Gauges and Gauge Transformations in 3-D Reconstruction from a Sequence of Images

Kenichi Kanatani
Department of Computer Science, Gunma University, Kiryu, Gunma 376-8515 Japan
Daniel D. Morris
Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 1523-3890, U.S.A.

ABSTRACT

This paper presents a consistent theory for describing indeterminacy and uncertainty of 3-D reconstruction from a sequence of images. First, we present a group-theoretical analysis of gauge transformations and gauges by invoking the Lie group theory. We then discuss how to evaluate the reliability of the solution that has indeterminacy. We also extend the Cramer-Rao lower bound to problems with internal indeterminacy. Finally, we describe the free-gauge approach and the normal form of a covariance matrix that is independent of particular gauges.

Keywords: 3-D reconstruction, Lie group theory, gauge transformation, normal form, reliability evaluation, Cramer-Rao lower bound.

1. INTRODUCTION

Numerous methods have been proposed for 3-D reconstruction from a sequence of images, but the reliability of the solution has typically been evaluated simply by direct first order analysis of the recovered parameters [10]. This, however, hides an important fact: *3-D reconstruction from images necessarily involves inherent indeterminacy* [8].

There are two sources for indeterminacy in parameterized 3-D shape reconstruction. One is the inherent physical indeterminacy caused by a loss of information during projection onto a camera; e.g., scale is indeterminate for a perspective camera. The second source is the over-parameterization of the problem: more than the minimal number of parameters are included for simplicity as well as symmetry of representation; e.g., a shape model with a vector representing each point implicitly specifies the absolute orientation and translation of the object even though these may be indeterminate.

These indeterminacies can be removed by imposing additional normalization constraints; e.g., we can fix the coordinate origin to a particular point of the object and normalize the size of something to unit length. But then the fixed or normalized parameters have no uncertainty by definition, and the uncertainty of all parameters is altered by merely changing the normalization conditions.

Thus, we cannot speak of uncertainty of a particular parameter in absolute terms unless it is *invariant* to normalization conditions, but this fact has been paid relatively little attention [8]. We call transformations of parameters caused by changing normalization conditions *gauge transformations* and particular choices of them *gauges* [5, 6, 7]. In the following, we give a formal description of the 3-D (Euclidean) reconstruction problem and present a group-theoretical analysis of gauge transformations and gauges by invoking the Lie group theory [1]. We then extend the Cramer-Rao lower bound to problems with internal indeterminacy. Finally, we describe the *free-gauge approach* and the *normal form* of a covariance matrix that is independent of particular gauges.

2. CAMERA IMAGING GEOMETRY

Suppose we track N rigidly moving feature points over M images. Let $(x_{\kappa\alpha}, y_{\kappa\alpha})$ be the image coordinates of the αth point in the κth frame. Here, we adopt the *camera-centered* description, assuming that an object is moving in the scene relative to a stationary camera, but the mathematical structure is the same if we view the camera as moving and taking images of a stationary object in the scene.

We identify the camera coordinate system with a global reference frame and define an arbitrary object coordinate system fixed to the object. Let $\boldsymbol{t}_\kappa$ and $\{\boldsymbol{i}_\kappa, \boldsymbol{j}_\kappa, \boldsymbol{k}_\kappa\}$ be, respectively, its origin and orthonormal basis vectors in the κth frame. Let $(s_{\alpha 1}, s_{\alpha 2}, s_{\alpha 3})$ be the coordinates of the αth feature point with respect to the object coordinate system. We define

$$\boldsymbol{R}_\kappa = \begin{pmatrix} \boldsymbol{i}_\kappa & \boldsymbol{j}_\kappa & \boldsymbol{k}_\kappa \end{pmatrix}, \qquad \boldsymbol{s}_\alpha = \begin{pmatrix} s_{\alpha 1} \\ s_{\alpha 2} \\ s_{\alpha 3} \end{pmatrix}. \tag{1}$$

Since $\{\boldsymbol{i}_\kappa, \boldsymbol{j}_\kappa, \boldsymbol{k}_\kappa\}$ are an orthonormal basis, $\boldsymbol{R}_\kappa$ is a rotation matrix. We call $\{\boldsymbol{R}_\kappa, \boldsymbol{t}_\kappa\}$ the *motion parameters* and $\{\boldsymbol{s}_\alpha\}$ the *shape vectors*.

Suppose a point (X, Y, Z) with respect to the camera coordinate system is projected to a point (x, y) on the image plane. We write this mapping as Π: $\mathcal{R}^3 \to \mathcal{R}^2$ ($\mathcal{R}$ denotes the set of real numbers) and call it the *camera model*. The mapping Π may be different from frame to frame. A point with object coordinates $(s_{\alpha 1}, s_{\alpha 2}, s_{\alpha 3})$ is a point $\boldsymbol{R}_\kappa \boldsymbol{s}_\alpha + \boldsymbol{t}_\kappa$ with respect to the camera coordinate system in the κth frame. Hence, the image coordinates $(x_{\alpha\kappa}, y_{\alpha\kappa})$ are expressed as a function of $\boldsymbol{t}_\kappa$, $\boldsymbol{R}_\kappa$, and $\boldsymbol{s}_\alpha$ in the form

$$\begin{pmatrix} x_{\alpha\kappa} \\ y_{\alpha\kappa} \end{pmatrix} = \Pi_\kappa[\boldsymbol{R}_\kappa \boldsymbol{s}_\alpha + \boldsymbol{t}_\kappa], \tag{2}$$

where Π_κ is the camera model for the κth frame. We call the mapping $\mathcal{P}_\kappa$: $SO(3) \times \mathcal{R}^3 \times \mathcal{R}^3 \to \mathcal{R}^2$ thus defined the *projection model* for the κth frame, where $SO(3)$ is the three-dimensional special orthogonal group (the group of 3-D rotations).

Example 1. For an orthographic camera, a point (X, Y, Z) with respect to the camera coordinate system is projected onto an image point (x, y) such that $x = X$ and $y = Y$. The projection model is given by

$$\begin{pmatrix} x_{\alpha\kappa} \\ y_{\alpha\kappa} \end{pmatrix} = \Pi(\boldsymbol{R}_\kappa \boldsymbol{s}_\alpha + \boldsymbol{t}_\kappa), \tag{3}$$

where $\Pi = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}$.

3. PARAMETER SPACE

We stack the two-dimensional vector given by eq. (2) for all the M frames and define a $2M$-dimensional vector $\boldsymbol{p}_\alpha$; the history of the projection of the αth feature point is identified with a single point $\boldsymbol{p}_\alpha$ in a $2M$-dimensional space $\mathcal{R}^{2M}$.

Let $\boldsymbol{\theta}$ be the vector that represents the shape vectors $\{\boldsymbol{s}_\alpha\}$, the motion parameters $\{\boldsymbol{R}_\kappa,\ \boldsymbol{t}_\kappa\}$, and the camera models for all the frames $\{\Pi_\kappa\}$. We define the projection model of the "$2M$-dimensional camera" in the form

$$\boldsymbol{p}_\alpha = \mathcal{P}_\alpha[\boldsymbol{\theta}]. \tag{4}$$

The domain $\mathcal{T}$ of the vector $\boldsymbol{\theta}$ is called the *parameter space*.

Example 2. For the orthographic camera model, the vector $\boldsymbol{\theta}$ has dimension $6M + 3N$, since the rotation matrix $\boldsymbol{R}_\kappa$ and the translation vector $\boldsymbol{t}_\kappa$ both have three degrees of freedom. The parameter space $\mathcal{T}$ is homeomorphic to the direct product of $\mathcal{R}^{3M+3N}$ and M copies of $SO(3)$:

$$\mathcal{T} \simeq \overbrace{SO(3) \times \cdots \times SO(3)}^{M} \times \mathcal{R}^{3M+3N}. \tag{5}$$

4. STATISTICAL NOISE MODEL

Let $\bar{\boldsymbol{p}}_\alpha$ be the true value of $\boldsymbol{p}_\alpha$ in the absence of noise, and write $\boldsymbol{p}_\alpha = \bar{\boldsymbol{p}}_\alpha + \Delta\boldsymbol{p}_\alpha$. We assume that the noise term $\Delta\boldsymbol{p}_\alpha$ is a Gaussian random variable; it may not be independent for different images, and its distribution may be different from point to point. We define the covariance matrices of $\{\boldsymbol{p}_\alpha\}$ by

$$V[\boldsymbol{p}_\alpha, \boldsymbol{p}_\beta] = E[\Delta\boldsymbol{p}_\alpha \Delta\boldsymbol{p}_\beta^\top], \tag{6}$$

where $E[\,\cdot\,]$ denotes expectation. We assume that $V[\boldsymbol{p}_\alpha, \boldsymbol{p}_\beta]$ is known *up to scale* and decompose it into the known *normalized covariance matrix* $V_0[\boldsymbol{p}_\alpha, \boldsymbol{p}_\beta]$ and an unknown *noise level* ϵ as follows [2, 3]:

$$V[\boldsymbol{p}_\alpha, \boldsymbol{p}_\beta] = \epsilon^2 V_0[\boldsymbol{p}_\alpha, \boldsymbol{p}_\beta]. \tag{7}$$

Define $2M \times 2M$ matrices $\boldsymbol{W}_{\alpha\beta}$, $\alpha, \beta = 1, ..., N$, by

$$\begin{pmatrix} \boldsymbol{W}_{11} & \cdots & \boldsymbol{W}_{1N} \\ \vdots & & \vdots \\ \boldsymbol{W}_{N1} & \cdots & \boldsymbol{W}_{NN} \end{pmatrix} = \begin{pmatrix} V_0[\boldsymbol{p}_1, \boldsymbol{p}_1] & \cdots & V_0[\boldsymbol{p}_1, \boldsymbol{p}_N] \\ \vdots & & \vdots \\ V_0[\boldsymbol{p}_N, \boldsymbol{p}_1] & \cdots & V_0[\boldsymbol{p}_N, \boldsymbol{p}_N] \end{pmatrix}^{-1}. \tag{8}$$

The *maximum likelihood solution*, often referred to as *bundle adjustment*, is obtained by minimizing

$$J(\boldsymbol{\theta}) = \sum_{\alpha,\beta=1}^{N} (\boldsymbol{p}_\alpha - \mathcal{P}_\alpha[\boldsymbol{\theta}], \boldsymbol{W}_{\alpha\beta}(\boldsymbol{p}_\beta - \mathcal{P}_\beta[\boldsymbol{\theta}]), \tag{9}$$

where and throughout this paper we denote the inner product of vectors $\boldsymbol{a}$ and $\boldsymbol{b}$ by $(\boldsymbol{a}, \boldsymbol{b})$.

5. GAUGE TRANSFORMATIONS

Generally, the solution that minimizes $J(\boldsymbol{\theta})$ is not unique. There are two sources of indeterminacy:

Frame indifference: Absolute translation and orientation of an object cannot be determined uniquely.

Projection insensitivity: Different 3-D configurations can be projected to the same image.

If there exists a transformation g of the parameter space $\mathcal{T}$ such that

$$J(\boldsymbol{\theta}) = J(g\boldsymbol{\theta}), \qquad \forall\boldsymbol{\theta} \in \mathcal{T}, \tag{10}$$

we call g a *gauge transformation*. The set of all such transformations forms a group $\mathcal{G}$, which we call the *group of gauge transformations* or the *gauge group* for short.

Example 3. If the object coordinate system is rotated by $\boldsymbol{R}$ around its origin and then translated by $\boldsymbol{t} = (t_1, t_2, t_3)^\top$, where $\boldsymbol{t}$ and $\boldsymbol{R}$ are defined with respect to the original object coordinate system, the shape vectors $\{\boldsymbol{s}_\alpha\}$ and the motion parameters $\{\boldsymbol{R}_\kappa,\ \boldsymbol{t}_\kappa\}$ are transformed in the form

$$\boldsymbol{s}'_\alpha = \boldsymbol{R}^\top(\boldsymbol{s}_\alpha - \boldsymbol{t}), \qquad \boldsymbol{R}'_\kappa = \boldsymbol{R}_\kappa\boldsymbol{R}, \qquad \boldsymbol{t}'_\kappa = \boldsymbol{R}_\kappa\boldsymbol{t} + \boldsymbol{t}_\kappa. \tag{11}$$

Example 4. For the orthographic camera model, we see from eq. (3) that if $\{\boldsymbol{s}_\alpha\}$ and $\{\boldsymbol{R}_\kappa,\ \boldsymbol{t}_\kappa\}$ are a solution, so are $\{\boldsymbol{s}_\alpha\}$ and $\{\boldsymbol{R}_\kappa,\ \boldsymbol{t}_\kappa + d_\kappa\boldsymbol{k}\}$ for all $d_\kappa \in \mathcal{R}$, where and throughout this paper we let $\boldsymbol{k} = (0, 0, 1)^\top$. We also see from eq. (3) that if $\{\boldsymbol{s}_\alpha\}$ and $\{\boldsymbol{R}_\kappa,\ \boldsymbol{t}_\kappa\}$ are a solution, so are $\{-\boldsymbol{s}_\alpha\}$ and $\{\boldsymbol{J}\boldsymbol{R}_\kappa, -\boldsymbol{J}\boldsymbol{t}_\kappa\}$, where $\boldsymbol{J} = \mathrm{diag}(-1, -1, 1)$ ($180°$ rotation around the Z-axis); the two solutions are *mirror images* of each other.

6. INFINITESIMAL GENERATORS

Let $\gamma \in \mathcal{G}$ be an element close to the identity element of $\mathcal{G}$. We call such an element an *infinitesimal gauge transformation*. An infinitesimal gauge transformation γ maps $\boldsymbol{\theta}$ in the form

$$\gamma\boldsymbol{\theta} = \boldsymbol{\theta} + \boldsymbol{D}(\boldsymbol{\theta}) + \cdots, \tag{12}$$

where $\boldsymbol{D}(\,\cdot\,)$ is a vector operator called the *infinitesimal gauge generator* of γ. The set of infinitesimal gauge generators becomes a linear space[1], which we denote by $\mathcal{D}_\theta(\mathcal{T})$. It can be proved that for every element $\boldsymbol{D}(\,\cdot\,) \in \mathcal{D}_\theta(\mathcal{G})$ there exists an infinitesimal gauge transformation γ that has $\boldsymbol{D}(\,\cdot\,)$ as its generator [1].

Example 5. The rotation around a unit vector $\boldsymbol{l}$ by a small angle $\Delta\Omega$ is written to a first approximation in the form $\Delta\boldsymbol{\Omega} \times \boldsymbol{I}$, where $\Delta\boldsymbol{\Omega} = \Delta\Omega\boldsymbol{l}$. Throughout this paper, the product $\boldsymbol{a} \times \boldsymbol{A}$ of a three-dimensional vector $\boldsymbol{a}$ and a 3×3 matrix is defined to be the matrix whose columns are the vector products of $\boldsymbol{a}$ with each of the three columns of $\boldsymbol{A}$. For the orthographic camera model, an infinitesimal gauge transform acts on $\{\boldsymbol{s}_\alpha\}$ and $\{\boldsymbol{R}_\kappa,\ \boldsymbol{t}_\kappa\}$ as follows:

$$\boldsymbol{D}(\boldsymbol{s}_\alpha) = -\Delta\boldsymbol{\Omega} \times \boldsymbol{s}_\alpha - \Delta\boldsymbol{t},$$
$$\boldsymbol{D}(\boldsymbol{R}_\kappa) = -\boldsymbol{R}_\kappa \times \Delta\boldsymbol{\Omega}, \qquad \boldsymbol{D}(\boldsymbol{t}_\kappa) = \boldsymbol{R}_\kappa\Delta\boldsymbol{t} + \Delta d_\kappa\boldsymbol{k}. \tag{13}$$

Thus, the operator $\boldsymbol{D}(\,\cdot\,)$ is linearly parameterized by $\Delta\boldsymbol{\Omega}$, $\Delta\boldsymbol{t}$, and $\{\Delta d_\kappa\}$, confirming that the linear space $\mathcal{D}(\mathcal{T})$ is $(M+6)$-dimensional.

7. GAUGES

We say that two values $\boldsymbol{\theta}$ and $\boldsymbol{\theta}'$ are *geometrically equivalent* and write $\boldsymbol{\theta} \sim \boldsymbol{\theta}'$ if there exists an element $g \in \mathcal{G}$ such that $\boldsymbol{\theta}' = g\boldsymbol{\theta}$. This implies that the true parameter space is not $\mathcal{T}$ but the *quotient space* $\mathcal{T}/\mathcal{G}$ of $\mathcal{T}$ with respect to this equivalence relation. Let $\mathcal{T}_\theta$ be an element of $\mathcal{T}/\mathcal{G}$, i.e., a subset of $\mathcal{T}$ consisting of all elements geometrically equivalent to $\boldsymbol{\theta}$:

$$\mathcal{T}_\theta = \{\boldsymbol{\theta}' \,|\, \boldsymbol{\theta}' \sim \boldsymbol{\theta}\}\Big(= \{g\boldsymbol{\theta} \,|\, g \in \mathcal{G}\}\Big). \tag{14}$$

[1]By introducing a product called *commutator*, we can make this linear space an algebra, which is isomomorphic to the *Lie algebra* $\mathcal{L}(\mathcal{G})$ of $\mathcal{G}$.

If the linear space $\mathcal{D}_\theta(\mathcal{T})$ is r-dimensional, $\mathcal{T}_\theta$ is an r-dimensional submanifold of $\mathcal{T}$ called the *leaf* associated with $\boldsymbol{\theta}$. The parameter space $\mathcal{T}$ is filled with such leaves. A space filled with leaves is called a *foliation* or a *foliated manifold*. The parameter space $\mathcal{T}$ is a typical foliation.

A natural way to choose a unique value of $\boldsymbol{\theta}$ for each leaf is to assign r equations

$$c_1(\boldsymbol{\theta}) = 0, \quad ..., \quad c_r(\boldsymbol{\theta}) = 0. \tag{15}$$

We call these equations a *gauge condition*, or a *gauge* for short, if the following are satisfied:

1. They are algebraically independent, defining a submanifold $\mathcal{C}$ of *codimension* r in $\mathcal{T}$; we call $\mathcal{C}$ the *gauge manifold*.
2. The gauge manifold $\mathcal{C}$ intersects *all* leaves $\mathcal{T}_\theta$ *transversally*, with each *connected component* at a single point.
3. For any $\boldsymbol{\theta} \in \mathcal{T}_\theta$ and $\boldsymbol{\theta}_\mathcal{C} = \mathcal{C} \cap \mathcal{T}_\theta$, there exists a *unique* element $g \in \mathcal{G}$ such that $\boldsymbol{\theta}_\mathcal{C} = g\boldsymbol{\theta}$.

Hereafter, we use the terms "gauge" and "gauge manifold" interchangeably, denoting both by $\mathcal{C}$. Introducing a gauge arbitrarily, we can find a solution $\boldsymbol{\theta}$ that minimizes eq. (9) uniquely for each connected component of the leaf $\mathcal{T}_\theta$, but we cannot distinguish two solutions that belong to disjoint components[2].

Example 6. For the orthographic camera model, we can uniquely specify the solution as follows. We choose the origin of the object coordinate system at the centroid of the feature points and align the axis orientation to the first frame of the camera coordinate system and let the Z-components of the translation $\boldsymbol{t}_\kappa$ be all zeros:

$$\sum_{\alpha=1}^{N} \boldsymbol{s}_\alpha = \boldsymbol{0}, \qquad \boldsymbol{R}_1 = \boldsymbol{I}, \qquad (\boldsymbol{k}, \boldsymbol{t}_\kappa) = 0. \tag{16}$$

We call this condition the *standard gauge*. Similarly, we define the standard gauge for the perspective camera model as follows:

$$\boldsymbol{R}_1 = \boldsymbol{I}, \qquad \boldsymbol{t}_1 = \boldsymbol{0}, \qquad \|\boldsymbol{t}_2\| = 1. \tag{17}$$

Example 7. If we differentiate eq. (9) with respect to $\boldsymbol{t}_\kappa$, set the result zero, and solve the resulting equation, we can express $\boldsymbol{t}_\kappa$ in terms of the remaining parameters. Substituting it to $J(\boldsymbol{\theta})$, we obtain a function of the remaining parameters; we call this description the *reduced model*. Similarly, we can eliminate $\{\boldsymbol{s}_\alpha\}$ and express $J(\boldsymbol{\theta})$ in terms of the motion parameters $\{\boldsymbol{R}_\kappa, \boldsymbol{t}_\kappa\}$ alone; we call this description the *epipolar model*.

8. ESTIMATORS

Our goal is to construct a function $\hat{\boldsymbol{\theta}}(\{\boldsymbol{p}_\alpha\})$ of the data $\{\boldsymbol{p}_\alpha\}$ that approximates $\boldsymbol{\theta}$ as closely as possible. Such a mapping $\hat{\boldsymbol{\theta}}$: $\mathcal{R}^{2MN} \to \mathcal{T}$ is called an *estimator* of $\boldsymbol{\theta}$. The solution $\hat{\boldsymbol{\theta}}$ that minimizes eq. (9) is called the *maximum likelihood estimator*. However, we cannot speak of the "true solution" $\boldsymbol{\theta}$; we can only speak of the "true leaf" $\mathcal{T}_\theta$, since any two elements of it are geometrically equivalent. We can make the solution unique by imposing a gauge $\mathcal{C}$; let us denote the resulting solution $\boldsymbol{\theta}_\mathcal{C}$. An estimator determined so that a gauge $\mathcal{C}$ is satisfied is denoted by $\hat{\boldsymbol{\theta}}_\mathcal{C}$. We assume that if noise does not exist, the estimator $\hat{\boldsymbol{\theta}}_\mathcal{C}$ coincides with the value $\boldsymbol{\theta}_\mathcal{C} \in \mathcal{T}_{\boldsymbol{\theta}} \cap \mathcal{C}$ that satisfies that gauge and belongs to the true leaf.

Let $T_{\theta_\mathcal{C}}(\mathcal{T})$ be the *tangent space* to $\mathcal{T}$ at $\boldsymbol{\theta}_\mathcal{C}$. This is an n-dimensional linear space, where, and in the remaining of this paper, n is the dimension of the parameter vector $\boldsymbol{\theta}$ and r is the dimension of the linear space $\mathcal{D}_\theta(\mathcal{T})$. Let $T_{\theta_\mathcal{C}}(\mathcal{C})$ be the tangent space to the gauge manifold $\mathcal{C}$ at $\boldsymbol{\theta}_\mathcal{C}$; it is an $(n-r)$-dimensional subspace of $T_{\theta_\mathcal{C}}(\mathcal{T})$. We assume that noise is very small so that the distribution of $\hat{\boldsymbol{\theta}}_\mathcal{C}$ is localized within a small region around $\boldsymbol{\theta}_\mathcal{C}$ and that the distribution of $\hat{\boldsymbol{\theta}}_\mathcal{C}$ over the gauge manifold $\mathcal{C}$ can be identified with the distribution over the tangent space $T_{\theta_\mathcal{C}}(\mathcal{C})$. Under this assumption, we define the *covariance matrix* of the estimator $\hat{\boldsymbol{\theta}}_\mathcal{C}$ by

$$V[\hat{\boldsymbol{\theta}}_\mathcal{C}] = E[(\hat{\boldsymbol{\theta}}_\mathcal{C} - \boldsymbol{\theta}_\mathcal{C})(\hat{\boldsymbol{\theta}}_\mathcal{C} - \boldsymbol{\theta}_\mathcal{C})^\top]. \tag{18}$$

This is a $n \times n$ singular matrix of rank $n-r$; its range is the tangent space $T_{\theta_\mathcal{C}}(\mathcal{C})$. Since the covariance matrix is defined in the tangent space $T_{\theta_\mathcal{C}}(\mathcal{T})$, it suffices to introduce coordinates to $T_{\theta_\mathcal{C}}(\mathcal{T})$ *locally* rather than to the parameter space $\mathcal{T}$ globally.

Example 8. For the reduced orthographic model, let $\{\hat{\boldsymbol{s}}_\alpha\}$ and $\{\hat{\boldsymbol{R}}_\kappa\}$ be estimators of $\{\boldsymbol{s}_\alpha\}$ and $\{\boldsymbol{R}_\kappa\}$, respectively, defined under gauge $\mathcal{C}$. The deviation of the shape vector $\boldsymbol{s}_\alpha$ is measured by $\Delta\boldsymbol{s}_\alpha = \hat{\boldsymbol{s}}_\alpha - \boldsymbol{s}_\alpha$. The relative rotation from $\boldsymbol{R}_\kappa$ to $\hat{\boldsymbol{R}}_\kappa$ is $\hat{\boldsymbol{R}}_\kappa\boldsymbol{R}_\kappa^\top$, which is a small rotation. Let $\boldsymbol{l}_\kappa$ be its axis (unit vector) and $\Delta\Omega_\kappa$ the angle of rotation around it. We put $\Delta\boldsymbol{\Omega}_\kappa = \Delta\Omega_\kappa\boldsymbol{l}_\kappa$ and take $\{\Delta\boldsymbol{s}_\alpha\}$ and $\{\Delta\boldsymbol{\Omega}_\kappa\}$ as *local coordinates* of $T_{\theta_\mathcal{C}}(\mathcal{T})$. In terms of them, the covariance matrix of $\{\hat{\boldsymbol{s}}_\alpha\}$ and $\{\hat{\boldsymbol{R}}_\kappa\}$ is defined as follows:

$$V[\hat{\boldsymbol{\theta}}_\mathcal{C}] = \begin{pmatrix} E[\Delta\boldsymbol{s}_1\Delta\boldsymbol{s}_1^\top] & \cdots & E[\Delta\boldsymbol{s}_1\Delta\boldsymbol{\Omega}_M^\top] \\ \vdots & & \vdots \\ E[\Delta\boldsymbol{\Omega}_M\Delta\boldsymbol{s}_1^\top] & \cdots & E[\Delta\boldsymbol{\Omega}_M\Delta\boldsymbol{\Omega}_M^\top] \end{pmatrix}. \tag{19}$$

9. GAUGE INVARIANTS

A function I: $\mathcal{T} \to \mathcal{R}$ is said to be a *gauge invariant* if

$$I(g\boldsymbol{\theta}) = I(g'\boldsymbol{\theta}) \tag{20}$$

for all g and g' that belong to the *same connected component* of the gauge group $\mathcal{G}$. It follows that a function I is a gauge invariant *if and only if it is invariant to infinitesimal gauge transformations*. Let γ be an infinitesimal gauge transformation, and $\boldsymbol{D}(\,\cdot\,)$ its generator. Gauge invariance implies

$$\begin{aligned} I(\boldsymbol{\theta}) &= I(\gamma\boldsymbol{\theta}) = I(\boldsymbol{\theta} + \boldsymbol{D}(\boldsymbol{\theta}) + \cdots) \\ &= I(\boldsymbol{\theta}) + (\nabla_\theta I, \boldsymbol{D}(\boldsymbol{\theta})) + \cdots. \end{aligned} \tag{21}$$

Let $\{\boldsymbol{D}_1(\,\cdot\,), ..., \boldsymbol{D}_r(\,\cdot\,)\}$ be an arbitrary basis of the linear space $\mathcal{D}_\theta(\mathcal{T})$. The corresponding vectors $\boldsymbol{D}_1(\boldsymbol{\theta})$, ..., $\boldsymbol{D}_r(\boldsymbol{\theta})$ are the basis of the tangent space $T_\theta(\mathcal{T}_\theta)$ to the leaf $\mathcal{T}_\theta$ associated with $\boldsymbol{\theta}$. Eq. (21) implies that $(\nabla_\theta I, \boldsymbol{D}_i(\boldsymbol{\theta})) = 0$, $i = 1, ..., r$. Thus, we have the following (the superscript $\perp$ denotes orthogonal complement):

Theorem 1 *A function I: $\mathcal{T} \to \mathcal{R}$ is a gauge invariant if and only if*

$$\nabla_\theta I \in T_\theta(\mathcal{T}_\theta)^\perp. \tag{22}$$

Let $\hat{\boldsymbol{\theta}}_\mathcal{C}$ and $\hat{\boldsymbol{\theta}}_{\mathcal{C}'}$ be estimators of $\boldsymbol{\theta}$ with respect to different gauges $\mathcal{C}$ and $\mathcal{C}'$, respectively, and let $V[\hat{\boldsymbol{\theta}}_\mathcal{C}]$ and $V[\hat{\boldsymbol{\theta}}_{\mathcal{C}'}]$ be their respective covariance matrices.

[2]This occurs, for example, for the two mirror image solutions described in Example 4

We say that the covariance matrices $V[\hat{\boldsymbol{\theta}}_{\mathcal{C}}]$ and $V[\hat{\boldsymbol{\theta}}_{\mathcal{C}'}]$ are *geometrically equivalent* if the uncertainty of any gauge invariant is the same whichever covariance matrix we use. Let $I(\boldsymbol{\theta})$ be a gauge invariant. Eq. (21) implies that its corresponding variance is to a first approximation

$$\begin{aligned} V[I] &= E[(\nabla_\theta I|_{\theta_{\mathcal{C}}}, \Delta\boldsymbol{\theta}_{\mathcal{C}})^2] \\ &= (\nabla_\theta I|_{\theta_{\mathcal{C}}}, E[\Delta\boldsymbol{\theta}_{\mathcal{C}}\Delta\boldsymbol{\theta}_{\mathcal{C}}^\top]\nabla_\theta I|_{\theta_{\mathcal{C}}}) \\ &= (\nabla_\theta I|_{\theta_{\mathcal{C}}}, V[\hat{\boldsymbol{\theta}}_{\mathcal{C}}]\nabla_\theta I|_{\theta_{\mathcal{C}}}). \end{aligned} \tag{23}$$

Let $\boldsymbol{\theta}_{\mathcal{C}}$ and $\boldsymbol{\theta}_{\mathcal{C}'}$ be the true values, i.e., the values we would have in the absence of noise, of $\hat{\boldsymbol{\theta}}_{\mathcal{C}}$ and $\hat{\boldsymbol{\theta}}_{\mathcal{C}'}$, respectively. If $\boldsymbol{\theta}_{\mathcal{C}}$ and $\boldsymbol{\theta}_{\mathcal{C}'}$ are geometrically equivalent, there exists a gauge transformation $g \in \mathcal{G}$ such that $\boldsymbol{\theta}_{\mathcal{C}} = g\boldsymbol{\theta}_{\mathcal{C}'}$. This transformation of $\mathcal{T}$ induces a linear mapping from the tangent space $T_{\theta_{\mathcal{C}'}}(\mathcal{C}')$ at $\boldsymbol{\theta}_{\mathcal{C}'}$ to the tangent space $T_{\theta_{\mathcal{C}}}(\mathcal{C})$ at $\boldsymbol{\theta}_{\mathcal{C}}$. We denote it symbolically by $\partial\boldsymbol{\theta}_{\mathcal{C}}/\partial\boldsymbol{\theta}_{\mathcal{C}'}$ and call it the *Jacobian matrix* from $T_{\theta_{\mathcal{C}'}}(\mathcal{C}')$ to $T_{\theta_{\mathcal{C}}}(\mathcal{C})$. Then, $\hat{\boldsymbol{\theta}}_{\mathcal{C}'} = \boldsymbol{\theta}_{\mathcal{C}'} + \Delta\boldsymbol{\theta}_{\mathcal{C}'}$ is geometrically equivalent $\hat{\boldsymbol{\theta}}_{\mathcal{C}} = \boldsymbol{\theta}_{\mathcal{C}} + (\partial\boldsymbol{\theta}_{\mathcal{C}}/\partial\boldsymbol{\theta}_{\mathcal{C}'})\Delta\boldsymbol{\theta}_{\mathcal{C}'}$. If $I(\boldsymbol{\theta})$ is a gauge invariant, its variance can also be computed from this in the form

$$V[I] = (\nabla_\theta I|_{\theta_{\mathcal{C}}}, \frac{\partial\boldsymbol{\theta}_{\mathcal{C}}}{\partial\boldsymbol{\theta}_{\mathcal{C}'}} V[\hat{\boldsymbol{\theta}}_{\mathcal{C}'}]\Big(\frac{\partial\boldsymbol{\theta}_{\mathcal{C}}}{\partial\boldsymbol{\theta}_{\mathcal{C}'}}\Big)^\top \nabla_\theta I|_{\theta_{\mathcal{C}}}). \tag{24}$$

Theorem 1 implies that $V[\hat{\boldsymbol{\theta}}_{\mathcal{C}}]$ and $V[\hat{\boldsymbol{\theta}}_{\mathcal{C}'}]$ are geometrically equivalent if and only if

$$(\boldsymbol{u}, \Big(V[\hat{\boldsymbol{\theta}}_{\mathcal{C}}] - \frac{\partial\boldsymbol{\theta}_{\mathcal{C}}}{\partial\boldsymbol{\theta}_{\mathcal{C}'}} V[\hat{\boldsymbol{\theta}}_{\mathcal{C}'}]\Big(\frac{\partial\boldsymbol{\theta}_{\mathcal{C}}}{\partial\boldsymbol{\theta}_{\mathcal{C}'}}\Big)^\top\Big)\boldsymbol{u}) = 0, \tag{25}$$

for all $\boldsymbol{u} \in T_{\theta_{\mathcal{C}}}(\mathcal{T}_\theta)^\perp$. This means that $T_{\theta_{\mathcal{C}}}(\mathcal{T}_\theta)^\perp$ is the null space of the symmetric matrix $V[\hat{\boldsymbol{\theta}}_{\mathcal{C}}] - \Big(\partial\boldsymbol{\theta}_{\mathcal{C}}/\partial\boldsymbol{\theta}_{\mathcal{C}'}\Big)V[\hat{\boldsymbol{\theta}}_{\mathcal{C}'}]\Big(\partial\boldsymbol{\theta}_{\mathcal{C}}/\partial\boldsymbol{\theta}_{\mathcal{C}'}\Big)^\top$. In other words, it has range $T_{\theta_{\mathcal{C}}}(\mathcal{T}_\theta)$, in which all eigenvectors for nonzero eigenvalues lie. We denote this relation by $V[\hat{\boldsymbol{\theta}}_{\mathcal{C}}] \equiv V[\hat{\boldsymbol{\theta}}_{\mathcal{C}'}] \bmod \mathcal{T}_\theta$. Thus, we have

Theorem 2 *Covariance matrices $V[\hat{\boldsymbol{\theta}}_{\mathcal{C}}]$ and $V[\hat{\boldsymbol{\theta}}_{\mathcal{C}'}]$ defined for different gauges are geometrically equivalent if and only if*

$$V[\hat{\boldsymbol{\theta}}_{\mathcal{C}}] \equiv V[\hat{\boldsymbol{\theta}}_{\mathcal{C}'}] \mod \mathcal{T}_\theta. \tag{26}$$

Example 9. For the reduced orthographic model, the parameter space $\mathcal{T}$ is locally parameterized by $\{\{\Delta\boldsymbol{s}_\alpha\}, \{\Delta\boldsymbol{\Omega}_\kappa\}\}$. Hence, the Jacobian matrix $\partial\boldsymbol{\theta}'/\partial\boldsymbol{\theta}$ is defined as a linear mapping from $\{\{\Delta\boldsymbol{s}_\alpha\}, \{\Delta\boldsymbol{\Omega}_\kappa\}\}$ at $\{\{\boldsymbol{s}_\alpha\}, \{\boldsymbol{R}_\kappa\}\}$ to $\{\{\Delta\boldsymbol{s}'_\alpha\}, \{\Delta\boldsymbol{\Omega}'_\kappa\}\}$ at $\{\{\boldsymbol{s}'_\alpha\}, \{\boldsymbol{R}'_\kappa\}\}$ given the gauge transformation $\{\boldsymbol{R}, \boldsymbol{t}\}$ that maps $\{\{\boldsymbol{s}_\alpha\}, \{\boldsymbol{R}_\kappa\}\}$ to $\{\{\boldsymbol{s}'_\alpha\}, \{\boldsymbol{R}'_\kappa\}\}$. From eqs. (11), we obtain

$$\Delta\boldsymbol{s}'_\alpha = \boldsymbol{R}^\top\Delta\boldsymbol{s}_\alpha, \qquad \Delta\boldsymbol{\Omega}'_\kappa = \Delta\boldsymbol{\Omega}_\kappa. \tag{27}$$

Hence, the Jacobian matrix has the form

$$\frac{\partial\boldsymbol{\theta}'}{\partial\boldsymbol{\theta}} = \mathrm{diag}(\boldsymbol{R}^\top, ..., \boldsymbol{R}^\top, \boldsymbol{I}, ..., \boldsymbol{I}), \tag{28}$$

where $\mathrm{diag}(\boldsymbol{A}, ..., \boldsymbol{B})$ denotes the block diagonal matrix with $\boldsymbol{A}$, ..., $\boldsymbol{B}$ as its diagonal blocks in that order.

10. EQUIVALENCE OF ESTIMATORS

Let $\hat{\boldsymbol{\theta}}_{\mathcal{C}}$ be an estimator of $\boldsymbol{\theta}$ with respect to a gauge $\mathcal{C}$. To a first approximation, the deviation $\Delta\boldsymbol{\theta}_{\mathcal{C}} = \hat{\boldsymbol{\theta}}_{\mathcal{C}} - \boldsymbol{\theta}_{\mathcal{G}}$ can be identified with an element in $T_{\theta_{\mathcal{C}}}(\mathcal{C})$. Let $\Delta\boldsymbol{\theta} \in T_{\theta_{\mathcal{C}}}(\mathcal{T})$ be an arbitrary vector. To a first approximation, $\boldsymbol{\theta}_{\mathcal{C}} + \Delta\boldsymbol{\theta}$ is geometrically equivalent to $\hat{\boldsymbol{\theta}}_{\mathcal{C}} = \boldsymbol{\theta}_{\mathcal{C}} + \Delta\boldsymbol{\theta}_{\mathcal{C}}$ if and only if $\Delta\boldsymbol{\theta} - \Delta\boldsymbol{\theta}_{\mathcal{C}} \in T_{\theta_{\mathcal{C}}}(\mathcal{T}_\theta)$. Since $\{\boldsymbol{D}_1(\boldsymbol{\theta}_{\mathcal{C}}), ..., \boldsymbol{D}_r(\boldsymbol{\theta}_{\mathcal{C}})\}$ is the basis of $T_{\theta_{\mathcal{C}}}(\mathcal{T}_\theta)$, this condition is equivalent to the existence of r numbers x_1, ..., x_r such that $\Delta\boldsymbol{\theta}_{\mathcal{C}} = \Delta\boldsymbol{\theta} + \sum_{i=1}^r x_i\boldsymbol{D}_i(\boldsymbol{\theta}_{\mathcal{C}})$ or

$$\Delta\boldsymbol{\theta}_{\mathcal{C}} = \Delta\boldsymbol{\theta} + \boldsymbol{U}_{\theta_{\mathcal{C}}}\boldsymbol{x}, \tag{29}$$

where we let $\boldsymbol{x} = (x_1, ..., x_r)^\top$ and define

$$\boldsymbol{U}_{\theta_{\mathcal{C}}} = \begin{pmatrix} \boldsymbol{D}_1(\boldsymbol{\theta}_{\mathcal{C}}) & \cdots & \boldsymbol{D}_r(\boldsymbol{\theta}_{\mathcal{C}}) \end{pmatrix}. \tag{30}$$

If the gauge $\mathcal{C}$ is defined by r equations $c_1(\boldsymbol{\theta}) = 0$, ..., $c_r(\boldsymbol{\theta}) = 0$, the tangent space $T_{\theta_{\mathcal{C}}}(\mathcal{C})$ is the orthogonal complement of the linear space spanned by $\{\nabla_\theta c_1|_{\theta_{\mathcal{C}}}, ... \nabla_\theta c_r|_{\theta_{\mathcal{C}}}\}$. It follows that $(\nabla_\theta c_i|_{\theta_{\mathcal{C}}}, \Delta\boldsymbol{\theta}_{\mathcal{C}}) = 0$ for i = 1, ..., r. From eq. (29), we have

$$\boldsymbol{V}_{\theta_{\mathcal{C}}}^\top\Delta\boldsymbol{\theta} + \boldsymbol{V}^\top\boldsymbol{U}_{\theta_{\mathcal{C}}}\boldsymbol{x} = \boldsymbol{0}, \tag{31}$$

where we have defined

$$\boldsymbol{V}_{\theta_{\mathcal{C}}} = \begin{pmatrix} \nabla_\theta c_1|_{\theta_{\mathcal{C}}} & \cdots & \nabla_\theta c_r|_{\theta_{\mathcal{C}}} \end{pmatrix}. \tag{32}$$

From eqs. (29) and (31), we obtain

$$\Delta\boldsymbol{\theta}_{\mathcal{C}} = \Delta\boldsymbol{\theta} - \boldsymbol{U}_{\theta_{\mathcal{C}}}(\boldsymbol{V}_{\theta_{\mathcal{C}}}^\top\boldsymbol{U}_{\theta_{\mathcal{C}}})^{-1}\boldsymbol{V}_{\theta_{\mathcal{C}}}^\top\Delta\boldsymbol{\theta} = \boldsymbol{Q}_{\theta_{\mathcal{C}}}^{\mathcal{C}}\boldsymbol{\theta}, \tag{33}$$

where

$$\boldsymbol{Q}_{\theta_{\mathcal{C}}}^{\mathcal{C}} = \boldsymbol{I} - \boldsymbol{U}_{\theta_{\mathcal{C}}}(\boldsymbol{V}_{\theta_{\mathcal{C}}}^\top\boldsymbol{U}_{\theta_{\mathcal{C}}})^{-1}\boldsymbol{V}_{\theta_{\mathcal{C}}}^\top, \tag{34}$$

which is an (oblique) *projection matrix* onto $T_{\theta_{\mathcal{C}}}(\mathcal{C})$ along $T_{\theta_{\mathcal{C}}}(\mathcal{T}_\theta)$. Thus, we have

Theorem 3 *An estimator $\hat{\boldsymbol{\theta}}_{\mathcal{C}}$ is geometrically equivalent to $\boldsymbol{\theta}_{\mathcal{C}} + \Delta\boldsymbol{\theta}$ to a first approximation if and only if*

$$\hat{\boldsymbol{\theta}}_{\mathcal{C}} = \boldsymbol{\theta}_{\mathcal{C}} + \boldsymbol{Q}_{\theta_{\mathcal{C}}}^{\mathcal{C}}\Delta\boldsymbol{\theta}. \tag{35}$$

Example 10. For the reduced orthographic model, the standard gauge (16) is expressed in the local coordinates $\{\Delta\boldsymbol{s}_\alpha\}$ and $\{\Delta\boldsymbol{\Omega}_\kappa\}$ in the form

$$\sum_{\alpha=1}^N \Delta\boldsymbol{s}_\alpha = \boldsymbol{0}, \qquad \Delta\boldsymbol{\Omega}_1 = \boldsymbol{0}. \tag{36}$$

This means that the six basis vectors of the tangent space to the standard gauge manifold $\mathcal{C}$ are the columns of the matrix

$$\boldsymbol{V}_\theta = \begin{pmatrix} \boldsymbol{I} & \cdots & \boldsymbol{I} & \boldsymbol{O} & \boldsymbol{O} & \cdots \boldsymbol{O} \\ \boldsymbol{O} & \cdots & \boldsymbol{O} & \boldsymbol{I} & \boldsymbol{O} & \cdots \boldsymbol{O} \end{pmatrix}^\top. \tag{37}$$

From eqs. (13), we can express the infinitesimal generator in terms of the local coordinates $\{\Delta\boldsymbol{s}_\alpha\}$ and $\{\Delta\boldsymbol{\Omega}_\kappa\}$ in the form

$$\begin{aligned} \boldsymbol{D}(\boldsymbol{\theta}) &= \boldsymbol{U}_\theta \begin{pmatrix} \Delta\boldsymbol{\Omega} \\ \Delta\boldsymbol{t} \end{pmatrix}, \\ \boldsymbol{U}_\theta &= \begin{pmatrix} \boldsymbol{s}_1\times\boldsymbol{I} & \cdots & \boldsymbol{s}_N\times\boldsymbol{I} & \boldsymbol{R}_1 & \cdots & \boldsymbol{R}_M \\ -\boldsymbol{I} & \cdots & -\boldsymbol{I} & \boldsymbol{O} & \cdots & \boldsymbol{O} \end{pmatrix}^\top. \end{aligned} \tag{38}$$

If $\{\boldsymbol{s}_\alpha\}$ and $\{\boldsymbol{R}_\kappa\}$ satisfy the standard gauge (16), we have

$$\boldsymbol{Q}_\theta^{\mathcal{C}} = \boldsymbol{I} - \begin{pmatrix} \boldsymbol{I}/N & \cdots & \boldsymbol{I}/N & \boldsymbol{s}_1\times\boldsymbol{I} & \boldsymbol{O} & \cdots & \boldsymbol{O} \\ \vdots & & \vdots & \vdots & \vdots & & \vdots \\ \boldsymbol{I}/N & \cdots & \boldsymbol{I}/N & \boldsymbol{s}_N\times\boldsymbol{I} & \boldsymbol{O} & \cdots & \boldsymbol{O} \\ \boldsymbol{O} & \cdots & \boldsymbol{O} & \boldsymbol{I} & \boldsymbol{O} & \cdots & \boldsymbol{O} \\ \boldsymbol{O} & \cdots & \boldsymbol{O} & \boldsymbol{R}_2 & \boldsymbol{O} & \cdots & \boldsymbol{O} \\ \vdots & & \vdots & \vdots & \vdots & \ldots & \vdots \\ \boldsymbol{O} & \cdots & \boldsymbol{O} & \boldsymbol{R}_M & \boldsymbol{O} & \cdots & \boldsymbol{O} \end{pmatrix}. \tag{39}$$

11. COVARIANCE TRANSFORMATIONS

Let $\hat{\boldsymbol{\theta}}_{\mathcal{C}}$ be an estimator of $\boldsymbol{\theta}$ with respect to gauge $\mathcal{C}$, and $V[\hat{\boldsymbol{\theta}}_{\mathcal{C}}]$ its covariance matrix. Let $\hat{\boldsymbol{\theta}}_{\mathcal{C}'}$ be a geometrically equivalent estimator of $\boldsymbol{\theta}$ with respect to another gauge $\mathcal{C}'$, and write

$$\hat{\boldsymbol{\theta}}_{\mathcal{C}} = \boldsymbol{\theta}_{\mathcal{C}} + \Delta\boldsymbol{\theta}_{\mathcal{C}}, \qquad \hat{\boldsymbol{\theta}}_{\mathcal{C}'} = \boldsymbol{\theta}_{\mathcal{C}'} + \Delta\boldsymbol{\theta}_{\mathcal{C}'}. \tag{40}$$

Since $\hat{\boldsymbol{\theta}}_{\mathcal{C}}$ is geometrically equivalent to $\boldsymbol{\theta}_{\mathcal{C}'} + \left(\partial\boldsymbol{\theta}_{\mathcal{C}'}/\partial\boldsymbol{\theta}'_{\mathcal{C}}\right)\Delta\boldsymbol{\theta}_{\mathcal{C}}$, Theorem 3 implies that $\hat{\boldsymbol{\theta}}_{\mathcal{C}}$ and $\hat{\boldsymbol{\theta}}_{\mathcal{C}'}$ are geometrically equivalent if and only if $\Delta\boldsymbol{\theta}_{\mathcal{C}'} = \boldsymbol{Q}^{\mathcal{C}'}_{\theta_{\mathcal{C}'}}(\partial\boldsymbol{\theta}_{\mathcal{C}'}/\partial\boldsymbol{\theta}_{\mathcal{C}})\Delta\boldsymbol{\theta}_{\mathcal{C}}$. Thus, we obtain the following rule of transformation of the covariance matrix induced by a change of the gauge:

Theorem 4 *Let $V[\hat{\boldsymbol{\theta}}_{\mathcal{C}}]$ be the covariance matrix of estimator $\hat{\boldsymbol{\theta}}_{\mathcal{C}}$ under gauge $\mathcal{C}$. For another gauge $\mathcal{C}'$, the covariance matrix of the corresponding estimator $\hat{\boldsymbol{\theta}}_{\mathcal{C}'}$ is given by*

$$V[\hat{\boldsymbol{\theta}}_{\mathcal{C}'}] = \boldsymbol{Q}^{\mathcal{C}'}_{\theta_{\mathcal{C}'}} \frac{\partial\boldsymbol{\theta}_{\mathcal{C}'}}{\partial\boldsymbol{\theta}_{\mathcal{C}}} V[\hat{\boldsymbol{\theta}}_{\mathcal{C}}] \left(\frac{\partial\boldsymbol{\theta}_{\mathcal{C}'}}{\partial\boldsymbol{\theta}_{\mathcal{C}}}\right)^{\top} \boldsymbol{Q}^{\mathcal{C}'\top}_{\theta_{\mathcal{C}'}}. \tag{41}$$

12. UNBIASED ESTIMATORS

Let $p(\{\boldsymbol{p}_\alpha\};\boldsymbol{\theta})$ be the probability density of the data $\{\boldsymbol{p}_\alpha\}$ parameterized by $\boldsymbol{\theta}$. We define its *score* $\boldsymbol{l}_\theta$ by

$$\boldsymbol{l}_\theta = \nabla_\theta \log p. \tag{42}$$

The *Fisher information matrix* is defined by

$$\boldsymbol{J}_\theta = E[\boldsymbol{l}_\theta \boldsymbol{l}_\theta^\top], \tag{43}$$

which is a symmetric positive semi-definite matrix. Because $p(\{\boldsymbol{p}_\alpha\};\boldsymbol{\theta})$ is a gauge invariant, $\log p$ is also a gauge invariant. According to Theorem 1, we have $\nabla_\theta \log p \in T_\theta(\mathcal{T}_\theta)^\perp$ and hence $\boldsymbol{l}_\theta \in T_\theta(\mathcal{T}_\theta)^\perp$. It follows that the Fisher information matrix $\boldsymbol{J}_{\theta_{\mathcal{C}}}$ has the range $T_\theta(\mathcal{T}_\theta)^\perp$ and the null space $T_\theta(\mathcal{T}_\theta)$.

We have so far assumed that an estimator $\hat{\boldsymbol{\theta}}_{\mathcal{C}}$ under gauge $\mathcal{C}$ takes the "true value" $\boldsymbol{\theta}_{\mathcal{C}} \in \mathcal{T}_{\boldsymbol{\theta}} \cap \mathcal{C}$ under that gauge in the absence of noise. In other words, we are implicitly assuming that all estimators are "unbiased". Formally, an estimator $\hat{\boldsymbol{\theta}}$ under gauge $\mathcal{C}$ is *unbiased* if

$$E[\hat{\boldsymbol{\theta}}_{\mathcal{C}} - \boldsymbol{\theta}_{\mathcal{C}}] = \boldsymbol{0}, \tag{44}$$

where we are assuming that the deviation $\hat{\boldsymbol{\theta}}_{\mathcal{C}} - \boldsymbol{\theta}_{\mathcal{C}}$ is so small and localized that it can be identified with an element in the tangent space $T_{\theta_{\mathcal{C}}}(\mathcal{C})$. Eq. (44) should hold for any value $\boldsymbol{\theta}_{\mathcal{U}} \in \mathcal{C}$. In particular, eq. (44) should be invariant to all perturbations in the form $\boldsymbol{\theta}_{\mathcal{C}} \to \boldsymbol{\theta}_{\mathcal{C}} + \delta\boldsymbol{\theta}_{\mathcal{C}}$, $\delta\boldsymbol{\theta}_{\mathcal{C}} \in T_{\theta_{\mathcal{C}}}(\mathcal{C})$. Hence, the corresponding first order variations of eq. (44) should vanish, which implies

$$\left(E[(\hat{\boldsymbol{\theta}}_{\mathcal{C}} - \boldsymbol{\theta}_{\mathcal{C}})\boldsymbol{l}_{\theta_{\mathcal{C}}}^\top] - \boldsymbol{I}\right)\delta\boldsymbol{\theta}_{\mathcal{C}} = \boldsymbol{0} \tag{45}$$

for an arbitrary vector $\delta\boldsymbol{\theta}_{\mathcal{C}} \in T_{\theta_{\mathcal{C}}}(\mathcal{C})$. This is equivalent to

$$E[(\hat{\boldsymbol{\theta}}_{\mathcal{C}} - \boldsymbol{\theta}_{\mathcal{C}})\boldsymbol{l}_{\theta_{\mathcal{C}}}^\top] = \boldsymbol{Q}^{\mathcal{C}}_{\theta_{\mathcal{C}}}. \tag{46}$$

13. CRAMER-RAO INEQUALITY

From identity (46), obtain

$$E[\begin{pmatrix}\hat{\boldsymbol{\theta}}_{\mathcal{C}}-\boldsymbol{\theta}_{\mathcal{C}} \\ \boldsymbol{l}_{\theta_{\mathcal{C}}}\end{pmatrix}\begin{pmatrix}\hat{\boldsymbol{\theta}}_{\mathcal{C}}-\boldsymbol{\theta}_{\mathcal{C}} \\ \boldsymbol{l}_{\theta_{\mathcal{C}}}\end{pmatrix}^{\top}] = \begin{pmatrix} V[\hat{\boldsymbol{\theta}}_{\mathcal{C}}] & \boldsymbol{Q}^{\mathcal{C}}_{\theta_{\mathcal{C}}} \\ \boldsymbol{Q}^{\mathcal{C}\top}_{\theta_{\mathcal{C}}} & \boldsymbol{Q}^{\mathcal{C}\top}_{\theta_{\mathcal{C}}}\boldsymbol{J}_{\theta_{\mathcal{C}}}\boldsymbol{Q}^{\mathcal{C}}_{\theta_{\mathcal{C}}} \end{pmatrix}. \tag{47}$$

This is a symmetric and positive semi-definite matrix by construction. Hence, the following is also a symmetric and positive semi-definite matrix:

$$\begin{pmatrix} \boldsymbol{Q}^{\mathcal{C}}_{\theta_{\mathcal{C}}} & -\boldsymbol{Q}^{\mathcal{C}}_{\theta_{\mathcal{C}}}\boldsymbol{J}^-_{\theta_{\mathcal{C}}} \\ \boldsymbol{O} & \boldsymbol{Q}^{\mathcal{C}}_{\theta_{\mathcal{C}}}\boldsymbol{J}^-_{\theta_{\mathcal{C}}} \end{pmatrix} \begin{pmatrix} V[\hat{\boldsymbol{\theta}}_{\mathcal{C}}] & \boldsymbol{Q}^{\mathcal{C}}_{\theta_{\mathcal{C}}} \\ \boldsymbol{Q}^{\mathcal{C}\top}_{\theta_{\mathcal{C}}} & \boldsymbol{J}_{\theta_{\mathcal{C}}} \end{pmatrix} \begin{pmatrix} \boldsymbol{Q}^{\mathcal{C}\top}_{\theta_{\mathcal{C}}} & \boldsymbol{O} \\ -\boldsymbol{J}^-_{\theta_{\mathcal{C}}}\boldsymbol{Q}^{\mathcal{C}\top}_{\theta_{\mathcal{C}}} & \boldsymbol{J}^-_{\theta_{\mathcal{C}}}\boldsymbol{Q}^{\mathcal{C}\top}_{\theta_{\mathcal{C}}} \end{pmatrix}. \tag{48}$$

Here, the superscript "$-$" denotes (Moore-Penrose) generalized inverse. If we note that $V[\hat{\boldsymbol{\theta}}_{\mathcal{C}}]$ has the range $T_{\theta_{\mathcal{C}}}(\mathcal{C})$ and hence $\boldsymbol{Q}^{\mathcal{C}}_{\theta_{\mathcal{C}}}V[\hat{\boldsymbol{\theta}}_{\mathcal{C}}]\boldsymbol{Q}^{\mathcal{C}\top}_{\theta_{\mathcal{C}}} = V[\hat{\boldsymbol{\theta}}_{\mathcal{C}}]$ and use the identity $\boldsymbol{J}^-_{\theta_{\mathcal{C}}}\boldsymbol{J}_{\theta_{\mathcal{C}}}\boldsymbol{J}^-_{\theta_{\mathcal{C}}} = \boldsymbol{J}^-_{\theta_{\mathcal{C}}}$, we see that the above matrix equals

$$\begin{pmatrix} V[\hat{\boldsymbol{\theta}}_{\mathcal{C}}] - \boldsymbol{Q}^{\mathcal{C}}_{\theta_{\mathcal{C}}}\boldsymbol{J}^-_{\theta_{\mathcal{C}}}\boldsymbol{Q}^{\mathcal{C}\top}_{\theta_{\mathcal{C}}} & \boldsymbol{O} \\ \boldsymbol{O} & \boldsymbol{Q}^{\mathcal{C}}_{\theta_{\mathcal{C}}}\boldsymbol{J}^-_{\theta_{\mathcal{C}}}\boldsymbol{Q}^{\mathcal{C}\top}_{\theta_{\mathcal{C}}} \end{pmatrix}. \tag{49}$$

Since this is a symmetric positive semi-definite matrix, so is $\boldsymbol{Q}^{\mathcal{C}}_{\theta_{\mathcal{C}}}\boldsymbol{J}^-_{\theta_{\mathcal{C}}}\boldsymbol{Q}^{\mathcal{C}\top}_{\theta_{\mathcal{C}}}$. Thus, we obtain the following theorem:

Theorem 5 *For an estimator $\hat{\boldsymbol{\theta}}_{\mathcal{C}}$ under gauge $\mathcal{C}$, the following inequality holds:*

$$V[\hat{\boldsymbol{\theta}}_{\mathcal{C}}] \succ \boldsymbol{Q}^{\mathcal{C}}_{\theta_{\mathcal{C}}}\left(\boldsymbol{J}_{\theta_{\mathcal{C}}}\right)^-_{n-r}\boldsymbol{Q}^{\mathcal{C}\top}_{\theta_{\mathcal{C}}}. \tag{50}$$

Here, the inequality $\boldsymbol{A} \succ \boldsymbol{B}$ for symmetric matrices $\boldsymbol{A}$ and $\boldsymbol{B}$ means that $\boldsymbol{A} - \boldsymbol{B}$ is a symmetric positive semi-definite matrix. Eq. (50) is an extension of the so called *Cramer-Rao inequality* [2, 4] to statistical estimation with internal indeterminacy. Let us call the right-hand side the *Cramer-Rao lower bound under gauge* $\mathcal{C}$.

14. FREE-GAUGE APPROACH

The Taylor expansion of the function J in the neighborhood of the true value $\bar{\boldsymbol{\theta}}_{\mathcal{C}}$ under gauge $\mathcal{C}$ has the form

$$J(\bar{\boldsymbol{\theta}}_{\mathcal{C}} + \Delta\boldsymbol{\theta}) = \bar{J} + (\nabla_\theta \bar{J}, \Delta\boldsymbol{\theta}) + \frac{1}{2}(\Delta\boldsymbol{\theta}, \nabla^2_\theta \bar{J}\Delta\boldsymbol{\theta}) + \cdots. \tag{51}$$

The bars over J in the expressions $\bar{J}$, $\nabla_\theta\bar{J}$, and $\nabla^2_\theta\bar{J}$ mean that these expressions are evaluated at the true value $\boldsymbol{\theta} = \bar{\boldsymbol{\theta}}_{\mathcal{C}}$.

Ignoring terms of order $O(\epsilon^3)$ or higher, differentiating eq. (51) with respect to $\Delta\boldsymbol{\theta}$, letting the result be zero, and substituting an estimate $\boldsymbol{\theta}_C$ for the true value $\bar{\boldsymbol{\theta}}_C$, we obtain

$$\nabla_\theta J + \nabla^2_\theta J \Delta\boldsymbol{\theta} = \boldsymbol{0}. \tag{52}$$

Since the rank of $\nabla^2_\theta J$ is r, this equation has infinitely many solutions. Among them, we can choose an arbitrary one. Let us choose the one that has minimum norm $\|\Delta\boldsymbol{\theta}\|$. This is given by simply computing the (Moore-Penrose) generalized inverse:

$$\Delta\boldsymbol{\theta} = -\left(\nabla^2_\theta J\right)^-_{n-r}\nabla_\theta J. \tag{53}$$

The subscript "$n-r$" means that the smallest r eigenvalues are replaced by zeros [2]. If we want to enforce a gauge $\mathcal{C}$, we replace $\Delta\boldsymbol{\theta}$ by $\boldsymbol{Q}^{\mathcal{C}}_\theta\Delta\boldsymbol{\theta}$ according to Theorem 3. However, this enforcement is not necessary at each step of the iterations. We can alternatively repeat the above update until the solution converges and then transform the solution so that it satisfies the gauge $\mathcal{C}$. We call this approach, first discussed by Triggs [9], the *free-gauge approach*.

15. NORMAL FORM

So far, the covariance matrix of an estimators of $\boldsymbol{\theta}$ is defined for a particular gauge. However, it can be defined *independently* of gauges. Let $\boldsymbol{\theta} \in \mathcal{T}_\theta$ be an arbitrary true value, and $\hat{\boldsymbol{\theta}}$ its estimator under an arbitrary gauge. Since $\boldsymbol{\theta}$ and $\hat{\boldsymbol{\theta}}$ may have very different values because of the gauge freedom, we cannot compare them directly. So, we apply an appropriate gauge transformation g such that $g\hat{\boldsymbol{\theta}}$ becomes the "closest" to $\boldsymbol{\theta}$ and then evaluate $E[(g\hat{\boldsymbol{\theta}}-\boldsymbol{\theta})(g\hat{\boldsymbol{\theta}}-\boldsymbol{\theta})^\top]$.

The formal description of this procedure is as follows: Define a gauge manifold $\mathcal{C}^*$ such that it passes through $\boldsymbol{\theta}$ and the tangent space $T_\theta(\mathcal{C}^*)$ to $\mathcal{C}^*$ coincides with the *orthogonal complement* $T_\theta(\mathcal{T}_\theta)^\perp$ of the tangent space $T_\theta(\mathcal{T}_\theta)$ spanned by the infinitesimal generators $\{\boldsymbol{D}_i(\boldsymbol{\theta})\}$ of the linear space $\mathcal{D}_\theta(\mathcal{T})$ of the gauge transformation group $\mathcal{G}$. We call the covariance matrix $V[\hat{\boldsymbol{\theta}}^*_{\mathcal{C}}]$ of the estimator $\hat{\boldsymbol{\theta}}_{\mathcal{C}^*}$ under that gauge the *normal form* and denoted it by $V^*[\hat{\boldsymbol{\theta}}]$. Note that the normal form is defined for *all* values that belong to the leaf $\mathcal{T}_\theta$ of true values.

Let $\hat{\boldsymbol{\theta}}_{\mathcal{C}}$ be an estimator equivalent to $\hat{\boldsymbol{\theta}}_{\mathcal{C}^*}$ defined under another gauge $\mathcal{C}$ that also passes through $\boldsymbol{\theta}$. Let $V[\hat{\boldsymbol{\theta}}_{\mathcal{C}}]$ be its covariance matrix. According to Theorem 4, the normal form $V^*[\hat{\boldsymbol{\theta}}]$ is given by

$$V^*[\hat{\boldsymbol{\theta}}] = \boldsymbol{P}_\theta V[\hat{\boldsymbol{\theta}}_{\mathcal{C}}]\boldsymbol{P}_\theta^\top, \tag{54}$$

where $\boldsymbol{P}_\theta$ is the (orthogonal) *projection matrix* onto $T_\theta(\mathcal{T}_\theta)^\perp$. It has the expression

$$\boldsymbol{P}_\theta = \boldsymbol{I} - \boldsymbol{U}_\theta(\boldsymbol{U}_\theta^\top\boldsymbol{U}_\theta)^{-1}\boldsymbol{U}_\theta^\top, \tag{55}$$

where $\boldsymbol{U}_\theta$ is the $n\times n$ matrix consisting of $\boldsymbol{D}_1(\boldsymbol{\theta})$,..., $\boldsymbol{D}_r(\boldsymbol{\theta})$ as its columns (see the first of eqs. (30)). Conversely, $V[\hat{\boldsymbol{\theta}}_{\mathcal{C}}]$ is given in terms of its normal form $V^*[\hat{\boldsymbol{\theta}}]$ in the form

$$V[\hat{\boldsymbol{\theta}}_{\mathcal{C}}] = \boldsymbol{Q}_\theta^{\mathcal{C}} V^*[\hat{\boldsymbol{\theta}}]\boldsymbol{Q}_\theta^{\mathcal{C}\top}. \tag{56}$$

In general, we have

Theorem 6 *Let $\hat{\boldsymbol{\theta}}_{\mathcal{C}}$ be an estimator defined under a gauge $\mathcal{C}$, and let $\boldsymbol{\theta}_{\mathcal{C}} = \mathcal{T}_{\boldsymbol{\theta}}\cap\mathcal{C}$ be the value of $\boldsymbol{\theta}_{\mathcal{C}}$ that satisfies that gauge and belongs to the true leaf. Let $V[\hat{\boldsymbol{\theta}}_{\mathcal{C}}]$ be its covariance matrix. Its normal form at an arbitrary point $\boldsymbol{\theta}\in\mathcal{T}_\theta$ is*

$$V^*[\hat{\boldsymbol{\theta}}] = \boldsymbol{P}_\theta\frac{\partial\boldsymbol{\theta}}{\partial\boldsymbol{\theta}_{\mathcal{C}}}V[\hat{\boldsymbol{\theta}}_{\mathcal{C}}]\Big(\frac{\partial\boldsymbol{\theta}}{\partial\boldsymbol{\theta}_{\mathcal{C}}}\Big)^\top\boldsymbol{P}_\theta^\top. \tag{57}$$

Conversely, $V[\hat{\boldsymbol{\theta}}_{\mathcal{C}}]$ is given in terms of its normal form at at $\boldsymbol{\theta}\in\mathcal{T}_\theta$ in the form

$$V[\hat{\boldsymbol{\theta}}_{\mathcal{C}}] = \boldsymbol{Q}_{\theta_{\mathcal{C}}}^{\mathcal{C}}\frac{\partial\boldsymbol{\theta}_{\mathcal{C}}}{\partial\boldsymbol{\theta}}V^*[\hat{\boldsymbol{\theta}}]\Big(\frac{\partial\boldsymbol{\theta}_{\mathcal{C}}}{\partial\boldsymbol{\theta}}\Big)^\top\boldsymbol{Q}_{\theta_{\mathcal{C}}}^{\mathcal{C}\top}. \tag{58}$$

Recall that the Cramer-Rao lower bound given by eq. (50) is defined for every true value $\boldsymbol{\theta}$ that belongs to the true leaf for any given gauge $\mathcal{C}$. Suppose eq. (50) holds for $\boldsymbol{\theta}$. If the gauge $\mathcal{C}$ is such that $T_\theta(\mathcal{C}) = T_\theta(\mathcal{T}_\theta)^\perp$, eq. (50) can be read as giving a lower bound on the covariance matrix of $\hat{\boldsymbol{\theta}}$ in the normal form. But if $T_\theta(\mathcal{C}) = T_\theta(\mathcal{T}_\theta)^\perp$, then $\boldsymbol{Q}_{\theta_{\mathcal{C}}}^{\mathcal{C}} = \boldsymbol{P}_\theta$. On the other hand, the Fisher information matrix $\boldsymbol{J}_\theta$ has the range $T_\theta(\mathcal{T}_\theta)^\perp$, so we obtain the following Cramer-Rao inequality:

Theorem 7 *For an estimator $\hat{\boldsymbol{\theta}}$, the following inequality holds for the normal form of its covariance matrix:*

$$V^*[\hat{\boldsymbol{\theta}}] \succ \Big(\boldsymbol{J}_\theta\Big)^-_{n-r}. \tag{59}$$

Since the maximum likelihood estimator attains the Cramer-Rao lower bound to a first approximation, we can approximately evaluate the Cramer-Rao lower bound by evaluating $\boldsymbol{J}_\theta$ at the computed estimate $\hat{\boldsymbol{\theta}}$. If we ignore terms of order $O(\epsilon^3)$ or higher in computing the maximum likelihood solution, it is easy to see that the Fisher information matrix takes the form

$$\boldsymbol{J}_\theta|_{\hat{\theta}} = \frac{1}{2}\nabla_\theta^2 J|_{\hat{\theta}}. \tag{60}$$

It follows that the free-gauge approach allows us to obtain not only an maximum likelihood solution $\hat{\boldsymbol{\theta}}$ but also its covariance matrix in the normal form as follows:

$$V^*[\hat{\boldsymbol{\theta}}] = -2\Big(\nabla_\theta^2 J|_{\hat{\theta}}\Big)^-_{n-r}. \tag{61}$$

The free-gauge approach is advantageous when the gauge manifold $\mathcal{C}$ is "nearly parallel" to the leaves $\mathcal{T}_\theta$, $\boldsymbol{\theta}\in\mathcal{T}$, since the numerical solution must travel along a long path to reach the solution leaf $\mathcal{T}_{\bar{\theta}}$ if the trajectory is constrained to be in $\mathcal{C}$. In contrast, the free-gauge approach allows trajectories to intersect the individual leaves $\mathcal{T}_\theta$, $\boldsymbol{\theta}\in\mathcal{T}$, "orthogonally", so the true leaf $\mathcal{T}_{\bar{\theta}}$ is reached along the shortest path [5, 6].

16. CONCLUDING REMARKS

We have presented a consistent theory for describing indeterminacy and uncertainty of 3-D reconstruction from a sequence of images. First, we presented a group-theoretical analysis of gauge transformations and gauges by invoking the Lie group theory. We have extended the Cramer-Rao lower bound to problems with internal indeterminacy. Finally, we described the free-gauge approach and the normal form of a covariance matrix that is independent of particular gauges. Applications of our theory to real images are given in [7].

Acknowledgments. The authors thank Bill Triggs of INRIA Rône-Alpes, France, and Philip McLauchlan of the University of Surrey, U.K., for helpful discussions. This work was in part supported by the Ministry of Education, Science, Sports and Culture, Japan under a Grant in Aid for Scientific Research C(2) (No. 11680377).

References

[1] K. Kanatani, *Group-Theoretical Methods in Image Understanding*, Springer, Berlin, 1990.

[2] K. Kanatani, *Statistical Optimization for Geometric Computation: Theory and Practice*, Elsevier, Amsterdam, 1996.

[3] K. Kanatani, Statistical optimization and geometric inference in computer vision, *Phi. Trans. R. Soc. Lond.*, A**356**-1740 (1998), 1303–1320.

[4] K. Kanatani, Cramer-Rao lower bounds for curve fitting, *Graphical Models Image Process.* **60**-2 (1998), 93–99.

[5] P. F. McLauchlan, Gauge invariance in projective 3D reconstruction, *IEEE Workshop on Multi-View Modeling and Analysis of Visual Scenes*, June 1999, Fort Collins, CO, U.S.A..

[6] P. F. McLauchlan, Coordinate-frame independence in optimization algorithms for 3D vision, *IEEE Workshop on Vision Algorithms: Theory and Practice*, September 1999, Corfu, Greece.

[7] D. D. Morris, K. Kanatani and T. Kanade, Uncertainty modeling for optimal structure from motion, *IEEE Workshop on Vision Algorithms: Theory and Practice*, September 1999, Corfu, Greece.

[8] R. Szeliski and S. B. Kang, Shape ambiguities in structure from motion, *IEEE Trans. Patt. Anal. Mach. Intell.*, **19**-5 (1997), pp. 506–512.

[9] B. Triggs, Optimal estimation of matching constraint, in R. Koch and L. Van Gool (eds.), *3D Structure from Multiple Images of Large-Scale Environments*, Lecture Notes in Computer Science, 1506, Springer, 1998, pp. 63–77.

[10] J. Weng, N. Ahuja and T. S. Huang, Optimal motion and structure estimation, *3D IEEE Trans. Patt. Anal. Mach. Intell.*, **15**-9 (1993), 864–884.