

Uncertainty Modeling and Geometric Inference

Kenichi Kanatani

Department of Computer Science, Okayama University, Okayama 700-8530 Japan
kanatani@suri.it.okayama-u.ac.jp

14.1 Introduction

Statistical inference from images is one of the key components of computer vision research today. Traditionally, statistical methods have been used for recognition and classification purposes. Recently, however, there are many studies of statistical analysis for *geometric inference* based on geometric primitives such as points and lines extracted by image processing operations.

However, the term “statistical” has somewhat a different meaning for such geometric inference problems than for the traditional recognition and classification purposes. This difference has often been overlooked, causing controversies over the validity of the statistical approach to geometric problems in general. In Sect. , we take a close look at this problem, tracing back the origin of feature uncertainty to image processing operations. In Sect. , we discuss the implications of asymptotic analysis in reference to geometric fitting and geometric model selection. In Sect. , we point out that a correspondence exists between the standard statistical analysis and the geometric inference problem. We also compare the capability of the geometric AIC and the geometric MDL in detecting degeneracy. In Sect. , we review recent progress in geometric fitting techniques for linear constraints, describing the FNS method, the HEIV method, the renormalization method, and other related techniques. In Sect. , we discuss the Neyman–Scott problem and semiparametric models in relation to geometric inference. Sect. presents our concluding remarks.

14.2 What Is Geometric Inference?

14.2.1 Ensembles for Geometric Inference

The goal of statistical methods is not to study the properties of observed data themselves but to infer the properties of the *ensemble* from which we




Fig. 1. **a** A feature point in an image of a building. **b** Its enlargement and the uncertainty of the feature location

regard the observed data as sampled. The ensemble may be a collection of existing entities (e.g., the entire population), but often it is a hypothetical set of conceivable possibilities. When a statistical method is employed, the underlying ensemble is often taken for granted. However, this issue is very crucial for geometric inference based on feature points.

Suppose, for example, we extract feature points, such as corners of walls and windows, from an image of a building and want to test if they are collinear. The reason why we need a statistical method is that the extracted feature positions have uncertainty. So, we have to judge the extracted feature points as collinear if they are sufficiently aligned. We can also evaluate the degree of uncertainty of the fitted line by propagating the uncertainty of the individual points. What is the ensemble that underlies this type of inference?

This question reduces to the question of why the uncertainty of the feature points occurs at all. After all, statistical methods are not necessary if the data are exact. Using a statistical method means regarding the current feature position as sampled from a set of its possible positions. But where else could it be if not in the current position?

14.2.2 Uncertainty of Feature Extraction

Many algorithms have been proposed for extracting feature points, including the Harris operator [12] and SUSAN [46], and their performance has been extensively compared [4, 41, 45]. However, if we use, for example, the Harris operator to extract a particular corner of a particular building image, the output is unique (Fig. 1). No matter how many times we repeat the extraction, we obtain the same point because no external disturbances exist and the internal parameters (e.g., thresholds for judgment) are unchanged. It follows that the current position is the sole possibility. How can we find it elsewhere?

If we closely examine the situation, we are compelled to conclude that other possibilities should exist because the extracted position is not necessarily correct. But if it is not correct, why did we extract it? Why didn't we extract the correct position in the first place? The answer is: *we cannot*.

14.2.3 Image Processing for Computer Vision

The reason why there exist so many feature extraction algorithms, none of them being definitive, is that they are aiming at an intrinsically impossible task. If we were to extract a point around which, say, the intensity varies to the largest degree in such and such a measure, the algorithm would be unique; variations may exist in intermediate steps, but the final output should be the same.

However, what we want is not “image properties” but “3-D properties” such as corners of a building, but the way a 3-D property is translated into an image property is intrinsically heuristic. As a result, as many algorithms can exist as the number of heuristics for its 2-D interpretation. If we specify a particular 3-D feature to extract, say a corner of a window, its appearance in the image is not unique. It is affected by many properties of the scene, including the details of its 3-D shape, the viewing orientation, the illumination condition, and the light reflectance properties of the material. A slight variation of any of them can result in a substantial difference in the image.

Theoretically, exact extraction would be possible if all the properties of the scene were exactly known, but to infer them from images is the very task of computer vision. It follows that we must make a guess in the image processing stage. For the current image, some guesses may be correct, but others may be wrong. The exact feature position could be found only by an (nonexisting) “ideal” algorithm that could guess everything correctly.

This observation allows us to interpret the “possible feature positions” to be *the positions that would be located by different (nonideal) algorithms based on different guesses*. It follows that the set of hypothetical positions should be associated with *the set of hypothetical algorithms*. The current position is regarded as produced by an algorithm sampled from it. This explains why one always obtains the same position no matter how many times one repeats extraction using that algorithm. To obtain a different position, one has to sample another algorithm.

Remark 1. We may view the statistical ensemble in the following way. If we repeat the *same* experiment, the result should always be the same. But if we declare that the experiment is the “same” if such and such are the same while other things can vary, then those variable conditions define the ensemble. The conventional view is to regard the experiment as the same if the *3-D scene* we are viewing is the same while other properties, such as the lighting condition, can vary. Then, the resulting image would be different for each (hypothetical) experiment, so one would obtain a different output each time, using the same image processing algorithm. The expected spread of the outputs measures the robustness of that algorithm. Here, however, we are viewing the experiment as the same *if the image is the same*. Then, we could obtain different results only by sampling other algorithms. The expected spread of the outputs measures

the uncertainty of feature detection from *that image*. We take this view because we are analyzing the reliability of geometric inference from a particular image, while the conventional view is suitable for assessing the robustness of a *particular algorithm*.

14.2.4 Covariance Matrix of a Feature Point

The performance of feature point extraction depends on the image properties around that point. If, for example, we want to extract a point in a region with an almost homogeneous intensity, the resulting position may be ambiguous whatever algorithm is used. In other words, the positions that potential algorithms would extract should have a large spread. If, on the other hand, the intensity greatly varies around that point, any algorithm could easily locate it accurately, meaning that the positions that the hypothetical algorithms would extract should have a strong peak. It follows that we may introduce for each feature point its *covariance matrix* that measures the spread of its potential positions.

Let $V[p_\alpha]$ be the covariance matrix of the α th feature point p_α . The above argument implies that we can estimate the qualitative characteristics of uncertainty but not its absolute magnitude. So, we write the covariance matrix $V[p_\alpha]$ in the form

$$V[p_\alpha] = \varepsilon^2 V_0[p_\alpha], \quad (1)$$

where ε is an unknown magnitude of uncertainty, which we call the *noise level*. The matrix $V_0[p_\alpha]$, which we call the (*scale*) *normalized covariance matrix*, describes the relative magnitude and the dependence on orientations.

Remark 2. The decomposition of $V[p_\alpha]$ into ε^2 and $V_0[p_\alpha]$ involves scale ambiguity. We assume that the decomposition is made unique by an appropriate scale normalization such as $\text{tr}V_0[p_\alpha] = 2$. However, the subsequent analysis does not depend on particular normalizations, so we do not explicitly specify it except that it should be done in such a way that ε is much smaller than the data themselves. Note that mathematically, modeling the covariance matrix by a common scale factor ε^2 and the individual matrix part $V_0[p_\alpha]$ is rather restrictive. However, this model is sufficient for most practical applications, as we describe in the following.

14.2.5 Covariance Matrix Estimation

If the intensity variations around p_α are almost the same in all directions, we can think of the probability distribution as isotropic, a typical equiprobability line, known as the *uncertainty ellipses*, being a circle (Fig. 1b). On the other hand, if p_α is on an object boundary, distinguishing it from nearby points

should be difficult whatever algorithm is used, so its covariance matrix should have an elongated uncertainty ellipse along that boundary.

However, existing feature extraction algorithms are usually designed to output those points that have large image variations around them, so points in a region with an almost homogeneous intensity or on object boundaries are rarely chosen. As a result, the covariance matrix of a feature point extracted by such an algorithm can be regarded as nearly isotropic. This has also been confirmed by experiments [26], justifying the use of the identity as the normalized covariance matrix $V_0[p_\alpha]$.

Remark 3. The intensity variations around different feature points are usually unrelated, so their uncertainty can be regarded as statistically independent. However, if we track feature points over consecutive video frames, it has been observed that the uncertainty has strong correlations over the frames [47].

Remark 4. Many interactive applications require humans to extract feature points by manipulating a mouse. Extraction by a human is also an “algorithm”, and it has been shown by experiments that humans are likely to choose “easy-to-see” points such as isolated points and intersections, avoiding points in a region with an almost homogeneous intensity or on object boundaries [26]. In this sense, the statistical characteristics of human extraction are very similar to machine extraction. This is no surprise if we recall that image processing for computer vision is essentially a heuristic that simulates human perception. It has also been reported that strong microscopic correlations exist when humans manually select corresponding feature points over multiple images [34].

14.2.6 Image Quality and Uncertainty

The uncertainty of feature points has often been identified with “image noise”, giving a misleading impression as if the feature locations were perturbed by random intensity fluctuations. Of course, we may obtain better results using higher-quality images whatever algorithm is used. However, the task of computer vision is not to analyze image properties but to study the 3-D properties of the scene. As long as the image properties and the 3-D properties do not correspond one to one, any image processing inevitably entails some degree of uncertainty, however high the image quality may be, and the result must be interpreted statistically. The underlying ensemble is the set of hypothetical (inherently imperfect) algorithms of image processing. Yet, the performance of image processing algorithms has often been evaluated by adding *independent Gaussian noise* to individual pixels.

Remark 5. This also applies to *edge detection*, whose goal is to find the boundaries of 3-D objects in the scene. In reality, all existing algorithms seek *edges*,




Fig. 2. **a** For the standard statistical analysis, it is desired that the accuracy increases rapidly as the number of experiments $n \rightarrow \infty$, because admissible accuracy can be reached with a smaller number of experiments. **b** For geometric inference, it is desired that the accuracy increases rapidly as the noise level $\varepsilon \rightarrow 0$, because larger data uncertainty can be tolerated for admissible accuracy

i.e., lines and curves across which the intensity changes discontinuously. Yet, this is regarded by many as an objective image processing task, and the detection performance is often evaluated by adding independent Gaussian noise to individual pixels. From the above considerations, we conclude that edge detection is also a heuristic, and hence no definitive algorithm will ever be found.

14.3 Asymptotic Analysis for Geometric Inference

14.3.1 What Is Asymptotic Analysis?

As stated earlier, *statistical estimation* refers to estimating the properties of an ensemble from a finite number of samples, assuming some knowledge, or a *model*, about the ensemble. If the uncertainty originates from external conditions, as in experiments in physics, the estimation accuracy can be increased by controlling the measurement devices and environments. For internal uncertainty, on the other hand, there is no way of increasing the accuracy except by repeating the experiment and doing statistical inference. However, repeating experiments usually entails costs, and in practice the number of experiments is often limited.

Taking account of this, statisticians usually evaluate the performance of estimation *asymptotically*, analyzing the growth in accuracy as the number n of experiments increases. This is justified because a method whose accuracy increases more rapidly as $n \rightarrow \infty$ can reach admissible accuracy *with fewer experiments* (Fig. 2a).

In contrast, the ensemble for geometric inference is, as we have seen, the set of potential feature positions that could be located if other (hypothetical) algorithms were used. As noted earlier, however, we can choose only *one* sample from the ensemble as long as we use a particular image processing

algorithm. In other words, the number n of experiments is 1. Then, how can we evaluate the performance of statistical estimation?

Evidently, we want a method whose accuracy is sufficiently high *even for large data uncertainty*. This implies that we need to analyze the growth in accuracy as the noise level ε decreases, because a method whose accuracy increases more rapidly as $\varepsilon \rightarrow 0$ can tolerate larger data uncertainty for admissible accuracy (Fig. 2b).

14.3.2 Geometric Fitting

We now illustrate the above consideration in more specific terms. Let $\{p_\alpha\}$, $\alpha = 1, \dots, N$, be the extracted feature points. Suppose each point should satisfy a parameterized constraint

$$F(p_\alpha, \mathbf{u}) = 0 \quad (2)$$

when no uncertainty exists. In the presence of uncertainty, Eq. (2) may not hold exactly. Our task is to estimate the parameter \mathbf{u} from observed positions $\{p_\alpha\}$ in the presence of uncertainty.

A typical problem of this form is to fit a line or a curve to given N points in the image, but this can be straightforwardly extended to multiple images. For example, if a point (x_α, y_α) in one image corresponds to a point (x'_α, y'_α) in another, we can regard them as a single point p_α in a four-dimensional joint space with coordinates $(x_\alpha, y_\alpha, x'_\alpha, y'_\alpha)$. If the camera imaging geometry is modeled as perspective projection, constraint (2) corresponds to the *epipolar equation*; the parameter \mathbf{u} is the *fundamental matrix* [13]. This is discussed in more detail in Sect. .

14.3.2.1 General Geometric Fitting

The above problem can be stated in abstract terms as *geometric fitting* as follows. We view a feature point in the image plane or a set of feature points in the joint space as an m -dimensional vector \mathbf{x} ; we call it a “datum”. Let $\{\mathbf{x}_\alpha\}$, $\alpha = 1, \dots, N$, be observed data. Their true values $\{\bar{\mathbf{x}}_\alpha\}$ are supposed to satisfy r constraint equations

$$F^{(k)}(\bar{\mathbf{x}}_\alpha, \mathbf{u}) = 0, \quad k = 1, \dots, r, \quad (3)$$

parameterized by a p -dimensional vector \mathbf{u} . We call Eq. (3) the (*geometric model*). The domain \mathcal{X} of the data $\{\mathbf{x}_\alpha\}$ is called the *data space*; the domain \mathcal{U} of the parameter \mathbf{u} is called the *parameter space*. The number r of the constraint equations is called the *rank* of the constraint. The r equations $F^{(k)}(\mathbf{x}, \mathbf{u}) = 0$, $k = 1, \dots, r$, are assumed to be mutually independent, defining a manifold \mathcal{S} of codimension r parameterized by \mathbf{u} in the data space \mathcal{X} . Equation (3) requires that the true values $\{\bar{\mathbf{x}}_\alpha\}$ be all in the manifold \mathcal{S} . Our task is to estimate the parameter \mathbf{u} from the noisy data $\{\mathbf{x}_\alpha\}$ (Fig. 3a).




Fig. 3. **a** Fitting a manifold \mathcal{S} to the data $\{\mathbf{x}_\alpha\}$. **b** Estimating $\{\bar{\mathbf{x}}_\alpha\}$ and \mathbf{u} by minimizing the sum of squared Mahalanobis distance with respect to the normalized covariance matrices $\mathbf{V}_0[\mathbf{x}_\alpha]$

14.3.2.2 Maximum Likelihood Estimation

Let

$$\mathbf{V}[\mathbf{x}_\alpha] = \varepsilon^2 \mathbf{V}_0[\mathbf{x}_\alpha] \tag{4}$$

be the covariance matrix of \mathbf{x}_α , where ε and $\mathbf{V}_0[\mathbf{x}_\alpha]$ are the noise level and the normalized covariance matrix, respectively. If the distribution of uncertainty is Gaussian, which we assume hereafter, the probability density of the data $\{\mathbf{x}_\alpha\}$ is given by

$$P(\{\mathbf{x}_\alpha\}) = C \prod_{\alpha=1}^N e^{-(\mathbf{x}_\alpha - \bar{\mathbf{x}}_\alpha, \mathbf{V}[\mathbf{x}_\alpha]^{-1}(\mathbf{x}_\alpha - \bar{\mathbf{x}}_\alpha))/2}, \tag{5}$$

where C is a normalization constant. Throughout this chapter, we denote the inner product of vectors \mathbf{a} and \mathbf{b} by (\mathbf{a}, \mathbf{b}) .

Maximum likelihood estimation (MLE) is finding the values of $\{\bar{\mathbf{x}}_\alpha\}$ and \mathbf{u} that maximize the *likelihood*, i.e., Eq. (6) into which the data $\{\mathbf{x}_\alpha\}$ are substituted, or equivalently minimize the sum of the squared *Mahalanobis distances* in the form

$$J = \sum_{\alpha=1}^N (\mathbf{x}_\alpha - \bar{\mathbf{x}}_\alpha, \mathbf{V}_0[\mathbf{x}_\alpha]^{-1}(\mathbf{x}_\alpha - \bar{\mathbf{x}}_\alpha)) \tag{6}$$

subject to the constraint (3) (Fig. 3b). The solution is called the *maximum likelihood (ML) estimator*. If the uncertainty is small, which we assume hereafter, constraint (3) can be eliminated by introducing Lagrange multipliers and applying first-order approximation. After some manipulations, we obtain the following form [14]:

$$J = \sum_{\alpha=1}^N \sum_{k,l=1}^r W_\alpha^{(kl)} F^{(k)}(\mathbf{x}_\alpha, \mathbf{u}) F^{(l)}(\mathbf{x}_\alpha, \mathbf{u}). \tag{7}$$

Here, $W_\alpha^{(kl)}$ is the (kl) element of the inverse of the $r \times r$ matrix whose (kl) element is $(\nabla_{\mathbf{x}} F_\alpha^{(k)}, \mathbf{V}_0[\mathbf{x}_\alpha] \nabla_{\mathbf{x}} F_\alpha^{(l)})$. We symbolically write

$$\left(W_\alpha^{(kl)} \right) = \left((\nabla_{\mathbf{x}} F_\alpha^{(k)}, \mathbf{V}_0[\mathbf{x}_\alpha] \nabla_{\mathbf{x}} F_\alpha^{(l)}) \right)^{-1}, \tag{8}$$

where $\nabla_{\mathbf{x}} F^{(k)}$ is the gradient of the function $F^{(k)}$ with respect to \mathbf{x} . The subscript α means that $\mathbf{x} = \mathbf{x}_\alpha$ is substituted.

Remark 6. The data $\{\mathbf{x}_\alpha\}$ may be subject to some constraints. For example, each \mathbf{x}_α may be a unit vector. The above formulation still holds if the inverse $\mathbf{V}_0[\mathbf{x}_\alpha]^{-1}$ in Eq. (6) is replaced by the Moore–Penrose generalized (or pseudo) inverse $\mathbf{V}_0[\mathbf{x}_\alpha]^-$ [14]. Similarly, the r constraints in Eq. (3) may be redundant, say only $r' (< r)$ of them are independent. The above formulation still holds if the inverse in Eq. (8) is replaced by the generalized inverse of rank r' with all but r' largest eigenvalues replaced by zero [14].

14.3.2.3 Accuracy of the ML Estimator

It can be shown [14] that the covariance matrix of the ML estimator $\hat{\mathbf{u}}$ has the form

$$\mathbf{V}[\hat{\mathbf{u}}] = \varepsilon^2 \mathbf{M}(\hat{\mathbf{u}})^{-1} + O(\varepsilon^4), \tag{9}$$

where

$$\mathbf{M}(\mathbf{u}) = \sum_{\alpha=1}^N \sum_{k,l=1}^r W_\alpha^{(kl)} \nabla_{\mathbf{u}} F_\alpha^{(k)} \nabla_{\mathbf{u}} F_\alpha^{(l)\top}. \tag{10}$$

Here, $\nabla_{\mathbf{u}} F^{(k)}$ is the gradient of the function $F^{(k)}$ with respect to \mathbf{u} . The subscript α means that $\mathbf{x} = \mathbf{x}_\alpha$ is substituted.

Remark 7. It can be proved that no other estimators could reduce the covariance matrix further than Eq. (9) except for the higher-order term $O(\varepsilon^4)$ [14, 17]. The ML estimator is optimal in this sense. Recall that we are focusing on the asymptotic analysis for $\varepsilon \rightarrow 0$. Thus, what we call the “ML estimator” should be understood to be a first approximation to the true ML estimator for small ε .

Remark 8. The p -dimensional parameter vector \mathbf{u} may be constrained. For example, it may be a unit vector. If it has only $p' (< p)$ degrees of freedom, the parameter space \mathcal{U} is a p' -dimensional manifold in \mathcal{R}^p . In this case, the matrix $\mathbf{M}(\mathbf{u})$ in Eq. (9) is replaced by $\mathbf{P}_{\mathbf{u}} \mathbf{M}(\mathbf{u}) \mathbf{P}_{\mathbf{u}}$, where $\mathbf{P}_{\mathbf{u}}$ is the projection matrix onto the tangent space to the parameter space \mathcal{U} at \mathbf{u} [14]. The inverse $\mathbf{M}(\hat{\mathbf{u}})^{-1}$ in Eq. (9) is replaced by the generalized inverse $\mathbf{M}(\hat{\mathbf{u}})^{-1}$ of rank p' [14].

14.3.3 Geometric Model Selection

Geometric fitting is to estimate the parameter \mathbf{u} of a given model. If we have multiple candidate models

$$F_1^{(k)}(\bar{\mathbf{x}}_\alpha, \mathbf{u}_1) = 0, \quad F_2^{(k)}(\bar{\mathbf{x}}_\alpha, \mathbf{u}_2) = 0, \quad F_3^{(k)}(\bar{\mathbf{x}}_\alpha, \mathbf{u}_3) = 0, \quad \dots, \quad (11)$$

from which we are to select an appropriate one for the observed data $\{\mathbf{x}_\alpha\}$, the problem is *geometric model selection* [14, 16, 18].

Suppose, for example, we want to fit a curve to given points in two dimensions. If they are almost collinear, a straight line may fit fairly well, but a quadratic curve may fit better, and a cubic curve even better. Which curve should we fit? A naive idea is to compare the *residual (sum of squares)*, i.e., the minimum value \hat{J} of J in Eq. (6); we select the one that has the smallest residual \hat{J} . This does not work, however, because the ML estimator $\hat{\mathbf{u}}$ is so determined as to minimize the residual \hat{J} , and the residual \hat{J} can be made arbitrarily smaller if the model is equipped with more parameters to adjust. So, the only conclusion would be to fit a curve of a sufficiently high degree passing through all the points.

14.3.3.1 Geometric AIC

The above observation leads to the idea of compensating for the negative bias of the residual caused by substituting the ML estimator. This is the principle of the *Akaike information criterion (AIC)* [1], which is derived from the asymptotic behavior of the *Kullback–Leibler information (or divergence)* as the number n of experiments goes to infinity. Doing a similar analysis to Akaike's and examining the asymptotic behavior as the noise level ε goes to zero, we can obtain the following *geometric AIC* [14, 15]:

$$\text{G-AIC} = \hat{J} + 2(Nd + p)\varepsilon^2 + O(\varepsilon^4). \quad (12)$$

Here, d is the dimension of the manifold \mathcal{S} defined by the constraint (3) in the data space \mathcal{X} , and p is the dimension of \mathbf{u} (i.e., the number of unknowns). The model for which Eq. (12) is the smallest is regarded as the best. The derivation of Eq. (12) is based on the following facts [14, 15]:

- The ML estimator $\hat{\mathbf{u}}$ converges to its true value as $\varepsilon \rightarrow 0$.
- The ML estimator $\hat{\mathbf{u}}$ obeys a Gaussian distribution under linear constraints, because the noise is assumed to be Gaussian. For nonlinear constraints, linear approximation can be justified in the neighborhood of the solution if ε is sufficiently small.
- A quadratic form in standardized Gaussian random variables is subject to a χ^2 distribution, whose expectation is equal to its degree of freedom.

14.3.3.2 Geometric MDL

Another well-known criterion for model selection is Rissanen's *minimum description length (MDL)* [42, 43, 44], which measures the goodness of a model by the minimum information theoretic code length of the data and the model. The basic idea is simple, but the following difficulties must be resolved to apply it in practice:

- Encoding a problem involving real numbers requires an infinitely long code length.
- The probability density, from which a minimum length code can be obtained, involves unknown parameters.
- The exact form of the minimum code length is very difficult to compute.

Rissanen [42, 43, 44] avoided these difficulties by quantizing the real numbers in a way that does not depend on individual models and substituting the ML estimators for the parameters. They, too, are real numbers, so they are also quantized. The quantization width is so chosen as to minimize the total description length (*two-stage encoding*). The resulting code length is evaluated asymptotically as the data length n goes to infinity. If we analyze the asymptotic behavior of encoding the geometric fitting problem as the noise level ε goes to zero, we obtain the following *geometric MDL* [20]:

$$\text{G-MDL} = \hat{J} - (Nd + p)\varepsilon^2 \log\left(\frac{\varepsilon}{L}\right)^2 + O(\varepsilon^2). \quad (13)$$

Here, L is a reference length chosen so that its ratio to the magnitude of data is $O(1)$, e.g., L can be taken to be the image size for feature point data. Its exact determination requires an a priori distribution that specifies where the data are likely to appear (we discuss this more in Sect.), but it has been observed that the model selection is not very much affected by L as long as it is within the same order of magnitude [20].

14.4 Standard Statistical Analysis vs. Geometric Inference

We now point out that a correspondence exists between the standard statistical analysis and the geometric inference problem. We also compare the capability of the geometric AIC and the geometric MDL in detecting degeneracy.

14.4.1 Standard Statistical Analysis

The asymptotic analysis in Sect. bears a strong resemblance to the standard statistical estimation problem: after observing n data $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n$, we

want to estimate the parameter $\boldsymbol{\theta}$ of the probability density $P(\mathbf{x}|\boldsymbol{\theta})$ called the (*stochastic*) *model*, according to which each datum is assumed to be sampled independently.

Maximum likelihood estimation (MLE) is to find the value $\boldsymbol{\theta}$ that maximizes $\prod_{i=1}^n P(\mathbf{x}_i|\boldsymbol{\theta})$, or equivalently minimizes its negative logarithm $-\sum_{i=1}^n \log P(\mathbf{x}_i|\boldsymbol{\theta})$. It can be shown that the covariance matrix $V[\hat{\boldsymbol{\theta}}]$ of the resulting ML estimator $\hat{\boldsymbol{\theta}}$ converges, under a mild condition, to \mathbf{O} as the number n of experiments goes to infinity (*consistency*) in the form

$$V[\hat{\boldsymbol{\theta}}] = \mathbf{I}(\boldsymbol{\theta})^{-1} + O\left(\frac{1}{n^2}\right), \quad (14)$$

where we define the *Fisher information matrix* $\mathbf{I}(\boldsymbol{\theta})$ by

$$\mathbf{I}(\boldsymbol{\theta}) = nE[(\nabla_{\boldsymbol{\theta}} \log P(\mathbf{x}|\boldsymbol{\theta}))(\nabla_{\boldsymbol{\theta}} \log P(\mathbf{x}|\boldsymbol{\theta}))^{\top}]. \quad (15)$$

The operation $E[\cdot]$ denotes expectation with respect to the density $P(\mathbf{x}|\boldsymbol{\theta})$. The first term in the right-hand side of Eq. (14) is called the *Cramer–Rao lower bound (CRLB)*, describing the minimum degree of fluctuations in all estimators. Thus, the ML estimator is optimal if n is sufficiently large (*asymptotic efficiency*).

If we have multiple candidate models

$$P_1(\mathbf{x}|\boldsymbol{\theta}_1), \quad P_2(\mathbf{x}|\boldsymbol{\theta}_2), \quad P_3(\mathbf{x}|\boldsymbol{\theta}_3), \quad \dots, \quad (16)$$

from which we are to select an appropriate one for the observations $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n$, the problem is (*stochastic*) *model selection*. Akaike's AIC has the following form:

$$\text{AIC} = -2 \sum_{i=1}^N \log P(\mathbf{x}_i|\hat{\boldsymbol{\theta}}) + 2k + O\left(\frac{1}{n}\right). \quad (17)$$

The model for which this quantity is the smallest is regarded as the best. The derivation of Eq. (17) is based on the following facts [1]:

- The maximum likelihood estimator $\hat{\boldsymbol{\theta}}$ converges to its true value as $n \rightarrow \infty$ (the *law of large numbers*).
- The maximum likelihood estimator $\hat{\boldsymbol{\theta}}$ asymptotically obeys a Gaussian distribution as $n \rightarrow \infty$ (the *central limit theorem*).
- A quadratic form in standardized Gaussian random variables is subject to a χ^2 distribution, whose expectation is equal to its degree of freedom.

Rissanen's MDL has the following form [43, 44]:

$$\text{MDL} = -\sum_{i=1}^n \log P(\mathbf{x}_i|\hat{\boldsymbol{\theta}}) + \frac{k}{2} \log \frac{n}{2\pi} + \log \int_{\mathcal{T}} \sqrt{|\mathbf{I}(\boldsymbol{\theta})|} d\boldsymbol{\theta} + O(1). \quad (18)$$

Here, $\hat{\boldsymbol{\theta}}$ is the ML estimator; the symbol $O(1)$ denotes terms of order 0 in n in the limit $n \rightarrow \infty$. In order that the integration in the right-hand side of

Eq. (18) exists, the domain \mathcal{T} of the parameter $\boldsymbol{\theta}$ must be compact. In other words, we must specify in the k -dimensional space of $\boldsymbol{\theta}$ a finite region \mathcal{T} in which the true value of $\boldsymbol{\theta}$ is likely to exist. This is nothing but the *Bayesian* standpoint that requires a prior distribution for the parameter to estimate. If it is not known, we must introduce an appropriate expedient to suppress an explicit dependence on the prior. Such an expedient is also necessary for the geometric MDL, i.e., the introduction of the reference length L in Eq. (18).

14.4.2 Dual Interpretations of Asymptotic Analysis

Thus, we have seen that the limit $n \rightarrow \infty$ for the standard statistical analysis corresponds to the limit $\varepsilon \rightarrow 0$ for geometric inference. For example, the covariance matrix of the ML estimator agrees with the Cramer–Rao lower bound up to $O(1/n^2)$ for $n \rightarrow \infty$ (see Eq. (14)), while for geometric inference it agrees with the lower bound up to $O(\varepsilon^4)$ for $\varepsilon \rightarrow 0$ (see Eq. (9)). It follows that $1/\sqrt{n}$ for the standard statistical analysis plays the same role as ε for geometric inference.

The same correspondence exists for model selection, too. The unknowns for geometric inference are the p parameters of the constraint plus the N true positions specified by the d coordinates of the d -dimensional manifold \mathcal{S} defined by the constraint. If Eq. (12) is divided by ε^2 , we have $\hat{J}/\varepsilon^2 + 2(Nd + p) + O(\varepsilon^2)$, which is (-2 times the logarithmic likelihood) + 2(the number of unknowns), the same form as Akaike’s AIC given by Eq. (17). The same holds for Eq. (13), which corresponds to Rissanen’s MDL given by Eq. (18) if ε is replaced by $1/\sqrt{n}$ [20].

This correspondence can be interpreted as follows. Since the underlying ensemble is hypothetical, we can actually observe only one sample as long as a particular algorithm is used. Suppose we hypothetically sample n different algorithms to find n different positions. The optimal estimate of the true position under the Gaussian model is their sample mean. The covariance matrix of the sample mean is $1/n$ times that of the individual samples. Hence, this hypothetical estimation is equivalent to dividing the noise level ε in Eq. (4) by \sqrt{n} .

In fact, there were attempts to generate a hypothetical *ensemble of algorithms* by randomly varying the internal parameters (e.g., the thresholds for judgments), not adding random noise to the image [5, 6]. Then, one can compute their means and covariance matrix. Such a process as a whole can be regarded as one operation that effectively achieves higher accuracy.

Thus, the asymptotic analysis for $\varepsilon \rightarrow 0$ is equivalent to the asymptotic analysis for $n \rightarrow \infty$, where n is the number of hypothetical observations. As a result, the expression $\cdots + O(1/\sqrt{n^k})$ in the standard statistical analysis turns into $\cdots + O(\varepsilon^k)$ in geometric inference.

14.4.3 Noise Level Estimation

In order to use the geometric AIC or the geometric MDL, we need to know the noise level ε . If not known, it must be estimated. Here arises a sharp contrast between the standard statistical analysis and our geometric inference.

For the standard statistical analysis, the noise magnitude is a *model parameter*, because “noise” is defined to be *the random effects that cannot be accounted for by the assumed model*. Hence, the noise magnitude should be estimated, if not known, *according to the assumed model*. For geometric inference, on the other hand, the noise level ε is *a constant that reflects the uncertainty of feature detection*. So, it should be estimated *independently of individual models*.

If we know the true model, it can be estimated from the residual \hat{J} using the knowledge that \hat{J}/ε^2 is subject to a χ^2 distribution with $rN - p$ degrees of freedom in the first-order [14]. Specifically, we obtain an unbiased estimator of ε^2 in the form

$$\hat{\varepsilon}^2 = \frac{\hat{J}}{rN - p}. \quad (19)$$

The validity of this formula has been confirmed by many simulations.

One may wonder if model selection is necessary at all when the true model is known. In practice, however, a typical situation where model selection is called for is *degeneracy detection*. In 3-D analysis from images, for example, the constraint (3) corresponds to our knowledge about the scene such as rigidity of motion. However, the computation fails if degeneracy occurs (e.g., the motion is zero). Even if exact degeneracy does not occur, the computation may become numerically unstable in near-degeneracy conditions. In such a case, the computation can be stabilized by switching to a model that describes the degeneracy [16, 21, 24, 25, 31, 39, 53].

Degeneracy means *addition* of new constraints, such as some quantity being zero. It follows that the manifold \mathcal{S} degenerates into a submanifold \mathcal{S}' of it. Since the general model still holds irrespective of the degeneracy, i.e., $\mathcal{S}' \subset \mathcal{S}$, we can estimate the noise level ε from the residual \hat{J} of the general model \mathcal{S} using Eq. (19).

Remark 9. Equation (19) can be intuitively understood as follows. Recall that \hat{J} is the sum of the square distances from $\{\mathbf{x}_\alpha\}$ to the manifold $\hat{\mathcal{S}}$ defined by the constraint $F^{(k)}(\mathbf{x}, \mathbf{u}) = 0$, $k = 1, \dots, r$. Since $\hat{\mathcal{S}}$ has codimension r (the dimension of the orthogonal directions to it), the residual \hat{J} should have expectation $rN\varepsilon^2$. However, $\hat{\mathcal{S}}$ is fitted by adjusting its p -dimensional parameter \mathbf{u} , so the expectation of \hat{J} reduces to $(rN - p)\varepsilon^2$.

Remark 10. It may appear that the residual \hat{J} of the general model cannot be stably computed in the presence of degeneracy. However, what is unstable is *model specification*, not the residual. For example, if we fit a planar surface to




Fig. 4. Fitting a space line and a plane to points in space

almost collinear points in 3-D, it is difficult to specify the fitted plane stably; the solution is very susceptible to noise. Yet, the residual is stably computed, since unique specification of the fit is difficult *because all the candidates have almost the same residual*.

Remark 11. Note that the noise level estimation from the general model \mathcal{S} by Eq. (19) is still valid even if degeneracy occurs, because degeneracy means shrinkage of the model manifold \mathcal{S}' *within* \mathcal{S} , which does not affect the data deviations in the “orthogonal” directions (in the Mahalanobis sense) to \mathcal{S} that account for the residual \hat{J} .

14.4.4 Comparing the Geometric AIC and the Geometric MDL

We now illustrate the different characteristics of the geometric AIC and the geometric MDL in detecting degeneracy. Consider a rectangular region $[0, 10] \times [-1, 1]$ on the x - y -plane in the x - y - z -space. We randomly take 11 points in it and magnify the region A times in the y -direction. Adding Gaussian noise of mean 0 and variance ε^2 to the x , y , and z coordinates of each point independently, we fit a space line and a plane in a statistically optimal manner (Fig. 4). The rectangular region degenerates into a line segment as $A \rightarrow 0$.

A space line is a one-dimensional model with four degrees of freedom; a plane is a two-dimensional model with three degrees of freedom. Their geometric AIC and geometric MDL are

$$\begin{aligned} \text{G-AIC}_l &= \hat{J}_l + 2(N+4)\varepsilon^2, & \text{G-AIC}_p &= \hat{J}_p + 2(2N+3)\varepsilon^2, \\ \text{G-MDL}_l &= \hat{J}_l - (N+4)\varepsilon^2 \log\left(\frac{\varepsilon}{L}\right)^2, & \text{G-MDL}_p &= \hat{J}_p - (2N+3)\varepsilon^2 \log\left(\frac{\varepsilon}{L}\right)^2, \end{aligned} \tag{20}$$

where the subscripts l and p refer to lines and planes, respectively. For each A , we compare the geometric AIC and the geometric MDL of the fitted line and plane and choose the one that has the smaller value. We used the reference length $L = 1$.

Figure 5a shows the percentage of choosing a line for $\varepsilon = 0.01$ after 1000 independent trials for each A . If there were no noise, it should be 0% for A

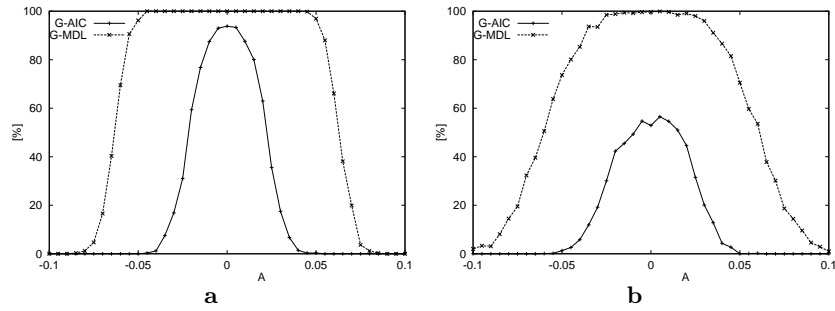


Fig. 5. The rate (%) of detecting a space line by the geometric AIC (solid lines with +) and the geometric MDL (dotted lines with \times) with **a** the true noise level and **b** the estimated noise level

$\neq 0$ and 100% for $A = 0$. In the presence of noise, the geometric AIC has a high capability of distinguishing a line from a plane, but it judges a line to be a plane with some probability. In contrast, the geometric MDL judges a line to be a line almost 100%, but it judges a plane to be a line over a wide range of A .

In Fig. 5a, we used the true value of ε^2 . Figure 5b shows the corresponding result using its estimate obtained from the general plane model by Eq. (19). We observe somewhat degraded but similar performance characteristics.

Thus, we can observe that the geometric AIC has a higher capability for detecting degeneracy than the geometric MDL, but the general model is chosen with some probability when the true model is degenerate. In contrast, the percentage for the geometric MDL to detect degeneracy when the true model is really degenerate approaches 100% as the noise decreases. This is exactly the dual statement to the well-known fact, called the *consistency of the MDL*, that the percentage for Rissanen's MDL to identify the true model converges to 100% in the limit of an infinite number of observations. Rissanen's MDL is regarded by many as superior to Akaike's AIC because the latter lacks this property.

At the cost of this consistency, however, the geometric MDL regards a wide range of nondegenerate models as degenerate. This is no surprise, since the penalty $-(Nd + p)\varepsilon^2 \log(\varepsilon/L)^2$ for the geometric MDL in Eq. (13) is heavier than the penalty $2(Nd + p)\varepsilon^2$ for the geometric AIC in Eq. (12). As a result, the geometric AIC is more faithful to the data than the geometric MDL, which is more likely to choose a degenerate model. This contrast has also been observed in many applications [23, 31].

Remark 12. Despite the fundamental difference of geometric model selection from the standard (stochastic) model selection, many attempts have been made in the past to apply Akaike's AIC and their variants to computer vision problems based on the asymptotic analysis of $n \rightarrow \infty$, where the interpretation of n is different from problem to problem [48, 49, 50, 51, 52]. Rissanen's MDL

is also used in computer vision applications. Its use may be justified if the problem has the standard form of linear/nonlinear regression [3, 32]. Often, however, the solution having a shorter description length was chosen with a rather arbitrary definition of the complexity [11, 27, 33].

Remark 13. Note that one cannot compare different model selection criteria in general terms, because each is based on its own logic. Not only that, one cannot prove that a particular criterion works at all. In fact, although Akaike's AIC and Rissanen's MDL are based on rigorous mathematics, there is no guarantee that they work well in practice. The mathematical rigor is in their *reduction* from their starting principles (the Kullback–Leibler information and the minimum description length principle), which are beyond proof. What one can tell is which criterion is more suitable for a particular application when used in a particular manner. The geometric AIC and the geometric MDL have shown to be effective in many computer vision applications [19, 22, 23, 24, 25, 31, 39, 53], but other criteria may be better in other applications. The important thing is, however, to understand the underlying logic of each criterion.

14.5 Linear Geometric Fitting

Now, we consider a special type of geometric fitting problem that most frequently arises in computer vision applications: the constraint is linear in both data and unknowns. We systematically review existing methods.

14.5.1 Linear Constraints

In many geometric inference problems of computer vision, the constraint (3) has the form

$$(\boldsymbol{\xi}(\bar{\mathbf{x}}_\alpha), \mathbf{u}) = 0, \quad (21)$$

where $\boldsymbol{\xi}(\cdot)$ is generally a nonlinear mapping from an m -dimensional vector to a p -dimensional vector. Evidently, the magnitude of \mathbf{u} is unconstrained, so we normalize it to a unit vector: $\|\mathbf{u}\| = 1$.

Example 1. Suppose we are given N points $\{(x_\alpha, y_\alpha)\}$, $\alpha = 1, \dots, N$, in two dimensions. Their true positions $\{(\bar{x}_\alpha, \bar{y}_\alpha)\}$ are assumed to be on a *conic* (a circle, an ellipse, a parabola, a hyperbola, or their degeneracy). Our task is to estimate the curve from the noisy data $\{(x_\alpha, y_\alpha)\}$. The constraint on $\{(\bar{x}_\alpha, \bar{y}_\alpha)\}$ is

$$A\bar{x}_\alpha^2 + 2B\bar{x}_\alpha\bar{y}_\alpha + C\bar{y}_\alpha^2 + 2(D\bar{x}_\alpha + E\bar{y}_\alpha) + F = 0 \quad (22)$$

for some coefficients A, B, \dots, D , not all being zero. This constraint reduces to Eq. (21) if we put

$$\boldsymbol{\xi}(x, y) = (x^2 \ 2xy \ y^2 \ 2x \ 2y \ 1)^\top, \quad \mathbf{u} = (A \ B \ C \ D \ E \ F)^\top. \quad (23)$$

The data space \mathcal{X} is a two-dimensional manifold in the six-dimensional space \mathcal{R}^6 ; the parameter space \mathcal{U} is the five-dimensional unit sphere S^6 centered on the origin of \mathcal{R}^6 .

Example 2. Suppose N points in a 3-D scene are projected to (x_α, y_α) in the first image and (x'_α, y'_α) in the second, $\alpha = 1, \dots, N$. If the camera imaging geometry is perspective projection, there exists a matrix \mathbf{F} of determinant 0 such that

$$\begin{pmatrix} \bar{x}_\alpha \\ \bar{y}_\alpha \\ 1 \end{pmatrix}, \mathbf{F} \begin{pmatrix} \bar{x}'_\alpha \\ \bar{y}'_\alpha \\ 1 \end{pmatrix} = 0, \quad (24)$$

which is called the *epipolar equation* [13]. The matrix \mathbf{F} is known as the *fundamental matrix*. For 3-D reconstruction from the images, we need to estimate the fundamental matrix \mathbf{F} from the noisy data $\{(x_\alpha, y_\alpha)\}$ and $\{(x'_\alpha, y'_\alpha)\}$. Equation (24) reduces to Eq. (21) if we put

$$\begin{aligned} \boldsymbol{\xi}(x, y, x', y') &= (xx' \ xy' \ x \ yx' \ yy' \ y \ x' \ y' \ 1)^\top, \\ \mathbf{u} &= (F_{11} \ F_{12} \ F_{13} \ F_{21} \ F_{22} \ F_{23} \ F_{31} \ F_{32} \ F_{33})^\top. \end{aligned} \quad (25)$$

The data space \mathcal{X} is a four-dimensional manifold in the nine-dimensional space \mathcal{R}^9 ; the parameter space \mathcal{U} is a seven-dimensional manifold defined by $\det \mathbf{F} = 0$ and $\|\mathbf{F}\| = 1$, where the matrix norm is define by $\|\mathbf{F}\| = \sqrt{\sum_{i,j=1}^3 F_{ij}^2}$.

For the linear constraint (21), the function J in Eq. (7) reduces to

$$J = \sum_{\alpha=1}^N \frac{(\boldsymbol{\xi}_\alpha, \mathbf{u})^2}{(\mathbf{u}, \mathbf{V}_0[\boldsymbol{\xi}_\alpha] \mathbf{u})}, \quad (26)$$

where $\mathbf{V}_0[\boldsymbol{\xi}_\alpha]$ is the normalized covariance matrix of $\boldsymbol{\xi}_\alpha$; we use the abbreviation $\boldsymbol{\xi}_\alpha = \boldsymbol{\xi}(\mathbf{x}_\alpha)$. The matrix $\mathbf{V}_0[\boldsymbol{\xi}_\alpha]$ can be expressed to a first approximation in the form

$$\mathbf{V}_0[\boldsymbol{\xi}_\alpha] = \nabla_{\mathbf{x}} \boldsymbol{\xi}|_{\mathbf{x}=\mathbf{x}_\alpha}^\top \mathbf{V}_0[\mathbf{x}_\alpha] \nabla_{\mathbf{x}} \boldsymbol{\xi}|_{\mathbf{x}=\mathbf{x}_\alpha}, \quad (27)$$

where $\nabla_{\mathbf{x}} \boldsymbol{\xi}$ is the $m \times p$ Jacobian matrix of $\boldsymbol{\xi}(\mathbf{x})$:

$$\nabla_{\mathbf{x}} \boldsymbol{\xi} = \begin{pmatrix} \partial \xi_1 / \partial x_1 & \cdots & \partial \xi_p / \partial x_1 \\ \vdots & & \vdots \\ \partial \xi_1 / \partial x_m & \cdots & \partial \xi_p / \partial x_m \end{pmatrix}. \quad (28)$$

The covariance matrix $\mathbf{V}[\hat{\mathbf{u}}]$ of the ML estimator $\hat{\mathbf{u}}$ given by Eq. (9) now reads

$$V[\hat{\mathbf{u}}] = \varepsilon^2 \left(\sum_{\alpha=1}^N \frac{\mathbf{P}_{\mathbf{u}} \boldsymbol{\xi}_{\alpha} \boldsymbol{\xi}_{\alpha}^{\top} \mathbf{P}_{\mathbf{u}}}{(\mathbf{u}, V_0[\boldsymbol{\xi}_{\alpha}] \mathbf{u})} \right)^{-} + O(\varepsilon^4), \quad (29)$$

where the superscript $-$ denotes the Moore–Penrose generalized inverse. The matrix $\mathbf{P}_{\mathbf{u}}$ denotes projection onto the tangent space to the parameter space \mathcal{U} at \mathbf{u} (cf. Remark 8). Since the leading term is the lower bound on the covariance matrix of any estimation (Remark 7), the ML estimator is optimal up to higher-order terms in ε .

Remark 14. Since we are focusing on the asymptotic analysis for $\varepsilon \rightarrow 0$, what we call the “ML estimator” is a first approximation to the true ML estimator for small ε (Remark 7). Note that if the parameter \mathbf{u} is not constrained, the generalized inverse in Eq. (29) can be replaced by the usual inverse, and the projection matrix $\mathbf{P}_{\mathbf{u}}$ is not necessary. However, \mathbf{u} is at least constrained to be a unit vector, and often additional constraints exist, e.g., $\det \mathbf{F} = 0$ on the fundamental matrix \mathbf{F} . If no constraints exist other than $\|\mathbf{u}\| = 1$, the covariance matrix $V[\hat{\mathbf{u}}]$ has rank $p - 1$, and its null space is in the direction of \mathbf{u} . The projection matrix $\mathbf{P}_{\mathbf{u}}$ in this case is

$$\mathbf{P}_{\mathbf{u}} = \mathbf{I} - \mathbf{u} \mathbf{u}^{\top}. \quad (30)$$

14.5.2 Least-Squares Method

If \mathbf{u} is constrained, the minimization of Eq. (26) should be carried out subject to the constraint, but this is very difficult in many cases. A practical approach to this is to ignore all the constraints except the normalization $\|\mathbf{u}\| = 1$ and do minimization over the $(p - 1)$ -dimensional sphere S^{p-1} in \mathcal{R}^p . This expedient is motivated by the fact that if the data $\{\mathbf{x}_{\alpha}\}$ are exact, the solution should automatically satisfy the remaining constraints. It follows that if the data uncertainty is very small, which we always assume, the resulting solution $\hat{\mathbf{u}}$ should satisfy all the constraints up to higher-order terms in ε .

However, the minimization of Eq. (26) is still nonlinear even if all constraints other than $\|\mathbf{u}\| = 1$ are ignored. The simplest approach is to solve Eq. (21) directly by (total) least squares, minimizing

$$J_{LS} = \sum_{\alpha=1}^N (\boldsymbol{\xi}_{\alpha}, \mathbf{u})^2. \quad (31)$$

If we define the second-order moment matrix

$$\mathbf{M} = \sum_{\alpha=1}^N \boldsymbol{\xi}_{\alpha} \boldsymbol{\xi}_{\alpha}^{\top}, \quad (32)$$

Eq. (31) is rewritten as

$$J_{LS} = (\mathbf{u}, \mathbf{M}\mathbf{u}). \quad (33)$$

The unit vector \mathbf{u} that minimizes this is the unit eigenvector of \mathbf{M} for the smallest eigenvalue. The resulting *least-squares (LS) solution* $\hat{\mathbf{u}}_{LS}$ is a very crude approximation to the ML estimator $\hat{\mathbf{u}}$. However, because of the ease of the computation, it is often used as an initial guess for computing the ML estimator $\hat{\mathbf{u}}$ by iterations.

14.5.3 Naive Method

If we define

$$\mathbf{M}(\mathbf{u}) = \sum_{\alpha=1}^N \frac{\boldsymbol{\xi}_\alpha \boldsymbol{\xi}_\alpha^\top}{(\mathbf{u}, \mathbf{V}_0[\boldsymbol{\xi}_\alpha] \mathbf{u})}, \quad (34)$$

Eq. (26) is written as

$$J = (\mathbf{u}, \mathbf{M}(\mathbf{u})\mathbf{u}). \quad (35)$$

This inspires the following iterations for computing the ML estimator:

1. Guess an appropriate initial value \mathbf{u}_0 , say the LS solution $\hat{\mathbf{u}}_{LS}$.
2. Assuming that \mathbf{u}_{i-1} is obtained (initially $i = 1$), let \mathbf{u}_i be the unit eigenvector of $\mathbf{M}(\mathbf{u}_{i-1})$ for the smallest eigenvalue.
3. Return \mathbf{u}_i if \mathbf{u}_i is sufficiently close to \mathbf{u}_{i-1} except for the sign. Otherwise, let $\mathbf{u}_{i-1} \leftarrow -\mathbf{u}_i$, and go back to step 2.

This scheme does not work, however, because the resulting solution $\hat{\mathbf{u}}$ is the value \mathbf{u} that minimizes $(\mathbf{u}, \mathbf{M}(\hat{\mathbf{u}})\mathbf{u})$, not $(\mathbf{u}, \mathbf{M}(\mathbf{u})\mathbf{u})$. In other words,

$$(\hat{\mathbf{u}}, \mathbf{M}(\hat{\mathbf{u}})\hat{\mathbf{u}}) < (\hat{\mathbf{u}} + \Delta\mathbf{u}, \mathbf{M}(\hat{\mathbf{u}})(\hat{\mathbf{u}} + \Delta\mathbf{u})) \quad (36)$$

for any nonzero perturbation $\Delta\mathbf{u}$, but not

$$(\hat{\mathbf{u}}, \mathbf{M}(\hat{\mathbf{u}})\hat{\mathbf{u}}) < (\hat{\mathbf{u}} + \Delta\mathbf{u}, \mathbf{M}(\hat{\mathbf{u}} + \Delta\mathbf{u})(\hat{\mathbf{u}} + \Delta\mathbf{u})). \quad (37)$$

A detailed analysis shows that $\hat{\mathbf{u}}$ is *biased* by $O(\varepsilon^2)$ [14]. Namely, if the fluctuations of the data $\{\mathbf{x}_\alpha\}$ are centered on their true values $\{\bar{\mathbf{x}}_\alpha\}$, the corresponding fluctuations of $\hat{\mathbf{u}}$ are around a value different from its true value by $O(\varepsilon^2)$. This causes inadmissible errors in many practical applications.

14.5.4 FNS Method

If the constraint on \mathbf{u} is ignored, the solution that minimizes Eq. (26) is obtained by solving $\nabla_{\mathbf{u}}J = \mathbf{0}$. Since

$$\nabla_{\mathbf{u}}J = \sum_{\alpha=1}^N \frac{2(\boldsymbol{\xi}_\alpha, \mathbf{u})\boldsymbol{\xi}_\alpha}{(\mathbf{u}, \mathbf{V}_0[\boldsymbol{\xi}_\alpha] \mathbf{u})} - \sum_{\alpha=1}^N \frac{2(\boldsymbol{\xi}_\alpha, \mathbf{u})^2 \mathbf{V}_0[\boldsymbol{\xi}_\alpha] \mathbf{u}}{(\mathbf{u}, \mathbf{V}_0[\boldsymbol{\xi}_\alpha] \mathbf{u})^2}, \quad (38)$$

the equation $\nabla_{\mathbf{u}}J = \mathbf{0}$ is written in the form

$$\mathbf{X}(\mathbf{u})\mathbf{u} = \mathbf{0}, \tag{39}$$

where

$$\mathbf{X}(\mathbf{u}) = \sum_{\alpha=1}^N \frac{\boldsymbol{\xi}_\alpha \boldsymbol{\xi}_\alpha^\top}{(\mathbf{u}, \mathbf{V}_0[\boldsymbol{\xi}_\alpha]\mathbf{u})} - \sum_{\alpha=1}^N \frac{(\boldsymbol{\xi}_\alpha, \mathbf{u})^2 \mathbf{V}_0[\boldsymbol{\xi}_\alpha]}{(\mathbf{u}, \mathbf{V}_0[\boldsymbol{\xi}_\alpha]\mathbf{u})^2}. \tag{40}$$

From this, we have the following scheme for solving Eq. (39):

1. Guess an appropriate initial value \mathbf{u}_0 , say the LS solution $\hat{\mathbf{u}}_{\text{LS}}$.
2. Assuming that \mathbf{u}_{i-1} is obtained (initially $i = 1$), solve the eigenvalue problem

$$\mathbf{X}(\mathbf{u}_{i-1})\mathbf{u} = \lambda\mathbf{u}. \tag{41}$$

Let \mathbf{u}_i be the unit eigenvector for the eigenvalue λ closest to 0.

3. Return \mathbf{u}_i if \mathbf{u}_i is sufficiently close to \mathbf{u}_{i-1} except for the sign. Otherwise, let $\mathbf{u}_{i-1} \leftarrow \mathbf{u}_i$, and go back to step 2.

The resulting solution $\hat{\mathbf{u}}$ satisfies Eq. (39). In fact, the value $\hat{\mathbf{u}}$ produced by the above iterations should satisfy

$$\mathbf{X}(\hat{\mathbf{u}})\hat{\mathbf{u}} = \lambda\hat{\mathbf{u}} \tag{42}$$

for some λ . Taking the inner product of $\hat{\mathbf{u}}$ and both sides, we have

$$(\hat{\mathbf{u}}, \mathbf{X}(\hat{\mathbf{u}})\hat{\mathbf{u}}) = \lambda. \tag{43}$$

Equation (40) implies that

$$(\hat{\mathbf{u}}, \mathbf{X}(\hat{\mathbf{u}})\hat{\mathbf{u}}) = \sum_{\alpha=1}^N \frac{(\hat{\mathbf{u}}, \boldsymbol{\xi}_\alpha)^2}{(\hat{\mathbf{u}}, \mathbf{V}_0[\boldsymbol{\xi}_\alpha]\hat{\mathbf{u}})} - \sum_{\alpha=1}^N \frac{(\boldsymbol{\xi}_\alpha, \hat{\mathbf{u}})^2 (\hat{\mathbf{u}}, \mathbf{V}_0[\boldsymbol{\xi}_\alpha]\hat{\mathbf{u}})}{(\hat{\mathbf{u}}, \mathbf{V}_0[\boldsymbol{\xi}_\alpha]\hat{\mathbf{u}})^2} = 0, \tag{44}$$

meaning that $\lambda = 0$. Thus, $\hat{\mathbf{u}}$ is indeed the solution of Eq. (39). This method was proposed by Chojnacki et al. [7] and is called the *fundamental numerical scheme (FNS) method*. Usually, the iterations converge very quickly.

Remark 15. Equation (44) is a consequence of the fact that the right-hand side of Eq. (26) is a *homogeneous function of degree 0* in \mathbf{u} . Since multiplying \mathbf{u} by any nonzero constant does not change the value of J , the gradient $\nabla_{\mathbf{u}}J$ is necessarily orthogonal to \mathbf{u} . Thus, $(\mathbf{u}, \nabla_{\mathbf{u}}J) = 2(\mathbf{u}, \mathbf{X}(\mathbf{u})\mathbf{u})$ is identically 0.

14.5.5 HEIV Method

Equation (39) can also be written as

$$\mathbf{M}(\mathbf{u})\mathbf{u} = \mathbf{L}(\mathbf{u})\mathbf{u}, \tag{45}$$

where

$$\mathbf{M}(\mathbf{u}) = \sum_{\alpha=1}^N \frac{\boldsymbol{\xi}_\alpha \boldsymbol{\xi}_\alpha^\top}{(\mathbf{u}, \mathbf{V}_0[\boldsymbol{\xi}_\alpha] \mathbf{u})}, \quad \mathbf{L}(\mathbf{u}) = \sum_{\alpha=1}^N \frac{(\boldsymbol{\xi}_\alpha, \mathbf{u})^2 \mathbf{V}_0[\boldsymbol{\xi}_\alpha]}{(\mathbf{u}, \mathbf{V}_0[\boldsymbol{\xi}_\alpha] \mathbf{u})^2}. \quad (46)$$

This implies the following scheme:

1. Guess an appropriate initial value \mathbf{u}_0 , say the LS solution $\hat{\mathbf{u}}_{\text{LS}}$.
2. Assuming that \mathbf{u}_{i-1} is obtained (initially $i = 1$), solve the generalized eigenvalue problem

$$\mathbf{M}(\mathbf{u}_{i-1}) \mathbf{u} = \lambda \mathbf{L}(\mathbf{u}_{i-1}) \mathbf{u}. \quad (47)$$

Let \mathbf{u}_i be the generalized eigenvector for the generalized eigenvalue closest to 1. The norm of \mathbf{u}_i is normalized to be

$$(\mathbf{u}_i, \mathbf{L}(\mathbf{u}_{i-1}) \mathbf{u}_i) = 1. \quad (48)$$

3. Return \mathbf{u}_i if \mathbf{u}_i is sufficiently close to \mathbf{u}_{i-1} except for the sign. Otherwise, let $\mathbf{u}_{i-1} \leftarrow \mathbf{u}_i$, and go back to step 2.

The resulting solution $\hat{\mathbf{u}}$ should satisfy

$$\mathbf{M}(\hat{\mathbf{u}}) \hat{\mathbf{u}} = \lambda \mathbf{L}(\hat{\mathbf{u}}) \hat{\mathbf{u}}, \quad (49)$$

for some λ . Taking the inner product of $\hat{\mathbf{u}}$ and both sides, we have

$$(\hat{\mathbf{u}}, \mathbf{M}(\hat{\mathbf{u}}) \hat{\mathbf{u}}) = \lambda, \quad (50)$$

because of the normalization convention given in Eq. (48), which implies from the second of Eqs. (46) that

$$1 = (\hat{\mathbf{u}}, \mathbf{L}(\hat{\mathbf{u}}) \hat{\mathbf{u}}) = \sum_{\alpha=1}^N \frac{(\boldsymbol{\xi}_\alpha, \mathbf{u})^2 (\hat{\mathbf{u}}, \mathbf{V}_0[\boldsymbol{\xi}_\alpha] \hat{\mathbf{u}})}{(\mathbf{u}, \mathbf{V}_0[\boldsymbol{\xi}_\alpha] \mathbf{u})^2} = \sum_{\alpha=1}^N \frac{(\boldsymbol{\xi}_\alpha, \mathbf{u})^2}{(\mathbf{u}, \mathbf{V}_0[\boldsymbol{\xi}_\alpha] \mathbf{u})}. \quad (51)$$

From the first of Eqs. (46), we see that

$$(\hat{\mathbf{u}}, \mathbf{M}(\hat{\mathbf{u}}) \hat{\mathbf{u}}) = \sum_{\alpha=1}^N \frac{(\hat{\mathbf{u}}, \boldsymbol{\xi}_\alpha)^2}{(\mathbf{u}, \mathbf{V}_0[\boldsymbol{\xi}_\alpha] \mathbf{u})} = 1, \quad (52)$$

meaning that $\lambda = 1$. Thus, $\hat{\mathbf{u}}$ is indeed the solution of Eq. (45). However, the matrix $\mathbf{L}(\mathbf{u})$ is usually singular, because the matrix $\mathbf{V}_0[\mathbf{x}_\alpha]$ in the second of Eqs. (46) is likely to degenerate. This is easily seen from Eq. (27): the dimension p of $\boldsymbol{\xi}_\alpha$ is generally larger than the dimension m of \mathbf{x}_α . Hence, the generalized eigenvalue problem in Eq. (47) needs to be reduced to subproblems of smaller dimensions. The reduced form (we omit the details, see [9]) was proposed by Leedan and Meer [28] and Matei and Meer [30] and called the *heteroscedastic errors-in-variables (HEIV) method*.

14.5.6 Renormalization Method

The reason why the solution of the naive method of Sect. is biased is that the matrix $\mathbf{M}(\mathbf{u})$ in Eq. (34) is biased. If we decompose the datum $\boldsymbol{\xi}_\alpha$ into its true value $\bar{\boldsymbol{\xi}}_\alpha$ and the noise term $\Delta\boldsymbol{\xi}_\alpha$, the expectation of Eq. (34) is

$$\begin{aligned} E[\boldsymbol{\xi}_\alpha \boldsymbol{\xi}_\alpha^\top] &= E[(\bar{\boldsymbol{\xi}}_\alpha + \Delta\boldsymbol{\xi}_\alpha)(\bar{\boldsymbol{\xi}}_\alpha + \Delta\boldsymbol{\xi}_\alpha)^\top] = E[\bar{\boldsymbol{\xi}}_\alpha \bar{\boldsymbol{\xi}}_\alpha^\top] + E[\bar{\boldsymbol{\xi}}_\alpha \Delta\boldsymbol{\xi}_\alpha^\top] \\ &\quad + E[\Delta\boldsymbol{\xi}_\alpha \bar{\boldsymbol{\xi}}_\alpha^\top] + E[\Delta\boldsymbol{\xi}_\alpha \Delta\boldsymbol{\xi}_\alpha^\top] = \bar{\boldsymbol{\xi}}_\alpha \bar{\boldsymbol{\xi}}_\alpha^\top + V_0[\boldsymbol{\xi}_\alpha]. \end{aligned} \quad (53)$$

Thus,

$$E[\mathbf{M}(\mathbf{u})] = \bar{\mathbf{M}}(\mathbf{u}) + \varepsilon^2 \mathbf{N}(\mathbf{u}) + O(\varepsilon^4), \quad (54)$$

where $\bar{\mathbf{M}}(\mathbf{u})$ is the value of $\mathbf{M}(\mathbf{u})$ evaluated using the true values $\{\bar{\boldsymbol{\xi}}_\alpha\}$ and

$$\mathbf{N}(\mathbf{u}) = \sum_{\beta=1}^N \frac{V_0[\boldsymbol{\xi}_\beta]}{(\mathbf{u}, V_0[\boldsymbol{\xi}_\beta] \mathbf{u})}. \quad (55)$$

Equation (54) implies that an unbiased solution can be obtained if the matrix $\mathbf{M}(\mathbf{u})$ in Eq. (35) is replaced by

$$\hat{\mathbf{M}}(\mathbf{u}) = \mathbf{M}(\mathbf{u}) - \varepsilon^2 \mathbf{N}(\mathbf{u}). \quad (56)$$

The square noise level ε^2 is unknown, but if we note that the smallest eigenvalue of $\bar{\mathbf{M}}(\mathbf{u})$ is 0, we can estimate ε^2 so that the smallest eigenvalue of $\hat{\mathbf{M}}(\mathbf{u})$ is 0. Thus, we obtain the following scheme:

1. Guess an appropriate initial value \mathbf{u}_0 , say the LS solution $\hat{\mathbf{u}}_{\text{LS}}$, and let $c_0 = 0$.
2. Assuming that \mathbf{u}_{i-1} and c_{i-1} are obtained (initially $i = 1$), solve the eigenvalue problem

$$(\mathbf{M}(\mathbf{u}_{i-1}) - c_{i-1} \mathbf{N}(\mathbf{u}_{i-1})) \mathbf{u} = \lambda \mathbf{u}. \quad (57)$$

Let \mathbf{u}_i be the unit eigenvector for the smallest eigenvalue λ .

3. Return \mathbf{u}_i if λ is sufficiently close to 0. Otherwise, let

$$c_i = c_{i-1} + \frac{\lambda}{(\mathbf{u}_i, \mathbf{N}(\mathbf{u}_{i-1}) \mathbf{u}_i)}. \quad (58)$$

4. Let $\mathbf{u}_{i-1} \leftarrow \mathbf{u}_i$, and go back to step 2.

Equations (57) and (58) imply that if c_i is close to 0 we have

$$(\mathbf{M}(\mathbf{u}_{i-1}) - c_i \mathbf{N}(\mathbf{u}_{i-1})) \mathbf{u}_i = \mathbf{0}. \quad (59)$$

In fact, the inner product of \mathbf{u}_i and the left-hand side is

$$\begin{aligned}
 & (\mathbf{u}_i, (\mathbf{M}(\mathbf{u}_{i-1}) - c_i \mathbf{N}(\mathbf{u}_{i-1})) \mathbf{u}_i) \\
 &= (\mathbf{u}_i, (\mathbf{M}(\mathbf{u}_{i-1}) - c_{i-1} \mathbf{N}(\mathbf{u}_{i-1})) \mathbf{u}_i) - \frac{\lambda(\mathbf{u}_i, \mathbf{N}(\mathbf{u}_{i-1}) \mathbf{u}_i)}{(\mathbf{u}_i, \mathbf{N}(\mathbf{u}_{i-1}) \mathbf{u}_i)} \\
 &= \lambda - \lambda = 0.
 \end{aligned} \tag{60}$$

If c_i is close to 0, the matrix $\mathbf{M}(\mathbf{u}_{i-1}) - c_i \mathbf{N}(\mathbf{u}_{i-1})$ is positive semidefinite, so Eq. (60) implies that \mathbf{u}_i is included in the null space of $\mathbf{M}(\mathbf{u}_{i-1}) - c_i \mathbf{N}(\mathbf{u}_{i-1})$, proving Eq. (59). Hence, the solution satisfies

$$(\mathbf{M}(\hat{\mathbf{u}}) - c \mathbf{N}(\hat{\mathbf{u}})) \hat{\mathbf{u}} = \mathbf{0}, \tag{61}$$

and c gives an estimate of ε^2 . This scheme was proposed by Kanatani [14] and called *renormalization*.

Remark 16. Historically, this method was proposed first; the HEIV and FNS methods were proposed as refinements to it. However, the renormalization solution and the HEIV/FNS solution (FNS and HEIV produce the same value) are both optimal in the sense that their covariance matrices differ only in the term $O(\varepsilon^4)$ in Eq. (29) [14]. This is confirmed by numerical simulations [7, 8, 9].

Remark 17. Renormalization tries to eliminate the bias term in Eq. (54) by “subtraction” in the form of Eq. (56). An alternative strategy would be to remove the bias by “division”. In fact, if we let $\tilde{\mathbf{M}}(\mathbf{u}) = \mathbf{N}(\mathbf{u})^{-1/2} \mathbf{M}(\mathbf{u}) \mathbf{N}(\mathbf{u})^{-1/2}$ (the negative square root is defined by replacing all its eigenvalues λ by $1/\sqrt{\lambda}$ in the canonical form), $E[\tilde{\mathbf{M}}(\mathbf{u})]$ and $\tilde{\mathbf{M}}(\mathbf{u})$ share the same eigenvectors up to $O(\varepsilon^4)$. If $\tilde{\mathbf{u}}$ is an eigenvector of $\tilde{\mathbf{M}}(\mathbf{u})$, the corresponding eigenvector of $\mathbf{M}(\mathbf{u})$ is $\mathbf{N}(\mathbf{u})^{-1/2} \tilde{\mathbf{u}}$. This implies that an unbiased solution is obtained by applying the naive method of Sect. to $\tilde{\mathbf{M}}(\mathbf{u})$. This strategy is known as *equilibration* or *whitening*. However, the matrix $\mathbf{N}(\mathbf{u})$ is often singular due to the degeneracy of $V_0[\boldsymbol{\xi}_\alpha]$ (cf. Sect.), so $\mathbf{N}(\mathbf{u})^{-1/2}$ cannot be computed. Still, it has been applied to a few problems for which $\mathbf{N}(\mathbf{u})$ does not degenerate [29, 35, 36].

14.5.7 Optimal Correction

In deriving the FNS, HEIV, and renormalization methods, we ignored all constraints on \mathbf{u} except $\|\mathbf{u}\| = 1$. Let the remaining constraints be

$$\phi^{(k)}(\mathbf{u}) = 0, \quad k = 1, \dots, r. \tag{62}$$

From Eq. (29), the normalized covariance of the ML estimator $\hat{\mathbf{u}}$ is given by

$$V_0[\hat{\mathbf{u}}] = \left(\mathbf{P}_{\hat{\mathbf{u}}} \mathbf{M}(\hat{\mathbf{u}}) \mathbf{P}_{\hat{\mathbf{u}}} \right)^{-}, \tag{63}$$

where $\mathbf{M}(\mathbf{u})$ is defined in Eq. (34) (or in Eqs. (46)). The maximum likelihood solution of \mathbf{u} that satisfies the constraint (62) is obtained to a first approximation by minimizing

$$J = (\hat{\mathbf{u}} - \mathbf{u}, \mathbf{V}_0[\hat{\mathbf{u}}]^{-1}(\hat{\mathbf{u}} - \mathbf{u})) \tag{64}$$

subject to Eq. (62). Introducing Lagrange multipliers and first-order approximation, we obtain the following solution [14]:

$$\mathbf{u}^* = \hat{\mathbf{u}} - \mathbf{V}_0[\hat{\mathbf{u}}] \sum_{k,l=1}^r w^{(kl)} \hat{\phi}^{(k)} \nabla_{\mathbf{u}} \hat{\phi}^{(l)}. \tag{65}$$

Here, $w^{(kl)}$ is the (kl) element of the inverse of the $r \times r$ matrix whose (kl) element is $(\nabla_{\mathbf{u}} \hat{\phi}^{(k)}, \mathbf{V}_0[\hat{\mathbf{u}}] \nabla_{\mathbf{u}} \hat{\phi}^{(l)})$, i.e.,

$$(w^{(kl)}) = \left((\nabla_{\mathbf{u}} \hat{\phi}^{(k)}, \mathbf{V}_0[\hat{\mathbf{u}}] \nabla_{\mathbf{u}} \hat{\phi}^{(l)}) \right)^{-1}. \tag{66}$$

The hat means that the ML estimator $\hat{\mathbf{u}}$ is substituted for \mathbf{u} . The normalized covariance matrix of the corrected value \mathbf{u}^* of Eq. (65) is

$$\mathbf{V}_0[\mathbf{u}^*] = \mathbf{V}_0[\hat{\mathbf{u}}] - \sum_{k,l=1}^r w^{(kl)} (\mathbf{V}_0[\hat{\mathbf{u}}] \nabla_{\mathbf{u}} \hat{\phi}^{(k)}) (\mathbf{V}_0[\hat{\mathbf{u}}] \nabla_{\mathbf{u}} \hat{\phi}^{(l)})^T \tag{67}$$

up to $O(\varepsilon^2)$ [14]. For a single constraint, Eqs. (65) and (67) reduce to

$$\mathbf{u}^* = \hat{\mathbf{u}} - \frac{\hat{\phi} \mathbf{V}_0[\hat{\mathbf{u}}] \nabla_{\mathbf{u}} \hat{\phi}}{(\nabla_{\mathbf{u}} \hat{\phi}, \mathbf{V}_0[\hat{\mathbf{u}}] \nabla_{\mathbf{u}} \hat{\phi})}, \tag{68}$$

$$\mathbf{V}_0[\mathbf{u}^*] = \mathbf{V}_0[\hat{\mathbf{u}}] - \frac{(\mathbf{V}_0[\hat{\mathbf{u}}] \nabla_{\mathbf{u}} \hat{\phi})(\mathbf{V}_0[\hat{\mathbf{u}}] \nabla_{\mathbf{u}} \hat{\phi})^T}{(\nabla_{\mathbf{u}} \hat{\phi}, \mathbf{V}_0[\hat{\mathbf{u}}] \nabla_{\mathbf{u}} \hat{\phi})}. \tag{69}$$

Remark 18. If the r constraints in Eq. (62) are redundant, say only r' ($< r$) of them are independent, the inverse in Eq. (66) is replaced by the generalized inverse of rank r' (cf. Remark 6).

Remark 19. If all the r constraints in Eq. (62) are independent, the rank of the matrix $\mathbf{V}_0[\mathbf{u}^*]$ given by Eq. (65) is smaller than $\mathbf{V}_0[\hat{\mathbf{u}}]$ by r . Intuitively, the ellipsoid that represents the uncertainty of \mathbf{u} in \mathcal{R}^p “collapses” in the r directions in which the constraint (62) is violated, while it keeps its shape in the directions orthogonal to them. Hence, the optimality of the ML estimator is not affected by doing this type of posterior correction [14].

Remark 20. Equation (65) enforces all the constraints only to a first approximation, so $\phi^{(k)}(\mathbf{u}^*)$, $k = 1, \dots, r$, may not exactly be 0, and \mathbf{u}^* may not

exactly be a unit vector. Such higher-order discrepancies can be eliminated by iterating Eqs. (68) and (69) in the form

$$\mathbf{u}^* \leftarrow N\left[\hat{\mathbf{u}} - \frac{\hat{\phi} \mathbf{V}_0[\hat{\mathbf{u}}] \nabla_{\mathbf{u}} \hat{\phi}}{(\nabla_{\mathbf{u}} \hat{\phi}, \mathbf{V}_0[\hat{\mathbf{u}}] \nabla_{\mathbf{u}} \hat{\phi})}\right], \quad (70)$$

$$\mathbf{V}_0[\mathbf{u}^*] \leftarrow \mathbf{P}_{\mathbf{u}^*} \left(\mathbf{V}_0[\hat{\mathbf{u}}] - \frac{(\mathbf{V}_0[\hat{\mathbf{u}}] \nabla_{\mathbf{u}} \hat{\phi})(\mathbf{V}_0[\hat{\mathbf{u}}] \nabla_{\mathbf{u}} \hat{\phi})^\top}{(\nabla_{\mathbf{u}} \hat{\phi}, \mathbf{V}_0[\hat{\mathbf{u}}] \nabla_{\mathbf{u}} \hat{\phi})} \right) \mathbf{P}_{\mathbf{u}^*}, \quad (71)$$

where $N[\cdot]$ denotes normalization to a unit vector ($N[\mathbf{v}] = \mathbf{v}/\|\mathbf{v}\|$), and $\mathbf{P}_{\mathbf{u}^*}$ is the projection matrix defined by Eq. (30). Equation (71) makes the null space of the $\mathbf{V}_0[\mathbf{u}^*]$ exactly compatible with \mathbf{u}^* .

14.6 Nuisance Parameters and Semiparametric Model

Finally, we discuss some new topics related to the use of statistical methods for geometric inference.

14.6.1 Asymptotic Parameters

The number n that appears in the standard statistical analysis is the *number of experiments*. It is also called the *number of trials*, the *number of observations*, and the *number of samples*. Evidently, the properties of the ensemble are revealed more precisely as more data are sampled from it.

However, the number n is often called the *number of data*, which has caused considerable confusion. For example, if we observe a 100-dimensional vector datum in one experiment, one may think that the “number of data” is 100, but this is wrong: the number n of experiments is 1. We are observing 1 sample from an ensemble of 100-dimensional vectors.

For character recognition, the underlying ensemble is the set of possible character images, and the learning process concerns the number n of training steps necessary to establish satisfactory responses. This is independent of the dimension N of the vector that represents each character. The learning performance is evaluated asymptotically as $n \rightarrow \infty$, not $N \rightarrow \infty$.

For geometric inference, however, many researchers have taken the dimension of the data as the “number of data” perhaps because the ensemble is hypothetical and one cannot sample more than one datum from it. However, if we extract, for example, 50 feature points, they constitute a 100-dimensional vector consisting of their x and y coordinates. If no other information, such as the image intensity, is used, the image is completely characterized by that vector. Applying a statistical method means regarding it as a sample from a hypothetical ensemble of 100-dimensional vectors.

14.6.2 Neyman–Scott Problem

In the past, many computer vision researchers have analyzed the asymptotic behavior as $N \rightarrow \infty$ without explicitly mentioning what the underlying ensemble is. This is perhaps motivated by a similar formulation in the statistical literature. Suppose, for example, a rodlike structure lies on the ground in the distance. We emit a laser beam toward it and estimate its position and orientation by observing the reflection of the beam, which is contaminated by noise. We assume that the laser beam can be emitted in any orientation any number of times, but the emission orientation is measured with noise. The task is to estimate the position and orientation of the structure as accurately as possible by emitting as small a number of beams as possible. Naturally, the estimation performance should be evaluated in the asymptotic limit $n \rightarrow \infty$ with respect to the number n of emissions.

The underlying ensemble is the set of all response times for all possible directions of emission. Usually, we are interested in the position and orientation of the structure but not the exact orientation of each emission, so the variables for the former are called the *structural parameters*, which are fixed in number, while the latter are called the *nuisance parameters*, which increase indefinitely as the number n of experiments increases [2]. Such a formulation is called the *Neyman–Scott problem* [37]. Since the constraint is an implicit function in the form of Eq. (3), we are considering an *errors-in-variables model* [10]. If we linearize the constraint by changing variables, the noise characteristics differs for each data component, so the problem is *heteroscedastic* [28].

To solve this problem, one can introduce a parametric model for the distribution of possible laser emission orientations, regarding the actual emissions as random samples from it. This formulation is called a *semiparametric model* [2]. An optimal solution can be obtained by finding a good *estimating function* [2, 40].

14.6.3 Semiparametric Model for Geometric Inference

Since the semiparametric model has something different from the geometric inference problem described in Sect. , a detailed analysis is required for examining if application of a semiparametric model to geometric inference will yield a desirable result [38, 40]. In any event, one should explicitly state what kind of ensemble (or ensemble of ensembles) is assumed before doing statistical analysis.

This is not merely a conceptual issue. It also affects the performance evaluation of simulation experiments. In doing a simulation, one can freely change the number N of feature points and the noise level ε . If the accuracy of method A is higher than method B for particular values of N and ε , one cannot conclude that method A is superior to method B, because opposite results may

come out for other values of N and ε . Here, we have two alternatives for performance evaluation: fixing ε and varying N to see if admissible accuracy is attained for a smaller number of feature point; fixing N and varying ε to see if larger data uncertainty can be tolerated for admissible accuracy. These two types of evaluation have different meanings. Our conclusion is that the results of one type of evaluation cannot directly be compared with the results of the other.

14.7 Conclusions

We have investigated the meaning of “statistical methods” for geometric inference based on image feature points. Tracing back the origin of feature uncertainty to image processing operations, we discussed the implications of asymptotic analysis in reference to geometric fitting and geometric model selection. We pointed out that a correspondence exists between the standard statistical analysis and the geometric inference problem. We also compared the capability of the geometric AIC and the geometric MDL in detecting degeneracy. Next, we reviewed recent progress in geometric fitting techniques for linear constraints, describing the FNS method, the HEIV method, the renormalization method, and other related techniques. Finally, we discussed the Neyman–Scott problem and semiparametric models in relation to geometric inference.

From these discussions, we conclude that applications of statistical methods require careful consideration about the nature of the problem in question and that different statistical theories are necessary for different classes of problems. In this sense, there is much room for new statistical theories to emerge as the scope of computer vision research expands. The important thing is, however, to always make clear the underlying hypotheses and assumptions, and to not simply use the methods in the statistical literature.

References

1. Akaike H. (1977) A new look at the statistical model identification. *IEEE Trans. Autom. Control* **16**:716–723
2. Amari S., Kawanabe M. (1997) Information geometry of estimating functions in semiparametric statistical models. *Bernoulli* **3**:29–54
3. Bubna K., Stewart C.V. (2000) Model selection techniques and merging rules for range data segmentation algorithms. *Comput. Vision Image Understand.* **80**:215–245
4. Chabat F., Yang G.Z., Hansell D.M. (1999) A corner orientation detector. *Image Vision Comput.* **17**:761–769
5. Cho K., Meer P. (1997) Image segmentation from consensus information. *Comput. Vision Image Understand.* **68**:72–89

6. Cho K., Meer P., Cabrera J. (1997) Performance assessment through bootstrap. *IEEE Trans. Patt. Anal. Mach. Intell.* **19**:1185–1198
7. Chojnacki W., Brooks M.J., van den Hengel A., Gawley D. (2000) On the fitting of surfaces to data with covariances. *IEEE Trans. Patt. Anal. Mach. Intell.* **22**:1294–1303
8. Chojnacki W., Brooks M.J., van den Hengel A. (2001) Rationalising the renormalisation method of Kanatani. *J. Math. Imaging Vision* **14**:21–38
9. Chojnacki W., Brooks M.J., van den Hengel A., Gawley D. (2004) From FNS to HEIV: A link between two vision parameter estimation methods. *IEEE Trans. Patt. Anal. Mach. Intell.* **26**:264–268
10. Fuller W.A. (1987) *Measurement Error Models*, Wiley, New York
11. Gu H., Shirai Y., Asada M. (1996) MDL-based segmentation and motion modeling in a long sequence of scene with multiple independently moving objects *IEEE Trans. Patt. Anal. Mach. Intell.* **18**:58–64
12. Harris C., Stephens M. (1988) A combined corner and edge detector. In: *Proceedings of Fourth Alvey Vision Conference*. Manchester, UK, pp. 147–151
13. Hartley R., Zisserman A. (2000) *Multiple View Geometry in Computer Vision*. Cambridge University Press, Cambridge, UK
14. Kanatani K. (1996) *Statistical Optimization for Geometric Computation: Theory and Practice*. Elsevier, Amsterdam
15. Kanatani K. (1998) Geometric information criterion for model selection. *Int. J. Comput. Vision* **26**:171–189
16. Kanatani K. (1998) Statistical optimization and geometric inference in computer vision. *Phil. Trans. Roy. Soc. Lond.* **A356**:1303–1320
17. Kanatani K. (1988) Cramer–Rao lower bounds for curve fitting. *Graphical Models Image Process.* **60**:93–99
18. Kanatani K. (2000) Model selection criteria for geometric inference. In: Bab-Hadiashar A., Suter D. (eds.) *Data Segmentation and Model Selection for Computer Vision: A Statistical Approach*, Springer, Berlin Heidelberg New York, pp. 91–115
19. Kanatani K. (2001) Motion segmentation by subspace separation and model selection. In: *Proceedings of Eighth International Conference on Computer Vision, vol. 2*. Vancouver, Canada, pp. 301–306
20. Kanatani K. (2002) Model selection for geometric inference, plenary talk. In: *Proceedings of Fifth Asian Conference on Computer Vision, vol. 1*. Melbourne, Australia, pp. xxi–xxxii
21. Kanatani K. (2002) Motion segmentation by subspace separation: Model selection and reliability evaluation. *Int. J. Image Graphics* **2**:179–197
22. Kanatani K. (2002) Evaluation and selection of models for motion segmentation. In: *Proceedings of Seventh European Conference on Computer Vision, vol. 3*. Copenhagen, Denmark, pp. 335–349
23. Kanatani K., Matsunaga C. (2002) Estimating the number of independent motions for multibody motion segmentation. In: *Proceedings of Fifth Asian Conference on Computer Vision, vol. 1*. Melbourne, Australia, pp. 7–12
24. Kanazawa Y., Kanatani K. (1997) Infinity and planarity test for stereo vision. *IEICE Trans. Inf. & Syst.* **E80-D**:774–779
25. Kanazawa Y., Kanatani K. Stabilizing image mosaicing by model selection. In: Pollefeys M., Van Gool L., Zisserman A., Fitzgibbon A. (eds.) *3D Structure from Images–SMILE 2000. LNCS, vol. 2018*. Springer, Berlin Heidelberg New York, 2001, pp. 35–51

26. Kanazawa Y., Kanatani K. (2001) Do we really have to consider covariance matrices for image features? In: *Proceedings of Eighth International Conference on Computer Vision, vol. 2*. Vancouver, Canada, pp. 586–591
27. Leclerc Y.G. (1989) Constructing simple stable descriptions for image partitioning. *Int. J. Comput. Vision* **3**:73–102
28. Leedan Y., Meer P. (2000) Heteroscedastic regression in computer vision: Problems with bilinear constraint. *Int. J. Comput. Vision*. **37**:127–150
29. MacLean W.J. (1999) Removal of translation bias when using subspace methods. In: *Proceedings of Seventh International Conference on Computer Vision, vol. 2*. Kerkyra, Greece, pp. 753–758
30. Matei B., Meer P. (2000) A generalized method for errors-in-variables problem in computer vision. In: *Proceedings of Fifteenth International Conference on Pattern Recognition, vol. 2*. Barcelona, Spain, pp. 18–25
31. Matsunaga C., Kanatani K. (2000) Calibration of a moving camera using a planar pattern: Optimal computation, reliability evaluation and stabilization by model selection. In: *Proceedings of Sixth European Conference on Computer Vision, vol. 2*. Dublin, Ireland, pp. 595–609
32. Maxwell B.A. (2000) Segmentation and interpretation of multicolored objects with highlights. *Comput. Vision Image Understand.* **77**:1–24
33. Maybank S.J., Sturm P.F. (1999) MDL, collineations and the fundamental matrix. In: *Proceedings of Tenth British Machine Vision Conference*. Nottingham, UK, pp. 53–62
34. Morris D.D., Kanatani K., Kanade T. (2001) Gauge fixing for accurate 3D estimation. In: *Proceedings of IEEE Conference on Computer Vision Pattern Recognition, vol. 2*. Kauai, Hawaii, pp. 343–350
35. Mühlich M., Mester R. (1998) The role of total least squares in motion analysis. In: *Proceedings of Fifth European Conference on Computer Vision, vol. 2*. Freiburg, Germany, pp. 305–321
36. Mühlich M., Mester R. (2001) A considerable improvement in pure parameter estimation using TLS and equilibration. *Patt. Recog. Lett.* **22**:1181–1189
37. Neyman J., Scott E.L. (1948) Consistent estimates based on partially consistent observations. *Econometrica* **16**:1–32
38. Ohta N. (2003) Motion parameter estimation from optical flow without nuisance parameters. In: *Third International Workshop on Statistical and Computational Theory of Vision*. Nice, France:
<http://www.stat.ucla.edu/~yuille/meetings/2003-workshop.php>. Cited 12 October 2003
39. Ohta N., Kanatani K. (1998) Moving object detection from optical flow without empirical thresholds. *IEICE Trans. Inf. & Syst.* **E81-D**:243–245
40. Okatani T., Deguchi K. (2003) Toward a statistically optimal method for estimating geometric relations from noisy data: Cases of linear relations. In: *Proceedings of IEEE Conference on Computer Vision Pattern Recognition, vol. 1*. Madison, WI, pp. 432–439
41. Reissfeld D., Wolfson H., Yeshurun Y. (1995) Context-free attentional operators: The generalized symmetry transform. *Int. J. Comput. Vision* **14**:119–130
42. Rissanen J. (1984) Universal coding, information, prediction and estimation. *IEEE Trans. Inform. Theory* **30**:629–636
43. Rissanen J. (1989) *Stochastic Complexity in Statistical Inquiry*. World Scientific, Singapore

44. Rissanen J. (1996) Fisher information and stochastic complexity, *IEEE Trans. Inform. Theory* **42**:40–47
45. Schmid C., Mohr R., Bauckhage C. (2000) Evaluation of interest point detectors. *Int J. Comput. Vision* **37**:151–172
46. Smith S.M., Brady J.M. (1997) SUSAN—A new approach to low level image processing. *Int. J. Comput. Vision* **23**:45–78
47. Sugaya Y., Kanatani K. (2003) Outlier removal for feature tracking by subspace separation. *IEICE Trans. Inf. & Syst.* **E86-D**:1095–1102
48. Torr P.H.S. (1997) An assessment of information criteria for motion model selection. In: *Proceedings of IEEE Conference on Computer Vision Pattern Recognition*. Puerto Rico, pp. 47–53
49. Torr P.H.S. (1998) Geometric motion segmentation and model selection. *Phil. Trans. Roy. Soc. Lond.* **A356**:1321–1340
50. Torr P.H.S. (2002) Bayesian model estimation and selection for epipolar geometry and generic manifold fitting. *Int. J. Comput. Vision* **50**:35–61
51. Torr P.H.S., FitzGibbon A., Zisserman A. (1998) Maintaining multiple motion model hypotheses through many views to recover matching and structure. In: *Proceedings of Sixth International Conference on Computer Vision*. Bombay, India, pp. 485–492
52. Torr P.H.S., Zisserman A. (2000) Concerning Bayesian motion segmentation, model averaging, matching and the trifocal tensor. In: *Proceedings of Sixth European Conference on Computer Vision, vol. 1*. Dublin, Ireland, pp. 511–528
53. Triono I., Ohta N., Kanatani K. (1998) Automatic recognition of regular figures by geometric AIC, *IEICE Trans. Inf. & Syst.* **E81-D**:246–248