3D Model Accuracy and Gauge Fixing

CMU-RI-TR-00-32

Daniel D. Morris   Kenichi Kanatani†   Takeo Kanade

Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

†Department of Computer Science
Gunma University
Kiryu, Gunma 376-8515, Japan

December 2000

Abstract

In computer vision we often estimate a 3D model of an object, and its uncertainty, up to an unknown scale factor. This unknown scale factor means that we cannot directly infer positions and lengths on the model, nor their uncertainties. In order to make quantitative distance measurements on this model, we must obtain the unknown scale factor. However, how we determine the scale factor will affect the accuracy of the rescaled model. In this paper we study the problem of estimating the absolute scale of an object and its uncertainty, starting from a 3D reconstruction up to a scale factor, and a reference length. The theory derived here can be used both to correctly transform a model covariance matrix under rescaling, and to select a good length on the object from which to obtain the scale, and so maximize accuracy.

Contents

1 Introduction

Methods for imaging objects have been used since ancient civilizations to make quantitative measurements of the world. The classic example is trigonometry, where elevation angles and distance measurements are used to calculate the height of a difficult-to-measure object. Modern computer vision provides tools for combining many measurements to obtain more sophisticated and accurate estimates of the 3D world [4, 10] including, for instance, ancient archeological structures [11]. However, many image-only methods estimate 3D only up to a scale factor, and so do not give direct length estimates. An additional step is needed to estimate the scale factor of the 3D world. In this work, we investigate how the accuracy of 3D estimates are affected by this unknown scale factor.

A covariance matrix is a widely used measure for the uncertainty of a model. We assume that our image-based 3D shape modeling algorithm returns a shape estimate and a covariance matrix, but that there is a global unknown scale factor that must be determined before the model and the real object can be made the same size. This unknown scale factor in 3D estimation has caused no great concern in the computer vision field. It is thought that we can measure a length of the real object, find the unknown scale factor, and then use this as a change-of-variables transformation to rescale both the model and the covariance as illustrated in Figure 1. This widely held assumption, however, is false. In section 2, we show that this in general gives the wrong covariance matrix, and that other covariances, like those illustrated in Figure 2 may result from rescaling to object.

The goal of this paper is to find the correct transformation of a covariance matrix that rescales it and also takes the measurement information into account. The aspect that makes this problem interesting, and not simply a change-of-variables, is the scale indeterminacy. Recently, gauge theory [8, 9, 10, 13] has been developed in the computer vision context to correctly treat indeterminacies in problems like this. This work on gauge theory has claimed that only gauge invariant properties are properly described by a model and its covariance matrix when it has indeterminacies or gauge freedoms. However, in this paper we will show that with the additional information obtained by a measurement, we can obtain an exact covariance matrix for the shape with no indeterminacy. We derive the correct transformation for obtaining the true covariance matrix and we illustrate this on a real example. Finally we derive some guidelines to maximize the final accuracy.

It is a surprise, that how the scale factor is determined, is critically important to covariance analysis. Previous work that dealt with shape uncertainties ignored the unknown scale factor [14, 12] or stated that how it is determined is unimportant [5]. But, as will be shown in this paper, the method by which the model scale is determined will significantly affect the accuracy of the model.

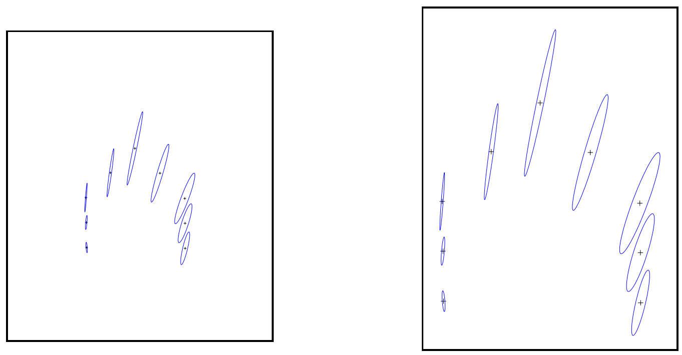

(a) (b)

Figure 1: (a) An object in the plane is shown by feature points, and its uncertainty is represented by the ellipses. (b) A standard change-of-variables is used to rescale the object and its covariance. We show in this paper that when the object is known only up to a scale factor, this in general is *not* the correct way to rescale the covariance of the object.

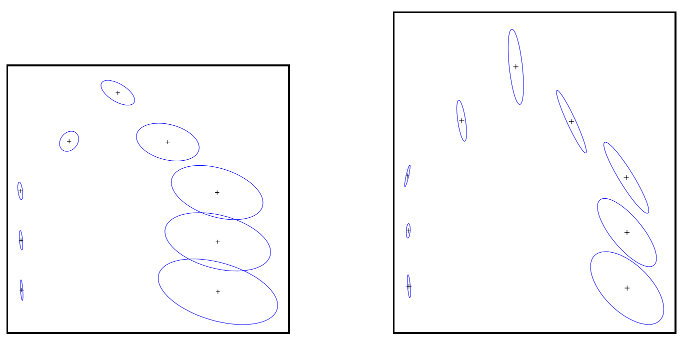

(a) (b)

Figure 2: If the object in Figure 1 is known only up to a scale factor, then there are many ways to rescale it which will not affect the final scale, but will change the covariance. Here two different methods are used to rescale the object, and pictured are the resulting covariances.

2 Problem Statement

Let us represent the shape of our object with a vector, $\boldsymbol{s}$, that contains all of the x, y, and z coordinates of all the individual feature points:

$$\boldsymbol{s} = \begin{pmatrix} \boldsymbol{s}_1 \\ \vdots \\ \boldsymbol{s}_N \end{pmatrix}, \quad \text{where} \quad \boldsymbol{s}_i = \begin{pmatrix} x_i \\ y_i \\ z_i \end{pmatrix}. \tag{1}$$

Here, $\boldsymbol{s}_i$ denotes the 3D position of the ith feature point. Let us assume further that the total shape, $\boldsymbol{s}$, is known only up to an unknown scale factor, and that its uncertainty, $\Delta\boldsymbol{s}$, has zero mean and a covariance matrix, $\boldsymbol{V}_{\boldsymbol{s}}$:

$$\boldsymbol{V}_{\boldsymbol{s}} = E[\Delta\boldsymbol{s}\,\Delta\boldsymbol{s}^{\top}]\,, \tag{2}$$

where $E[\cdot]$ denotes the expectation.

Since our model is known only up to a scale factor, we perform an additional measurement in order to specify the model exactly. Say we measure the distance between feature points, i and j, on the real object and find that its value is d'. We conclude that for our model, $\boldsymbol{s}$, to correspond to the true object, it must be rescaled by a scale factor given by:

$$a = \frac{d'}{||\boldsymbol{s}_i - \boldsymbol{s}_j||}\,. \tag{3}$$

The new model is then:

$$\boldsymbol{s}' = a\boldsymbol{s}\,. \tag{4}$$

Our question is: Now that we know the scale, what is the covariance, $\boldsymbol{V}_{\boldsymbol{s}'}$, of the rescaled model?

2.1 Naive Solution

A direct answer may be obtained as follows. Perturbations of the new model are given by:

$$\Delta\boldsymbol{s}' = a\Delta\boldsymbol{s}\,, \tag{5}$$

and so its covariance is:

$$\begin{aligned} \boldsymbol{V}_{\boldsymbol{s}'} &= E[\Delta\boldsymbol{s}'\,\Delta\boldsymbol{s}'^{\top}] \\ &= a^2\boldsymbol{V}_{\boldsymbol{s}}\,. \end{aligned} \tag{6}$$

Thus we would conclude that we can rescale the covariance with the square of the scale factor. This is exactly a change-of-variables transformation. It is correct way to transform the covariance matrix if we are changing the measurement units. Hence one might suppose this is how we should rescale a covariance matrix in general, and so, for example, obtain the result in Figure 1.

2.2 Contradiction

However, it is easy to see that there is something wrong with equation (6). We have measured the distance, d', between points i and j on our rescaled model, $\boldsymbol{s}'$. Assuming no errors in the measurement of d',[1] then this distance has no uncertainty and so must have zero variance. But in our new model, $\boldsymbol{s}'$, we have: $d' = ||\boldsymbol{s}'_i - \boldsymbol{s}'_j||$, and so perturbations of d' are expressed:

$$\Delta d' = \nabla_{\boldsymbol{s}'} d'^{\top} \Delta \boldsymbol{s}' \,. \tag{7}$$

Then, using the covariance in equation (6), the variance of d' would be:

$$\sigma^2_{d'} = \nabla_{\boldsymbol{s}'} d'^{\top} \boldsymbol{V}_{\boldsymbol{s}'} \nabla_{\boldsymbol{s}'} d' \,, \tag{8}$$

where we have defined the gradient of a scalar d' to be:

$$\nabla_{\boldsymbol{s}'} d' = \begin{pmatrix} \partial d' / \partial s'_{x1} \\ \vdots \\ \partial d' / \partial s'_{zN}) \end{pmatrix} \,. \tag{9}$$

This calculated variance, $\sigma^2_{d'}$, could be zero for a special covariance, $\boldsymbol{V}_{\boldsymbol{s}}$, and choice of features, but in general it will not be zero. By assumption, however, this distance has zero variance, and so there is a contradiction. We conclude that even though we have the correct scale factor a, unless the variance, $\sigma^2_{d'}$, between the points we chose is zero, simply scaling the covariance by a^2, as in equation (6), is the wrong transformation.

3 Shape up to a scale factor

When the shape of an object is recovered by an image-based method such as Structure from Motion, its scale is arbitrary. The object may have been larger and further away from the camera, or smaller and closer to the camera without changing the measurements at all. This indeterminacy is called a *gauge freedom*. In this section we will describe the basics of gauge theory and how invariants are able to capture the essential information of a model and its covariance. A more detailed description of gauge theory can be found in [7, 8, 10]. We will use this to show how the measurement of a length of an object should fix the scale and transform the covariance.

3.1 Gauge Orbits

When a model space, $\mathcal{S}$, is defined only up to an unknown transformation, such as scale, we say the space is filled with *gauge orbits*. A gauge orbit is a manifold within the model space, $\boldsymbol{G}_{\boldsymbol{s}} \in \mathcal{S}$, where all points on the gauge orbit are the same under the model interpretation.[2] We say all the points on a gauge orbit are *geometrically*

[1] We show in section 3.6 how to incorporate measurement uncertainties, but for simplicity we do not consider that here.

[2] Mathematicians call such an orbit a *leaf*, and a space filled with leafs a *foliation*.

equivalent. A *gauge transformation*, g, takes one point, $\boldsymbol{s} \in \boldsymbol{G_s}$ to another point on the orbit: $\boldsymbol{s}' \in \boldsymbol{G_s}$ and is expressed as:

$$\boldsymbol{s}' = g(\boldsymbol{s})\,, \qquad \forall g \in \mathcal{G} \tag{10}$$

where $\mathcal{G}$ is the group of gauge transformations. Figure 3(a) illustrates a point on its gauge orbit. In our case scale acts as a linear gauge transformation, $g(\boldsymbol{s}) = a\boldsymbol{s}$, and points on the gauge orbit correspond to models that are identical except by a scale factor.

3.2 Constraints and Covariance Subspace

The covariance matrix provides a first order perturbation analysis around a point in the parameter space. This error probability is illustrated by ellipses representing contours of equal probability in Figure 3(a). However, our shape belongs not to a single point, but to an orbit of points, all geometrically equivalent. Thus it does not make sense that ellipses with different probability could intersect an orbit, all of whose points are equivalent and so have equal probability. In order to speak, then, of perturbation analysis we need to add a set of constraints that restricts the shape to a unique point on the orbit. In our case, the gauge freedom is one dimensional, and so a single constraint on the shape will suffice. Denote this as:

$$c(\boldsymbol{s}) = 0\,. \tag{11}$$

This defines a *gauge manifold*, or simply *gauge*, $\mathcal{C}$, of points satisfying this constraint and having dimension $3N-1$. Enforcing a constraint like this, we call *gauge fixing* or *choosing a gauge*. Perturbations must be in the tangent space of this manifold, as illustrated in Figure 3(b). Thus the rank of the covariance matrix, $\boldsymbol{V_s}$, that governs these perturbations must be at most $3N-1$. Derivations of how to obtain this covariance can be found in references [7, 8, 10].

3.3 Problem Reformulation

We started with a shape, $\boldsymbol{s}$, and covariance, $\boldsymbol{V_s}$, in an unknown gauge, $\mathcal{C}$. We made a measurement, d', and will use this as a constraint on our model:

$$d' - ||\boldsymbol{s}'_i - \boldsymbol{s}'_j|| = 0\,, \tag{12}$$

and as such it defines a new gauge, $\mathcal{C}'$. We obtained a rescaled shape, $\boldsymbol{s}' = a\boldsymbol{s}$, that belongs to the same gauge orbit, $\boldsymbol{G_s}$, and also satisfies the constraint of this new gauge, $\mathcal{C}'$. Now we want the covariance, $\boldsymbol{V_{s'}}$, at $\boldsymbol{s}'$ that lies in the tangent space of $\mathcal{C}'$. To obtain this covariance, we will derive a geometric equivalence relationship using gauge invariants.

3.4 Perturbations of Invariants

Now we consider gauge invariants as a way to extract the true information from the shape and its covariance. A *gauge invariant* is a function of the model parameters,

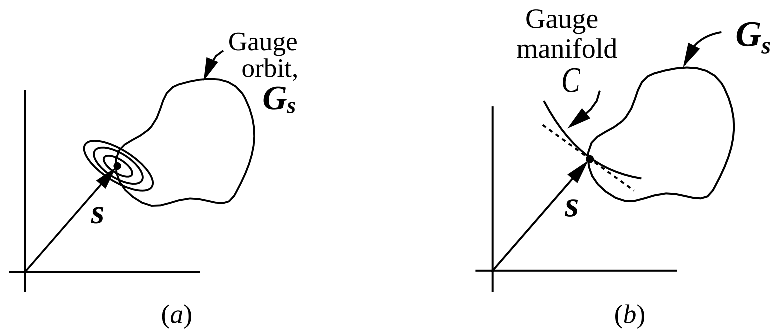

Figure 3: Plot (a) shows the contradiction we obtain when we use a full rank covariance to describe perturbations of a point on a gauge orbit. Ellipses of different probability intersect the gauge orbit, all of whose points are equivalent and hence have equal probability. Instead, we must choose a gauge by imposing a constraint that reduces our parameter space to the gauge manifold, $\mathcal{C}$, which intersects the gauge orbit at a single point, as shown in (b). Perturbations are now restricted to the tangent plane to the gauge manifold, shown by the dashed line. (These plots are two dimensional for clarity only. In our case, the vector, $\boldsymbol{s}$, is $3N$ dimensional, the gauge orbit, $\boldsymbol{G_s}$, is one dimensional, and the gauge, $\mathcal{C}$, is $3N-1$ dimensional.)

$I(\boldsymbol{s})$, that is the same for all points on the gauge orbit:

$$I(\boldsymbol{s}) = I(a\boldsymbol{s}) \qquad \forall a \neq 0\,. \tag{13}$$

Since our model is defined "only up to a scale factor," it must be gauge invariants that are the true quantities described by the shape. Properties that are invariant to scale include angles and ratios of lengths.

Since an invariant remains constant on a gauge orbit, its gradient along the gauge orbit must be zero. This means:

$$\nabla_{\boldsymbol{s}} I \in T[\boldsymbol{G_s}]^{\perp} \tag{14}$$

where $T[\boldsymbol{G_s}]$ is the tangent space to the gauge orbit at $\boldsymbol{s}$, and "$\perp$" refers to the orthogonal complement. This is true for all invariants.

We can now analyze the transformation of perturbations in one gauge, $\mathcal{C}$, to another gauge, $\mathcal{C}'$. The result will be used to transform the covariances.

A perturbation of an invariant $I(\boldsymbol{s})$ can be expressed as:

$$\Delta I(\boldsymbol{s}) = \nabla_{\boldsymbol{s}} I(\boldsymbol{s})^{\top}\, \Delta \boldsymbol{s}\,. \tag{15}$$

But since $I(\boldsymbol{s}) = I(\boldsymbol{s}')$, where $\boldsymbol{s}' = a\boldsymbol{s}$, equivalently:

$$\begin{aligned} \Delta I(\boldsymbol{s}) &= \nabla_{\boldsymbol{s}'} I(\boldsymbol{s}')^{\top}\, \Delta \boldsymbol{s}' = \nabla_{\boldsymbol{s}} I(\boldsymbol{s})^{\top} \frac{\partial \boldsymbol{s}}{\partial \boldsymbol{s}'}^{\top} \Delta \boldsymbol{s}' \\ &= \nabla_{\boldsymbol{s}} I(\boldsymbol{s})^{\top} \frac{1}{a} \Delta \boldsymbol{s}'\,, \end{aligned} \tag{16}$$

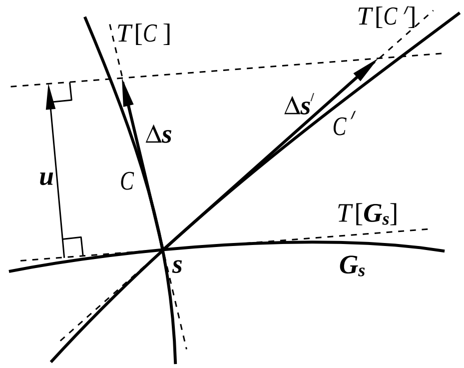

Figure 4: Perturbations in the tangent planes to two gauges, $\mathcal{C}$ and $\mathcal{C}'$, are shown when $a = 1$. These perturbations are geometrically equivalent when their difference, $\Delta \boldsymbol{s} - \Delta \boldsymbol{s}'/a$, is orthogonal to all vectors $\boldsymbol{u} \in T[\boldsymbol{G_s}]^\perp$.

where we represent the Jacobian matrix for a change-of-variables as:

$$\frac{\partial \boldsymbol{s}}{\partial \boldsymbol{s}'} = \frac{1}{a}\boldsymbol{I} \tag{17}$$

and $\boldsymbol{I}$ is the identity matrix. Now from equations (15) and (16) we have:

$$\nabla_{\boldsymbol{s}} I(\boldsymbol{s})^\top (\Delta \boldsymbol{s} - \frac{1}{a}\Delta \boldsymbol{s}') = 0\,. \tag{18}$$

Both perturbations, $\Delta \boldsymbol{s}$ and $\Delta \boldsymbol{s}'$, give the same perturbation of the invariant I. We say two perturbations are *geometrically equivalent* if they give the same perturbation of any invariant. We now state the following theorem governing this equivalence:

Geometric Equivalence

Theorem 3.1. *Two perturbations, $\Delta \boldsymbol{s} \in T[\mathcal{C}]$ and $\Delta \boldsymbol{s}' \in T[\mathcal{C}']$, where $\boldsymbol{s}' = a\boldsymbol{s}$, are geometrically equivalent if and only if:*

$$\boldsymbol{u}^\top (\Delta \boldsymbol{s} - \frac{1}{a}\Delta \boldsymbol{s}') = 0 \quad \forall \boldsymbol{u} \in T[\boldsymbol{G_s}]^\perp\,. \tag{19}$$

Proof. From equation (14) we know that *if* equation (19) is true, then equation (18) follows for all invariants, and so the perturbations are geometrically equivalent. Furthermore, we can construct a set of invariants whose gradients span $T[\boldsymbol{G_s}]^\perp$, for example: $I_i(\boldsymbol{s}) = s_i/s_1$, for $i = 2, \ldots, 3N$, where s_i is the ith element of $\boldsymbol{s}$. Thus if equation (18) is true for all invariants, equation must (19) follow. □

Geometric equivalence is illustrated pictorially in Figure 4. Here we see that to be geometrically equivalent means that the difference between the perturbations, $\Delta \boldsymbol{s} - \Delta \boldsymbol{s}'/a$, lies in the tangent space to the gauge orbit, $T[\boldsymbol{G_s}]$.

3.5 Gauge Fixing

Our initial covariance, $\boldsymbol{V_s}$, describes perturbations, $\Delta \boldsymbol{s}$, tangent to the unknown gauge, $\mathcal{C}$. We now want to find geometrically equivalent perturbations, $\Delta \boldsymbol{s}'$, tangent to our new gauge, $\mathcal{C}'$.

One solution to the geometric equivalence relationship, in equation (19), is a simple rescaling of the perturbation $\Delta \boldsymbol{s}' = a\Delta \boldsymbol{s}$. But this is not the only solution, and in general the resulting perturbation will not be tangent to the new gauge, $\mathcal{C}'$. The full solution is:

$$\Delta \boldsymbol{s}' = a\Delta \boldsymbol{s} + \boldsymbol{b} \tag{20}$$

where $\boldsymbol{b}$ is any vector in the tangent to the gauge orbit: $T[\boldsymbol{G_s}]$. The gauge orbit is defined as $\boldsymbol{s}' = a\boldsymbol{s}$, for variations in a. Its tangent direction at $\boldsymbol{s}'$ is given by:

$$\frac{\partial \boldsymbol{s}'}{\partial a} = \boldsymbol{s}\,. \tag{21}$$

Thus we can write $\boldsymbol{b} = x\boldsymbol{s}$ with some unknown coefficient x, which we will now solve for.

Let $\boldsymbol{v}$ be a vector orthogonal to the gauge tangent space, $\boldsymbol{v} \in T[\mathcal{C}']^{\perp}$. One such vector is given by the gradient of the constraint:

$$\boldsymbol{v} \;\; = \;\; \nabla_{\boldsymbol{s}'} c(\boldsymbol{s}') \tag{22}$$

where the constraint $c(\boldsymbol{s}') = 0$ is defined by equation (12). We know $\boldsymbol{v}$ is orthogonal to $\Delta \boldsymbol{s}'$, and so can write:

$$\boldsymbol{v}^{\top} \Delta \boldsymbol{s}' = 0\,. \tag{23}$$

Applying this to equation (20), letting $\boldsymbol{b} = x\boldsymbol{s}$, and solving for x, we obtain: $x = -(\boldsymbol{v}^{\top}\boldsymbol{s})^{-1}\boldsymbol{v}^{\top} a\Delta \boldsymbol{s}$. Then substituting this into equation (20) we get:

$$\begin{aligned} \Delta \boldsymbol{s}' &= a\Delta \boldsymbol{s} - \boldsymbol{s}(\boldsymbol{v}^{\top}\boldsymbol{s})^{-1}\boldsymbol{v}^{\top} a\Delta \boldsymbol{s} \\ &= a\boldsymbol{Q}\Delta \boldsymbol{s}\,, \end{aligned} \tag{24}$$

where

$$\boldsymbol{Q} = \boldsymbol{I} - \frac{\boldsymbol{s}\boldsymbol{v}^{\top}}{\boldsymbol{v}^{\top}\boldsymbol{s}}\,. \tag{25}$$

The matrix $a\boldsymbol{Q}$ is an oblique projection operator as illustrated in Figure 5. It takes any perturbation $\Delta \boldsymbol{s}$ to a geometrically equivalent perturbation $\Delta \boldsymbol{s}'$ in a new gauge $\mathcal{C}'$.

We can now state our result for the transformation of a covariance matrix when we measure one of the lengths on the object:

$$\begin{aligned} \boldsymbol{V_{s'}} &= E[\Delta \boldsymbol{s}'\, \Delta \boldsymbol{s}'^{\top}] \\ &= a^2 \boldsymbol{Q}\boldsymbol{V_s}\boldsymbol{Q}^{\top}\,. \end{aligned} \tag{26}$$

The resulting matrix, $\boldsymbol{V_{s'}}$ is geometrically equivalent to the original covariance, $\boldsymbol{V_s}$, and also it satisfies our measurement constraint to first order. It thus gives the true uncertainty of the shape after we make measurement d'.

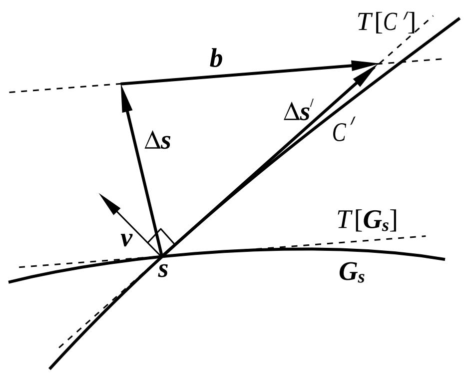

Figure 5: An oblique projection onto the tangent space of the gauge manifold $\mathcal{C}'$: $\Delta \boldsymbol{s}' = a\Delta \boldsymbol{s} + \boldsymbol{b}$ $= a\boldsymbol{Q}\Delta \boldsymbol{s}$, here shown when $a = 1$.

3.6 Measurement Uncertainty

We can now ask what happens when there is uncertainty in the measurement of length d' on the object. We expand our basic scale equation (4) including a perturbation term, Δa, on scale itself:

$$\Delta \boldsymbol{s}' = a\boldsymbol{Q}\boldsymbol{s} + \boldsymbol{s}\Delta a\,. \tag{27}$$

The second term is simplified as follows. The scale is $a = d'/d$ where $d = ||\boldsymbol{s}_i - \boldsymbol{s}_j||$, and so $\Delta a = \Delta d'/d$. If d' is measured with standard deviation σ_m, then substituting for Δa into equation (27), and finding the covariance we obtain:

$$\boldsymbol{V}_{\boldsymbol{s}'} = a^2 \boldsymbol{Q}\boldsymbol{V}_{\boldsymbol{s}}\boldsymbol{Q}^\top + \sigma_m^2 \frac{\boldsymbol{s}\boldsymbol{s}^\top}{d^2}\,. \tag{28}$$

We see that the measurement component to the variance is inversely weighted by the length on the object that is measured. Thus measuring longer lengths reduces this component of the error.

3.7 Choosing a Good Gauge

The choice of gauge affects the final accuracy of our results. It is natural to want to choose a gauge that maximizes the accuracy of the model. In this section we consider two measures of final accuracy and how gauges should be chosen to optimize these.

Optimizing Overall Accuracy

One measure for accuracy is the trace of $\boldsymbol{V}_{\boldsymbol{s}'}$, which is the sum of the individual 3D feature-point variances. Reducing this corresponds to improving the overall accuracy, but ignoring cross-correlation effects. If $\boldsymbol{V}_{\boldsymbol{s}}$ includes off-diagonal elements, the analysis becomes complicated, and so we approximate it with uniform and uncorrelated noise: $\boldsymbol{V}_{\boldsymbol{s}} = \sigma_0^2 \boldsymbol{I}$. Now we ask: If we can measure the distance, $d' = ||\boldsymbol{s}'_i - \boldsymbol{s}'_j||$, between any two points i and j, which two points should be used to minimize the trace: $\mathrm{Tr}[\boldsymbol{V}_{\boldsymbol{s}'}]$?

The vector $\boldsymbol{v}$, from equation (22), is zero except for the ith and jth elements which are: $(\boldsymbol{s}_i - \boldsymbol{s}_j)/d$ and $(\boldsymbol{s}_j - \boldsymbol{s}_i)/d$ respectively. From this we get: $\boldsymbol{v}^\top \boldsymbol{s} = d$ and $\boldsymbol{v}^\top \boldsymbol{v} = 2$, and $\boldsymbol{s}\boldsymbol{v}^\top$ is a square matrix whose only two, non-zero, diagonal blocks are: $\boldsymbol{s}_i(\boldsymbol{s}_i - \boldsymbol{s}_j)^\top$ and $\boldsymbol{s}_j(\boldsymbol{s}_j - \boldsymbol{s}_i)^\top$. Using these, and substituting for $\boldsymbol{Q}$ and $\boldsymbol{V}_{\boldsymbol{s}}$ in equation (28), we obtain:

$$\mathrm{Tr}[\boldsymbol{V}_{\boldsymbol{s}'}] = (3N-2)a^2\sigma_0^2 + (2a^2\sigma_0^2 + \sigma_m^2)\frac{||\boldsymbol{s}||^2}{d^2} \, . \tag{29}$$

We see that given constant measurement error, σ_m, the longer the length on the model we choose, d, the smaller the total uncertainty.

Optimizing One-Length Accuracy

For real models, there is typically strong correlation between features and so the approximation that $\boldsymbol{V}_{\boldsymbol{s}}$ is uncorrelated may be poor. Also, we may be interested in the accuracy of only part of the model, and in some cases just one length. Let us say that our goal is to estimate a certain length, $e' = ||\boldsymbol{s}'_k - \boldsymbol{s}'_l||$, with the greatest accuracy, and that we can set the scale factor by measuring another length, d'. What qualities should d' have to minimize the variance of e'?

Let $\boldsymbol{d} = (e \;\; d)^\top$ and $\boldsymbol{d}' = a\boldsymbol{d}$, and so we can write:

$$\begin{aligned} \boldsymbol{V}_{\boldsymbol{d}} &= \nabla_{\boldsymbol{s}}\boldsymbol{d}^\top \boldsymbol{V}_{\boldsymbol{s}} \nabla_{\boldsymbol{s}}\boldsymbol{d} & (30) \\ &\equiv \begin{pmatrix} \sigma_e^2 & \sigma_{ed} \\ \sigma_{de} & \sigma_d^2 \end{pmatrix} \, . & (31) \end{aligned}$$

If we measure d' with variance σ_m^2, we obtain from equation (28):

$$\boldsymbol{V}_{\boldsymbol{d}'} = a^2 \boldsymbol{Q}\boldsymbol{V}_{\boldsymbol{d}}\boldsymbol{Q}^\top + \sigma_m^2 \frac{\boldsymbol{d}\boldsymbol{d}^\top}{d^2} \tag{32}$$

where $\partial\boldsymbol{d}'/\partial a = \boldsymbol{d}$, $\boldsymbol{v} = (0 \;\; 1)^\top$, and so

$$\boldsymbol{Q} = \begin{pmatrix} 1 & -e/d \\ 0 & 0 \end{pmatrix} \, . \tag{33}$$

We want to minimize the variance of e', which can be obtained from equation (32) as:

$$\sigma_e^{2\prime} = a^2(\sigma_e^2 - 2\frac{e}{d}\sigma_{ed} + \left(\frac{e}{d}\right)^2 \sigma_d^2) + \left(\frac{e}{d}\right)^2 \sigma_m^2 \, . \tag{34}$$

This is a quadratic in the ratio e/d, and its minimum has two cases. The first is when the noise is uncorrelated or anti-correlated, ($\sigma_{ed} \leq 0$). The variance, $\sigma_e^{2\prime}$, is reduced when the ratio, e/d, is reduced. Thus given constant variances, the longer the length we measure, d', the more accurate our estimate for e' is. The second case is when the noise is positively correlated, $\sigma_{ed} > 0$. The ratio that minimizes $\sigma_e^{2\prime}$ is then:

$$\frac{e}{d} = \frac{a^2\sigma_{ed}}{a^2\sigma_d^2 + \sigma_m^2} \, . \tag{35}$$

If the noise is perfectly correlated, namely $\sigma_{ed} = \sigma_e\sigma_d$, and $\sigma_m^2 = 0$, and the ratio, e/d, is given by equation (35), then the length e' will be perfectly estimated with zero covariance.

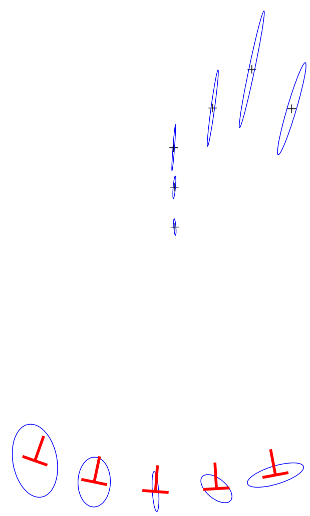

Figure 6: From five images, an object shape and camera motion are estimated, along with a covariance matrix. The "⊥" symbol shows the camera positions and orientations. The ellipses illustrate the covariance of each point. The correlation between features is not shown here, but it is significant.

4 Results

First we illustrate with a synthetic example the significance of gauge fixing on shape covariances. Figure 6 illustrates the shape and motion estimated using a Structure from Motion algorithm [3, 4, 10, 11], along with the covariance of all the features and camera positions. This shape is known up to a scale factor,[3] and so to obtain an exact model estimate we must rescale it. If we are given a length on the object we can rescale it, and also transform the covariance accordingly. Figure 7 shows the effect on the shape covariance when different lengths are used to rescale the shape.

We now illustrate on a real 3D object how model accuracy changes when different object measurements are used to fix the scale. We start with a set of hand registered features in an image sequence as shown in Figure 8. A Structure from Motion algorithm was used to obtain the 3D shape, $\boldsymbol{s}$, and its covariance, $\boldsymbol{V_s}$, up to an unknown scale factor as shown in Figure 9.

We want to know what the diagonal length of the television screen is and how accurately can we know it. Since we can recover 3D shape only up to a scale factor, the shape alone will not determine this length. We first have to find the overall scale of the recovered shape. Now it may be that we know the length of another object in the scene, in which case this could be used to obtain the scale. For the purpose of this experiment, we will look at a number of different objects in the scene whose lengths we know, and find out which gives us the best estimate of the TV diagonal length.

[3] Here we assume overall translation and rotation are known.

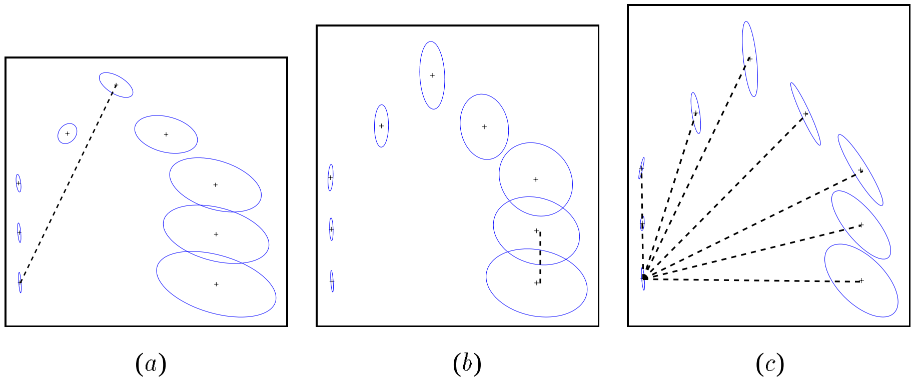

Figure 7: In (a) and (b) the object is rescaled by measuring the distance between the points joined by the dashed line. In (c) the scale is determined using the average distance of all the features to the bottom–left feature. In each case the covariance is appropriately transformed as in equation (26). These plots illustrate the effects of gauge fixing on covariances.

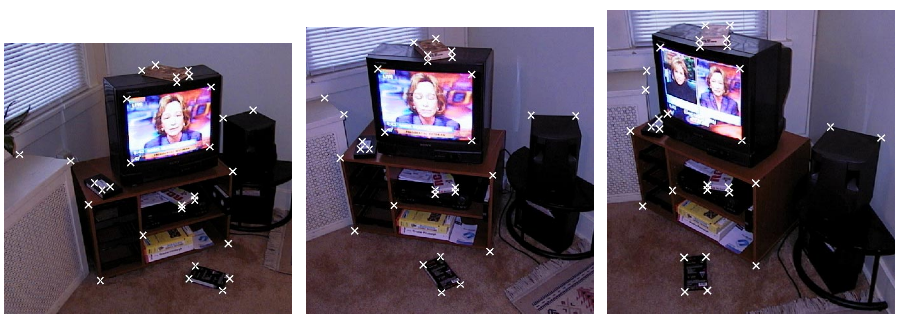

Figure 8: Three images from a seven-image sequence with hand-registered features.

Consider, then, the 13 lines, measured in 3D, and shown in Figure 10. We made 13 separate estimates of the TV diagonal length, in each case using one of these lengths to fix our gauge, $\mathcal{C}'$, and obtain the scale. Figure 11 shows these estimates, along with their predicted standard deviations, σ'_e, obtained using equation (28). We assumed no error in the measurement of the 3D line lengths, $\sigma_m = 0$. We note that the actual error corresponds well with the uncertainty given by the predicted error.

We notice, from Figure 11, that there is a large variation in uncertainty of the TV diagonal, depending on which line is used to fix the scale. In general, using the longer lines lead to better estimates, and the shortest lines, 5 and 11, gave the greatest uncertainty. The puzzle is line 8, which gives a more accurate estimate than line 3, which is longer. This is explained by Figures 12 and 13 which show that line 8 is more strongly correlated with the TV diagonal length, and that the ratio with the TV diagonal length, given by equation (35), is closer to its optimal value for line 8 than for line 3.

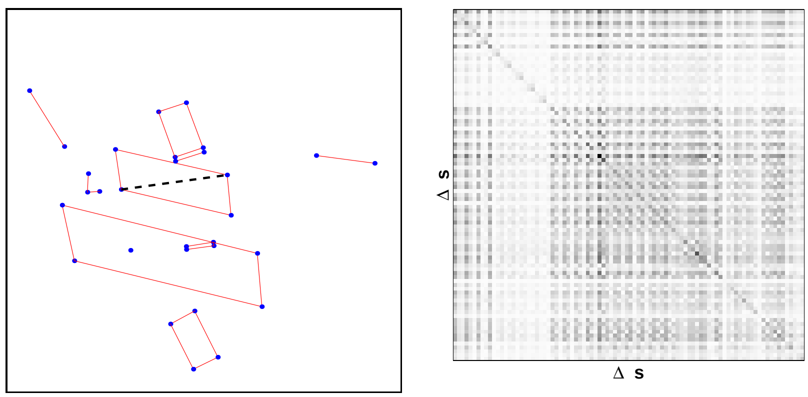

(a) (b)

Figure 9: (a) The 3D reconstruction of point features, s, from the image sequence in Figure 8 using a Structure from Motion algorithm. Lines are drawn for clarity only. (b) The covariance, $\boldsymbol{V_s}$, of this shape calculated in an arbitrary gauge. Image features were assumed to have uniform, identical noise, and whose magnitude we solved for. The plot shows strong correlation between features.

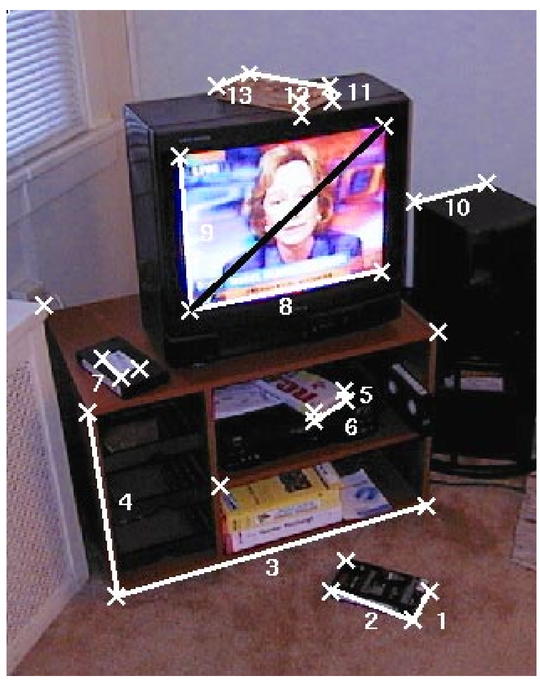

Figure 10: We want to estimate the diagonal length of the TV shown by the black line. To do this we need to know the length of one of the white lines.

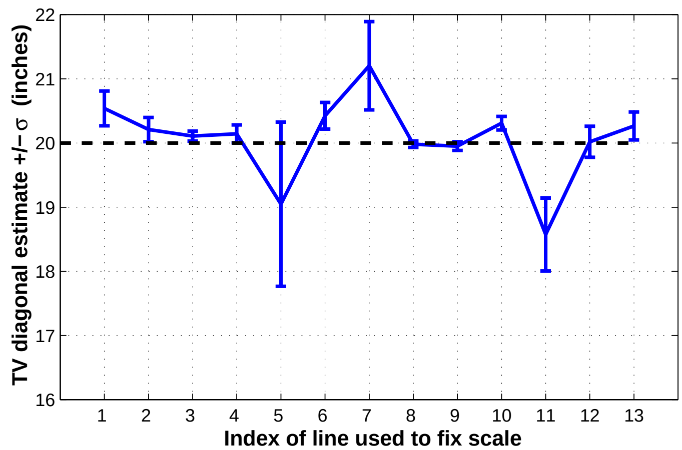

Figure 11: The true length of the TV diagonal is 20 inches. This plot shows the estimate we obtained of this length using each of the different lines to fix the scale. The number on the abscissa corresponds to the line in Figure 10. The uncertainty, σ'_e, given by the error bars, varies greatly depending on which line was used; the largest uncertainty with standard deviation of 1.3 inches is obtained using line 5, and the smallest with standard deviation of 0.05 inches is obtained using line 8. Line 8 provides a better constraint than line 3, despite line 3 being longer, for two reasons: it is more strongly correlated with the TV diagonal as shown in Figure 12, and the ratio of its length with the TV diagonal length is closer to the optimal value in equation (35), as shown by Figure 13.

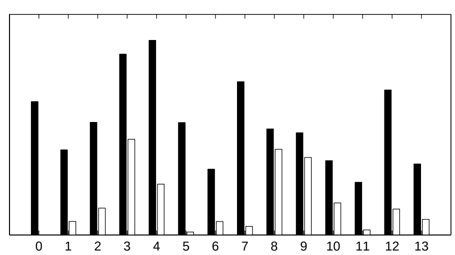

Figure 12: The black bars show the standard deviation, σ_i, for each line length in Figure 10, calculated from the unscaled shape covariance, $\boldsymbol{V_s}$. Bar 0 corresponds to the TV diagonal. The scale is arbitrary and so left unmarked. The white bars indicate the proportion of cross-correlation, and are calculated as: $r_{0i}\sigma_i$, where $r_{0i} = \sigma_{0i}/\sigma_0\sigma_i$ is the correlation coefficient ($-1 \le r_{0i} \le 1$). When perfectly correlated with the TV diagonal, the white bar will equal the black bar. Hence we see that lines 8 and 9 are most strongly correlated with the TV diagonal.

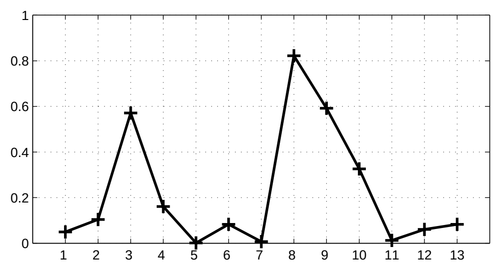

Figure 13: Given the corresponding variance and cross-correlation plotted in Figure 12, we can use equation (35) to calculate what the optimal length would be for each line. By "optimal" we mean it would minimize the variance of the TV diagonal estimate. Here we plot the ratio of the actual line length to the optimal line length for each line. Line 8 is closest to its optimal length. This is a better measure to use than length in selecting a line to fix the scale. While all of these lines are shorter than their optimal lengths, we found in other experiments that lines can also be longer than their optimal lengths.

5 Conclusion

We have derived and demonstrated an unexpected consequence of fixing the scale of an object known only up to a scale factor. The covariance of the resulting shape is dramatically affected by how the scale is determined. We assume the scale is determined by the distance between two points on the object. Simply rescaling the covariance matrix with the square of the scale factor leaves the indeterminacy in the matrix and does not account for the measurement. Instead, the covariance must be transformed so that the model uncertainty between the measured points is correctly distributed over the rest of the points. We used gauge theory to derive the covariance transformation, when the scale is fixed by making a measurement, and showed how choosing different gauges will affect the final accuracy.

A qualitative analysis of our results leads to three conclusions: (a) When there is little or no correlation between feature points and lengths, accuracy is best when the gauge is fixed by measuring the largest distance on the object. (b) The effect of measurement error on the 3D lengths is reduced for longer lengths. (c) When there is significant positive correlation between lengths in the model, then measuring a longer line is no longer better. Instead, the stronger the correlation, and the closer the ratio, of measured to estimated length, is to the ratio in equation (35), the better the final accuracy.

In some cases, our shape covariance may have further gauge freedoms such as an over-all translation and rotation. The same gauge principles still apply and our work can be easily extended. But if we consider only lengths, rather than absolute 3D positions, we can use the fact that lengths are invariant to translation and rotation, and so we can directly apply the theory developed here.

References

[1] K. B. Atkinson, *Close Range Photogrammetry and Machine Vision*, Whittles Publishing, Caithness, Scotland, (1996).

[2] R. M. Haralick, Propagating covariance in computer vision, *Int. J. of Patt. Recog. and AI.*, **10**(5), pp. 561–572, (1996).

[3] R. Hartley, Euclidean reconstruction from uncalibrated views, in *Proc. DARPA–ESPRIT Workshop App. Invariants in Comp. Vision*, pp. 187–202, Azores, Portugal, (October 1993).

[4] A. Heyden & K. Astrom, Euclidean reconstruction from image sequences with varying and unknown focal length and principal point, in *Proc. Comp. Vision Patt. Recog.*, pp. 438–443, Peurto Rico, (June 1997).

[5] J. Inigo Thomas, A. Hanson & J. Oliensis, Refining 3D reconstructions: A theoretical and experimental study of the effect of cross-correlations, *CVGIP*, **60**(3), pp. 359–370, (November 1994).

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

[7] K. Kanatani & D. D. Morris, Gauges and gauge transformations in 3-D reconstruction from a sequence of images, in *Proc. Fourth Asian Conf. Comp. Vision*, volume 2, pp. 1046–1051, (2000).

[8] K. Kanatani & D. D. Morris, Gauge and gauge transformations for uncertainty description of geomtric structure with indeterminacy, *IEEE Transactions on Information Theory*, (2001), to appear.

[9] P. F. McLauchlan, Gauge independence in optimization algorithms for 3D vision, in *Vision Algorithms: Theory and Practice*, pp. 183–199, Springer, Berlin, (2000).

[10] D. D. Morris, K. Kanatani & T. Kanade, Uncertainty modeling for optimal structure from motion, in *Vision Algorithms: Theory and Practice*, pp. 200–217, Springer, Berlin, (2000).

[11] M. Pollefeys, R. Koch & L. Van Gool, Self-calibration and metric reconstruction in spite of varying and unknown internal camera parameters, *Int. J. Comp. Vision*, **32**(1), pp. 7–26, (1999).

[12] R. Szeliski & S. B. Kang, Shape ambiguities in structure from motion, *IEEE Trans. Patt. Anal. Mach. Intelligence*, **19**(5), pp. 506–512, (May 1997).

[13] B. Triggs, P. F. McLauchlan, R. I. Hartley & A. W. Fitzibbon, Bundle adjustment – a modern synthesis, in *Vision Algorithms: Theory and Practice*, pp. 298–375, Springer, Berlin, (2000).

[14] J. Weng, N. Ahuja & T. Huang, Optimal motion and structure estimation, *IEEE Trans. Patt. Anal. Mach. Intelligence*, **15**(9), pp. 864–884, (1993).