-region method for low-rank tensor completion*, Numer. Linear Algebra Appl., 25 (2018).
- [40] A. HEYDEN AND K. ÅSTRÖM, *Algebraic properties of multilinear constraints*, Math. Meth. Appl. Sci., 20 (1997), pp. 1135–1162.
- [41] N. HIGHAM, *Accuracy and Stability of Numerical Algorithms*, SIAM, Philadelphia, PA, 1996.
- [42] M. W. HIRSCH, *Differential Topology*, no. 33 in Graduate Text in Mathematics, Springer-Verlag, 1976.
- [43] S. HOLTZ, T. ROHWEDDER, AND R. SCHNEIDER, *On manifolds of tensors of fixed TT-rank*, Numer. Math., 120 (2012), pp. 701–731.
- [44] J. HUMPHERYS, T. J. JARVIS, AND E. J. EVANS, *Foundations of Applied Mathematics*, SIAM, 2017.
- [45] D. KRESSNER, M. STEINLECHNER, AND B. VANDEREYCKEN, *Low-rank tensor completion by Riemannian optimization*, BIT Numer. Math., 54 (2014), pp. 447–468.
- [46] Z. KÚKELOVÁ, *Algebraic Methods in Computer Vision*, PhD thesis, Czech Technical University in Prague, 2013.
- [47] J. M. LEE, *Riemannian Manifolds: Introduction to Curvature*, Springer-Verlag, 1997.
- [48] J. M. LEE, *Introduction to Smooth Manifolds*, Springer, New York, USA, 2 ed., 2013.
- [49] Y. LIU AND W. WANG, *A revisit to least squares orthogonal distance fitting of parametric curves and surfaces*, in Advances in Geometric Modeling and Processing, F. Chen and B. Jüttler, eds., Lecture Notes in Computer Science, 2008.
- [50] MATLAB, *R2017b*. Natick, Massachusetts, 2017.
- [51] S. MAYBANK, *Theory of Reconstruction from Image Motion*, vol. 28 of Springer Series in Information Sciences, Springer-Verlag, Berlin, 1993.
- [52] B. O’NEILL, *Semi-Riemannian Geometry*, Academic Press, 1983.
- [53] ———, *Elementary Differential Geometry*, Elsevier, revised second edition ed., 2001.
- [54] J. PEÑA AND V. ROSCHINA, *A data-independent distance to infeasibility for linear conic systems*, SIAM J. Optim., 30 (2020), pp. 1049–1066.
- [55] P. PETERSEN, *Riemannian Geometry*, vol. 171 of Graduate Texts in Mathematics, Springer, New York, second ed., 2006.
- [56] I. R. PORTEOUS, *The normal singularities of a submanifold*, J. Differ. Geom., 5 (1971).
- [57] M. PSENKA AND N. BOUMAL, *Second-order optimization for tensors with fixed tensor-trains rank*, in OPT2020: 12th Annual Workshop on Optimization for Machine Learning, 2020.
- [58] J. RENEGAR, *Incorporating condition measures into the complexity theory of linear programming*, SIAM J. Optim., 5 (1995), pp. 506–524.
- [59] ———, *Linear programming, complexity theory and elementary functional analysis*, Math. Program., 70 (1995), pp. 279–351.
- [60] J. R. RICE, *A theory of condition*, SIAM J. Numer. Anal., 3 (1966), pp. 287–310.
- [61] M. SPIVAK, *A Comprehensive Introduction to Differential Geometry*, vol. 2, Publish or Perish, Inc., 1999.
- [62] M. STEINLECHNER, *Riemannian optimization for high-dimensional tensor completion*, SIAM J. Sci. Comput., 38 (2016), pp. S461–S484.
- [63] A. SUNG-JOON, *Geometric fitting of parametric curves and surfaces*, J. Inf. Process. Syst., 4 (2008), pp. 153–158.
- [64] R. THOM, *Sur la theorie des enveloppes*, J. Math. Pures Appl., 41 (1962), pp. 177–192.
- [65] A. USCHMAJEW, *Zur Theorie der Niedrigrangapproximation in Tensorprodukten von Hilberträumen*, PhD thesis, TU Berlin, 2013.
- [66] B. VANDEREYCKEN, *Low-rank matrix completion by Riemannian optimization*, SIAM J. Optim., 23 (2013), pp. 1214–1236.
- [67] J. R. VERA, *Ill-posedness and the complexity of deciding existence of solutions to linear programs*, SIAM J. Optim., 6 (1996), pp. 549–569.
- [68] ———, *Geometric measures of convex sets and bounds on problem sensitivity and robustness for conic linear optimization*, Math. Program., 147 (2014), pp. 47–79.
- [69] VISUAL GEOMETRY GROUP, UNIVERSITY OF OXFORD, *Model house data set*, Last accessed 28 september 2020. Access online at <https://www.robots.ox.ac.uk/~vgg/data/data-mview.html>.
- [70] H. WEYL, *On the volume of tubes*, Amer. J. Math., 2 (1939), pp. 461–472.
- [71] J. H. WILKINSON, *Rounding Errors in Algebraic Processes*, Prentice-Hall, Englewood Cliffs, New Jersey, 1963.
- [72] ———, *The Algebraic Eigenvalue Problem*, Oxford University Press, Amen House, London, United Kingdom, 1965.
- [73] H. WOŹNIAKOWSKI, *Numerical stability for solving nonlinear equations*, Numer. Math., 27 (1976/77), pp. 373–390.