Covariance Propagation for Guided Matching

Benjamin Ochoa[1,2] and Serge Belongie[2]

[1] BAE Systems, 10920 Technology Place, San Diego, CA 92127, USA
[2] University of California, San Diego, La Jolla, CA 92093, USA
{bochoa, sjb}@ucsd.edu

Abstract. We present a general approach and analytical method for determining a search region for use in guided matching under projective mappings. Our approach is based on the propagation of covariance through a first-order approximation of the error model to define the boundary of the search region for a specified probability and we provide an analytical expression for the Jacobian matrix used in the covariance propagation calculation. The resulting closed-form expression is easy to implement and generalizes to n dimensions. We apply our method to point-to-point mapping under a planar homography, point-to-line mapping under a fundamental matrix, and mosaic construction from video in the case of the video looping back on itself.

1 Introduction

In this work we address the problem of determining a search region used for establishing feature correspondences over multiple views given an estimate of the projective mapping that relates these views. Corresponding features are defined as the set of features that are the images of the same pre-image feature. Consider two different cameras imaging a scene. A 3D point in the scene is imaged as a 2D point in the image plane of each of the cameras. The image point in one of the cameras corresponds to the image point in the other camera and both image points correspond to the pre-image 3D scene point. Several projective models (e.g., the fundamental matrix) have been developed in computer vision that allow features to be mapped between views without explicit knowledge of the 3D scene structure. However, in the presence of noise or uncertainty, the mapped feature may not be coincident with the true corresponding feature and a search must be performed to locate the true correspondence. In the absence of uncertainty information, the true correspondence may be located anywhere in the image, assuming the pre-image feature was imaged by the camera.

Guided matching methods are often used to reduce the size of the search region from the entire image to a region expected to contain the corresponding feature. One simple guided matching method is to specify a search region bounded at a fixed distance from the mapped feature. Although easy to implement, this simple method generally yields either an undersized region, which may not include the true correspondence, or an oversized region, which may include features that are similar to the true correspondence, increasing the potential of a false

match. Our approach uses covariance propagation to define the search region for projective mappings. The search region is bounded by a specified probability that the region contains the feature. This approach can be used, for example, to propagate the spatial covariance of a homogeneous point through a planar projective (homography) or epipolar (fundamental matrix) transformation, where the transformation may additionally have an associated covariance. We present an expression that allows for the determination of the covariance of the mapping of the point as a function of the above covariances.

Several approaches to propagating uncertainty in structure from motion have been proposed appealing to different statistical techniques. One approach is to use Monte Carlo methods, which is highly general but computationally expensive. Alternatively, analytical frameworks have been developed by Kanatani [1], Förstner [2], and others (e.g., [3], [4]). However, when these mappings are well approximated locally by an affine transform, a first-order model has proved to be sufficient [5]. This linearized approximation of the error model is commonly used in computer vision and is the approach adopted in this paper. Central to this approach to uncertainty propagation is the Jacobian matrix of the mapping. One method for estimating the Jacobian is to perform numerical differentiation using forward differencing, which in practice may yield a Jacobian matrix of unexpected rank due to numerical inaccuracies. Alternatively, one can derive specialized analytical expressions on a per-mapping basis, e.g., planar homography [6] and fundamental matrix [7], [8]. We derive a novel analytical expression for the Jacobian applicable to all projective mappings, obviating the need for specialized expressions. The resulting closed-form expression is general and easy to implement. The same expression can be generalized to n dimensions and can also be applied to other projective mappings such as composition of homographies.

This paper is structured as follows. We describe covariance propagation for projective mappings under the linearized error model in Section 2. In Section 3, we discuss guided matching in terms of uncertainty bounds with particular detail given to 2D points and lines. Experimental results for point-to-point mapping under a planar homography and point-to-line mapping under a fundamental matrix are given in Section 4. In Section 5, we apply our approach to mosaic construction from video in the case when the image sequence loops back on itself. Finally, we give our conclusions in Section 6.

Notation In this paper, if a capital letter is used to denote a matrix, then the vector denoted by the corresponding lower case letter is composed of the entries of the matrix by

$$\mathtt{A} \in \mathbb{R}^{m \times n} \Leftrightarrow \mathtt{A} = \begin{bmatrix} \mathbf{a}_1^\top \\ \mathbf{a}_2^\top \\ \vdots \\ \mathbf{a}_m^\top \end{bmatrix}, \mathbf{a} = \begin{pmatrix} \mathbf{a}_1 \\ \mathbf{a}_2 \\ \vdots \\ \mathbf{a}_m \end{pmatrix} \in \mathbb{R}^{mn}$$

where $\mathbf{a}_i^\top \in \mathbb{R}^n$ is the ith row of $\mathtt{A}$ (i.e., $\mathbf{a} = \mathrm{vec}(\mathtt{A}^\top)$).

2 Nonlinear Propagation of Covariance

Let $\mathbf{x} \in \mathbb{R}^n$ be a random vector with mean $\boldsymbol{\mu}_\mathbf{x}$ and covariance matrix $\Sigma_\mathbf{x}$, and let $f : \mathbb{R}^n \to \mathbb{R}^m$ be a nonlinear function. Up to first-order approximation, $\mathbf{y} = f(\mathbf{x}) \approx f(\boldsymbol{\mu}_\mathbf{x}) + \mathtt{J}(\mathbf{x} - \boldsymbol{\mu}_\mathbf{x})$, where $\mathtt{J} \in \mathbb{R}^{m \times n}$ is the Jacobian matrix $\partial f / \partial \mathbf{x}$ evaluated at $\boldsymbol{\mu}_\mathbf{x}$. If f is approximately affine in the region about the mean of the distribution, then this approximation is reasonable and the random vector $\mathbf{y} \in \mathbb{R}^m$ has mean $\boldsymbol{\mu}_\mathbf{y} \approx f(\boldsymbol{\mu}_\mathbf{x})$ and covariance $\Sigma_\mathbf{y} \approx \mathtt{J}\Sigma_\mathbf{x}\mathtt{J}^\top$.

If $\mathbf{x}$ is composed of two random vectors $\mathbf{a}$ and $\mathbf{b}$ such that $\mathbf{x} = (\mathbf{a}^\top, \mathbf{b}^\top)^\top$, then

$$\Sigma_\mathbf{y} \approx \begin{bmatrix} \mathtt{J}_\mathbf{a} & \mathtt{J}_\mathbf{b} \end{bmatrix} \begin{bmatrix} \Sigma_\mathbf{a} & \Sigma_\mathbf{ab} \\ \Sigma_\mathbf{ba} & \Sigma_\mathbf{b} \end{bmatrix} \begin{bmatrix} \mathtt{J}_\mathbf{a}^\top \\ \mathtt{J}_\mathbf{b}^\top \end{bmatrix} \tag{1}$$

where $\mathtt{J}_\mathbf{a} = \partial \mathbf{y} / \partial \mathbf{a}$ and $\mathtt{J}_\mathbf{b} = \partial \mathbf{y} / \partial \mathbf{b}$.

2.1 Projective Mappings

Consider a homogeneous 2D point $\mathbf{x}$ represented by the vector $(x, y, w)^\top \in \mathbb{R}^3$. The vector $s(x, y, w)^\top$, where s is any nonzero scalar, represents the same 2D point as $(x, y, w)^\top$. It follows that $(x, y, w)^\top \sim s(x, y, w)^\top$, where $\sim$ denotes equality up to a nonzero scale factor.

Due to the use of homogeneous representations, several projective mappings are only determined up to scale [9]. Examples include point imaging $\mathbf{x} \sim \mathtt{P}\mathbf{X}$, where $\mathtt{P} \in \mathbb{R}^{3 \times 4}$ is the projective camera that maps the homogeneous 3D point $\mathbf{X}$ to a homogeneous 2D point $\mathbf{x}$; projective transformation accumulation $\mathtt{H}_{a,c} \sim \mathtt{H}_{b,c}\mathtt{H}_{a,b}$, where $\mathtt{H}_{a,c}, \mathtt{H}_{b,c}, \mathtt{H}_{a,b} \in \mathbb{R}^{(n+1)\times(n+1)}$ are n-dimensional projective transformations and $\mathtt{H}_{a,c}$ represents the transformation from a to c; and point-to-line mapping $\boldsymbol{\ell}' \sim \mathtt{F}\mathbf{x}$, where $\mathtt{F} \in \mathbb{R}^{3 \times 3}$ is the fundamental matrix, which maps a homogeneous 2D point $\mathbf{x}$ in one image to a homogeneous 2D line $\boldsymbol{\ell}'$ in another image. All of these mappings can be generalized as $\mathtt{C} \sim \mathtt{AB}$, where $\mathtt{C} \in \mathbb{R}^{m \times n}$, $\mathtt{A} \in \mathbb{R}^{m \times p}$, and $\mathtt{B} \in \mathbb{R}^{p \times n}$. However, because $\mathtt{C}$ is only determined up to scale, the entries of $\mathtt{C}$ may vary without bound. This poses an issue for covariance propagation and uncertainty analysis. It is usual to impose the constraint that $\|\mathtt{C}\| = 1$, where $\|\cdot\|$ denotes the Frobenius norm. Under this constraint, the generalized mapping is $\mathtt{C} = (\mathtt{AB})/\|\mathtt{AB}\|$ and the variance of the entries of $\mathtt{C}$ are constrained accordingly.

2.2 Jacobian Matrices

For the expression

$$\mathtt{C} = \frac{\mathtt{AB}}{\|\mathtt{AB}\|}$$

where $\mathtt{A} \in \mathbb{R}^{m \times p}$, $\mathtt{B} \in \mathbb{R}^{p \times n}$, and $\mathtt{C} \in \mathbb{R}^{m \times n}$, the matrices $\partial\mathbf{c}/\partial\mathbf{a}$ and $\partial\mathbf{c}/\partial\mathbf{b}$ may be computed analytically as

$$\mathtt{J}_{\mathbf{a}} = \frac{\partial \mathbf{c}}{\partial \mathbf{a}} = \frac{1}{\|\mathtt{AB}\|} \left[\left(\mathtt{I}_{m \times m} \otimes \mathtt{B}^{\top} \right) - \mathbf{c} \operatorname{vec}(\mathtt{BC}^{\top})^{\top} \right] \tag{2}$$

$$\mathtt{J}_{\mathbf{b}} = \frac{\partial \mathbf{c}}{\partial \mathbf{b}} = \frac{1}{\|\mathtt{AB}\|} \left[\left(\mathtt{A} \otimes \mathtt{I}_{n \times n} \right) - \mathbf{c} \operatorname{vec}(\mathtt{C}^{\top}\mathtt{A})^{\top} \right] \tag{3}$$

where $\otimes$ denotes the Kronecker product. The derivations of $\partial\mathbf{c}/\partial\mathbf{a}$ and $\partial\mathbf{c}/\partial\mathbf{b}$ are given in the appendix. Applying (1), these are the Jacobian matrices used to approximate the covariance $\Sigma_{\mathbf{c}}$.

3 Guided Matching

3.1 Search Region

If a Gaussian random vector $\mathbf{x}$ has mean $\boldsymbol{\mu}_{\mathbf{x}}$ and covariance $\Sigma_{\mathbf{x}}$, then the squared Mahalanobis distance between $\mathbf{x}$ and $\boldsymbol{\mu}_{\mathbf{x}}$ satisfies a χ^2_r distribution where r is the degrees of freedom of $\mathbf{x}$. It follows that a percentage α of all instances of $\mathbf{x}$ will satisfy the condition

$$(\mathbf{x} - \boldsymbol{\mu}_{\mathbf{x}})^{\top} \Sigma_{\mathbf{x}}^{+} (\mathbf{x} - \boldsymbol{\mu}_{\mathbf{x}}) \leq k^2 \tag{4}$$

where k^2 is the inverse of the chi-square cumulative distribution function with r degrees of freedom and probability α, and $\Sigma_{\mathbf{x}}^{+}$ is the pseudo-inverse of the covariance matrix $\Sigma_{\mathbf{x}}$ with rank r.

3.2 2D Points and Lines

A derivation of the uncertainty bounds for homogeneous 2D lines may be found in [7], [8], and [5]. In this section we derive uncertainty bounds for 2D points and state the bounds for 2D lines using the duality principle.

Let $\boldsymbol{\mu}_{\mathbf{x}}$ and $\Sigma_{\mathbf{x}}$ be the mean and covariance, respectively, of a homogeneous 2D point $\mathbf{x}$. The covariance matrix has rank 2, thereby constraining the 3-vector to 2 degrees of freedom. For a given k, we can determine if an instance of $\mathbf{x}$ is within the some bounds directly from (4). However, it is often the case that we want to determine if an *arbitrary* homogeneous point, not necessarily drawn from the distribution of $\mathbf{x}$, is within these bounds. This cannot be accomplished using (4) directly and instead must be determined geometrically as follows.

For a given k, the set of points with equal likelihood in the distribution of $\mathbf{x}$ is given by

$$(\mathbf{x} - \boldsymbol{\mu}_{\mathbf{x}})^{\top} \Sigma_{\mathbf{x}}^{+} (\mathbf{x} - \boldsymbol{\mu}_{\mathbf{x}}) = k^2 \tag{5}$$

For further analysis, we apply a change of coordinates such that

$$\Sigma'_{\mathbf{x}} = \mathtt{U}\Sigma_{\mathbf{x}}\mathtt{U}^{\top} = \begin{bmatrix} \tilde{\Sigma}'_{\mathbf{x}} & \mathbf{0} \\ \mathbf{0}^{\top} & 0 \end{bmatrix}$$

where $\tilde{\Sigma}'_{\mathbf{x}} \in \mathbb{R}^{2\times 2}$ is a nonsingular diagonal matrix, and $\boldsymbol{\mu}'_{\mathbf{x}} = \mathtt{U}\boldsymbol{\mu}_{\mathbf{x}} = (\tilde{\boldsymbol{\mu}}'^{\top}_{\mathbf{x}}, 1)^{\top}$ and $\mathbf{x}' = \mathtt{U}\mathbf{x} = (\tilde{\mathbf{x}}'^{\top}, 1)^{\top}$. The similarity $\mathtt{U} = s\mathtt{V}^{\top}$, where the orthogonal matrix $\mathtt{V}^{\top}$ is obtained from the eigen decomposition $\Sigma_{\mathbf{x}} = \mathtt{V}\mathtt{D}\mathtt{V}^{\top}$ and s is chosen such that the last entry in the 3-vector $s\mathtt{V}^{\top}\boldsymbol{\mu}_{\mathbf{x}}$ is equal to 1. The matrix $\mathtt{D} = \mathrm{diag}(\lambda_1, \lambda_2, 0)$ contains the eigenvalues of $\Sigma_{\mathbf{x}}$. Using this, we can show

$$\begin{aligned}(\mathbf{x} - \boldsymbol{\mu}_{\mathbf{x}})^{\top}\Sigma^{+}_{\mathbf{x}}(\mathbf{x} - \boldsymbol{\mu}_{\mathbf{x}}) &= k^2 \\ (\mathbf{x}' - \boldsymbol{\mu}'_{\mathbf{x}})^{\top}\Sigma'^{+}_{\mathbf{x}}(\mathbf{x}' - \boldsymbol{\mu}'_{\mathbf{x}}) &= k^2 \\ (\tilde{\mathbf{x}}' - \tilde{\boldsymbol{\mu}}'_{\mathbf{x}})^{\top}\tilde{\Sigma}'^{-1}_{\mathbf{x}}(\tilde{\mathbf{x}}' - \tilde{\boldsymbol{\mu}}'_{\mathbf{x}}) &= k^2\end{aligned} \tag{6}$$

which can be written more fully as

$$\tilde{\mathbf{x}}'^{\top}\tilde{\Sigma}'^{-1}_{\mathbf{x}}\tilde{\mathbf{x}}' - \tilde{\boldsymbol{\mu}}'^{\top}_{\mathbf{x}}\tilde{\Sigma}'^{-1}_{\mathbf{x}}\tilde{\mathbf{x}}' - \tilde{\mathbf{x}}'^{\top}\tilde{\Sigma}'^{-1}_{\mathbf{x}}\tilde{\boldsymbol{\mu}}'_{\mathbf{x}} + \tilde{\boldsymbol{\mu}}'^{\top}_{\mathbf{x}}\tilde{\Sigma}'^{-1}_{\mathbf{x}}\tilde{\boldsymbol{\mu}}'_{\mathbf{x}} - k^2 = 0$$

or in matrix form

$$\begin{pmatrix}\tilde{\mathbf{x}}'^{\top} & 1\end{pmatrix}\begin{bmatrix}\tilde{\Sigma}'^{-1}_{\mathbf{x}} & -\tilde{\Sigma}'^{-1}_{\mathbf{x}}\tilde{\boldsymbol{\mu}}'_{\mathbf{x}} \\ -\tilde{\boldsymbol{\mu}}'^{\top}_{\mathbf{x}}\tilde{\Sigma}'^{-1}_{\mathbf{x}} & \tilde{\boldsymbol{\mu}}'^{\top}_{\mathbf{x}}\tilde{\Sigma}'^{-1}_{\mathbf{x}}\tilde{\boldsymbol{\mu}}'_{\mathbf{x}} - k^2\end{bmatrix}\begin{pmatrix}\tilde{\mathbf{x}}' \\ 1\end{pmatrix} = 0$$

This is equivalent to

$$\begin{aligned}\begin{pmatrix}\tilde{\mathbf{x}}'^{\top} & 1\end{pmatrix}\begin{bmatrix}\tilde{\boldsymbol{\mu}}'_{\mathbf{x}}\tilde{\boldsymbol{\mu}}'^{\top}_{\mathbf{x}} - k^2\tilde{\Sigma}'_{\mathbf{x}} & \tilde{\boldsymbol{\mu}}'_{\mathbf{x}} \\ \tilde{\boldsymbol{\mu}}'^{\top}_{\mathbf{x}} & 1\end{bmatrix}^{-1}\begin{pmatrix}\tilde{\mathbf{x}}' \\ 1\end{pmatrix} &= 0 \\ \mathbf{x}'^{\top}\left[\boldsymbol{\mu}'_{\mathbf{x}}\boldsymbol{\mu}'^{\top}_{\mathbf{x}} - k^2\Sigma'_{\mathbf{x}}\right]^{-1}\mathbf{x}' &= 0\end{aligned}$$

which is the equation of a conic. The conic $\mathtt{C}' = [\boldsymbol{\mu}'_{\mathbf{x}}\boldsymbol{\mu}'^{\top}_{\mathbf{x}} - k^2\Sigma'_{\mathbf{x}}]^{-1}$ is formed by the points that satisfy (6). Transforming back to the original coordinate system, $\mathtt{C} = \mathtt{U}^{\top}\mathtt{C}'\mathtt{U}$, the set of equal-likelihood points that satisfy (5) form the homogeneous conic

$$\mathtt{C} = \left[\boldsymbol{\mu}_{\mathbf{x}}\boldsymbol{\mu}^{\top}_{\mathbf{x}} - k^2\Sigma_{\mathbf{x}}\right]^{-1} \tag{7}$$

representing an ellipse containing $\boldsymbol{\mu}_{\mathbf{x}}$. An arbitrary point $\mathbf{x}_0$ is on the interior of the ellipse if $\mathbf{x}_0^{\top}\mathtt{C}\mathbf{x}_0$ has the same sign as $\boldsymbol{\mu}^{\top}_{\mathbf{x}}\mathtt{C}\boldsymbol{\mu}_{\mathbf{x}}$.

Using the duality between points and lines, and conics and dual conics, the same approach is employed for homogeneous 2D lines. The set of equal-likelihood lines in the distribution of a random homogeneous line $\boldsymbol{\ell}$ with mean $\boldsymbol{\mu}_{\boldsymbol{\ell}}$ and covariance $\Sigma_{\boldsymbol{\ell}}$ satisfies

$$(\boldsymbol{\ell} - \boldsymbol{\mu}_{\boldsymbol{\ell}})^{\top}\Sigma^{+}_{\boldsymbol{\ell}}(\boldsymbol{\ell} - \boldsymbol{\mu}_{\boldsymbol{\ell}}) = k^2$$

for a given k. The set of lines form the homogeneous dual conic $\mathtt{C}^* = \left[\boldsymbol{\mu}_{\boldsymbol{\ell}}\boldsymbol{\mu}^{\top}_{\boldsymbol{\ell}} - k^2\Sigma_{\boldsymbol{\ell}}\right]^{-1}$, which is the the adjoint of the matrix $\mathtt{C}$. For a non-singular symmetric matrix $\mathtt{C} \sim (\mathtt{C}^*)^{-1}$, therefore the conic that forms the envelope of lines is given by

$$\mathtt{C} = \boldsymbol{\mu}_{\boldsymbol{\ell}}\boldsymbol{\mu}^{\top}_{\boldsymbol{\ell}} - k^2\Sigma_{\boldsymbol{\ell}} \tag{8}$$

This conic is a hyperbola with branches symmetric about $\boldsymbol{\mu}_{\boldsymbol{\ell}}$. An arbitrary point $\mathbf{x}_0$ lies inside the region between the two branches of the hyperbola if $\mathbf{x}_0^{\top}\mathtt{C}\mathbf{x}_0$ has the same sign as $\mathbf{x}^{\top}_{\boldsymbol{\ell}}\mathtt{C}\mathbf{x}_{\boldsymbol{\ell}}$, where $\mathbf{x}_{\boldsymbol{\ell}}$ is any point that lies on the line $\boldsymbol{\mu}_{\boldsymbol{\ell}}$. Two points $\mathbf{x}_1$ and $\mathbf{x}_2$ on the line $\boldsymbol{\mu}_{\boldsymbol{\ell}}$ may be determined by $\boldsymbol{\mu}^{\top}_{\boldsymbol{\ell}}\left[\mathbf{x}_1\ \mathbf{x}_2\right] = 0$, where the matrix $\left[\mathbf{x}_1\ \mathbf{x}_2\right]$ is the null space of $\boldsymbol{\mu}^{\top}_{\boldsymbol{\ell}}$. One of these points can be used to determine the sign of $\mathbf{x}^{\top}_{\boldsymbol{\ell}}\mathtt{C}\mathbf{x}_{\boldsymbol{\ell}}$.

3.3 Points and Hyperplanes in n Dimensions

Generalizing the above results from 2 to n dimensions is straightforward. A homogeneous n-dimensional point $\mathbf{X} \in \mathbb{R}^{(n+1)}$ with mean $\boldsymbol{\mu}_{\mathbf{X}}$ and covariance $\Sigma_{\mathbf{X}}$ of rank n is bounded by the homogeneous n-dimensional quadric

$$\mathtt{Q} = \left[\boldsymbol{\mu}_{\mathbf{X}}\boldsymbol{\mu}_{\mathbf{X}}^{\top} - k^2\Sigma_{\mathbf{X}}\right]^{-1}$$

representing an ellipsoid in n dimensions containing $\boldsymbol{\mu}_{\mathbf{X}}$. By duality, a homogeneous hyperplane $\boldsymbol{\pi}$ is bounded by the dual quadric $\mathtt{Q}^* = \left[\boldsymbol{\mu}_{\boldsymbol{\pi}}\boldsymbol{\mu}_{\boldsymbol{\pi}}^{\top} - k^2\Sigma_{\boldsymbol{\pi}}\right]^{-1}$ of the same dimension as the hyperplane, where $\boldsymbol{\mu}_{\boldsymbol{\pi}}$ and $\Sigma_{\boldsymbol{\pi}}$ are the mean and covariance, respectively, of the hyperplane. The hyperboloid bounding the uncertainty of the hyperplane is given by the quadric

$$\mathtt{Q} = \boldsymbol{\mu}_{\boldsymbol{\pi}}\boldsymbol{\mu}_{\boldsymbol{\pi}}^{\top} - k^2\Sigma_{\boldsymbol{\pi}}$$

which is symmetric about the hyperplane.

4 Two-View Geometry

In this section we apply our method to point-to-point mapping under a planar homography and point-to-line mapping under a fundamental matrix. The maximum likelihood estimate and its covariance is determined from image point correspondences by 2D block adjustment [10] and two-view bundle adjustment [11] for the planar homography and fundamental matrix, respectively. First, points are detected in each of the images using the Förstner operator [12]. For each point in image 1, its initial corresponding point is established by searching for the point in image 2 that has the highest local normalized cross-correlation value. The resulting set of initial point correspondences are used as input to RANSAC [13], which provides both a linear estimate of the model and its set of inlier point correspondences. Lastly, the reprojection error is minimized using a sparse implementation of the Levenberg-Marquardt algorithm [14]. We retrieve the covariance matrix of the parameters after minimization.

For analysis, we select a point $\mathbf{x}$ in image 1 that did not participate in block adjustment. The point $\mathbf{x} = (\tilde{\mathbf{x}}^{\top}, 1)^{\top}$ has covariance

$$\Sigma_{\mathbf{x}} = \begin{bmatrix} \tilde{\Sigma}_{\mathbf{x}} & \mathbf{0} \\ \mathbf{0}^{\top} & 0 \end{bmatrix}$$

where the inhomogeneous coordinate $\tilde{\mathbf{x}} = (\tilde{x}, \tilde{y})^{\top}$ has assumed covariance $\tilde{\Sigma}_{\mathbf{x}} = \mathtt{I}_{2\times 2}$.

Figure 1 shows the results of point-to-point mapping under the estimated planar homography. The mapped point is computed by

$$\mathbf{x}' = \frac{\mathtt{H}\mathbf{x}}{\|\mathtt{H}\mathbf{x}\|}$$

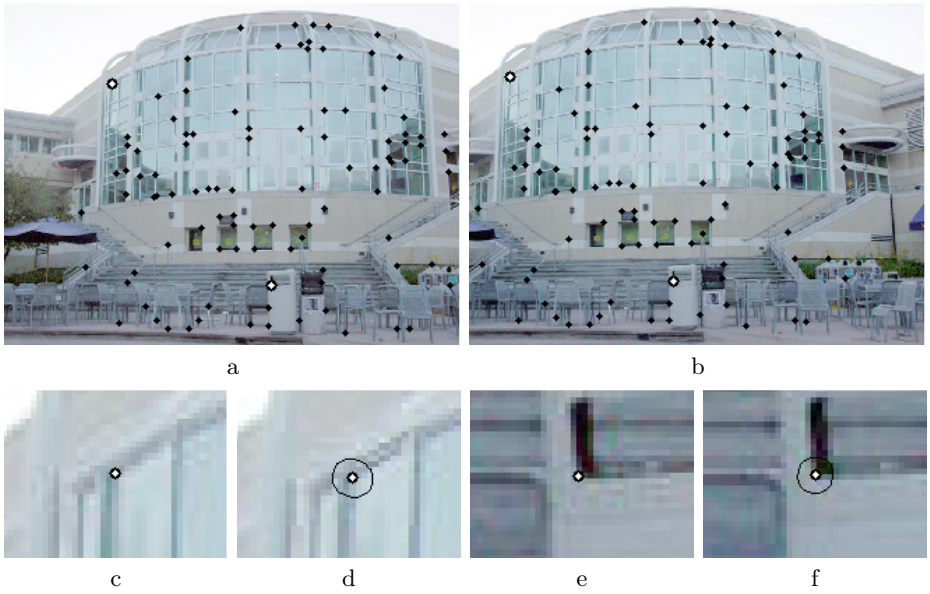

Fig. 1. Point-to-point mapping under a planar homography. (a) (b) The left and right images acquired from a camera undergoing pure rotation about its center with corresponding points used in 2D block adjustment (in black). There are 97 point correspondences. Two additional points have been selected in the left image (in white) and mapped to the right image (in white). The uncertainty ellipses associated with the mapped points are contained in the right image (in black). The ellipses correspond to a probability of 99%. (c) The left image zoomed in on the first selected point. (d) The right image zoomed in on the corresponding first mapped point. (e) The left image zoomed in on the second selected point. (f) The right image zoomed in on the corresponding second mapped point. Note that the eccentricity of the ellipse associated with the first point is slightly greater than that of the ellipse associated with the second point. This is because the second point is surrounded by points used in block adjustment, while the first point is not.

From (1), the covariance of $\mathbf{x}'$ is $\Sigma_{\mathbf{x}'} \approx \mathtt{J}_{\mathbf{h}}\Sigma_{\mathbf{h}}\mathtt{J}_{\mathbf{h}}^{\top} + \mathtt{J}_{\mathbf{x}}\Sigma_{\mathbf{x}}\mathtt{J}_{\mathbf{x}}^{\top}$, where $\mathtt{J}_{\mathbf{h}}$ and $\mathtt{J}_{\mathbf{x}}$ are computed from (2) and (3), respectively, and the associated uncertainty ellipse is computed from (7). Points that did participate in block adjustment are correlated to the estimated homography. If one of these points were selected, then the cross-covariance $\Sigma_{\mathbf{hx}}$ would be nonzero and the covariance of $\mathbf{x}'$ calculated as

$$\Sigma_{\mathbf{x}'} \approx \begin{bmatrix} \mathtt{J}_{\mathbf{h}} & \mathtt{J}_{\mathbf{x}} \end{bmatrix} \begin{bmatrix} \Sigma_{\mathbf{h}} & \Sigma_{\mathbf{hx}} \\ \Sigma_{\mathbf{xh}} & \Sigma_{\mathbf{x}} \end{bmatrix} \begin{bmatrix} \mathtt{J}_{\mathbf{h}}^{\top} \\ \mathtt{J}_{\mathbf{x}}^{\top} \end{bmatrix}$$

Similarly, results for point-to-line mapping under a fundamental matrix are shown in Figure 2. The line $\boldsymbol{\ell}'$ corresponding to the point $\mathbf{x}$ is computed by

$$\boldsymbol{\ell}' = \frac{\mathtt{F}\mathbf{x}}{\|\mathtt{F}\mathbf{x}\|}$$

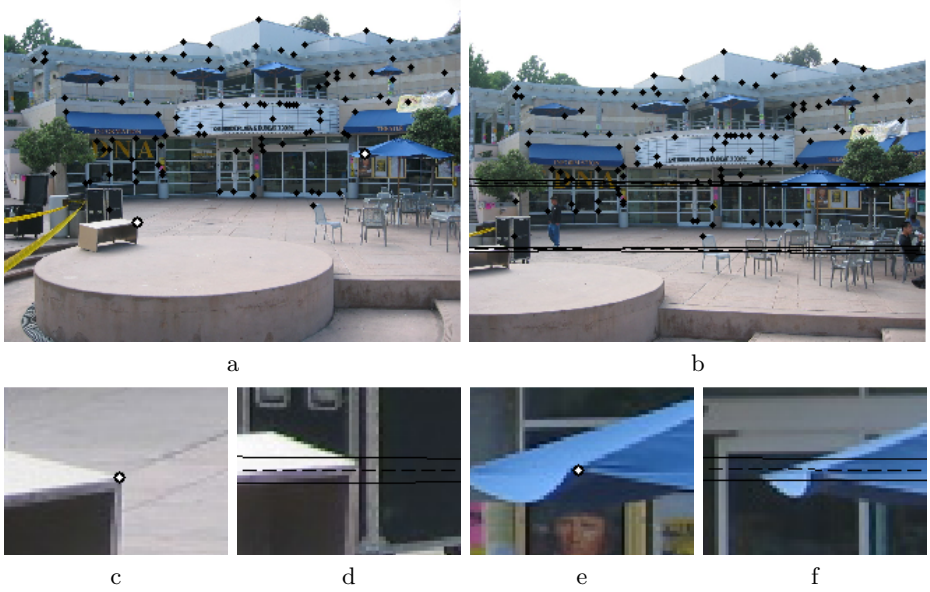

Fig. 2. Point-to-line mapping under a fundamental matrix. (a) (b) The left and right images with corresponding points used in bundle adjustment (in black). There are 101 point correspondences. Two additional points have been selected in the left image (in white) and mapped to lines in the right image (in dashed black). The uncertainty hyperbolas associated with the mapped lines are contained in the right image (in black). The hyperbolas correspond to a probability of 99%. (c) The left image zoomed in on the first selected point. (d) The right image zoomed in on the corresponding first mapped line. (e) The left image zoomed in on the second selected point. (f) The right image zoomed in on the corresponding second mapped line. Note that the mapped lines miss the corresponding points, but that the corresponding points are within the uncertainty bounds.

The covariance of $\boldsymbol{\ell}'$ is $\Sigma_{\boldsymbol{\ell}'} \approx \mathtt{J}_{\mathbf{f}}\Sigma_{\mathbf{f}}\mathtt{J}_{\mathbf{f}}^{\top} + \mathtt{J}_{\mathbf{x}}\Sigma_{\mathbf{x}}\mathtt{J}_{\mathbf{x}}^{\top}$ with associated uncertainty hyperbola given by (8).

5 Mosaic Construction from Video

This section describes use of our approach in the application of video mosaicing. More specifically, we apply our approach to the special case of the video looping back on itself, i.e., the sensor returns to image a region of the scene that it imaged at a previous time. We seek to determine the search region in the previously acquired frames that spatially overlap with the looped back frames but are not temporal neighbors with these frames. Mosaic construction from video is performed in a sequential manner as follows.

For each video frame, features are detected using the method described in [15]. This method detects windows of bidirectional texturedness, which are

good features to track in video. Nonmaxima suppression is applied to detected features to limit their number. A pyramidal implementation of Lucas-Kanade [16] determines the translation of each feature from the current frame to the previous one. Just as in Section 4, RANSAC is applied to the inter-frame correspondences, and the reprojection error is minimized using a sparse implementation of the Levenberg-Marquardt algorithm and the covariance matrix is retrieved. The homographies are accumulated such that the current frame n is mapped back to frame 1 of the video by

$$\mathtt{H}_{n,1} = \frac{\mathtt{H}_{n-1,1}\mathtt{H}_{n,n-1}}{\|\mathtt{H}_{n-1,1}\mathtt{H}_{n,n-1}\|} = \frac{\mathtt{AB}}{\|\mathtt{AB}\|}$$

and the covariance of $\mathtt{H}_{n,1}$ is approximately $\mathtt{J}_{\mathbf{a}}\Sigma_{\mathbf{a}}\mathtt{J}_{\mathbf{a}}^{\top} + \mathtt{J}_{\mathbf{b}}\Sigma_{\mathbf{b}}\mathtt{J}_{\mathbf{b}}^{\top}$.

As the homographies between successive frames $\mathtt{H}_{n,n-1}$ are accumulated, so are their uncertainties. It is expected that the uncertainty of $\mathtt{H}_{n,1}$ will increase with n, i.e., looped back frames will not align with previous frames containing images of the same region of the scene and this misalignment will increase as the time between these frames increase. Figure 3 illustrates the results of this approach.

6 Conclusions

In this paper we have presented an analytical method based on covariance propagation for determining the search region in guided matching. Our primary contribution is a general approach and accompanying analytical expression for the Jacobian matrix used in this approach. Not only does this obviate the need for specialized expressions, it can also be applied to other projective mappings, such as composition of homographies. We have shown that the uncertainty of homogeneous 2D points and lines are bounded by a conic and generalized this to homogeneous points and hyperplanes in n dimensions. Finally, we have applied this general approach to point-to-point mapping under a planar homography, point-to-line mapping under a fundamental matrix, and mosaic construction from video in the case of the video looping back on itself. This general approach can be used in other practical applications such as camera pose estimation over time and scene reconstruction from video.

Acknowledgments

The authors thank Sameer Agarwal, Josh Wills, and Kristin Branson at the University of California, San Diego, and Alan Sussman and John Dolloff at BAE Systems for helpful discussions. The original image pairs in figures 1 and 2 were provided by Kristin Branson and Neil McCurdy, respectively.

a b c

Fig. 3. Mosaic construction from video. (a) A planar mosaic sequentially constructed from a video containing 706 frames acquired from a camera undergoing pure rotation about its center. The video begins at the upper left corner of the face of the building and moves clockwise around the border of the face returning to the upper left corner. Note that the last frame is not aligned with the first frame due to uncertainty accumulation. (b) (c) The last and first frames (images) of the video. Three points (in red, yellow, and green) have been selected in the last image and mapped to the first image. The uncertainty ellipses associated with the mapped points are contained in the first image (also in red, yellow, and green). The ellipses correspond to probabilities of 50% and 99% for each mapped point. In this case, the corresponding points in the first image are contained in the ellipses corresponding to a probability of 50%.

References

1. Kanatani, K.: Statistical Optimization for Geometric Computation: Theory and Practice. Elsevier Science (1996)
2. Förstner, W.: Uncertainty and projective geometry. In Corrochano, E.B., ed.: Handbook of Geometric Computing: Applications in Pattern Recognition, Computer Vision, Neuralcomputing, and Robotics. Springer-Verlag (2005) 493–534
3. Seitz, S.M., Anandan, P.: Implicit representation and scene reconstruction from probability density functions. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (1999) 28–34
4. Triggs, B.: Joint feature distributions for image correspondence. In: Proceedings of the International Conference on Computer Vision. (2001) 201–108
5. Hartley, R., Zisserman, A.: Multiple View Geometry in Computer Vision. Cambridge University Press (2000)
6. Criminisi, A.: A plane measuring device. Image and Vision Computing **40** (1999) 625–634
7. Csurka, G., Zeller, C., Zhang, Z., Faugeras, O.D.: Characterizing the uncertainty of the fundamental matrix. Computer Vision and Image Understanding **68** (1997) 18–36
8. Zhang, Z.: Determining the epipolar geometry and its uncertainty: A review. International Journal of Computer Vision **27** (1998) 161–195
9. Semple, J.G., Kneebone, G.T.: Algebraic Projective Geometry. Oxford University Press (1952)
10. Capel, D., Zisserman, A.: Automated mosaicing with super-resolution zoom. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (1998) 885–891
11. Triggs, B., McLauchlan, P.F., Hartley, R.I., Fitzgibbon, A.W.: Bundle adjustment – A modern synthesis. In: Vision Algorithms: Theory and Practice. Springer-Verlag (2000)
12. Förstner, W., Gülch, E.: A fast operator for detection and precise location of distinct points, corners and centres of circular features. In: Proceedings of the ISPRS Conference on Fast Processing of Photogrammetric Data. (1987) 281–305
13. Fischler, M.A., Bolles, R.C.: Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. Communications of the ACM **24** (1981) 381–395
14. Hartley, R.I.: Euclidean reconstruction from uncalibrated views. In: Proceedings of the Workshop on Applications of Invariance in Computer Vision. (1994) 237–256
15. Shi, J., Tomasi, C.: Good features to track. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (1994) 593–600
16. Lucas, B.D., Kanade, T.: An iterative image registration technique with an application to stereo vision. In: Proceedings of the International Joint Conference on Artificial Intelligence. (1981) 674–679
17. Graham, A.: Kronecker Products and Matrix Calculus with Applications. Ellis Horwood Ltd. (1981)

Appendix

Assume $\mathtt{A} \in \mathbb{R}^{m \times p}$, $\mathtt{B} \in \mathbb{R}^{p \times n}$, and $\mathtt{C} \in \mathbb{R}^{m \times n}$. For the equation

$$\mathtt{C} = \frac{\mathtt{AB}}{\|\mathtt{AB}\|} = \frac{\mathtt{M}}{\|\mathtt{M}\|}$$

where $\|\cdot\|$ denotes the Frobenius norm, we seek the partial derivatives of the entries of $\mathtt{C}$ with respect to the entries of both $\mathtt{A}$ and $\mathtt{B}$ [17]. For clarity, the matrix product $\mathtt{AB}$ is denoted by $\mathtt{M}$, which allows for $\mathrm{vec}((\mathtt{AB})^\top)$ to be represented by $\mathbf{m}$. The partial derivative of $\mathbf{c}$ with respect to $\mathbf{a}$ is computed as

$$\frac{\partial \mathbf{c}}{\partial \mathbf{a}} = \frac{1}{\|\mathbf{m}\|^2}\left[\|\mathbf{m}\|\frac{\partial \mathbf{m}}{\partial \mathbf{a}} - \mathbf{m}\frac{\partial \|\mathbf{m}\|}{\partial \mathbf{a}}\right]$$

or equivalently

$$\frac{\partial \mathbf{c}}{\partial \mathbf{a}} = \frac{1}{\|\mathbf{m}\|}\left[\frac{\partial \mathbf{m}}{\partial \mathbf{a}} - \mathbf{c}\frac{\partial \|\mathbf{m}\|}{\partial \mathbf{a}}\right] \tag{9}$$

Similarly,

$$\frac{\partial \mathbf{c}}{\partial \mathbf{b}} = \frac{1}{\|\mathbf{m}\|}\left[\frac{\partial \mathbf{m}}{\partial \mathbf{b}} - \mathbf{c}\frac{\partial \|\mathbf{m}\|}{\partial \mathbf{b}}\right] \tag{10}$$

The partial derivative of $\mathbf{m}$ with respect to both $\mathbf{a}$ and $\mathbf{b}$ is given by

$$\frac{\partial \mathbf{m}}{\partial \mathbf{a}} = \mathtt{I}_{m\times m} \otimes \mathtt{B}^\top \text{and } \frac{\partial \mathbf{m}}{\partial \mathbf{b}} = \mathtt{A} \otimes \mathtt{I}_{n\times n}$$

where $\otimes$ denotes the Kronecker product. The partial derivative of $\|\mathtt{M}\|$ with respect to both $\mathtt{A}$ and $\mathtt{B}$ is

$$\frac{\partial \|\mathtt{M}\|}{\partial \mathtt{A}} = \mathtt{C}\mathtt{B}^\top \text{and } \frac{\partial \|\mathtt{M}\|}{\partial \mathtt{B}} = \mathtt{A}^\top\mathtt{C}$$

It follows that

$$\frac{\partial \|\mathbf{m}\|}{\partial \mathbf{a}} = \mathrm{vec}(\mathtt{B}\mathtt{C}^\top)^\top \text{and } \frac{\partial \|\mathbf{m}\|}{\partial \mathbf{b}} = \mathrm{vec}(\mathtt{C}^\top\mathtt{A})^\top$$

Substituting into (9) and (10) yields

$$\frac{\partial \mathbf{c}}{\partial \mathbf{a}} = \frac{1}{\|\mathtt{AB}\|}\left[\left(\mathtt{I}_{m\times m} \otimes \mathtt{B}^\top\right) - \mathbf{c}\,\mathrm{vec}(\mathtt{B}\mathtt{C}^\top)^\top\right]$$
$$\frac{\partial \mathbf{c}}{\partial \mathbf{b}} = \frac{1}{\|\mathtt{AB}\|}\left[\left(\mathtt{A} \otimes \mathtt{I}_{n\times n}\right) - \mathbf{c}\,\mathrm{vec}(\mathtt{C}^\top\mathtt{A})^\top\right]$$