

On intrinsic Cramér-Rao bounds for Riemannian submanifolds and quotient manifolds

Nicolas Boumal, *Member, IEEE*

Abstract

We study Cramér-Rao bounds (CRB's) for estimation problems on Riemannian manifolds. In (S.T. Smith, *Covariance, subspace, and intrinsic Cramér-Rao bounds*, IEEE TSP, 53(5):1610–1630, 2005), the author gives intrinsic CRB's in the form of matrix inequalities relating the covariance of estimators and the Fisher information of estimation problems. We focus on estimation problems whose parameter space $\bar{\mathcal{P}}$ is a Riemannian submanifold or a Riemannian quotient manifold of a parent space \mathcal{P} , that is, estimation problems on manifolds with either deterministic constraints or ambiguities. The CRB's in the aforementioned reference would be expressed w.r.t. bases of the tangent spaces to $\bar{\mathcal{P}}$. In some cases though, it is more convenient to express covariance and Fisher information w.r.t. bases of the tangent spaces to \mathcal{P} . We give CRB's w.r.t. such bases expressed in terms of the geodesic distances on the parameter space. The bounds are valid even for singular Fisher information matrices. In two examples, we show how the CRB's for synchronization problems (including a type of sensor network localization problem) differ in the presence or absence of anchors, leading to bounds for estimation on either submanifolds or quotient manifolds with very different interpretations.

Keywords: Cramér-Rao bounds, CRB, Riemannian manifolds, submanifolds, quotient manifolds, intrinsic bounds, estimation bounds, singular Fisher information matrix, singular FIM, graph Laplacian, sensor network localization, synchronization.

Copyright © 2012 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

N. Boumal is with the department of mathematical engineering, ICTEAM Institute, at the Université catholique de Louvain (Belgium).

Contact: nicolasboumal@gmail.com.

I. INTRODUCTION

We study Cramér-Rao bounds (CRB's) for estimation problems on Riemannian manifolds (loosely, nonlinear spaces with a smooth structure such as the sphere for example). In such problems, one would like to estimate a deterministic but unknown parameter θ belonging to a manifold \mathcal{P} , given a measurement y belonging to a probability space \mathcal{M} . The measurement y is a random variable whose probability density function is shaped by θ . It is because y is distributed differently for different θ 's that sampling (observing) y reveals information about θ . We assume some familiarity of the reader with differential geometry. The appendix briefly reviews the tools needed in this paper, fixes notation and provides references to classical introductory textbooks.

Estimation problems on manifolds arise naturally in camera networks pose estimation [1], angular synchronization [2], covariance matrix estimation and subspace estimation [3] and many other applications, see references therein. Cramér-Rao bounds relate the covariance matrix of estimators to the Fisher information matrix (FIM) of an estimation problem through matrix inequalities. The classical results deal with estimation on Euclidean spaces [4]. More recently, a number of authors have established similar bounds in the manifold setting, see [3], [5] and the many references therein. We focus on bounds for unbiased estimators.

More formally, let \mathcal{P} be a Riemannian manifold and \mathcal{M} be a probability space, i.e., a measurable space with nonnegative measure μ such that $\mu(\mathcal{M}) = 1$. We consider an estimation problem on the parameter space \mathcal{P} based on measurements in \mathcal{M} , such that the probability density function of the measurement given a parameter $\theta \in \mathcal{P}$ is $f(\cdot; \theta): \mathcal{M} \rightarrow \mathbb{R}$. Let $L: \mathcal{P} \rightarrow \mathbb{R}$ be the associated log-likelihood function

$$L(\theta) = \log f(y; \theta). \quad (1)$$

The related Fisher information form at θ , $\mathbf{F}: \mathbb{T}_\theta \mathcal{P} \times \mathbb{T}_\theta \mathcal{P} \rightarrow \mathbb{R}$, is defined as (all expectations are taken w.r.t. y):

$$\mathbf{F}[u, v] \triangleq \mathbb{E} \{DL(\theta)[u] \cdot DL(\theta)[v]\}, \quad (2)$$

where $\mathbb{T}_\theta \mathcal{P}$ is the tangent space to \mathcal{P} at θ and $DL(\theta)[u]$ denotes the directional derivative of L at θ along the tangent vector u (see the appendix). The bilinear form \mathbf{F} is symmetric, positive semidefinite. If it is positive definite—that is, $\mathbf{F}[u, u] > 0$ whenever $u \neq 0$ —the Cramér-Rao bounds in [3], which make use of the inverse of the matrix representing \mathbf{F} , apply.

In this paper, we consider estimation problems such that \mathbf{F} is *not necessarily* positive definite. Singularity of \mathbf{F} typically arises when the measurements are not sufficient to determine the parameter, i.e., ambiguities remain. For example, locating a point $p = (x, y, z)$ in space based solely on information about the bearing $p/\|p\|$ is impossible, since nothing is known about the distance between p and the origin. The Fisher information matrix of such a problem would only be positive *semidefinite*.

To resolve these ambiguities, one can proceed in at least two ways. Firstly, one can add constraints on θ , based on additional knowledge about the parameter. By restricting the parameter space to $\bar{\mathcal{P}} \subset \mathcal{P}$, a submanifold of \mathcal{P} , one may hope that the resulting estimation problem is well-posed. For example, if one knows the distance between p and the origin is 1, one should perform the estimation on the sphere $\bar{\mathcal{P}} = \mathbb{S}^2 = \{(x, y, z) : x^2 + y^2 + z^2 = 1\}$ rather than on $\mathcal{P} = \mathbb{R}^3$. Alternatively, one can recognize that the parameter space is made of equivalence classes, that is, sets of parameters that are equally valid estimators for they give rise to the same measurement distribution. In this scenario, one ends up with an estimation problem on a quotient manifold $\bar{\mathcal{P}} = \mathcal{P}/\sim$, where \sim is an equivalence relation on \mathcal{P} stating that $\theta, \theta' \in \mathcal{P}$ are equivalent if they give rise to the same distribution of the measurements. Continuing with our example, all points p with the same bearing $p/\|p\|$ would give rise to the same measurement distribution, hence are indistinguishable and should be grouped into an equivalence class.

The treatment of submanifolds hereafter may also be useful when the FIM is invertible. In that scenario, one is interested in studying the Cramér-Rao bounds of the original problem, and the effect on those bounds caused by incorporating additional knowledge about θ .

The direct way to address these issues would be to work on the smaller space $\bar{\mathcal{P}}$ directly, writing down Fisher information and covariance with respect to bases of the tangent spaces to $\bar{\mathcal{P}}$, leading to Cramér-Rao bounds according to [3]. However, we argue that the tangent spaces of \mathcal{P} sometimes make more sense to the user: that is why the problem was defined on \mathcal{P} rather than $\bar{\mathcal{P}}$ to begin with. Furthermore, when $\bar{\mathcal{P}}$ is a quotient manifold, its tangent spaces are rather abstract objects to work with. It is hence desirable to have equivalent Cramér-Rao bounds expressed as matrix inequalities w.r.t. bases of tangent spaces of \mathcal{P} instead. This is what the theorems in this communication achieve. The present work is based on [3] and derives its consequences for unbiased estimators in the presence of indeterminacies (ambiguities) or under additional constraints.

The case of constrained Cramér-Rao bounds, that is, estimation on Riemannian submanifolds of \mathbb{R}^d , has been studied extensively [6], [7], [8]. Notably, in [8], the authors describe $\bar{\mathcal{P}}$ through a set of equality constraints and they express the covariance in terms of distances in the embedding Euclidean space \mathbb{R}^d . In this paper, we more generally consider Riemannian submanifolds of any Riemannian manifold \mathcal{P} .

Furthermore, for the simple versions of the CRB's, only an orthogonal projector from the tangent spaces of \mathcal{P} to those of $\bar{\mathcal{P}}$ are required. More importantly, the covariance matrix in the proposed bounds is expressed in terms of the Riemannian, or geodesic, distance on $\bar{\mathcal{P}}$, which may be more natural for a number of applications.

The case of CRB's for estimation problems with singular FIM has also been investigated extensively [7], [9], [10]. The classical remedy is to use the Moore-Penrose pseudoinverse, hereafter referred to as the pseudoinverse, of the FIM instead of the inverse in the CRB (in what follows, we use the notation A^\dagger to denote the pseudoinverse of a matrix $A \in \mathbb{R}^{n \times m}$, that is, the unique matrix in $\mathbb{R}^{m \times n}$ such that $AA^\dagger A = A$, $A^\dagger AA^\dagger = A^\dagger$ and the matrices AA^\dagger and $A^\dagger A$ are symmetric). When the singularity is due to indeterminacies (a notion we make precise in Section III), Xavier and Barroso [10] showed a nice interpretation of the role of the pseudoinverse by recasting the estimation problem on a Riemannian quotient manifold $\bar{\mathcal{P}}$. In the latter reference, the authors give a geometric interpretation for the kernel of the FIM and propose a CRB-type bound they name IVLB [5] for the variance of unbiased estimators for such problems. In their bound, the possible curvature of $\bar{\mathcal{P}}$ is captured through a single number: an upperbound on the sectional curvatures of $\bar{\mathcal{P}}$. In comparison, since the present results are based on [3], the proposed bounds concern the whole covariance matrix (the trace of which coincides with the variance). The pseudoinverse of the FIM appears naturally through the same manipulations as in [10]. The additional curvature terms in the CRB (Section IV) take the whole Riemannian curvature tensor into account. This is especially useful when $\bar{\mathcal{P}}$ is flat or almost flat in most directions but has significant curvature in a few directions, which happens naturally for product spaces. In such scenarios, the IVLB tends to be overly optimistic, i.e., less restrictive—hence less informative—because it has to assume maximum curvature in all directions. In comparison, the bounds derived here based on [3] are able to capture complex curvature structures if need be.

A lot of work has been accomplished to gain other types of geometric insight into the Cramér-Rao bounds. Scharf and McWhorter [11] study the CRB for multi-parameter estimation problems where the parameters are partitioned into two groups. Their study shows how the angles between the subspaces spanned by the sensitivity vectors of the two groups are related to the CRB. Another trend of geometric interpretation comes from the field of information geometry, where one endows the parameter space with Fisher information as a metric, turning the parameter space (typically, \mathbb{R}^n) into a Riemannian manifold [12]. The metric “bends” space, so that the larger the Riemannian distance between two parameters, the easier it is for an estimator to distinguish them.

The CRB's presented in this communication hold *at large SNR*. The origin of this provision is double.

Firstly, the definition of covariance on a manifold uses the logarithmic map on that manifold, which is only locally well-defined. It is thus necessary to require the noise level to be low enough so that the estimator $\hat{\theta}$ of a parameter θ will, with high probability, belong to a neighborhood of θ where the logarithm is well-defined. This can be relevant even on flat manifolds such as the circle $SO(2)$ for example, which is compact. Secondly, on curved manifolds, the proof of the main theorem in [3] relies on truncated Taylor expansions. Those are legitimate only at large enough SNR so that typical errors are small compared to the scale at which curvature becomes a dominant feature.

Let $e = \{e_1, \dots, e_d\}$ be an orthonormal basis of $T_\theta \mathcal{P}$ w.r.t. the Riemannian metric $\langle \cdot, \cdot \rangle_\theta$. The Fisher information matrix of the estimation problem on \mathcal{P} w.r.t. the basis e is a $d \times d$ symmetric matrix defined by:

$$(F_e)_{ij} = \mathbf{F}[e_i, e_j] = \mathbb{E} \{DL(\theta)[e_i] \cdot DL(\theta)[e_j]\}, \quad (3)$$

where L is the log-likelihood function (1). The covariance matrix C_e w.r.t. the basis e will be defined separately for the submanifold (Section II) and the quotient manifold (Section III) cases, then F_e and C_e will be linked through matrix inequalities. At first, we will neglect curvature terms that may appear due to the possible curvature of $\bar{\mathcal{P}}$. This will result in simple statements (Theorems 2 and 3). These are practically useful because the curvature terms are often negligible at large SNR [3]. Then, we will establish the CRB's including curvature terms (Section IV). Finally, we will illustrate the usage of these theorems through two examples (Section V).

II. RIEMANNIAN SUBMANIFOLDS

Let us consider the constrained estimation problem on the space $\bar{\mathcal{P}} \subset \mathcal{P}$, a Riemannian submanifold of \mathcal{P} , such that $\theta \in \bar{\mathcal{P}}$ and for which the log-likelihood function $\bar{L} = L|_{\bar{\mathcal{P}}}$ is the restriction of L to $\bar{\mathcal{P}}$. This situation arises when one adds supplementary constraints on the parameter θ . For example, some of the target parameters are known or deterministically related. The Fisher information is simply the restriction $\bar{\mathbf{F}} = \mathbf{F}|_{T_\theta \bar{\mathcal{P}} \times T_\theta \bar{\mathcal{P}}}$. We assume $\bar{\mathbf{F}}$ is invertible, i.e., the added constraints fix possible ambiguities in the estimation problem. Figure 1 depicts the situation.

Let $\hat{\theta}$ be any unbiased estimator for the estimation problem, that is, $\hat{\theta}: \mathcal{M} \rightarrow \bar{\mathcal{P}}$ maps every possible realization of the measurement y to a parameter $\hat{\theta}(y)$ and has zero bias:

$$\forall \theta \in \mathcal{P}, \quad b(\theta) = \mathbb{E} \left\{ \text{Log}_\theta(\hat{\theta}(y)) \right\} = 0, \quad (4)$$

where $\text{Log}_\theta: \bar{\mathcal{P}} \rightarrow T_\theta \bar{\mathcal{P}}$ is the logarithmic map at θ on $\bar{\mathcal{P}}$ (see the appendix). For example, on a Euclidean space, $\text{Log}_\theta(\hat{\theta}(y)) = \hat{\theta}(y) - \theta$. For conciseness, we often write $\hat{\theta}$ to mean $\hat{\theta}(y)$. The covariance matrix




Fig. 1. $\bar{\mathcal{P}}$ is a Riemannian submanifold of \mathcal{P} . We consider estimation problems for which the parameter to estimate is θ , a point of $\bar{\mathcal{P}}$. In this drawing, for simplicity, we chose $\mathcal{P} = \mathbb{R}^2$. The vectors $e = (e_1, e_2)$ form an orthonormal basis of $T_\theta \mathcal{P} \equiv \mathbb{R}^2$, while $\bar{e} = (\bar{e}_1)$ is an orthonormal basis of the tangent space $T_\theta \bar{\mathcal{P}}$. The operator P_θ projects vectors of $T_\theta \mathcal{P}$ orthogonally onto $T_\theta \bar{\mathcal{P}}$. We express the Cramér-Rao bounds for such problems in terms of the basis e , which at times may be more convenient than defining a basis \bar{e} for each point θ .

of $\hat{\theta}$ w.r.t. the basis e is defined as:

$$(C_e)_{ij} = \mathbb{E} \left\{ \left\langle \text{Log}_\theta(\hat{\theta}), e_i \right\rangle_\theta \cdot \left\langle \text{Log}_\theta(\hat{\theta}), e_j \right\rangle_\theta \right\}, \quad (5)$$

where, as always in this paper, the expectation is taken w.r.t. the measurements $y \sim f(y; \theta)$. The goal is to link C_e and F_e through a matrix inequality.

Let $\bar{e} = \{\bar{e}_1, \dots, \bar{e}_{\bar{d}}\}$ be an orthonormal basis of $T_\theta \bar{\mathcal{P}} \subset T_\theta \mathcal{P}$ w.r.t. the Riemannian metric $\langle \cdot, \cdot \rangle_\theta$. Let E be the $\bar{d} \times d$ matrix such that $E_{ij} = \langle \bar{e}_i, e_j \rangle_\theta$. E is orthonormal: $EE^\top = I_{\bar{d}}$, but in general, $P_e \triangleq E^\top E \neq I_d$. Furthermore, let $P_\theta: T_\theta \mathcal{P} \rightarrow T_\theta \bar{\mathcal{P}}$ be the orthogonal projector onto $T_\theta \bar{\mathcal{P}}$. Clearly, P_e is the matrix representation of P_θ w.r.t. the basis e , that is: $\langle P_\theta e_i, e_j \rangle = (P_e)_{ij}$.

A direct application of the CRB's in [3] to the estimation problem on $\bar{\mathcal{P}}$ would link the covariance matrix $C_{\bar{e}}$ of $\hat{\theta}$ and the inverse Fisher information matrix $\bar{F}_{\bar{e}}^{-1}$ w.r.t. the basis \bar{e} . More precisely,

$$\begin{aligned} (C_{\bar{e}})_{ij} &= \mathbb{E} \left\{ \left\langle \text{Log}_\theta(\hat{\theta}), \bar{e}_i \right\rangle_\theta \cdot \left\langle \text{Log}_\theta(\hat{\theta}), \bar{e}_j \right\rangle_\theta \right\}, \\ (\bar{F}_{\bar{e}})_{ij} &= \mathbb{E} \left\{ D\bar{L}(\theta)[\bar{e}_i] \cdot D\bar{L}(\theta)[\bar{e}_j] \right\}, \\ C_{\bar{e}} &\succeq \bar{F}_{\bar{e}}^{-1} + \text{curvature terms}. \end{aligned} \quad (6)$$

We argue that it is sometimes convenient to work with C_e and F_e directly, to avoid the necessity to define and work with the basis \bar{e} . This is what the next theorem achieves, right after we establish a technical lemma.

Lemma 1. Let $E \in \mathbb{R}^{\bar{d} \times d}$, $A \in \mathbb{R}^{d \times d}$, $B \in \mathbb{R}^{\bar{d} \times \bar{d}}$, with $\bar{d} \leq d$, $A = A^\top$, $B = B^\top$ and $EE^\top = I_{\bar{d}}$, i.e., E is orthonormal. Further assume that $\ker E \subset \ker A$. Then,

$$EAE^\top \succeq B \quad \Rightarrow \quad A \succeq E^\top BE. \quad (7)$$

Proof: Since $\mathbb{R}^d = \text{im } E^\top \oplus \ker E$, for all $x \in \mathbb{R}^d$, there exist unique vectors $y \in \mathbb{R}^{\bar{d}}$ and $z \in \mathbb{R}^d$ such that $x = E^\top y + z$ and $Ez = 0$. It follows that:

$$\begin{aligned} x^\top Ax &= y^\top EAE^\top y + z^\top Az + 2y^\top EAz \\ (\text{since } Ez = 0 \Rightarrow Az = 0) &= y^\top EAE^\top y \\ (\text{since } EAE^\top \succeq B) &\geq y^\top By \\ (\text{since } Ex = EE^\top y + Ez = y) &= x^\top E^\top BE x. \end{aligned}$$

This holds for all x , hence $A \succeq E^\top BE$. ■

Theorem 2 (CRB on submanifolds). Given any unbiased estimator $\hat{\theta}$ for the estimation problem on the Riemannian submanifold $\bar{\mathcal{P}}$ with log-likelihood $\bar{L} = L|_{\bar{\mathcal{P}}}$ (1), at large SNR, the $d \times d$ covariance matrix C_e (5) and the $d \times d$ Fisher information matrix F_e (3) obey the matrix inequality (assuming $\text{rank}(P_e F_e P_e) = \bar{d}$):

$$C_e \succeq (P_e F_e P_e)^\dagger + \text{curvature terms}, \quad (8)$$

where the $d \times d$ matrix $P_e = E^\top E$ is the orthogonal projector from $\mathbb{T}_\theta \mathcal{P}$ to $\mathbb{T}_\theta \bar{\mathcal{P}}$ w.r.t. the basis e and \dagger denotes Moore-Penrose pseudoinversion. Furthermore, the spectrum of $(P_e F_e P_e)^\dagger$ is the spectrum of \bar{F}_e^{-1} with $d - \bar{d}$ additional zeroes. In particular, neglecting curvature terms:

$$\text{trace}(C_e) = \text{trace}(C_{\bar{e}}) \geq \text{trace}(\bar{F}_e^{-1}) = \text{trace}((P_e F_e P_e)^\dagger).$$

Proof: Since $\hat{\theta} \in \bar{\mathcal{P}}$, $\text{Log}_\theta(\hat{\theta}) \in \mathbb{T}_\theta \bar{\mathcal{P}}$. Consequently, for all $u \in \mathbb{T}_\theta \mathcal{P}$, $\langle \text{Log}_\theta(\hat{\theta}), u \rangle_\theta = \langle \text{Log}_\theta(\hat{\theta}), P_\theta u \rangle_\theta$, where $P_\theta u$ is the orthogonal projection of u on $\mathbb{T}_\theta \bar{\mathcal{P}}$. The orthogonal projection of the basis vector e_i on $\mathbb{T}_\theta \bar{\mathcal{P}}$ expands in the basis \bar{e} as $P_\theta e_i = \sum_j \langle \bar{e}_j, e_i \rangle_\theta \bar{e}_j = \sum_j E_{ji} \bar{e}_j$. Then, by bilinearity, $(C_e)_{ij} = \sum_{k,\ell} E_{ki} E_{\ell j} (C_{\bar{e}})_{k\ell}$. In matrix form,

$$C_e = E^\top C_{\bar{e}} E. \quad (9)$$

Since $EE^\top = I_{\bar{d}}$, it also holds that $C_{\bar{e}} = EC_e E^\top$. The vectors of \bar{e} expand in the basis e as $\bar{e}_i = \sum_j \langle \bar{e}_i, e_j \rangle_\theta e_j = \sum_j E_{ij} e_j$. By bilinearity again, $(\bar{F}_e)_{ij} = \sum_{k,\ell} E_{ik} E_{j\ell} (F_e)_{k\ell}$. In matrix form,

$$\bar{F}_e = EF_e E^\top. \quad (10)$$

Notice that the assumption $\text{rank}(P_e F_e P_e) = \bar{d}$ is equivalent to the assumption that \bar{F}_e is invertible. Then, substituting in (6), we find $EC_e E^\top \succeq (EF_e E^\top)^{-1}$. Since $\ker C_e = \ker(E^\top C_{\bar{e}} E) \supset \ker E$, Lemma 1 applies and it follows that (neglecting curvature terms):

$$C_e \succeq E^\top (EF_e E^\top)^{-1} E. \quad (11)$$

Finally, from the definition of pseudoinverse, it is easily checked that

$$E^\top (EF_e E^\top)^{-1} E = (E^\top EF_e E^\top E)^\dagger. \quad (12)$$

Since $P_e = E^\top E$, this concludes the proof of the main part.

We now establish the spectrum property. Since \bar{F}_e^{-1} is symmetric positive definite, there exist a diagonal matrix D and an orthogonal matrix U of size $\bar{d} \times \bar{d}$ such that $\bar{F}_e^{-1} = UDU^\top$. Hence,

$$(P_e F_e P_e)^\dagger = E^\top UDU^\top E = V \begin{pmatrix} D & \\ & 0 \end{pmatrix} V^\top, \quad (13)$$

with $V = \begin{pmatrix} E^\top U & (E^\top U)^\perp \end{pmatrix}$ a $d \times d$ orthogonal matrix. The trace property follows easily (neglecting curvature terms):

$$\text{trace}(C_e) = \text{trace}(E^\top C_{\bar{e}} E) = \text{trace}(C_{\bar{e}}) \geq \text{trace}(\bar{F}_e^{-1}) = \text{trace}((P_e F_e P_e)^\dagger). \quad (14)$$

■

The trace property is especially interesting, as it bounds the variance of the estimator $\hat{\theta}$, expressed w.r.t. the Riemannian distance dist on $\bar{\mathcal{P}}$:

$$\text{trace}(C_e) = \text{trace}(C_{\bar{e}}) = \mathbb{E} \left\{ \|\text{Log}_\theta(\hat{\theta})\|^2 \right\} = \mathbb{E} \left\{ \text{dist}^2(\theta, \hat{\theta}) \right\} \triangleq \text{var}_{\hat{\theta}}(\theta).$$

Here is one way of interpreting the bound (8). Expand the random error vector $\text{Log}_\theta(\hat{\theta}) = \sum_i x_i e_i$ with random coefficients x_i . From the definition, $(C_e)_{ii} = \mathbb{E} \{x_i^2\}$. Then, equation (8) implies $\mathbb{E} \{x_i^2\} \geq (P_e F_e P_e)_{ii}^\dagger$, which limits how well the i^{th} coordinate can be estimated. For example, when $\bar{\mathcal{P}}$ is Euclidean, $\text{Log}_\theta(\hat{\theta}) = \hat{\theta} - \theta$ and $\mathbb{E} \{x_i^2\} = \mathbb{E} \{(\hat{\theta}_i - \theta_i)^2\}$.

Notice that it is not necessary to explicitly construct a basis \bar{e} in order to use Theorem 2. Indeed, the orthogonal projector P_e is often easy to compute without requiring an explicit factorization as $E^\top E$. For example, the orthogonal projector from \mathbb{R}^3 onto the tangent space to the sphere \mathbb{S}^2 at θ , denoted $T_\theta \mathbb{S}^2$, w.r.t. the canonical basis of \mathbb{R}^3 is simply $P_e = I_3 - \theta\theta^\top$, where I_3 is the 3×3 identity matrix. This is fortunate since, because of the hairy ball theorem, it is impossible to define bases \bar{e} of $T_\theta \mathbb{S}^2$ for all θ in a smooth way, making it rather inconvenient to work with such bases.

III. RIEMANNIAN QUOTIENT MANIFOLDS

Whenever two parameters $\theta, \theta' \in \mathcal{P}$ give rise to the same measurement distribution, they are indistinguishable, in the sense that no argument based on the observed measurement can be used to favor one parameter over the other as estimator. This observation motivates the definition of the following equivalence relation (remember the definitions of f and L (1)):

$$\theta \sim \theta' \Leftrightarrow f(\cdot, \theta) \equiv f(\cdot, \theta') \text{ a.e. on } \mathcal{M}. \quad (15)$$

The quotient space $\bar{\mathcal{P}} = \mathcal{P}/\sim$ —that is, the set of equivalence classes—then becomes the natural parameter space on which the estimation should be performed. Figures 2 and 3, courtesy of the authors of [10], depict the concept of quotient manifold and of the related basic objects we introduce hereafter, namely submersions and horizontal/vertical spaces. See also the appendix.




Fig. 2. The parameter space \mathcal{P} is partitioned into equivalence classes, called fibers. The Riemannian submersion π maps each $\theta \in \mathcal{P}$ to its corresponding equivalence class $[\theta] \in \bar{\mathcal{P}}$. The space of equivalence classes is the quotient space $\bar{\mathcal{P}} = \mathcal{P}/\sim$, also a Riemannian manifold. *Figure courtesy of [10].*

We now consider the mapping π from \mathcal{P} to $\bar{\mathcal{P}}$, which maps each parameter θ to its equivalence class $[\theta]$,

$$\pi: \mathcal{P} \rightarrow \bar{\mathcal{P}}: \theta \mapsto \pi(\theta) = [\theta] \triangleq \{\theta' \in \mathcal{P} : \theta' \sim \theta\}, \quad (16)$$




Fig. 3. Each fiber $\pi(\theta) = [\theta]$ is a Riemannian submanifold of \mathcal{P} . The tangent space to a fiber at θ is the vertical space V_θ . The orthogonal complement of V_θ in $T_\theta\mathcal{P}$ is the horizontal space H_θ . The differential of π , noted $D\pi(\theta)$, is an isometry between H_θ and the abstract tangent space $T_{[\theta]}\bar{\mathcal{P}}$. This makes it convenient to represent abstract tangent vectors to $\bar{\mathcal{P}}$ as horizontal vectors. *Figure courtesy of [10].*

and concentrate on the case where π is a Riemannian submersion [13][14]. That is, $\bar{\mathcal{P}}$ is a Riemannian quotient manifold of \mathcal{P} . In particular, $[\theta]$ is a Riemannian submanifold of \mathcal{P} (a *fiber*). The log-likelihood function $\bar{L}: \bar{\mathcal{P}} \rightarrow \mathbb{R}$ is well-defined by $\bar{L}([\theta]) \triangleq L(\theta)$.

The tangent space to $[\theta]$ at θ , named the vertical space V_θ , is a subspace of the tangent space $T_\theta\mathcal{P}$. The orthogonal complement of the vertical space, named the horizontal space H_θ , is such that $T_\theta\mathcal{P} = H_\theta \oplus V_\theta$. The pushforward $D\pi(\theta): T_\theta\mathcal{P} \rightarrow T_{[\theta]}\bar{\mathcal{P}}$ of a Riemannian submersion induces a metric on the abstract tangent space $T_{[\theta]}\bar{\mathcal{P}}$:

$$\forall u, v \in H_\theta, \quad \langle D\pi(\theta)[u], D\pi(\theta)[v] \rangle_{[\theta]} \triangleq \langle u, v \rangle_\theta. \quad (17)$$

The definition of Riemannian submersion ensures that this is well-defined [14]. We mention two useful properties:

$$\ker D\pi(\theta) = V_\theta, \quad \text{and} \quad (18)$$

$$D\pi(\theta)|_{H_\theta}: H_\theta \rightarrow T_{[\theta]}\bar{\mathcal{P}} \text{ is an isometry.} \quad (19)$$

Let $[\hat{\theta}]: \mathcal{M} \rightarrow \bar{\mathcal{P}}$ be any unbiased estimator for the present problem. Define the covariance matrix of

$[\hat{\theta}]$ w.r.t. the basis e as:

$$(C_e)_{ij} = \mathbb{E} \left\{ \langle \xi, e_i \rangle_\theta \cdot \langle \xi, e_j \rangle_\theta \right\}, \text{ with}$$

$$\xi = (D\pi(\theta)|_{H_\theta})^{-1} [\text{Log}_{[\theta]}([\hat{\theta}])]. \quad (20)$$

The error vector ξ is the shortest horizontal vector at θ such that $\text{Exp}_\theta(\xi) \in [\hat{\theta}]$. The exponential map is the inverse of the logarithmic map, see the appendix. On a Euclidean space, $\text{Exp}_\theta(\xi) = \theta + \xi$.

Let $\bar{e} = (\bar{e}_1, \dots, \bar{e}_{\bar{d}})$ be an orthonormal basis of $T_{[\theta]}\bar{\mathcal{P}}$. A direct application of the CRB's in [3] to the estimation problem on $\bar{\mathcal{P}}$ would link the covariance matrix $C_{\bar{e}}$ of $[\hat{\theta}]$ and the inverse Fisher information matrix $\bar{F}_{\bar{e}}^{-1}$ w.r.t. the basis \bar{e} . More precisely,

$$(C_{\bar{e}})_{ij} = \mathbb{E} \left\{ \left\langle \text{Log}_{[\theta]}([\hat{\theta}]), \bar{e}_i \right\rangle_{[\theta]} \cdot \left\langle \text{Log}_{[\theta]}([\hat{\theta}]), \bar{e}_j \right\rangle_{[\theta]} \right\},$$

$$(\bar{F}_{\bar{e}})_{ij} = \mathbb{E} \left\{ D\bar{L}([\theta])[\bar{e}_i] \cdot D\bar{L}([\theta])[\bar{e}_j] \right\},$$

$$C_{\bar{e}} \succeq \bar{F}_{\bar{e}}^{-1} + \text{curvature terms}. \quad (21)$$

Since $T_{[\theta]}\bar{\mathcal{P}}$ is an abstract space, we argue that it is often convenient to work with the more concrete objects C_e and F_e instead.

Theorem 3 (CRB on quotient manifolds). *Given any unbiased estimator $[\hat{\theta}]$ for the estimation problem on the Riemannian quotient manifold $\bar{\mathcal{P}} = \mathcal{P}/\sim$ (15) with log-likelihood L (1), at large SNR, the $d \times d$ covariance matrix C_e (20) and the $d \times d$ Fisher information matrix F_e (3) obey the matrix inequality (assuming $\text{rank}(F_e) = \bar{d}$):*

$$C_e \succeq F_e^\dagger + \text{curvature terms}, \quad (22)$$

where \dagger denotes Moore-Penrose pseudoinversion. Furthermore, the spectrum of F_e^\dagger is the spectrum of $\bar{F}_{\bar{e}}^{-1}$ with $d - \bar{d}$ additional zeroes. In particular, neglecting curvature terms:

$$\text{trace}(C_e) = \text{trace}(C_{\bar{e}}) \geq \text{trace}(\bar{F}_{\bar{e}}^{-1}) = \text{trace}(F_e^\dagger). \quad (23)$$

Proof: It is convenient to introduce the orthonormal basis of H_θ related to \bar{e} as $\tilde{e} = (\tilde{e}_1, \dots, \tilde{e}_{\bar{d}})$, with $\bar{e}_i = D\pi(\theta)[\tilde{e}_i]$. The $\bar{d} \times d$ matrix E such that $E_{ij} = \langle \tilde{e}_i, e_j \rangle_\theta$ will prove useful. E is orthonormal: $EE^\top = I_{\bar{d}}$, but in general, $E^\top E \neq I_d$.

Let us denote the orthogonal projection of $u \in T_\theta\mathcal{P}$ onto the horizontal space H_θ as $PH_\theta u$. Since $\xi = (D\pi(\theta)|_{H_\theta})^{-1} [\text{Log}_{[\theta]}([\hat{\theta}])]$ is a horizontal vector, $\langle \xi, u \rangle_\theta = \langle \xi, PH_\theta u \rangle$. Furthermore, $D\pi(\theta)[PH_\theta u] =$

$D\pi(\theta)[u]$. Then, using the fact that $D\pi(\theta)|_{H_\theta}$ is an isometry, it follows that

$$(C_e)_{ij} = \mathbb{E} \left\{ \langle \xi, e_i \rangle_\theta \cdot \langle \xi, e_j \rangle_\theta \right\} \quad (24)$$

$$= \mathbb{E} \left\{ \langle \xi, PH_\theta e_i \rangle_\theta \cdot \langle \xi, PH_\theta e_j \rangle_\theta \right\} \quad (25)$$

$$= \mathbb{E} \left\{ \left\langle \text{Log}_{[\theta]}([\hat{\theta}]), D\pi(\theta)[e_i] \right\rangle_{[\theta]} \cdot \left\langle \text{Log}_{[\theta]}([\hat{\theta}]), D\pi(\theta)[e_j] \right\rangle_{[\theta]} \right\}. \quad (26)$$

The vector $D\pi(\theta)[e_i] \in T_{[\theta]}\bar{\mathcal{P}}$ expands in the basis \bar{e} as $D\pi(\theta)[e_i] = \sum_j E_{ji} \bar{e}_j$. Indeed,

$$\langle D\pi(\theta)[e_i], \bar{e}_j \rangle_{[\theta]} = \langle D\pi(\theta)[e_i], D\pi(\theta)[\bar{e}_j] \rangle_{[\theta]} = \langle e_i, \bar{e}_j \rangle_\theta.$$

It follows that $(C_e)_{ij} = \sum_{k,\ell} E_{ki} E_{\ell j} (C_{\bar{e}})_{k\ell}$. In matrix form:

$$C_e = E^\top C_{\bar{e}} E. \quad (27)$$

Since $EE^\top = I_{\bar{e}}$, it also holds that $C_{\bar{e}} = EC_e E^\top$.

We now similarly link F_e and $\bar{F}_{\bar{e}}$. In doing so, we exploit the fact that the gradient $\text{grad } L(\theta)$ is a horizontal vector. This stems from the fact that the log-likelihood function L is constant over fibers (equivalence classes).

$$\begin{aligned} (F_e)_{ij} &= \mathbb{E} \{ DL(\theta)[e_i] \cdot DL(\theta)[e_j] \} \\ &= \mathbb{E} \left\{ \langle \text{grad } L(\theta), e_i \rangle_\theta \cdot \langle \text{grad } L(\theta), e_j \rangle_\theta \right\} \\ &= \mathbb{E} \left\{ \langle \text{grad } L(\theta), PH_\theta e_i \rangle_\theta \cdot \langle \text{grad } L(\theta), PH_\theta e_j \rangle_\theta \right\} \\ &\text{(expand } PH_\theta e_i \text{ and } PH_\theta e_j \text{ in the basis } \bar{e}) \\ &= \sum_{k,\ell} E_{ki} E_{\ell j} \mathbb{E} \left\{ \langle \text{grad } L(\theta), \bar{e}_k \rangle_\theta \cdot \langle \text{grad } L(\theta), \bar{e}_\ell \rangle_\theta \right\} \\ &= \sum_{k,\ell} E_{ki} E_{\ell j} \mathbb{E} \left\{ \langle D\pi(\theta)[\text{grad } L(\theta)], \bar{e}_k \rangle_{[\theta]} \cdot \langle D\pi(\theta)[\text{grad } L(\theta)], \bar{e}_\ell \rangle_{[\theta]} \right\} \\ &= \sum_{k,\ell} E_{ki} E_{\ell j} \mathbb{E} \left\{ \langle \text{grad } \bar{L}([\theta]), \bar{e}_k \rangle_{[\theta]} \cdot \langle \text{grad } \bar{L}([\theta]), \bar{e}_\ell \rangle_{[\theta]} \right\} \\ &= \sum_{k,\ell} E_{ki} E_{\ell j} \mathbb{E} \left\{ D\bar{L}([\theta])[\bar{e}_k] \cdot D\bar{L}([\theta])[\bar{e}_\ell] \right\} \\ &= \sum_{k,\ell} E_{ki} E_{\ell j} (\bar{F}_{\bar{e}})_{k\ell}. \end{aligned}$$

In matrix form,

$$F_e = E^\top \bar{F}_{\bar{e}} E. \quad (28)$$

Notice that the assumption $\text{rank}(F_e) = \bar{d}$ is equivalent to the assumption that $\bar{F}_{\bar{e}}$ is invertible. The latter equation thus highlights that $\ker F_e = \ker E$, which makes sense since $\ker E$ is the vertical space V_θ (more precisely, it is the space of coordinate vectors of vertical vectors w.r.t. the basis e). Again, by orthonormality of E , it also holds that $\bar{F}_{\bar{e}} = EF_eE^\top$. Combining these rules, it follows that:

$$F_e = E^\top EF_eE^\top E. \quad (29)$$

Applying Lemma 1 to the inequality (21) and using arguments similar to the proof of Theorem 2 finally yields:

$$C_e \succeq F_e^\dagger + \text{curvature terms}, \quad (30)$$

since

$$E^\top (EF_eE^\top)^{-1} E = (E^\top EF_eE^\top E)^\dagger = F_e^\dagger. \quad (31)$$

The spectrum and trace properties follow directly, see proof of Theorem 2. ■

Again, there is no need to construct bases \tilde{e} or \bar{e} in order to use Theorem 3. Notice that it still holds that $\text{trace}(C_e) = \text{trace}(C_{\bar{e}}) = \mathbb{E} \{ \|\xi\|_\theta^2 \} = \mathbb{E} \{ \text{dist}^2([\theta], [\hat{\theta}]) \}$, where dist is the Riemannian distance on $\bar{\mathcal{P}}$, since $D\pi(\theta)|_{H_\theta}$ is an isometry.

IV. INCLUDING CURVATURE TERMS

The intrinsic Cramér-Rao bounds developed in [3] include special terms accounting for the possible curvature of the parameter space $\bar{\mathcal{P}}$ —for intuition on flatness and curvature, see the appendix; for references, see [13], [15]. As noted in [3, p. 1615], “*the proof of the CRB in Euclidean spaces relies on the fact that $\partial/\partial\theta(\hat{\theta} - \theta) = -I$. However, for arbitrary Riemannian manifolds, $\nabla \text{Log}_\theta(\hat{\theta}) = -I + \text{second- and higher-order terms involving the manifold’s sectional and Riemannian curvatures,}$ ” where ∇ denotes the Riemannian connection on $\bar{\mathcal{P}}$, i.e., a notion of derivatives of vector fields on Riemannian manifolds. In a nutshell: the origin of curvature terms in intrinsic CRB’s is the non-commutativity of derivatives of vector fields on curved spaces. As further noted [3, p. 1618], “*the significance of the curvature terms is an open question that depends on the specific application; however, as noted earlier, these terms become negligible for small errors and biases.*”*

The curvature terms vanish if $\bar{\mathcal{P}}$ is flat, that is, if it is locally isometric to a Euclidean space. In such cases, theorems 2 and 3 suffice. When $\bar{\mathcal{P}}$ is not flat, the curvature terms may nevertheless often

be neglected for high enough signal-to-noise ratio. The argument developed in [3] to that end concludes that neglecting the curvature terms is legitimate as soon as estimation errors obey

$$\text{dist}(\theta, \hat{\theta}) \ll \frac{1}{\sqrt{K_{\max}}}, \quad (32)$$

where K_{\max} is an upperbound on the absolute value of the sectional curvatures of $\bar{\mathcal{P}}$ at θ . Intuitively, this is the scale at which curvature plays a minor role.

Condition (32) involves an upperbound on the sectional curvature of $\bar{\mathcal{P}}$. As a consequence, it may be overly restrictive for parameter spaces which have small curvature in most directions, and large curvature in a few. An important class of such spaces consists in all product manifolds.

As an example, let us consider the problem of estimating $(\theta_1, \dots, \theta_N) \in \bar{\mathcal{P}} = \mathbb{S}^2 \times \dots \times \mathbb{S}^2$, the product of N spheres. $\bar{\mathcal{P}}$ has unit curvature along tangent 2-planes pertaining to a single sphere, but zero curvature along all 2-planes spanning exactly two distinct spheres. Of course, $K_{\max} = 1$. If estimating θ_i and θ_j , $i \neq j$, are two independent but identical tasks, one should expect the distribution of $\text{dist}(\theta_i, \hat{\theta}_i)$ to be independent of i . Consequently, $\text{dist}(\theta, \hat{\theta})$ grows as \sqrt{N} , whereas K_{\max} remains constant. Hence, condition (32) becomes increasingly restrictive with growing N . Of course, since the N tasks are independent and can be considered separately, the negligibility of the curvature terms should not depend on N , which brings the conclusion that simply describing the curvature of $\bar{\mathcal{P}}$ through K_{\max} may not be enough.

For such parameter spaces, it is necessary to explicitly compute the curvature terms in the intrinsic Cramér-Rao bounds, if only to show that they are indeed negligible at reasonable SNR. We now set out to give versions of theorems 2 and 3 including curvature terms, computable without constructing other bases than e , the basis of $T_\theta \mathcal{P}$. This will require the Riemannian curvature tensor of $\bar{\mathcal{P}}$. Useful references to look up/compute this tensor are [13, Lemma 3.39, Cor. 3.58, Thm 7.47, Cor. 11.10][15][16].

A. Curvature terms for submanifolds

The random error vector $X \triangleq \text{Log}_\theta(\hat{\theta})$ expands in the basis \bar{e} as $X = \sum_i \bar{x}_i \bar{e}_i$, with $\bar{x}_1, \dots, \bar{x}_{\bar{d}}$ random variables. Notice that $(C_{\bar{e}})_{ij} = \mathbb{E} \{ \langle X, \bar{e}_i \rangle_\theta \langle X, \bar{e}_j \rangle_\theta \} = \mathbb{E} \{ \bar{x}_i \bar{x}_j \}$. Let \bar{R} be the Riemannian curvature tensor of $\bar{\mathcal{P}}$. The mapping $(u, v, w, z) \in (T_\theta \bar{\mathcal{P}})^4 \mapsto \langle \bar{R}(u, v)w, z \rangle_\theta$ is linear in its four arguments [15].

Smith introduces the symmetric 2-form $\bar{\mathbf{R}}_{\mathbf{m}}: \mathbb{T}_\theta \bar{\mathcal{P}} \times \mathbb{T}_\theta \bar{\mathcal{P}} \rightarrow \mathbb{R}$ defined by [3, eq. (34)]:

$$\bar{\mathbf{R}}_{\mathbf{m}}[\bar{e}_i, \bar{e}_j] = \mathbb{E} \left\{ \langle \bar{R}(X, \bar{e}_i) \bar{e}_j, X \rangle_\theta \right\} \quad (33)$$

$$= \mathbb{E} \left\{ \sum_{k,\ell} \langle \bar{R}(\bar{e}_k, \bar{e}_i) \bar{e}_j, \bar{e}_\ell \rangle_\theta \bar{x}_k \bar{x}_\ell \right\} \quad (34)$$

$$= \sum_{k,\ell} \langle \bar{R}(\bar{e}_k, \bar{e}_i) \bar{e}_j, \bar{e}_\ell \rangle_\theta (C_{\bar{e}})_{k\ell}. \quad (35)$$

From the latter expression, it is apparent that the entries of the matrix associated to $\bar{\mathbf{R}}_{\mathbf{m}}$ are linear combinations of the entries of $C_{\bar{e}}$. Generalizing this to any matrix, the following linear map is defined:

$$\bar{R}_m: \mathbb{R}^{\bar{d} \times \bar{d}} \rightarrow \mathbb{R}^{\bar{d} \times \bar{d}}: M \mapsto \bar{R}_m(M), \text{ with}$$

$$(\bar{R}_m(M))_{ij} = \sum_{k,\ell} \langle \bar{R}(\bar{e}_k, \bar{e}_i) \bar{e}_j, \bar{e}_\ell \rangle_\theta M_{k\ell}. \quad (36)$$

At large SNR, the CRB with curvature terms in [3, Cor. 2] reads

$$C_{\bar{e}} \succeq \bar{F}_{\bar{e}}^{-1} - \frac{1}{3} (\bar{R}_m(\bar{F}_{\bar{e}}^{-1}) \bar{F}_{\bar{e}}^{-1} + \bar{F}_{\bar{e}}^{-1} \bar{R}_m(\bar{F}_{\bar{e}}^{-1})). \quad (37)$$

In order to provide an equivalent of (37) only referencing the basis e , we introduce the following symmetric 2-form on $\mathbb{T}_\theta \mathcal{P} \times \mathbb{T}_\theta \mathcal{P}$:

$$\mathbf{R}_{\mathbf{m}}[e_i, e_j] \triangleq \bar{\mathbf{R}}_{\mathbf{m}}[P_\theta e_i, P_\theta e_j]. \quad (38)$$

Notice that, since $X \in \mathbb{T}_\theta \bar{\mathcal{P}}$, we have $X = P_\theta X$. Expanding in the basis e , $X = \sum_i x_i e_i = \sum_i x_i P_\theta e_i$ with random variables x_1, \dots, x_d and $(C_e)_{ij} = \mathbb{E} \{x_i x_j\}$. It follows that:

$$\mathbf{R}_{\mathbf{m}}[e_i, e_j] = \mathbb{E} \left\{ \langle \bar{R}(X, P_\theta e_i) P_\theta e_j, X \rangle_\theta \right\} \quad (39)$$

$$= \sum_{k,\ell} \langle \bar{R}(P_\theta e_k, P_\theta e_i) P_\theta e_j, P_\theta e_\ell \rangle_\theta (C_e)_{k\ell}. \quad (40)$$

From there, we introduce the following linear map:

$$R_m: \mathbb{R}^{d \times d} \rightarrow \mathbb{R}^{d \times d}: M \mapsto R_m(M), \text{ with}$$

$$(R_m(M))_{ij} = \sum_{k,\ell} \langle \bar{R}(P_\theta e_k, P_\theta e_i) P_\theta e_j, P_\theta e_\ell \rangle_\theta M_{k\ell}. \quad (41)$$

Riemannian curvature is often specified by a formula for $\langle \bar{R}(u, v)v, u \rangle$. Hence the standard polarization identity for symmetric bilinear forms may be useful to compute $\mathbf{R}_{\mathbf{m}}$:

$$4\mathbf{R}_{\mathbf{m}}[e_i, e_j] = \mathbf{R}_{\mathbf{m}}[e_i + e_j, e_i + e_j] - \mathbf{R}_{\mathbf{m}}[e_i - e_j, e_i - e_j]. \quad (42)$$

We use the linear maps R_m and \bar{R}_m in the following theorem:

Theorem 4 (CRB on submanifolds, with curvature). *(Continued from Theorem 2). Including terms due to the possible curvature of $\bar{\mathcal{P}}$, at large SNR, the covariance matrix C_e (5) of any unbiased estimator $\hat{\theta}: \mathcal{M} \rightarrow \bar{\mathcal{P}}$ and the Fisher information matrix F_e (3) w.r.t. the orthonormal basis e of $\mathbb{T}_\theta \mathcal{P}$ obey the following matrix inequality (assuming $\text{rank}(P_e F_e P_e) = \bar{d}$):*

$$C_e \succeq \tilde{F}_e^\dagger - \frac{1}{3} \left(R_m(\tilde{F}_e^\dagger) \tilde{F}_e^\dagger + \tilde{F}_e^\dagger R_m(\tilde{F}_e^\dagger) \right), \quad (43)$$

where $\tilde{F}_e = P_e F_e P_e$ and $R_m: \mathbb{R}^{d \times d} \rightarrow \mathbb{R}^{d \times d}$ is as defined by (41).

Proof: We start from the CRB w.r.t. the basis \bar{e} (37):

$$C_{\bar{e}} \succeq \bar{F}_{\bar{e}}^{-1} - \frac{1}{3} \left(\bar{R}_m(\bar{F}_{\bar{e}}^{-1}) \bar{F}_{\bar{e}}^{-1} + \bar{F}_{\bar{e}}^{-1} \bar{R}_m(\bar{F}_{\bar{e}}^{-1}) \right). \quad (44)$$

By expanding the projections $P_\theta e_i = \sum_j \langle \bar{e}_j, e_i \rangle \bar{e}_j = \sum_j E_{ji} \bar{e}_j$ and exploiting the linearity of $\langle \bar{R}(u, v) w, z \rangle_\theta$ in its four arguments, the matrix relation below comes forth:

$$\forall M \in \mathbb{R}^{d \times d}, \quad R_m(M) = E^\top \bar{R}_m(EME^\top)E. \quad (45)$$

From the proof of Theorem 2, recall that $C_{\bar{e}} = EC_e E^\top$ and $\bar{F}_{\bar{e}}^{-1} = E(P_e F_e P_e)^\dagger E^\top$. The relation (45) yields $\bar{R}_m(\bar{F}_{\bar{e}}^{-1}) = ER_m((P_e F_e P_e)^\dagger)E^\top$. Substituting in the CRB gives:

$$EC_e E^\top \succeq E \left(\tilde{F}_e^\dagger - \frac{1}{3} \left(R_m(\tilde{F}_e^\dagger) \tilde{F}_e^\dagger + \tilde{F}_e^\dagger R_m(\tilde{F}_e^\dagger) \right) \right) E^\top,$$

where we used the fact that $R_m(M)P_e = P_e R_m(M) = R_m(M)$, which is easily established from (45). Lemma 1 applies and concludes the proof, since $P_e(P_e F_e P_e)^\dagger P_e = (P_e F_e P_e)^\dagger$. ■

B. Curvature terms for quotient manifolds

We follow the same line of thought as for submanifolds. The random error vector $X \triangleq \text{Log}_{[\theta]}([\hat{\theta}])$ expands in the basis \bar{e} as $X = \sum_i \bar{x}_i \bar{e}_i$, with $\bar{x}_1, \dots, \bar{x}_{\bar{d}}$ random variables and $(C_{\bar{e}})_{ij} = \mathbb{E} \{ \bar{x}_i \bar{x}_j \}$. Let \bar{R} be the Riemannian curvature tensor of $\bar{\mathcal{P}}$. We consider $\bar{\mathbf{R}}_m: \mathbb{T}_{[\theta]} \bar{\mathcal{P}} \times \mathbb{T}_{[\theta]} \bar{\mathcal{P}} \rightarrow \mathbb{R}$ defined by:

$$\bar{\mathbf{R}}_m[\bar{e}_i, \bar{e}_j] = \mathbb{E} \left\{ \langle \bar{R}(X, \bar{e}_i) \bar{e}_j, X \rangle_{[\theta]} \right\} \quad (46)$$

$$= \sum_{k, \ell} \langle \bar{R}(\bar{e}_k, \bar{e}_i) \bar{e}_j, \bar{e}_\ell \rangle_{[\theta]} (C_{\bar{e}})_{k\ell}. \quad (47)$$

A linear map on $\bar{d} \times \bar{d}$ matrices follows:

$$\bar{R}_m: \mathbb{R}^{\bar{d} \times \bar{d}} \rightarrow \mathbb{R}^{\bar{d} \times \bar{d}}: M \mapsto \bar{R}_m(M), \quad \text{with}$$

$$(\bar{R}_m(M))_{ij} = \sum_{k, \ell} \langle \bar{R}(\bar{e}_k, \bar{e}_i) \bar{e}_j, \bar{e}_\ell \rangle_{[\theta]} M_{k\ell}. \quad (48)$$

Again, at large SNR, the CRB (37) holds. To express it only referencing the basis e , we introduce the following symmetric 2-form:

$$\mathbf{R}_m[e_i, e_j] \triangleq \bar{\mathbf{R}}_m[\mathrm{D}\pi(\theta)[e_i], \mathrm{D}\pi(\theta)[e_j]]. \quad (49)$$

Let $\xi = (\mathrm{D}\pi(\theta)|_{\mathbb{H}_\theta})^{-1}[X]$ be the unique horizontal vector such that $\mathrm{D}\pi(\theta)[\xi] = X$. Expanding ξ in the basis e as $\xi = \sum_i x_i e_i$, we find $X = \sum_i x_i \mathrm{D}\pi(\theta)[e_i]$ with random variables x_1, \dots, x_d and $(C_e)_{ij} = \mathbb{E}\{x_i x_j\}$. It follows that:

$$\mathbf{R}_m[e_i, e_j] = \sum_{k, \ell} \langle \bar{R}(\mathrm{D}\pi(\theta)[e_k], \mathrm{D}\pi(\theta)[e_i]) \mathrm{D}\pi(\theta)[e_j], \mathrm{D}\pi(\theta)[e_\ell] \rangle_{[\theta]} (C_e)_{k\ell}.$$

From there, we introduce the following linear map:

$$R_m: \mathbb{R}^{d \times d} \rightarrow \mathbb{R}^{d \times d}: M \mapsto R_m(M), \quad \text{with} \quad (50)$$

$$(R_m(M))_{ij} = \sum_{k, \ell} \langle \bar{R}(\mathrm{D}\pi(\theta)[e_k], \mathrm{D}\pi(\theta)[e_i]) \mathrm{D}\pi(\theta)[e_j], \mathrm{D}\pi(\theta)[e_\ell] \rangle_{[\theta]} M_{k\ell}.$$

Theorem 5 (CRB on quotient manifolds, with curvature). *(Continued from Theorem 3). Including terms due to the possible curvature of $\bar{\mathcal{P}}$, at large SNR, the covariance matrix C_e (20) of any unbiased estimator $\hat{\theta}: \mathcal{M} \rightarrow \bar{\mathcal{P}}$ and the Fisher information matrix F_e (3) w.r.t. the orthonormal basis e of $\mathbb{T}_\theta \mathcal{P}$ obey the following matrix inequality (assuming $\mathrm{rank}(F_e) = \bar{d}$):*

$$C_e \succeq F_e^\dagger - \frac{1}{3} \left(R_m(F_e^\dagger) F_e^\dagger + F_e^\dagger R_m(F_e^\dagger) \right), \quad (51)$$

where $R_m: \mathbb{R}^{d \times d} \rightarrow \mathbb{R}^{d \times d}$ is as defined by (50).

Proof: The proof is very similar to that of Theorem 4. We start from the CRB w.r.t. the basis \bar{e} (37). Expanding $\mathrm{D}\pi(\theta)[e_i] = \mathrm{D}\pi(\theta)[PH_\theta e_i] = \sum_j \langle \bar{e}_j, e_i \rangle \mathrm{D}\pi(\theta)[\bar{e}_j] = \sum_j E_{ji} \bar{e}_j$ and exploiting linearity of $\langle \bar{R}(\cdot, \cdot), \cdot \rangle_{[\theta]}$ in its four arguments, relation (45) is established for the operators \bar{R}_m (48) and R_m (50) too. From the proof of Theorem 3, recall that $C_{\bar{e}} = EC_e E^\top$ and $\bar{F}_{\bar{e}}^{-1} = EF_e^\dagger E^\top$. The relation (45) yields $\bar{R}_m(\bar{F}_{\bar{e}}^{-1}) = ER_m(F_e^\dagger)E^\top$. Substituting in the CRB gives:

$$EC_e E^\top \succeq E \left(F_e^\dagger - \frac{1}{3} \left(R_m(F_e^\dagger) F_e^\dagger + F_e^\dagger R_m(F_e^\dagger) \right) \right) E^\top,$$

where we used the fact that $R_m(M)P_e = P_e R_m(M) = R_m(M)$, which is easily established from (45). Lemma 1 applies and concludes the proof, since $P_e F_e^\dagger P_e = F_e^\dagger$. \blacksquare

V. EXAMPLES

We take a look at two examples of the family of synchronization problems [2]. In such problems, one considers a group \mathcal{G} and a set of N group elements $g_1, \dots, g_N \in \mathcal{G}$. The g_i 's are to be estimated based on noisy measurements of group element ratios $g_i g_j^{-1}$. When \mathcal{G} has a manifold structure, that is, when it is a Lie group, synchronization falls within the spectrum of estimation on manifolds. The first example is synchronization on the group of translations \mathbb{R}^n , which makes for a simple geometry and helps fix ideas. The second example is a simplified version of synchronization on $\text{SO}(3)$, the group of rotations in \mathbb{R}^3 . A more elaborate treatment of synchronization on $\text{SO}(n)$ is given in [17]. Synchronization problems illustrate how both theorems for submanifolds and quotient manifolds can apply to the same setting, with rich interpretation.

A. Synchronization of translations

Let $\theta = (\theta_1, \dots, \theta_N)$ be a vector of N unknown but deterministic points in \mathbb{R}^n . Those can be thought of as positions, states, opinions, etc. of N agents. Let us consider an undirected graph on N nodes with edge set \mathcal{E} , such that for each edge $\{i, j\} \in \mathcal{E}$ we have a noisy measurement of the relative state $h_{ij} = \theta_j - \theta_i + n_{ij}$, where the $n_{ij} \sim \mathcal{N}(0, \Sigma)$ are i.i.d. normally distributed noise vectors. By symmetry, $h_{ij} = -h_{ji}$, so $n_{ij} = -n_{ji}$. While it is important to assume independence of noise on distinct edges to keep the derivation simple, it is easy to relax the assumption that they have identical distributions. We assume identical distributions to keep notations simple.

The task is to estimate the θ_i 's from the h_{ij} 's, thus $\mathcal{P} = (\mathbb{R}^n)^N$, and we set out to derive Cramér-Rao bounds for this problem. An alternative way of obtaining this result can be found in [18]. Decentralized algorithms to execute this synchronization can be found there and in [19].

The log-likelihood function $L: \mathcal{P} \rightarrow \mathbb{R}$ reads, with $\hat{\theta} = (\hat{\theta}_1, \dots, \hat{\theta}_N)$ and $V_i = \{j : \{i, j\} \in \mathcal{E}\}$ the set of neighbors of node i and dropping additive constants:

$$L(\hat{\theta}) = \frac{1}{2} \sum_{i=1}^N \sum_{j \in V_i} -\frac{1}{2} (h_{ij} - \hat{\theta}_j + \hat{\theta}_i)^\top \Sigma^{-1} (h_{ij} - \hat{\theta}_j + \hat{\theta}_i).$$

In order to compute the Fisher information matrix for this problem, we need to pick an orthonormal basis of $\text{T}_\theta \mathcal{P} \equiv \mathcal{P}$. We choose the basis such that the first n vectors correspond to the canonical basis for the first copy of \mathbb{R}^n in \mathcal{P} , the next n vectors correspond to the canonical basis for the second copy of \mathbb{R}^n in \mathcal{P} , etc., totaling nN orthonormal basis vectors. The gradient of $L(\hat{\theta})$ w.r.t. $\hat{\theta}_i$ in this basis is

the following vector in \mathbb{R}^n :

$$\text{grad}_i L(\hat{\boldsymbol{\theta}}) = \sum_{j \in V_i} \Sigma^{-1} (h_{ij} - \hat{\theta}_j + \hat{\theta}_i). \quad (52)$$

Hence, $\text{grad}_i L(\boldsymbol{\theta}) = \sum_{j \in V_i} \Sigma^{-1} n_{ij}$. The FIM F (3) is formed of $N \times N$ blocks of size $n \times n$. Due to independence of the n_{ij} 's and $n_{ij} = -n_{ji}$,

$$\mathbb{E} \left\{ (\Sigma^{-1} n_{ij})(\Sigma^{-1} n_{k\ell})^\top \right\} = \Sigma^{-1} \mathbb{E} \left\{ n_{ij} n_{k\ell}^\top \right\} \Sigma^{-1} = \begin{cases} \Sigma^{-1} & \text{if } (i, j) = (k, \ell), \\ -\Sigma^{-1} & \text{if } (i, j) = (\ell, k), \\ 0 & \text{otherwise.} \end{cases} \quad (53)$$

Hence, the (i, j) th block of the FIM is given by:

$$F_{ij} = \mathbb{E} \left\{ \text{grad}_i L(\boldsymbol{\theta}) \cdot \text{grad}_j L(\boldsymbol{\theta})^\top \right\} = \begin{cases} |V_i| \Sigma^{-1} & \text{if } i = j, \\ -\Sigma^{-1} & \text{if } \{i, j\} \in \mathcal{E}, \\ 0 & \text{otherwise.} \end{cases} \quad (54)$$

The structure of the graph Laplacian is apparent. Let $D = \text{diag}(|V_1|, \dots, |V_N|)$ be the degree matrix and let A be the adjacency matrix of the measurement graph. The Laplacian $\mathcal{L} = D - A$ is tied to the FIM via:

$$F = \mathcal{L} \otimes \Sigma^{-1}, \quad (55)$$

where \otimes denotes the Kronecker product.

Of course, since we only have relative measurements, we can only hope to recover the θ_i 's up to a global translation. And indeed, for every translation vector $t \in \mathbb{R}^n$, we have $L(\hat{\boldsymbol{\theta}}) = L(\hat{\boldsymbol{\theta}} + t)$, where $\hat{\boldsymbol{\theta}} + t \triangleq (\hat{\theta}_1 + t, \dots, \hat{\theta}_N + t)$. That is, all $\hat{\boldsymbol{\theta}} + t$ induce the same distribution of the measurements h_{ij} , and are thus indistinguishable. This is the root of the rank deficiency of the FIM. Surely, if the graph is connected, the all-ones vector $\mathbf{1}_N$ forms a basis of $\ker \mathcal{L}$. Consequently, $\ker F$ consists of all vectors of the form $\mathbf{1}_N \otimes t$, with arbitrary $t \in \mathbb{R}^n$. Naturally, these correspond to global translations by t .

To resolve this ambiguity, we can either add constraints, most naturally in the form of anchors, or work on the quotient space.

a) With anchors: Let us consider $A \subset \{1, \dots, N\}$, $A \neq \emptyset$, such that all θ_i with $i \in A$ are known; these are anchors. The resulting parameter space $\bar{\mathcal{P}} = \{\hat{\boldsymbol{\theta}} \in \mathcal{P} : \hat{\theta}_i = \theta_i \forall i \in A\}$ is a Riemannian submanifold of \mathcal{P} . The orthogonal projector from $T_{\boldsymbol{\theta}} \mathcal{P}$ to $T_{\boldsymbol{\theta}} \bar{\mathcal{P}}$ simply sets all components of a tangent vector corresponding to anchored nodes to zero. Formally, $P = I_A \otimes I_n$, where I_A is a diagonal matrix

of size N whose i^{th} diagonal entry is 1 if $i \notin A$ and 0 otherwise. It follows that $PP^\top = I_A \mathcal{L} I_A \otimes \Sigma^{-1} = \mathcal{L}_A \otimes \Sigma^{-1}$, with the obvious definition for \mathcal{L}_A : the Laplacian with rows and columns corresponding to anchored nodes forced to zero. $\bar{\mathcal{P}}$ is Euclidean, hence it is flat and its curvature tensor vanishes identically. Theorem 2 yields the anchored CRB for the covariance matrix C of an unbiased estimator on $\bar{\mathcal{P}}$:

$$\mathbb{E} \left\{ (\hat{\boldsymbol{\theta}} - \boldsymbol{\theta})(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta})^\top \right\} \triangleq C \succeq \mathcal{L}_A^\dagger \otimes \Sigma. \quad (56)$$

We used the fact that Kronecker product and pseudoinversion commute [20, Fact 7.4.32]. This bound is easily interpreted in terms of individual nodes. Indeed, by definition, inequality (56) means that for all $x \in \mathbb{R}^{nN}$, $x^\top C x \geq x^\top (\mathcal{L}_A^\dagger \otimes \Sigma) x$. In particular, setting $x = e_i \otimes e_k$ with e_i the i^{th} canonical basis vector of \mathbb{R}^N and e_k the k^{th} canonical basis vector of \mathbb{R}^n , we have:

$$\mathbb{E} \left\{ (\hat{\theta}_i - \theta_i)_k^2 \right\} \geq (\mathcal{L}_A^\dagger)_{ii} \cdot \Sigma_{kk}. \quad (57)$$

Summing over $k = 1 \dots n$, this translates into a lower bound on the variance for estimating the state of node i :

$$\mathbb{E} \left\{ \|\hat{\theta}_i - \theta_i\|^2 \right\} \geq (\mathcal{L}_A^\dagger)_{ii} \cdot \text{trace}(\Sigma). \quad (58)$$

This puts forward the importance of the diagonal of \mathcal{L}_A^\dagger , which captures the topology of the measurement graph and the anchor placement. Taking traces on both sides of (56), we obtain an inequality for the total variance:

$$\mathbb{E} \left\{ \text{dist}^2(\hat{\boldsymbol{\theta}}, \boldsymbol{\theta}) \right\} = \mathbb{E} \left\{ \sum_{i=1}^N \|\hat{\theta}_i - \theta_i\|^2 \right\} \geq \text{trace}(\mathcal{L}_A^\dagger) \text{trace}(\Sigma).$$

Notice that it would have been simple to pick a new basis for $T_{\boldsymbol{\theta}}\bar{\mathcal{P}}$, but this would have required a renumbering of the rows and columns of the matrices appearing in the CRB. If the ambiguities are fixed not by adding anchors but, more generally, by adding one or more (for example) linear constraints of the form $a_1\theta_1 + \dots + a_N\theta_N = b$, it becomes less obvious how to pick a meaningful basis for $T_{\boldsymbol{\theta}}\bar{\mathcal{P}}$ without breaking symmetry. In comparison, the projection method used here will apply gracefully, preserving symmetry and row/column ordering in the CRB matrices.

b) Without anchors: If there are no anchors, perhaps because there is no meaningful reference to begin with, we work on the quotient space $\bar{\mathcal{P}} = \mathcal{P}/\sim$, where $\boldsymbol{\theta} \sim \boldsymbol{\theta}'$ iff there exists a translation vector $t \in \mathbb{R}^n$ such that $\boldsymbol{\theta} = \boldsymbol{\theta}' + t$. The distance between the equivalence classes $[\boldsymbol{\theta}]$ and $[\boldsymbol{\theta}']$ on $\bar{\mathcal{P}}$ is the distance between their best aligned members, that is:

$$\text{dist}^2([\boldsymbol{\theta}], [\boldsymbol{\theta}']) = \min_{t \in \mathbb{R}^n} \sum_{i=1}^N \|\theta_i + t - \theta'_i\|^2. \quad (59)$$

The optimal t is easily seen to be $t = \frac{1}{N} \sum_{i=1}^N \theta'_i - \theta_i$, which amounts to aligning the centers of mass of θ and θ' . Consequently, if we denote by θ_c the centered version of θ —i.e., θ translated such that its center of mass is at the origin—we find that:

$$\text{dist}^2([\theta], [\theta']) = \text{dist}^2(\theta_c, \theta'_c) = \sum_{i=1}^N \|\theta_{c,i} - \theta'_{c,i}\|^2. \quad (60)$$

It follows that $\text{dist}^2([\theta], [\theta']) = \text{dist}^2(\theta_c, \theta'_c)$, hence the mapping $[\theta] \mapsto \theta_c$ is an isometry between $\bar{\mathcal{P}}$ and a Euclidean space. We thus conclude that $\bar{\mathcal{P}}$ is a flat manifold and that its curvature tensor vanishes identically [15, Chap. 7]. Theorem 3 and the fact that Kronecker product and pseudoinversion commute [20, Fact 7.4.32] then yield:

$$\mathbb{E} \left\{ (\hat{\theta}_c - \theta_c)(\hat{\theta}_c - \theta_c)^\top \right\} \triangleq C \succeq \mathcal{L}^\dagger \otimes \Sigma, \text{ and} \quad (61)$$

$$\mathbb{E} \left\{ \sum_{i=1}^N \|\hat{\theta}_{c,i} - \theta_{c,i}\|^2 \right\} \geq \text{trace}(\mathcal{L}^\dagger) \text{trace}(\Sigma). \quad (62)$$

We now interpret the CRB (61). Because of the ambiguity in the anchor-free scenario, it does not make much sense to ask what the variance for estimating a specific state is going to be. Rather, one should establish bounds for the variance on estimating the relative state between two nodes, i and j . Let $x = (e_i - e_j) \otimes e_k$ with e_i, e_j the i^{th} and j^{th} canonical basis vectors of \mathbb{R}^N and e_k the k^{th} canonical basis vector of \mathbb{R}^n . Notice that x is a horizontal vector (its components sum to zero). Applying $x^\top \cdot x$ on both sides of (61) yields:

$$\mathbb{E} \left\{ \left((\hat{\theta}_i - \hat{\theta}_j) - (\theta_i - \theta_j) \right)_k^2 \right\} \geq (e_i - e_j)^\top \mathcal{L}^\dagger (e_i - e_j) \cdot \Sigma_{kk}.$$

Notice that there is no need to center $\hat{\theta}$ nor θ anymore, since the quantities involved are relative states. Summing over $k = 1 \dots n$ gives a lower-bound on the variance for estimating the relative state between node i and node j :

$$\mathbb{E} \left\{ \left\| (\hat{\theta}_i - \hat{\theta}_j) - (\theta_i - \theta_j) \right\|^2 \right\} \geq (e_i - e_j)^\top \mathcal{L}^\dagger (e_i - e_j) \cdot \text{trace}(\Sigma). \quad (63)$$

A nice interpretation is now possible. Indeed, the quantity $(e_i - e_j)^\top \mathcal{L}^\dagger (e_i - e_j)$ is well-known to correspond to the Euclidean commute time distance (ECTD) between nodes i and j [21]. It is small if many short paths connect the two nodes and if those paths have edges with large weights which, in our case, means measurements of high quality. Furthermore, the authors of [21] show how one can produce an embedding of the nodes in, say, the plane such that two nodes are close-by if the ECTD separating them is small. This is done via a projection akin to PCA and is an interesting visualization tool as it leads to a plot of the graph such that easily synchronizable nodes are clustered together.

Notice that the bound without anchors has a very different interpretation than that of the bound one would obtain by artificially fixing an arbitrary node. Notice also that, since we did not need to switch to a different basis to obtain the bounds, regardless of which anchors we did or did not choose, it is always the same rows and columns of the matrices in the CRB's that refer to a specific node, which is rather convenient.

The maximum-likelihood estimator in the absence of anchors is easily obtained as the minimum-norm solution to the problem $\max L(\hat{\theta})$ (which is concave, quadratic). This estimator is centered and we state without proof that it is efficient, i.e., its covariance is exactly $\mathcal{L}^\dagger \otimes \Sigma$. In the anchored case, the maximum-likelihood estimator is conveniently obtained via quadratic programming.

For the sake of simplicity, we considered a connected graph. In general, the graph might be disconnected, and there would then be more ambiguity. It is obvious that, in general, there is an \mathbb{R}^n ambiguity for each connected component that does not include an anchor. The CRB's presented here can easily be derived to take care of this more general situation: one simply needs to redefine the equivalence relation \sim accordingly. This in turn leads to a new quotient space with an appropriate notion of distance and covariance. The theorems established in this paper apply seamlessly to this more general scenario.

B. Synchronization of rotations

We now consider a second example, closely related to the first one. Let $\text{SO}(3) = \{R \in \mathbb{R}^{3 \times 3} : R^\top R = I_3, \det(R) = 1\}$ denote the set of rotation matrices in \mathbb{R}^3 . Let $\mathbf{R} = (R_1, \dots, R_N)$ denote a set of unknown but deterministic rotations: those are the parameters. The natural parameter space is thus $\mathcal{P} = \text{SO}(3)^N$. Consider an undirected graph on N nodes with edge set \mathcal{E} as in the previous example. For each edge $\{i, j\} \in \mathcal{E}$, $i < j$, we are given a measurement of the relative rotation between R_i and R_j :

$$H_{ij} = R_i R_j^\top + \sigma N_{ij}, \quad (64)$$

where $\sigma > 0$ and N_{ij} is a 3×3 matrix whose entries are i.i.d. sampled from $\mathcal{N}(0, 1)$. Let us also assume that the N_{ij} 's are independent from each other. The proposed noise model could be disputed. Indeed, measurements of relative rotations should be rotation matrices too. In contrast, in the proposed model, the H_{ij} 's are arbitrary 3×3 matrices. We conduct the calculations with this noise model here for it admits a concise exposition. A more compelling model is considered in [17].

The task is to estimate the rotations $(R_i)_{i=1, \dots, N}$ from the measurements $(H_{ij})_{\{i < j\} \in \mathcal{E}}$. We expect a singular FIM since the measurements only convey information about relative rotations, not absolute

rotations. Indeed, all sets of rotations of the form $\mathbf{R}Q \triangleq (R_1Q, \dots, R_NQ)$ for arbitrary $Q \in \text{SO}(3)$ lead to the same probability distribution function for the measurements.

We will need a few tools to work on $\text{SO}(3)$, see [22]. The tangent space at $R \in \text{SO}(3)$ is the set $\text{T}_R\text{SO}(3) = \{R\Omega : \Omega \in \mathfrak{so}(3)\}$, where $\mathfrak{so}(3) = \{\Omega \in \mathbb{R}^{3 \times 3} : \Omega^\top = -\Omega\}$ is the set of skew-symmetric matrices. It is endowed with the usual metric $\langle R\Omega, R\Omega' \rangle_R = \text{trace}(\Omega^\top \Omega')$, turning it into a Riemannian submanifold of $\mathbb{R}^{3 \times 3}$. The logarithmic map between two rotation matrices R and R' is given by the matrix logarithm $\log: \text{SO}(3) \rightarrow \mathfrak{so}(3)$ as $\text{Log}_R(R') = R \log(R^\top R')$. The Riemannian distance follows as

$$\text{dist}(R, R') = \left\| \log(R^\top R') \right\|_{\text{F}} = \sqrt{2}\alpha, \quad (65)$$

where $\alpha \in [0, \pi]$ is the angle by which the aligning rotation $R^\top R'$ rotates and $\|\cdot\|_{\text{F}}$ denotes the Frobenius norm: $\|A\|_{\text{F}}^2 = \text{trace}(A^\top A)$. By element-wise extension, the tangent space to \mathcal{P} at \mathbf{R} is $\text{T}_{\mathbf{R}}\mathcal{P} = \{\mathbf{R}\boldsymbol{\Omega} \triangleq (R_1\Omega_1, \dots, R_N\Omega_N) : \boldsymbol{\Omega} = (\Omega_1, \dots, \Omega_N) \in \mathfrak{so}(3)^N\}$. The inner product is $\langle \mathbf{R}\boldsymbol{\Omega}, \mathbf{R}\boldsymbol{\Omega}' \rangle_{\mathbf{R}} = \sum_{i=1}^N \text{trace}(\Omega_i^\top \Omega'_i)$.

The probability distribution function of the random matrix $\sigma N_{ij} = H_{ij} - R_i R_j^\top$ is proportional to $\exp(-\|N_{ij}\|_{\text{F}}^2 / (2\sigma^2))$. Hence, up to unimportant additive constants, the log-likelihood function at $\hat{\mathbf{R}} = (\hat{R}_1, \dots, \hat{R}_N) \in \mathcal{P}$ is:

$$L(\hat{\mathbf{R}}) = -\frac{1}{2\sigma^2} \sum_{\{i < j\} \in \mathcal{E}} \left\| H_{ij} - \hat{R}_i \hat{R}_j^\top \right\|_{\text{F}}^2. \quad (66)$$

It is easy to compute the directional derivative of the log-likelihood at \mathbf{R} along the tangent vector $\mathbf{R}\boldsymbol{\Omega}$:

$$DL(\mathbf{R})[\mathbf{R}\boldsymbol{\Omega}] = \frac{1}{\sigma} \sum_{\{i < j\} \in \mathcal{E}} \left\langle N_{ij}, R_i(\Omega_i - \Omega_j)R_j^\top \right\rangle_{\text{F}}. \quad (67)$$

Hence, using independence of the entries of the N_{ij} 's and the fact that the N_{ij} 's have zero-mean, the Fisher information metric at \mathbf{R} is given by:

$$\begin{aligned} \mathbf{F}(\mathbf{R})[\mathbf{R}\boldsymbol{\Omega}, \mathbf{R}\boldsymbol{\Omega}] &= \mathbb{E} \left\{ (DL(\mathbf{R})[\mathbf{R}\boldsymbol{\Omega}])^2 \right\} \\ &= \frac{1}{\sigma^2} \sum_{\{i < j\} \in \mathcal{E}} \mathbb{E} \left\{ \left\langle N_{ij}, R_i(\Omega_i - \Omega_j)R_j^\top \right\rangle_{\text{F}}^2 \right\} \\ &= \frac{1}{\sigma^2} \sum_{\{i < j\} \in \mathcal{E}} \left\| R_i(\Omega_i - \Omega_j)R_j^\top \right\|_{\text{F}}^2 \\ &= \frac{1}{\sigma^2} \sum_{\{i < j\} \in \mathcal{E}} \|\Omega_i - \Omega_j\|_{\text{F}}^2. \end{aligned} \quad (68)$$

The FIM is thus independent of the true rotations \mathbf{R} . From now on, we will write \mathbf{F} instead of $\mathbf{F}(\mathbf{R})$. To represent \mathbf{F} as a matrix, we need to pick an orthonormal basis of $\mathbb{T}_{\mathbf{R}}\mathcal{P}$. Let E_1, E_2, E_3 form an orthonormal basis of $\mathfrak{so}(3)$. This choice leads to an orthonormal basis for the tangent space $\mathbb{T}_{\mathbf{R}}\mathcal{P}$:

$$(e_{ik})_{i=1,\dots,N,k=1,\dots,3} \quad (69)$$

with $e_{ik} = (0, \dots, 0, R_i E_k, 0, \dots, 0)$ a zero vector except for the i^{th} component equal to $R_i E_k$. The matrix F which represents the quadratic form \mathbf{F} w.r.t. this basis is a $3N \times 3N$ symmetric matrix composed of 3×3 blocks. $F_{ij,k\ell}$ refers to the (k, ℓ) entry in the (i, j) block. The value of this entry,

$$F_{ij,k\ell} = \mathbf{F}[e_{ik}, e_{j\ell}], \quad (70)$$

is best obtained via the polarization identity (42) applied to equation (68). For example, let us consider $\{i < j\} \in \mathcal{E}$ and $k, \ell \in \{1, 2, 3\}$. Polarization yields:

$$F_{ij,k\ell} = \frac{1}{4\sigma^2} \left(\|E_k - E_\ell\|_{\mathbf{F}}^2 - \|E_k + E_\ell\|_{\mathbf{F}}^2 \right). \quad (71)$$

For $k = \ell$, this evaluates to $-1/\sigma^2$. For $k \neq \ell$, this evaluates to zero. Hence, the (i, j) block of F is $(-1/\sigma^2)I_3$ if i and j are adjacent nodes. Similar algebra establishes that if nodes i and j are not adjacent, the corresponding block is zero. Similarly, the i^{th} diagonal block of F has the form $(|V_i|/\sigma^2)I_3$, where $|V_i|$ is the degree of node i . We thus obtain the same Laplacian structure as in the first example:

$$F = \frac{1}{\sigma^2}(\mathcal{L} \otimes I_3). \quad (72)$$

F is indeed rank-deficient. As in the previous example, for connected graphs, we have $\ker F = \{\mathbf{1}_N \otimes t : t \in \mathbb{R}^3\}$. These vectors correspond to the rotation of all R_i 's in the same direction, which indeed leaves relative rotations unaffected.

c) With anchors: Let us consider $A \subset \{1, \dots, N\}$, $A \neq \emptyset$, such that all R_i with $i \in A$ are known; these are anchors. The resulting parameter space $\bar{\mathcal{P}} = \{\hat{\mathbf{R}} \in \mathcal{P} : \hat{R}_i = R_i \forall i \in A\}$ is a Riemannian submanifold of \mathcal{P} .

Let $\hat{\mathbf{R}}$ be an unbiased estimator for the anchored problem. As defined by equation (5), the covariance matrix of that estimator w.r.t. the basis of e_{ik} 's is a $3N \times 3N$ matrix with blocks of size 3×3 . Indexing entry (k, ℓ) in the block (i, j) as $C_{ij,k\ell}$, we have by definition:

$$C_{ij,k\ell} = \mathbb{E} \left\{ \langle X, e_{ik} \rangle_{\mathbf{R}} \langle X, e_{j\ell} \rangle_{\mathbf{R}} \right\}, \quad (73)$$

where $X = \text{Log}_{\mathbf{R}}(\hat{\mathbf{R}}) = \left(\text{Log}_{R_1}(\hat{R}_1), \dots, \text{Log}_{R_N}(\hat{R}_N) \right)$ is the (random) error vector. In particular, the trace of C is the expected squared norm of X , that is, the expected squared distance between the true rotations \mathbf{R} and the estimator $\hat{\mathbf{R}}$.

The orthogonal projector from $T_{\mathbf{R}}\mathcal{P}$ to $T_{\mathbf{R}}\bar{\mathcal{P}}$ sets all components of a tangent vector corresponding to anchored nodes to zero. Hence, as in the previous example, we define \mathcal{L}_A to be the Laplacian with rows and columns corresponding to anchored nodes forced to zero. Theorem 2 yields the anchored CRB:

$$C \succeq \sigma^2 \mathcal{L}_A^\dagger \otimes I_3 + \text{curvature terms.} \quad (74)$$

This bound lends itself to a similar interpretation as in the previous example.

At large SNR, we expect the curvature terms to become negligible. Unfortunately, invoking the simple argument (32) is not sufficient. Indeed, $SO(3)$ has constant sectional curvature of $1/8$. Hence, a (tight) upperbound on the absolute value of the sectional curvatures of $\bar{\mathcal{P}}$ is $K_{\max} = 1/8$. Requiring $\text{dist}(\mathbf{R}, \hat{\mathbf{R}}) \ll \sqrt{8}$ is too restrictive when a large number N of rotations have to be estimated. In [17], the curvature terms for $\bar{\mathcal{P}}$ are derived and it is shown that

$$R_m(\sigma^2 \mathcal{L}_A^\dagger \otimes I_3) = \frac{3\sigma^2}{4} \text{ddiag}(\mathcal{L}_A^\dagger) \otimes I_3, \quad (75)$$

where ddiag puts all off-diagonal entries of a matrix to zero. Hence, Theorem 4 yields the following CRB:

$$C \succeq \sigma^2 \mathcal{L}_A^\dagger \otimes I_3 + \frac{3\sigma^4}{4} \left(\text{ddiag}(\mathcal{L}_A^\dagger) \mathcal{L}_A^\dagger + \mathcal{L}_A^\dagger \text{ddiag}(\mathcal{L}_A^\dagger) \right) \otimes I_3. \quad (76)$$

As expected, for large SNR, that is, for small σ^2 , the curvature terms are negligible.

Taking traces on both sides of (74), we obtain an inequality for the total variance (neglecting curvature):

$$\mathbb{E} \left\{ \text{dist}^2(\hat{\mathbf{R}}, \mathbf{R}) \right\} = \mathbb{E} \left\{ \sum_{i=1}^N \text{dist}^2(\hat{R}_i, R_i) \right\} \geq 3\sigma^2 \text{trace}(\mathcal{L}_A^\dagger).$$

d) Without anchors: If there are no anchors, we work on the quotient space $\bar{\mathcal{P}} = \mathcal{P}/\sim$, where $\mathbf{R} \sim \mathbf{R}'$ iff there exists a rotation $Q \in SO(3)$ such that $\mathbf{R}' = \mathbf{R}Q \triangleq (R_1Q, \dots, R_NQ)$. The distance between the equivalence classes $[\mathbf{R}]$ and $[\mathbf{R}']$ on $\bar{\mathcal{P}}$ is the distance between their best aligned members:

$$\text{dist}^2([\mathbf{R}], [\mathbf{R}']) = \min_{Q \in SO(3)} \text{dist}^2(\mathbf{R}Q, \mathbf{R}') = \min_{Q \in SO(3)} \sum_{i=1}^N \left\| \log(Q^\top R_i^\top R'_i) \right\|_{\mathbb{F}}^2. \quad (77)$$

This is well defined since $SO(3)$ is compact. Computing the optimal Q amounts to computing a Karcher mean on $SO(3)$, a problem for which good algorithms are available [23], [24].

A simple application of Theorem 3 then yields the following bound on the variance of an unbiased estimator $\hat{\mathbf{R}}$ for the anchor-free synchronization problem:

$$\mathbb{E} \left\{ \text{dist}^2([\mathbf{R}], [\hat{\mathbf{R}}]) \right\} \geq 3\sigma^2 \text{trace}(\mathcal{L}^\dagger) + \text{curvature terms.}$$

A more complete study of the anchor-free case is conducted in [17]. Among other things, it is shown there that the curvature terms for the quotient space $\bar{\mathcal{P}}$ are on the order of $\mathcal{O}(\sigma^4(1 + 1/N))$. Hence, they remain negligible for large SNR, as expected.

VI. CONCLUSIONS

We proposed four theorems that are meant to ease the use of the intrinsic Cramér-Rao bounds developed in [3] when the actual parameter space is a Riemannian submanifold or a Riemannian quotient manifold of a (usually more natural) parent space. We showed on two simple examples how these theorems provide meaningful bounds for estimation problems with indeterminacies, whether these are dealt with by including prior knowledge in the form of constraints or by acknowledging the quotient nature of the parameter space. We also observed on these same examples that fixing indeterminacies by adding constraints results in different CRB's than if the quotient nature is acknowledged.

ACKNOWLEDGMENTS

NB gratefully acknowledges J. Xavier and V. Barroso for figures 2–3 as well as P.-A. Absil, A. Singer and the anonymous referees for their useful comments about the original manuscript. This paper presents research results of the Belgian Network DYSCO (Dynamical Systems, Control, and Optimization), funded by the Interuniversity Attraction Poles Programme, initiated by the Belgian State, Science Policy Office. The scientific responsibility rests with its authors. NB is an FNRS research fellow (Aspirant).

APPENDIX

In this appendix, we give an informal and fast-paced review of the elementary tools of differential geometry used in the present paper. A number of standard textbooks cover all the necessary concepts [13], [16], [25]. The monograph [14] features an algorithmically-oriented introduction to differential geometry and is available for free on the publisher's website. The footnotes in [3] can also constitute a refresher.

A *smooth manifold* is a set \mathcal{P} that locally “looks like” \mathbb{R}^d , that is, \mathcal{P} can locally be charted by a choice of smooth coordinates (x_1, \dots, x_d) , $x_i: \Omega \subset \mathcal{P} \rightarrow \mathbb{R}$, up to some natural consistency conditions for charts with overlapping domains. A trivial example is \mathbb{R}^d itself. A simple example is the sphere $\mathbb{S}^2 = \{\theta \in \mathbb{R}^3 : \theta^\top \theta = 1\}$. Locally, \mathbb{S}^2 can be mapped to an open set of \mathbb{R}^2 in a smooth way. The *dimension* of the manifold is the number of coordinates required: $\dim \mathcal{P} = d$. A curve on \mathcal{P} , $c: \mathbb{R} \rightarrow \mathcal{P}$, is smooth if it is smooth when expressed in the charts.

There exist many different manifolds of practical interest for signal processing applications. We cite just a few. The (compact) Stiefel manifold $\text{St}(n, p)$ is the set of $n \times p$ orthonormal matrices [14]. Special

cases include the sphere \mathbb{S}^{n-1} ($p = 1$) and the orthogonal group $O(n)$ ($p = n$). The Grassmann manifold $\text{Gr}(n, p)$ is the set of vector subspaces of dimension p embedded in \mathbb{R}^n [14]. The special orthogonal group $\text{SO}(n)$ is the set of $n \times n$ orthogonal matrices with determinant 1, which corresponds to the set of rotations in \mathbb{R}^n [22].

At each point $\theta \in \mathcal{P}$, a *tangent space* $T_\theta \mathcal{P}$ is defined, which is a vector space of dimension d . For embedded manifolds such as \mathbb{S}^2 , the tangent spaces correspond to the usual concept of tangent planes to \mathbb{S}^2 in \mathbb{R}^3 : $T_\theta \mathbb{S}^2 = \{u \in \mathbb{R}^3 : \theta^\top u = 0\}$. In a more general setting, it is desirable to define $T_\theta \mathcal{P}$ without referencing an embedding space. This is done by considering smooth curves on \mathcal{P} that pass through θ and grouping the curves that do so with the same “velocity” (derivative as expressed in the charts) into an equivalence class, which is then called a tangent vector.

Let $L: \mathcal{P} \rightarrow \mathbb{R}$ be a smooth scalar field over \mathcal{P} . The *directional derivative* of L at θ along a tangent vector u is defined by $DL(\theta)[u] = (L \circ c)'(0)$, the derivative of $L \circ c: \mathbb{R} \rightarrow \mathbb{R}$ at $t = 0$, where \circ denotes function composition and $c: \mathbb{R} \rightarrow \mathcal{P}$ is a smooth curve passing through θ at $t = 0$ with velocity u . For the sphere for example, this correspondence amounts to $c'(0) = u$, where c is seen as a function from \mathbb{R} to \mathbb{R}^3 . If L is smoothly defined in an open neighborhood of \mathbb{S}^2 in \mathbb{R}^3 , this reduces to a classical derivative: $DL(\theta)[u] = \lim_{t \rightarrow 0} (L(\theta + tu) - L(\theta))/t$.

Each tangent space may be endowed with a metric (an inner product) $\langle \cdot, \cdot \rangle_\theta: T_\theta \mathcal{P} \times T_\theta \mathcal{P} \rightarrow \mathbb{R}$. The associated norm is defined by $\|u\|_\theta^2 = \langle u, u \rangle_\theta$. When this metric varies continuously as a function of θ , \mathcal{P} is a *Riemannian manifold*. The gradient of a scalar field L at θ , noted $\text{grad } L(\theta)$, gives the steepest-ascent direction of L at θ . By definition, it is the unique tangent vector that satisfies the following:

$$\forall u \in T_\theta \mathcal{P}, \quad \langle \text{grad } L(\theta), u \rangle_\theta = DL(\theta)[u]. \quad (78)$$

Notice that, while the notion of directional derivative does not depend on a particular choice of metric, the notion of gradient does.

A curve $\gamma: I \subset \mathbb{R} \rightarrow \mathcal{P}$ is a *geodesic* if it has zero acceleration. A proper definition of acceleration along a curve requires the concept of *affine connection* ∇ , which we omit here. For each $\theta \in \mathcal{P}$ and $u \in T_\theta \mathcal{P}$, there exists a unique geodesic γ_u such that $\gamma_u(0) = \theta$ and $\gamma_u'(0) = u$. The *exponential map* at θ is the mapping $\text{Exp}_\theta: T_\theta \mathcal{P} \rightarrow \mathcal{P}$ defined by $\text{Exp}_\theta(u) = \gamma_u(1)$. The *logarithmic map* is the principal inverse of the exponential map, $\text{Log}_\theta: \mathcal{P} \rightarrow T_\theta \mathcal{P}$, that is, $u = \text{Log}_\theta(\theta')$ is the shortest tangent vector such that $\text{Exp}_\theta(u) = \theta'$. The notions of geodesic, exponential map and logarithmic map are usually only defined locally. For $\mathcal{P} = \mathbb{R}^d$ with the usual metric $\langle u, v \rangle_\theta = \text{trace}(u^\top v)$, geodesics are straight lines, $\text{Exp}_\theta(u) = \theta + u$ and $\text{Log}_\theta(\theta') = \theta' - \theta$. On the sphere \mathbb{S}^2 , geodesics are great circles, such as the

equator. The *Riemannian distance* or *geodesic distance* between two points $\theta, \theta' \in \mathcal{P}$, is the length of the shortest geodesic arc connecting the two points: $\text{dist}(\theta, \theta') = \|\text{Log}_\theta(\theta')\|_\theta$. On the unit sphere, the Riemannian distance is the arc length, or angle in radians, separating two points.

A *flat manifold* is a manifold that is locally isometric to a Euclidean space \mathbb{R}^d . For example, any one-dimensional smooth manifold is flat. Think of the circle $\mathbb{S}^1 = \{\theta \in \mathbb{R}^2 : \theta^\top \theta = 1\}$ as a Riemannian submanifold of \mathbb{R}^2 with its usual metric: it is always possible to cut out a bit of \mathbb{S}^1 and straighten it without distorting distances. Similarly, the torus $\mathbb{S}^1 \times \mathbb{S}^1$ is flat, being the product of two circles. This example may contrast with our intuition if we think of the torus as an embedded manifold of \mathbb{R}^3 : their topologies are the same, but their Riemannian structures are not. On the other hand, the sphere \mathbb{S}^2 is not flat: attempting to flatten an orange peel will necessarily result in tearing, stretching or compression. Given two tangent vectors $u, v \in T_\theta \mathcal{P}$ (not co-linear), the *sectional curvature* $K(u, v)$ is a real number that measures the curvature of \mathcal{P} along the 2-dimensional subspace of $T_\theta \mathcal{P}$ spanned by u and v . For example, on \mathbb{S}^2 , all sectional curvatures are equal to 1. The *Riemannian curvature tensor* R is another (but equivalent) means of quantifying the “non-flatness” of a manifold. It is related to the sectional curvatures via $K(u, v) = \langle R(u, v)v, u \rangle_\theta / (\|u\|_\theta^2 \|v\|_\theta^2 - \langle u, v \rangle_\theta^2)$. For flat manifolds, curvature is zero.

Let \mathcal{P} be a Riemannian manifold and let $\bar{\mathcal{P}}$ be a submanifold of \mathcal{P} , i.e., $\bar{\mathcal{P}}$ is embedded in \mathcal{P} , like the sphere is embedded in 3-space. The tangent space $T_\theta \bar{\mathcal{P}}$ is a vector subspace of $T_\theta \mathcal{P}$. $\bar{\mathcal{P}}$ is turned into a *Riemannian submanifold* of \mathcal{P} by giving it a Riemannian metric that is the restriction of the Riemannian metric of \mathcal{P} . For example, the classical metric on \mathbb{R}^3 is $\langle u, v \rangle = u^\top v$. By defining $\langle u, v \rangle_\theta = u^\top v$ on $T_\theta \mathbb{S}^2$, we equip \mathbb{S}^2 with a Riemannian submanifold structure. The orthogonal projector from $T_\theta \mathcal{P}$ to $T_\theta \bar{\mathcal{P}}$ is noted P_θ . For the sphere, it is given by $P_\theta u = (I - \theta \theta^\top)u$.

Let \mathcal{P} be a Riemannian manifold and let \sim be an equivalence relation on \mathcal{P} . Let $[\theta] = \{\theta' \in \mathcal{P} : \theta' \sim \theta\}$ be the equivalence class of θ and let $\bar{\mathcal{P}} = \{[\theta] : \theta \in \mathcal{P}\}$ be the set of equivalence classes. Let us further assume that $\bar{\mathcal{P}}$ is a smooth manifold, i.e., it can be smoothly charted. Then, the projection $\pi: \mathcal{P} \rightarrow \bar{\mathcal{P}}$ is a *submersion* if its differential at θ , $D\pi(\theta): T_\theta \mathcal{P} \rightarrow T_{[\theta]} \bar{\mathcal{P}}$, is a surjective map for all θ . In that case, $\bar{\mathcal{P}} = \mathcal{P}/\sim$ is a *quotient manifold*, because we “quotient out” the equivalence relation from \mathcal{P} . The equivalence classes $\pi(\theta)$ are termed the *fibers*; they are Riemannian submanifolds of \mathcal{P} . As such, they admit tangent spaces, called the *vertical spaces*: $V_\theta = T_\theta \pi(\theta) = \ker D\pi(\theta)$. The vertical space is a vector subspace of $T_\theta \mathcal{P}$. The orthogonal complement to V_θ in $T_\theta \mathcal{P}$ is called the *horizontal space*, $H_\theta = \{u \in T_\theta \mathcal{P} : \langle u, v \rangle_\theta = 0 \forall v \in V_\theta\}$. If furthermore $\bar{\mathcal{P}}$ is endowed with a Riemannian metric and

$D\pi(\theta)$ restricted to H_θ is an isometry, in the sense that

$$\forall \theta \in \mathcal{P}, \forall u, v \in H_\theta, \quad \langle D\pi(\theta)[u], D\pi(\theta)[v] \rangle_{[\theta]} = \langle u, v \rangle_\theta,$$

then π is a *Riemannian submersion* and $\bar{\mathcal{P}}$ is a *Riemannian quotient manifold* of \mathcal{P} . Being an isometry, $D\pi(\theta)$ establishes a natural correspondence between the abstract tangent vectors in $T_{[\theta]}\bar{\mathcal{P}}$ and the often more concrete horizontal vectors. The *horizontal lift* of a tangent vector $X \in T_{[\theta]}\bar{\mathcal{P}}$ to H_θ is the unique horizontal vector $\xi \in H_\theta$ that verifies $D\pi(\theta)[\xi] = X$. The definition of Riemannian submersion ensures that this is well-defined. The orthogonal projector from the total tangent space $T_\theta\mathcal{P}$ to H_θ is here noted PH_θ .

REFERENCES

- [1] R. Tron and R. Vidal, “Distributed image-based 3D localization of camera sensor networks,” in *Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on*. IEEE, 2009, pp. 901–908.
- [2] A. Singer, “Angular synchronization by eigenvectors and semidefinite programming,” *Applied and Computational Harmonic Analysis*, 2010.
- [3] S. Smith, “Covariance, subspace, and intrinsic Cramér-Rao bounds,” *Signal Processing, IEEE Transactions on*, vol. 53, no. 5, pp. 1610–1630, 2005.
- [4] C. Rao, “Information and accuracy attainable in the estimation of statistical parameters,” *Bulletin of the Calcutta Mathematical Society*, vol. 37, no. 3, pp. 81–91, 1945.
- [5] J. Xavier and V. Barroso, “Intrinsic Variance Lower Bound (IVLB): an extension of the Cramér-Rao bound to Riemannian manifolds,” in *Acoustics, Speech, and Signal Processing, 2005. Proceedings.(ICASSP’05). IEEE International Conference on*, vol. 5. IEEE, 2005, pp. 1033–1036.
- [6] J. Gorman and A. Hero, “Lower bounds for parametric estimation with constraints,” *Information Theory, IEEE Transactions on*, vol. 36, no. 6, pp. 1285–1301, nov 1990.
- [7] Z. Ben-Haim and Y. Eldar, “On the constrained Cramér-Rao bound with a singular Fisher information matrix,” *Signal Processing Letters, IEEE*, vol. 16, no. 6, pp. 453–456, june 2009.
- [8] P. Stoica and B. Ng, “On the Cramér-Rao bound under parametric constraints,” *Signal Processing Letters, IEEE*, vol. 5, no. 7, pp. 177–179, 1998.
- [9] P. Stoica and T. Marzetta, “Parameter estimation problems with singular information matrices,” *Signal Processing, IEEE Transactions on*, vol. 49, no. 1, pp. 87–90, 2001.
- [10] J. Xavier and V. Barroso, “The Riemannian geometry of certain parameter estimation problems with singular Fisher information matrices,” in *Acoustics, Speech, and Signal Processing, 2004. Proceedings.(ICASSP’04). IEEE International Conference on*, vol. 2. IEEE, 2004, pp. 1021–1024.

- [11] L. Scharf and L. McWhorter, “Geometry of the Cramer-Rao bound,” *Signal Processing*, vol. 31, no. 3, pp. 301 – 311, 1993. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/016516849390088R>
- [12] S. Amari and H. Nagaoka, *Methods of information geometry*. American Mathematical Society, 2007, vol. 191.
- [13] B. O’Neill, *Semi-Riemannian geometry: with applications to relativity*. Academic Press, 1983, vol. 103.
- [14] P.-A. Absil, R. Mahony, and R. Sepulchre, *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, 2008.
- [15] J. Lee, *Riemannian manifolds: An introduction to curvature*, ser. Graduate Texts in Mathematics. Springer, 1997, vol. 176.
- [16] I. Chavel, *Riemannian geometry: a modern introduction*, ser. Cambridge Tracts in Mathematics. Cambridge University Press, 1993, vol. 108.
- [17] N. Boumal, A. Singer, P.-A. Absil, and V. D. Blondel, “Cramér-Rao bounds for synchronization of rotations,” *arXiv preprint arXiv:1211.1621*, 2012.
- [18] S. Howard, D. Cochran, W. Moran, and F. Cohen, “Estimation and registration on graphs,” *Arxiv preprint arXiv:1010.2983*, 2010.
- [19] W. Russell, D. Klein, and J. Hespanha, “Optimal estimation on the graph cycle space,” *Signal Processing, IEEE Transactions on*, vol. 59, no. 6, pp. 2834–2846, 2011.
- [20] D. Bernstein, *Matrix mathematics: theory, facts, and formulas*. Princeton University Press, 2009.
- [21] M. Saerens, F. Fouss, L. Yen, and P. Dupont, “The principal components analysis of a graph, and its relationships to spectral clustering,” *Machine Learning: ECML 2004*, pp. 371–383, 2004.
- [22] D. Bump, *Lie groups*, ser. Graduate Texts in Mathematics. Springer, 2004, vol. 225.
- [23] Q. Rentmeesters and P.-A. Absil, “Algorithm comparison for Karcher mean computation of rotation matrices and diffusion tensors,” in *EUSIPCO 2011: 19th European Signal Processing Conference*, 2011, pp. 2229–2233.
- [24] J. Manton, “A globally convergent numerical algorithm for computing the centre of mass on compact lie groups,” in *Control, Automation, Robotics and Vision Conference, 2004. ICARCV 2004 8th*, vol. 3. IEEE, 2004, pp. 2211–2216.
- [25] W. Boothby, *An introduction to differentiable manifolds and Riemannian geometry*, ser. Pure and Applied Mathematics. Elsevier, 1986, vol. 120.



Nicolas Boumal received a bachelor and a masters of engineering from the Université catholique de Louvain (Belgium) in 2008 and 2010, specializing in applied mathematics. He is now pursuing a Ph.D. degree as a research fellow of the FNRS. His current focus is on optimization on manifolds, for which he develops the Manopt toolbox. Particular applications he is interested in include regression problems on manifolds, low-rank matrix completion and synchronization on Lie groups.