

Image Inpainting using Two -View Epipolar Geometry

Prasanna Rangarajan

Panos Papamichalis

Marc Christensen

Abstract

We take a fresh look at the problem of removing occluders in an image using inpainting. We examine a geometric method that utilizes a second image of the scene from a different viewpoint, to identify the occluded objects. We recover the missing intensities by using the geometric relationship between corresponding points in the two images. The relationship is generally specified by the “epipolar line constraint”, and degenerates to a projective homography under special circumstances. For this reason, we use the plane + parallax approach to estimate the epipolar geometry relating the two images. We fill-in missing pixels by copying information from the respective epipolar lines in the second image.

The success of the proposed framework hinges on the ability to estimate the “epipolar line constraint” from noisy correspondences. To this end, we analyze the uncertainty in estimating a homography from noisy correspondences. We rely on the knowledge of this uncertainty to identify parallax vectors best suited for estimating the epipolar geometry.

We do not make any explicit assumptions about the nature or the extent of camera motion, only requiring that the occluded objects are static and undergo limited perspective change. To illustrate the effectiveness of the proposed approach, we “inpaint” occluders in wide-baseline situations.

1. Introduction

The objective of *image inpainting* is to produce a modified image where the inpainted region merges seamlessly with the rest of the image, in a manner not obvious to the untrained eye. Applications of image inpainting include *occluding object removal, wire/rig removal, retouching of damaged paintings & photographs and error concealment*.

Traditional image inpainting methods[1, 2, 5, 12] infer missing intensities using basic principles like continuation of edges (*level lines*), and texture synthesis. The quality of the inpainting depends on the size of the missing region, the geometry of the occluded objects, and the fill-order¹. Small regions are inpainted effortlessly, while inpainting large regions is difficult and often produces unrealistic results.

The proposed approach extends the scope of current image inpainting methods, by using the *epipolar line constraint* to infer missing intensities from an additional image

¹fill-order refers to the order in which the missing pixels are inpainted.

(*exemplar image*). The use of an *exemplar image* permits us to inpaint large regions, regardless of the geometry of the occluded objects and the fill-order¹.

The use of an *exemplar image* in inpainting is a recent development. Kang et al. [11] recover the missing pixels by assuming that the two images are *affinely* related. Jain[9] assumes that the occluded objects are planar in nature, and recovers the missing intensities using a *projective homography*. Also, there are *spatio-temporal* techniques that attempt to recover missing pixels using several images. Please refer to [15] for a review of such techniques. In this paper, we focus our attention on image inpainting using a single *exemplar image*, for arbitrary camera motion.

The novelty of our approach lies in the use of camera *and-or* occluder motion to reveal the missing intensities. Motion in the occluder either *fully or partially* reveals the occluded objects in the *exemplar image*. In the case of static occluders, camera motion induces an apparent motion in the occluders, depending on their depth. As a result, previously occluded objects are revealed *fully or partially* in the *exemplar image*. In either situation, the occluded objects can experience significant foreshortening due to camera motion. As this makes the task of inpainting difficult, we require that occluded objects undergo limited perspective change.

From now on, we will refer to the *image to be inpainted* as *view-1* and the *exemplar image* as *view-2*.

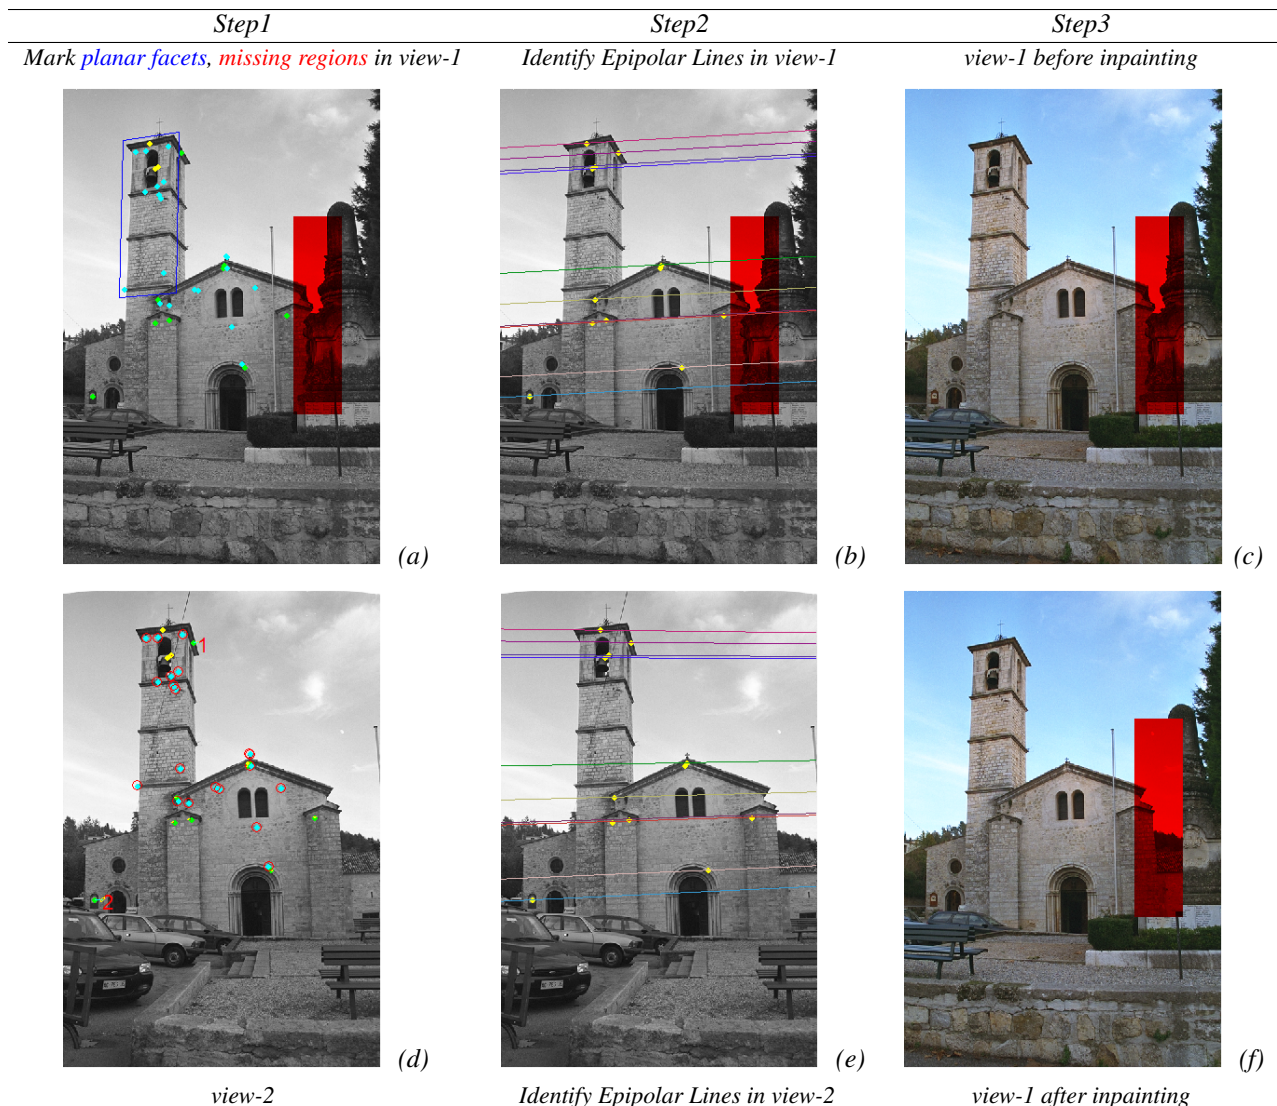
2. Proposed Approach

The challenge in the proposed inpainting framework, lies in accurately estimating the geometric relationship between *views-1 & 2*. The relationship is completely specified by a *homography* in these situations

- *occluded object/region belongs to planar facet or*
- *small baseline and distant occluders or*
- *small change in camera perspective between views-1 & 2*

We inpaint the missing region in *view-1*, by warping *view-2* into *view-1*, using the *inverse homography*. In all other situations, the “*epipolar line constraint*” describes the geometric relationship between corresponding points in *views-1 & 2*. We inpaint missing pixels in *view-1*, by copying intensities from the respective epipolar lines in *view-2*.

In either situation, we need to identify putative point correspondences in the two views. We start by identifying corners with *pixel precision* using the *fast corner detection* algorithm[14], and use *log-polar normalized cross correlation* (*LP-NCC*) [16] to match the neighborhood of each cor-



The objective is to inpaint parts of the pillar-like structure (transparent red region) in Fig1(a), using an image of the scene (Fig1(d)), from a different viewpoint. We inpaint the missing pixels in Fig1(a) by copying intensities from the respective epipolar lines in Fig1(e). Epipolar lines sharing the same color in Figs1(b),1(e) are in correspondence. The result is shown in Fig1(f).

The challenge is to accurately estimate the epipolar geometry relating Fig1(a) and Fig1(d), from a limited number of noisy point correspondences. The proposed approach relies on the uncertainty in estimating homographies and parallax vectors from noisy point correspondences, to accurately estimate the epipolar geometry.

Figure 1. Image Inpainting using Epipolar Geometry for the Valbonne church scene (courtesy INRIA Robotvis Group)

ner in view-1, with every corner in view-2. We prune the list of matches to obtain unique point correspondences. Unfortunately, these point correspondences are noisy. This problem plagues most *corner estimation* and *correspondence matching* algorithms. So, we proceed to estimate the *homography* or the *epipolar geometry* relating the two views, from noisy correspondences (Steps 1 & 2 in Fig.1). This is followed by the actual inpainting (Step 3 in Fig.1).

The success of the proposed inpainting framework hinges on the ability to estimate the epipolar geometry from noisy correspondences. So we start by analyzing the uncer-

tainty in estimating a homography from noisy correspondences, in Sec-3.2. In Sec-3.3.2, we propagate the uncertainty in estimating a homography to *parallax vectors*. In Sec-4 we examine the use of uncertainty analysis in inpainting, and provide examples of inpainting using the proposed approach in Sec-5.

3. Uncertainty Analysis

Different structures for the uncertainty in a homography, have been proposed by Criminisi[4], Suter[3], Kanatani[10], and Hartley [7]. Unfortunately, these struc-

tures cannot be used directly to find : *parallax vectors* and *planar facets* that are best suited for estimating epipolar geometry from noisy point correspondences. With this objective in mind, we re-examine the problem of determining the uncertainty in a homography. The utility of the proposed structure of uncertainty lies in the ease with which it can be extended to different normalization schemes, and propagated through projective entities such as *algebraic transfer error*, *parallax vectors*, *lines* and *epipoles*.

3.1. Notation

- **boldface** symbols are *vectors*, underlined symbols are *homogeneous vectors*, and CAPITAL letters represent matrices
- The ordered pair (x, y) represents a point in Euclidian 2-space (\mathbb{R}^2). The 3-vector $\underline{x} = (x, y, 1)^T$ is the corresponding point in *homogeneous coordinates* (Projective 2-space \mathbb{P}^2).
- $A(i, :)$ represents the i^{th} row of the matrix A
- $\text{diag}(\mathbf{v})$ is a diagonal matrix with entries of the vector \mathbf{v} along the principal diagonal
- I_d denotes a $d \times d$ identity matrix
- vec -operator[8] vectorizes a matrix by stacking its columns.
- $[\underline{e}]_{\underline{x}\underline{y}}$ denotes the use of skew-symmetric matrices to compute the vector product $\underline{e} \times \underline{y}$
- $\|\mathbf{v}\|$ denotes the l^2 -norm of the vector \mathbf{v}

3.2. Homography

A homography is a non-singular matrix that describes a linear mapping in *homogenous coordinates* (\mathbb{P}^2). If a homography H_p relates point correspondences $\{(x_i, y_i), (x'_i, y'_i)\}$, the following bilinear relationships hold

$$x'_i = \frac{h_{11}x_i + h_{12}y_i + h_{13}}{h_{31}x_i + h_{32}y_i + h_{33}} \quad y'_i = \frac{h_{21}x_i + h_{22}y_i + h_{23}}{h_{31}x_i + h_{32}y_i + h_{33}} \quad (1)$$

where h_{ij} are entries of the matrix H_p . *eq.(1)* suggests that only the ratio of the entries of H_p is significant. As there are only 8 independent ratios amongst the 9 elements of H_p , a homography has 8 degrees of freedom (d.o.f.). This forms the basis for estimating a homography from 4 or more point correspondences using the Direct Linear Transform algorithm[7].

3.2.1 Direct Linear Transform (DLT) algorithm

Using *kronecker products* and the *vec operator*[8], we rewrite *eq.(1)* in the form of the linear system

$$\underbrace{\begin{pmatrix} \left[\begin{array}{cc} 1 & 0 \\ 0 & 1 \end{array} \right] \otimes \underline{x}_i^T \\ A_{pi} \end{pmatrix}}_{A_{pi}} \underbrace{\begin{bmatrix} h_{11} \\ h_{12} \\ \dots \\ h_{33} \end{bmatrix}}_{\mathbf{h}_p = \text{vec}(H_p^T)} = \underbrace{\begin{bmatrix} 0 \\ 0 \\ \dots \\ 0 \end{bmatrix}}_{\mathbf{0}} \quad (2)$$

Given $M > 4$ sets of point correspondences related by a homography H_p , we estimate H_p by solving the following overdetermined linear system [7]

$$A_p \mathbf{h}_p = \mathbf{0}, \quad A_p = \begin{bmatrix} A_{p1}^T & \dots & A_{pM}^T \end{bmatrix}^T \quad (3)$$

In the absence of noise, A_p is rank-deficient and $\mathbf{h}_p \in$ nullspace of A_p . In the presence of noise A_p is almost always full rank. In either case, \mathbf{h}_p is estimated as the *right singular vector* corresponding to the *least singular value* of A_p^2 . It is easy to observe that \mathbf{h}_p has 8 d.o.f. as $\|\mathbf{h}_p\| = 1$.

There is always an uncertainty associated with estimating H_p using DLT, as the point correspondences used to estimate H_p are noisy. Typically, noise affects corners in both *views-1 & 2*. But, for the sake of analysis we assume that only corners $\{(x'_i, y'_i)\}$ in *view-2* are affected by noise. This assumption is valid only when we seek to find the corner in *view-2* that corresponds to a noisy corner in *view-1*.

3.2.2 Noise model

We assume that the noise $\boldsymbol{\varepsilon}'_i = (\boldsymbol{\varepsilon}'_{xi}, \boldsymbol{\varepsilon}'_{yi})^T$ that affects the corners (x'_i, y'_i) in *view-2* is i.i.d³ gaussian with zero mean and covariance $E\{\boldsymbol{\varepsilon}'_i \boldsymbol{\varepsilon}'_i{}^T\} = \text{diag}([\sigma_x^2, \sigma_y^2])$. This model implies that the error in localizing a corner in the x-direction is independent of the error in the y-direction.

Under these assumptions, we examine the effect of estimating a homography from noisy pixel coordinates, using the DLT algorithm.

3.2.3 Perturbation Analysis of the DLT algorithm

In the presence of noise, we estimate the noisy homography $\mathbf{h}_p + \delta_{h_p}$ as the *right singular vector* corresponding to the *least singular value* of the noisy matrix

$$A_p + N_p = \begin{bmatrix} \left[\begin{array}{cc} 1 & 0 \\ 0 & 1 \end{array} \right] \otimes \underline{x}_1^T \\ \dots \\ \left[\begin{array}{cc} 1 & 0 \\ 0 & 1 \end{array} \right] \otimes \underline{x}_M^T \end{bmatrix}, \quad \text{where} \quad N_p = [\dots N_p^T \dots]^T, \quad N_p^i = \begin{bmatrix} 0 & 0 & -\boldsymbol{\varepsilon}'_{xi} \\ 0 & 0 & -\boldsymbol{\varepsilon}'_{yi} \end{bmatrix} \otimes \underline{x}_i^T \quad (4)$$

For small noise levels⁴, the effect of estimating a homography from noisy correspondences is best explained by the theory of *first-order perturbation analysis* [4, 6]. Using this theory, we express the perturbation in the homography δ_{h_p} , and its associated uncertainty Λ_{h_p} in the compact form

$$\delta_{h_p} = A_p^\dagger (W \otimes I_2) \Xi \quad (5)$$

$$\Lambda_{h_p} = E\{\delta_{h_p} \delta_{h_p}^T\} = A_p^\dagger (W \otimes I_2) E\{\Xi \Xi^T\} (W^T \otimes I_2) A_p^{\dagger T} \quad (6)$$

$$= A_p^\dagger [W W^T \otimes \text{diag}([\sigma_x^2, \sigma_y^2])] A_p^{\dagger T} \quad (7)$$

where A_p^\dagger is the *pseudoinverse* of the matrix A_p , $W = \text{diag}([\omega_1 \dots \omega_M]) = \text{diag}([H_p(3, :)\underline{x}_1 \dots H_p(3, :)\underline{x}_M])$,

² \mathbf{h}_p is the solution to the *Total Least Squares* problem of *eq.(3)*

³ independent and identically distributed

⁴Quantifying "small" is difficult since it depends on the image size. A noise variance of 4 pixels in a 1024 x 768 image is \ll noise variance of 4 pixels in a 352 x 288 image.

$\Xi = [\varepsilon'_{x1} \ \varepsilon'_{y1} \ \dots \ \varepsilon'_{xM} \ \varepsilon'_{yM}]^T$. The matrix W is comprised of the *homogenous scale factors* $\{\omega_i\}$, while Ξ is comprised of the noise affecting correspondences in *view-2*.

Appendix-1 provides details on deriving the expressions for δ_{h_p} and Λ_{h_p} . Smaller entries in Λ_{h_p} correspond to a smaller uncertainty in estimating \mathbf{h}_p , and *vice-versa*.

The proposed structure of Λ_{h_p} (eq.(7)) has geometric appeal as explained in Sec-3.2.4, and is easily extended to *normalized DLT*[7] in Sec-3.2.5. In addition, Λ_{h_p} can be propagated to projective entities such as *algebraic transfer error* (Sec-3.2.6), *parallax vectors* (Sec-3.3.2), *lines* and *epipoles*. Further, the structure of Λ_{h_p} in eq.(7) allows us to implicitly estimate the unknown noise variances, as described in 3.2.6.

Another advantage of the proposed structure of Λ_{h_p} is that it can be readily expressed in alternative forms⁵, such as the structure of the covariance matrix proposed by Hartley [7] and Kanatani [10].

$$\Lambda_{h_p} = \left([(W^{-1} \otimes I_2) A_p]^T [E\{\Xi\Xi^T\}]^{-1} [(W^{-1} \otimes I_2) A_p] \right)^\dagger \quad (8)$$

3.2.4 Structure of the covariance matrix Λ_{h_p}

- Λ_{h_p} is singular, $\text{rank}(\Lambda_{h_p}) \leq 8$, and $\mathbf{h}_p = \text{null}(\Lambda_{h_p})$. The redundancy ($\|\mathbf{h}_p\| = 1$) in estimating a homography using DLT, forces Λ_{h_p} to be singular. As a result, there is no proper density function associated with the perturbation δ_{h_p} .
- The uncertainty in \mathbf{h}_p increases with noise variance. But for large noise variances the expression for Λ_{h_p} is inaccurate, as first-order perturbation analysis of the matrix $A_p + N_p$ is no longer sufficient.
- The uncertainty in \mathbf{h}_p increases with the *homogenous scale factors* $\{\omega_i\}$. By choosing point correspondences that are clustered around the fixed-point⁶ of the homography, we can reduce the uncertainty in \mathbf{h}_p . But, we run the risk of incorrectly estimating a projective transform as being nearly affine. Alternatively, we could arrive at a better estimate of the homography by choosing point correspondences that are spread over a larger area. But this comes at the expense of increased uncertainty (larger ω_i 's). This tradeoff is reminiscent of the bias-variance tradeoff in standard estimators.
- The uncertainty in \mathbf{h}_p depends on the orientation of the observed planar facet. The least uncertainty in \mathbf{h}_p results when the camera in view-2 is looking straight at the planar facet (fronto-parallel situation). The reason is that the least error in localizing corners in view-2 results, when camera-2 is in a fronto-parallel situation. Based on this observation, we choose the planar facet that most directly faces camera-

⁵ Using the property $(AB)^\dagger = B^\dagger A^\dagger$, we can rewrite eq.(6) as eq.(8)

⁶ $\underline{x}_i' = \underline{x}_i, \Rightarrow \omega_i = 1$

(reduces uncertainty) and has the largest area (reduces bias), to estimate the epipolar geometry.

Unfortunately, the DLT algorithm outlined in Sec-3.2.1 cannot be used to estimate the homography H_p , as the ordering of the *right singular vectors* of A_p is sensitive to small perturbations in the pixel coordinates. Hartley[7] points out that the sensitivity is due to the large condition number of the matrix $A_p^T A_p$, and observes that an anisotropic normalization of the pixel coordinates alleviates the problem.

Instead of estimating a homography from noisy pixel coordinates, we estimate the homography in normalized coordinates, using the *Normalized DLT algorithm*[7]. To be specific, the normalized homography $\mathbf{h}_p^n + \delta_{h_p}^n$ is estimated as the *right singular vector* corresponding to the *least singular value* of the normalized matrix

$$A_p^n + N_p^n = \begin{bmatrix} \left[\begin{array}{cc} 1 & 0 \\ 0 & 1 \end{array} - \begin{pmatrix} x_1^n' + \varepsilon_{x1}^n \\ y_1^n' + \varepsilon_{y1}^n \end{pmatrix} \right] \otimes \underline{x}_1^{nT} \\ \dots \\ \left[\begin{array}{cc} 1 & 0 \\ 0 & 1 \end{array} - \begin{pmatrix} x_M^n' + \varepsilon_{xM}^n \\ y_M^n' + \varepsilon_{yM}^n \end{pmatrix} \right] \otimes \underline{x}_M^{nT} \end{bmatrix} \quad (9)$$

$$\underline{x}_i^n = \frac{N_1}{N_1(3,3)} \underline{x}_i, \quad \left[\begin{array}{c} x_i^n' + \varepsilon_{xi}^n \\ y_i^n' + \varepsilon_{yi}^n \\ 1 \end{array} \right] = \frac{N_2}{N_2(3,3)} \underbrace{\left[\begin{array}{c} x_i' + \varepsilon_{xi}' \\ y_i' + \varepsilon_{yi}' \\ 1 \end{array} \right]}_{\hat{\underline{x}}_i} \quad (10)$$

$$N_1 = \text{ichol} \left(\sum_{i=1}^M \underline{x}_i \underline{x}_i^T \right), N_2 = \text{ichol} \left(\sum_{i=1}^M \hat{\underline{x}}_i \hat{\underline{x}}_i^T \right) \quad (11)$$

The *inverse cholesky factors*[6] N_1 & N_2 normalize pixel coordinates in *views-1* & 2 respectively. For small noise levels, the dependence of N_2 on the noise terms in eq.(11) is so negligible that we ignore it. The homography H_p in pixel coordinates is related to the homography H_p^n in normalized coordinates as

$$H_p = \frac{N_2^{-1} H_p^n N_1}{\|N_2^{-1} H_p^n N_1\|}, \quad \mathbf{h}_p = \frac{(N_2^{-1} \otimes N_1^T) \mathbf{h}_p^n}{\|(N_2^{-1} \otimes N_1^T) \mathbf{h}_p^n\|} = \frac{\mathbf{h}}{\|\mathbf{h}\|} \quad (12)$$

3.2.5 Perturbation Analysis of Normalized DLT

Along the lines of Sec-3.2.3, we can express the perturbation in the normalized homography ($\delta_{h_p}^n$) as

$$\delta_{h_p}^n = A_p^{n\dagger} (W^n \otimes I_2) \Xi^n \quad (13)$$

$$W^n = \text{diag}([H_p^n(3,:) \underline{x}_1^n \dots H_p^n(3,:) \underline{x}_M^n]), \quad \Xi^n = [\varepsilon_{x1}^n \ \varepsilon_{y1}^n \dots \varepsilon_{xM}^n \ \varepsilon_{yM}^n]^T$$

The corresponding uncertainty is specified as

$$\Lambda_{h_p}^n = A_p^{n\dagger} \left[W^n W^{nT} \otimes N_{2S} \text{diag}([\sigma_x^2, \sigma_y^2]) N_{2S}^T \right] A_p^{n\dagger T} \quad (14)$$

$$\text{where } N_{2S} = \begin{bmatrix} N_2(1,1) & N_2(1,2) \\ N_2(2,1) & N_2(2,2) \end{bmatrix}$$

Appendix-2 provides details on deriving $\Lambda_{h_p}^n$. Notice that the structure of $\Lambda_{h_p}^n$ in eq.(14) is similar to the structure of Λ_{h_p} in eq.(7), with the exception of the term N_{2S} .

However, *eq.(14)* only describes the uncertainty of the normalized homography \mathbf{h}_p^n . We need to propagate $\Lambda_{h_p}^n$ through *eq.(12)*, to determine Λ_{h_p} , in pixel coordinates. The process begins by propagating noise through \mathbf{h} in *eq.(12)*, and finally through the normalization $\frac{\mathbf{h}}{\|\mathbf{h}\|}$. The uncertainty in \mathbf{h}_p has the following structure

$$\Lambda_{h_p} = \frac{1}{\|\mathbf{h}\|} \left[I_9 - \frac{\mathbf{h}\mathbf{h}^T}{\|\mathbf{h}\|^2} \right] \Lambda_h \left[I_9 - \frac{\mathbf{h}\mathbf{h}^T}{\|\mathbf{h}\|^2} \right] \frac{1}{\|\mathbf{h}\|} \quad (15)$$

where $\mathbf{h} = (N_2^{-1} \otimes N_1^T) \mathbf{h}_p^n$, and $\Lambda_h = (N_2^{-1} \otimes N_1^T) \Lambda_{h_p^n} (N_2^{-T} \otimes N_1)$. Substituting for $\Lambda_{h_p^n}$ from *eq.(14)*, we find that

$$\Lambda_{h_p} = \mathcal{M}_h A_p^{n \dagger} \left[(W^n)^2 \otimes N_{2S} \text{diag}([\sigma_x^2, \sigma_y^2]) N_{2S}^T \right] A_p^{n \dagger T} \mathcal{M}_h^T$$

where $\mathcal{M}_h = \frac{1}{\|\mathbf{h}\|} \left[I_9 - \frac{\mathbf{h}\mathbf{h}^T}{\|\mathbf{h}\|^2} \right] (N_2^{-1} \otimes N_1^T)$ (16)

Alternatively, we can rewrite *eq.(16)* in terms of σ_x^2 & σ_y^2 as

$$\Lambda_{h_p} = \sigma_x^2 \Lambda_{h_p}^x + \sigma_y^2 \Lambda_{h_p}^y \quad (17)$$

$$\Lambda_{h_p}^x = \mathcal{M}_h A_p^{n \dagger} \left[(W^n)^2 \otimes N_{2S} \text{diag}([1, 0]) N_{2S}^T \right] A_p^{n \dagger T} \mathcal{M}_h^T$$

$$\Lambda_{h_p}^y = \mathcal{M}_h A_p^{n \dagger} \left[(W^n)^2 \otimes N_{2S} \text{diag}([0, 1]) N_{2S}^T \right] A_p^{n \dagger T} \mathcal{M}_h^T$$

In *Sec-3.2.6*, we use this structure of Λ_{h_p} to find σ_x^2 & σ_y^2 . Comments about the structure of the uncertainty Λ_{h_p} from *Sec-3.2.4* are still valid.

We conclude the study of uncertainty in estimating homographies using DLT, by noting that the estimated homography $\widehat{\mathbf{h}}_p$ is a random vector with

mean \mathbf{h}_p true homography defined in *eq.(12)*
 covariance Λ_{h_p} defined in *eq.(16)*

Note: There is no proper density function associated with $\widehat{\mathbf{h}}_p$, as Λ_{h_p} is singular.

3.2.6 Numerically estimating Λ_{h_p}

Provided we know the noise variances (σ_x^2, σ_y^2), we can find an estimate $\widehat{\Lambda}_{h_p}$ of the true covariance Λ_{h_p} , by replacing true quantities (A_p^n, H_p^n) in *eq.(17)*, with observed or estimated quantities ($A_p^n + N_p^n, \widehat{H}_p^n$).

In the absence of noise statistics, we estimate σ_x^2, σ_y^2 from the statistics of the algebraic transfer error in view-2. The algebraic transfer error for the i^{th} point correspondence is defined as

$$\mathbf{r}_i' = \begin{bmatrix} x_i' + \varepsilon_{x_i}' \\ y_i' + \varepsilon_{y_i}' \end{bmatrix} - \underbrace{\frac{1}{\widehat{H}_p(3, :)} \begin{bmatrix} \widehat{H}_p(1, :)} \\ \widehat{H}_p(2, :)}_{\mathbf{y}_i'} x_i' \quad (18)$$

⁷ From *eq.(12)* we find that a perturbation $\delta_{h_p}^n$ in \mathbf{h}_p^n , induces the perturbation $\delta_h = (N_2^{-1} \otimes N_1^T) \delta_{h_p}^n$, in \mathbf{h} . The corresponding uncertainty is $\Lambda_h = E\{\delta_h \delta_h^T\} = (N_2^{-1} \otimes N_1^T) \Lambda_{h_p^n} (N_2^{-T} \otimes N_1)$

\mathbf{r}_i' is a random vector with 2 sources of uncertainty

1. noisy point correspondences in view-2
 2. error in point transfer using the noisy homography \widehat{H}_p
- Assuming the transfer errors to be gaussian distributed⁸,

$$\mathbf{r}_i' \sim \mathcal{N}(\mathbf{0}, \text{diag}([\sigma_x^2, \sigma_y^2]) + \Lambda_{\mathbf{y}_i'}), \quad \Lambda_{\mathbf{y}_i'} = J_{h_p}^i \Lambda_{h_p} J_{h_p}^{i T} \quad (19)$$

$$J_{h_p}^i = \frac{1}{H_p(3, :)} x_i' \begin{bmatrix} H_p(3, :)} x_i' & 0 & -H_p(1, :)} x_i' \\ H_p(3, :)} x_i' & 0 & -H_p(2, :)} x_i' \end{bmatrix} \otimes x_i'^T \quad (20)$$

Please refer to [7] for details on deriving $\Lambda_{\mathbf{y}_i'}$. It can be shown that the matrix $\mathbf{R} = [\mathbf{r}_1' \dots \mathbf{r}_M']$ formed by stacking the algebraic transfer errors for the M point correspondences, has the following statistics

$$E\{\mathbf{R}\mathbf{R}^T\} = \sum_{i=1}^M E\{\mathbf{r}_i' \mathbf{r}_i'^T\} = M \begin{bmatrix} \sigma_x^2 & 0 \\ 0 & \sigma_y^2 \end{bmatrix} + \sum_{i=1}^M J_{h_p}^i \Lambda_{h_p} J_{h_p}^{i T}$$

$$= M \begin{bmatrix} \sigma_x^2 & 0 \\ 0 & \sigma_y^2 \end{bmatrix} + \sum_{i=1}^M \left[\sigma_x^2 J_{h_p}^i \Lambda_{h_p}^x J_{h_p}^{i T} + \sigma_y^2 J_{h_p}^i \Lambda_{h_p}^y J_{h_p}^{i T} \right] \quad (21)$$

The left-hand-side of *eq.(21)* can be estimated directly from the data. We find estimates $\widehat{\sigma}_x^2, \widehat{\sigma}_y^2$ of the true noise variances σ_x^2, σ_y^2 , by solving the linear system of *eq.(21)*, after these substitutions: $\widehat{J}_{h_p}^i, \widehat{\Lambda}_{h_p}^x, \widehat{\Lambda}_{h_p}^y \rightarrow J_{h_p}^i, \Lambda_{h_p}^x, \Lambda_{h_p}^y$.

3.2.7 Refining the homography \widehat{H}_p

We now examine how knowledge of $\widehat{\Lambda}_{h_p}$ could be used to refine \widehat{H}_p . Suppose, \mathcal{S} is the set of noisy correspondences in the two views, and \mathcal{S}_{h_p} is the subset with homography \widehat{H}_p . We refine \widehat{H}_p by using correspondences in $\mathcal{S} \setminus \mathcal{S}_{h_p}$ ⁹ that are consistent with \widehat{H}_p , while discarding inconsistent correspondences in \mathcal{S}_{h_p} . We declare the i^{th} point correspondence as consistent only when the Mahalanobis norm of its algebraic transfer error given by $d_{alg}[i]$ is < 5.9915 .

$$d_{alg}[i] = \mathbf{r}_i'^T \left(\begin{bmatrix} \widehat{\sigma}_x^2 & 0 \\ 0 & \widehat{\sigma}_y^2 \end{bmatrix} + \widehat{J}_{h_p} \widehat{\Lambda}_{h_p} \widehat{J}_{h_p}^T \right) \mathbf{r}_i' \quad (22)$$

where $\widehat{\sigma}_x^2, \widehat{\sigma}_y^2$ are the estimated noise variances, and \widehat{J}_{h_p} is defined in *eq.(20)*. The number 5.9915 specifies the Probability(misclassifying an inconsistent correspondence) $< 5\%$ [7]. This follows from the fact that $d_{alg}[i]$ is a χ_2^2 distributed random variable with 2 d.o.f. The proposed thresholding scheme is based on the thresholding scheme for sampson distances, utilized in estimating homographies using RANSAC[7].

The success of the proposed inpainting scheme depends on the ability to reliably estimate the epipolar geometry from noisy correspondences. So we examine the effect of noise on estimating the epipolar line constraint.

⁸ The gaussian assumption breaks down in the presence of outliers.

⁹The set difference $A \setminus B$ is defined as $A \setminus B = \{x : x \in A \ \& \ x \notin B\}$

3.3. Epipolar Geometry

Given two views of a static scene, the *epipolar line constraint*[7] describes a generic relationship between corresponding points¹⁰ in the two views. Suppose $\underline{x}_j = [x_j, y_j, 1]^T$ and $\underline{x}'_j = [x'_j, y'_j, 1]^T$ are corresponding points¹⁰ in *views-1* & *2* respectively. They satisfy the *epipolar line constraint*

$$\underbrace{\underline{x}'_j{}^T F \underline{x}_j}_{l_{j2}} = 0 \quad , \quad \underbrace{\underline{x}_j{}^T F^T \underline{x}'_j}_{l_{j1}} = 0 \quad (23)$$

where F is a singular matrix depending on the relative position and orientation of the cameras in the two views. The matrix F is referred to as the fundamental matrix. The *epipolar line constraint* derives its name from the fact that it describes a point-line incidence relation¹¹. The entities l_{j1}, l_{j2} are referred to as conjugate epipolar lines.

If \underline{x}_j is a missing pixel, its corresponding point in *view-2* must satisfy the *epipolar line constraint* of eq.(23). We exploit this knowledge to inpaint the missing pixel \underline{x}_j .

3.3.1 Estimating the Fundamental matrix

Given a planar facet that induces a homography H_p between the two views and the epipole¹² \underline{e}' in *view-2*, the fundamental matrix F is specified as

$$F = [\underline{e}']_x H_p \quad (24)$$

The epipole \underline{e}' is the intersection of all *epipolar lines* in *view-2*. As mentioned before, if $\underline{x}_j, \underline{x}'_j$ are corresponding points in the two views, \underline{x}'_j lies on the line $l_{j2} = F \underline{x}_j$. It can be shown[7] that the point $\tilde{\underline{x}}'_j \sim H_p \underline{x}_j$ ¹³, also lies on the line l_{j2} . This means that the line joining \underline{x}'_j and $\tilde{\underline{x}}'_j$ is a fragment of the epipolar line l_{j2} . Given two such fragments we can compute \underline{e}' , and the full epipolar geometry follows from eq.(24). The vector joining $\tilde{\underline{x}}'_j$ and \underline{x}'_j is referred to as the *parallax vector*. Consequently, this approach to estimating F is referred to as the *plane+parallax* method.

To reliably estimate the fundamental matrix using the *plane+parallax* method, we need to identify atleast two parallax vectors with little-or-no uncertainty in their orientation. This is a challenging task given noisy point correspondences in the two views. For short parallax vectors (\underline{x}'_j & $\tilde{\underline{x}}'_j$ are close), even small perturbations in the end-points significantly alters the orientation of the parallax vector. This suggests that the length of a parallax vector could be used to determine its reliability. But, it is likely that grossly mismatched correspondences are mistaken

¹⁰images of the same 3D world point in the two views

¹¹ (x_j, y_j) lies on line l_{j1} in *view-1*, and (x'_j, y'_j) lies on line l_{j2} in *view-2*

¹²The epipole \underline{e}' is the image of the 1st-pinhole as seen by *camera-2*.

¹³ $\tilde{\underline{x}}'_j$ represents the point \underline{x}_j transferred by the homography H_p

for reliable parallax vectors. So, we assume that grossly mismatched correspondences have been removed prior to inpainting. The sources of error affecting the orientation of a parallax vector are

1. noise in the correspondence \underline{x}'_j
2. noise in $\tilde{\underline{x}}'_j$ due to noise in the estimated homography \widehat{H}_p

We now examine how uncertainty in estimating H_p can be combined with the length of a parallax vector, to identify the two most reliable *parallax vectors*.

3.3.2 Reliability of parallax vectors

Given a noisy homography \widehat{H}_p , an estimate of its uncertainty $\widehat{\Lambda}_{h_p}$, noise variances $\widehat{\sigma}_x^2, \widehat{\sigma}_y^2$, and a noisy point correspondence $\{(x_j, y_j), (x'_j + \epsilon'_{x_j}, y'_j + \epsilon'_{y_j})\}$ that is *inconsistent*¹⁴ with the homography, we can construct an epipolar line segment by joining $\left(\frac{\widehat{H}_p(1,:)x_j}{\widehat{H}_p(3,:)x_j}, \frac{\widehat{H}_p(2,:)x_j}{\widehat{H}_p(3,:)x_j}\right)$ and $(x'_j + \epsilon'_{x_j}, y'_j + \epsilon'_{y_j})$. The length of this line segment is

$$d_j = \left\| \begin{bmatrix} x'_j + \epsilon'_{x_j} \\ y'_j + \epsilon'_{y_j} \end{bmatrix} - \frac{1}{\widehat{H}_p(3,:)x_j} \begin{bmatrix} \widehat{H}_p(1,:)x_j \\ \widehat{H}_p(2,:)x_j \end{bmatrix} \right\|^2 = \|\mathbf{r}_j\|^2 \quad (25)$$

It is easy to observe that d_j is a random variable. Using the noise model of Sec-3.2.2 and the covariance structure $(\Lambda_{y'_j})$ of \mathbf{y}'_j defined in eqs.(19) and (20), we can prove that

$$E\{d_j\} = \tilde{d}_j + \text{trace} \left(\begin{bmatrix} \sigma_x^2 & 0 \\ 0 & \sigma_y^2 \end{bmatrix} \right) + \text{trace}(\Lambda_{y'_j})$$

$$\tilde{d}_j = \left\| \begin{bmatrix} x'_j \\ y'_j \end{bmatrix} - \frac{1}{H_p(3,:)x_j} \begin{bmatrix} H_p(1,:)x_j \\ H_p(2,:)x_j \end{bmatrix} \right\|^2 \quad (26)$$

where \tilde{d}_j is the true length of the epipolar line segment. We define the reliability of a parallax vector as the ratio $\tilde{d}_j^{-1}(\sigma_x^2 + \sigma_y^2 + \text{trace}(\Lambda_{y'_j}))$. The ratio is minimum when the noise is minimum and the length of the line segment is maximum. Given 2 or more correspondences that are inconsistent with the homography \widehat{H}_p , we find the two *parallax vectors* with the highest reliability in the sense of

$$d_j^{-1} \left(\widehat{\sigma}_x^2 + \widehat{\sigma}_y^2 + \text{trace}(\widehat{J}_{h_p} \widehat{\Lambda}_{h_p} \widehat{J}_{h_p}^T) \right) \quad (27)$$

where d_j is the observed length of the epipolar line segment, defined in eq.(25). The chosen *parallax vectors* are intersected to find the epipole \underline{e}' in *view-2*. The full epipolar geometry follows from eq.(24).

4. Putting it all together

Fig.(1) summarizes the process of inpainting in the proposed approach. The user first selects images corresponding to *views-1* & *2*. We then proceed to identify putative correspondences in the two views (*colored dots in Figs1(a),1(d)*).

¹⁴ $d_{alg}[j] > 5.9915$, where $d_{alg}[j]$ is defined in eq.(22)

Next, the user outlines a hypothetical plane \mathcal{R} & the regions to be inpainted $\{\Omega_k\}$, in view-1. The hypothetical plane outlined in blue in Fig 1(a), is any region that includes correspondences from a planar facet. Using the set of correspondences $\in \mathcal{R}$, we estimate a homography \widehat{H}_p , its associated uncertainty $\widehat{\Lambda}_{h_p}$, and noise variances $\widehat{\sigma}_x^2, \widehat{\sigma}_y^2$. Using the strategy of Sec-3.2.7, we refine the homography \widehat{H}_p until the ratio $\frac{\widehat{\sigma}_x^2 + \widehat{\sigma}_y^2}{\#\text{consistent correspondences}}$ is a minimum. The result of the refinement is shown in Figs. 1(a), 1(d). The cyan points in Figs. 1(a), 1(d) are consistent correspondences, while the yellow & green points are the inconsistent correspondences $\in \mathcal{R}$ & $\notin \mathcal{R}$ respectively. We use the refined homography \widehat{H}_p to determine the two most reliable parallax vectors, as outlined in Sec-3.3.2. The best set of parallax vectors labeled 1 and 2 in Fig. 1(d) are intersected to find the epipole e' in view-2. The full epipolar geometry follows from eq.(24). If there is more than one planar facet in the scene, the best inpainting is achieved for the planar facet that most directly faces camera-2, as described in Sec-3.2.4.

4.1. Inpainting using the ‘‘epipolar line constraint’’

From Sec-3.3, we know that if x_m is a missing pixel in view-1, its corresponding point in view-2 satisfies the epipolar line constraint of eq.(23). Suppose $l_{m1} = [\text{null}(F)]_x x_m$, and $l_{m2} = Fx_m$ are conjugate epipolar lines in views-1 & 2, corresponding to x_m . In principle, we could inpaint x_m by searching for the best matching patch along the epipolar line l_{m2} , in view-2. But this is time consuming, computationally intensive, and not always correct. A simpler alternative to inferring missing pixels is globally stretching/shrinking l_{m2} until fragments of l_{m1} & l_{m2} are aligned.

We repeat the steps outlined below for each pair of conjugate epipolar lines $\{l_{m1}, l_{m2}\}$, in each region Ω_k .

1. Identify landmarks on l_{m1} & l_{m2} as points with the largest gradient magnitude along the epipolar line (\Rightarrow edge runs across the epipolar line)
2. Match the neighborhood around each landmark in view-1 with every landmark in view-2, using LP-NCC
3. Find at least 2 uniquely matching landmarks on l_{m1} & l_{m2}
4. Rotate l_{m1} & l_{m2} around the respective epipoles
5. Find the 1-d affine transform A_m that aligns l_{m1} & l_{m2} , using the coordinates of the matching landmarks
6. Transform the points $(x_m, y_m) \in (l_{m1} \cap \Omega_k)$ using A_m , to find points on l_{m2} with the missing intensities

Any error in identifying landmarks introduces errors in the 1-d affine transform and hence the inpainting. Failure to identify landmarks results in the inability to inpaint a specific conjugate line pair.

Contrary to our assumption, conjugate epipolar lines can be aligned using a locally affine transform. We are currently investigating the use of elastic registration[13] techniques to locally align a pair of conjugate epipolar lines.

5. Results

The effectiveness of the proposed approach is best demonstrated in Fig. 1, where we attempt to inpaint parts of a pillar like structure. It is difficult to accurately estimate the epipolar geometry in this example, as we have a limited number of correspondences, and short parallax vectors. A quick look at Fig. 1(d) reveals that the proposed inpainting scheme has succeeded in inpainting the pillar like structure. The quality of the inpainting can be judged by the accuracy with which corners on the roof, and vertical structures within the missing region have been reconstructed.

We provide two more examples of the proposed inpainting, wherein traditional image inpainting methods would have failed. In the first example (Figs. 2(a)-2(d)) we attempt to recover the missing detail in the bookshelf, while in the second example (Figs. 2(e)-2(h)) we inpaint a lamp & the window panes behind it. A quick look at Fig. 2(a), 2(h), reveals that we have successfully recovered the missing detail in the bookshelf of Fig. 2(a), and inpainted the lamp & the missing window panes in Fig. 2(h). Notice that we succeed in recovering the arch like structure on the window panes, in Fig. 2(h). But, a closer examination of Figs. 2(d), 2(h) reveals that the some of the objects are scaled incorrectly. This is a drawback of assuming that a pair of conjugate epipolar lines can be aligned globally using a 1-d affine transform.

6. Conclusions

We extend the scope of image inpainting by using the epipolar line constraint to infer the missing pixels from an exemplar image. We overcome the difficulty in accurately estimating the epipolar geometry from noisy point correspondences, by using the proposed structure of uncertainty in a homography and the proposed measure of the reliability of a parallax vector.

- The quality of the inpainting can be improved further, by
1. using elastic matching to align conjugate epipolar lines
 2. non-linear refinement[7] of the estimated homography
 3. using additional views to fully inpaint the occluding objects.

Appendix-1. Perturbation in DLT homography

Using first-order perturbation analysis[6] we can express the perturbation in the DLT homography (δ_{h_p}) in terms of the unperturbed singular values & right singular vectors $\{\sigma_j, \mathbf{v}_j\}_{j=1}^9$ of A_p .

$$\delta_{h_p} = - \sum_{j=1}^8 \left(\frac{\mathbf{v}_j \mathbf{v}_j^T}{\sigma_j^2} [A_p^T N_p \mathbf{h}_p + N_p^T A_p \mathbf{h}_p] \right) \quad (A1)$$

We simplify eq.(A1) by noting that $A_p \mathbf{h}_p = \mathbf{0}$, and

- o $A_p^T = \sum_{k=1}^9 \sigma_k \mathbf{v}_k \mathbf{u}_k^T$ singular value decomposition of A_p
- o $\mathbf{v}_m^T \mathbf{v}_n = \delta[m-n]$ orthogonality of right singular vectors

$$\delta_{h_p} = - \left(\sum_{j=1}^8 \frac{\mathbf{v}_j \mathbf{u}_j^T}{\sigma_j} \right) N_p \mathbf{h}_p = -A_p^\dagger N_p \mathbf{h}_p \quad (A2)$$

where A_p^\dagger is the pseudoinverse [6] of the matrix A_p . Using eq.(4) we can express the term $N_p \mathbf{h}_p$ as

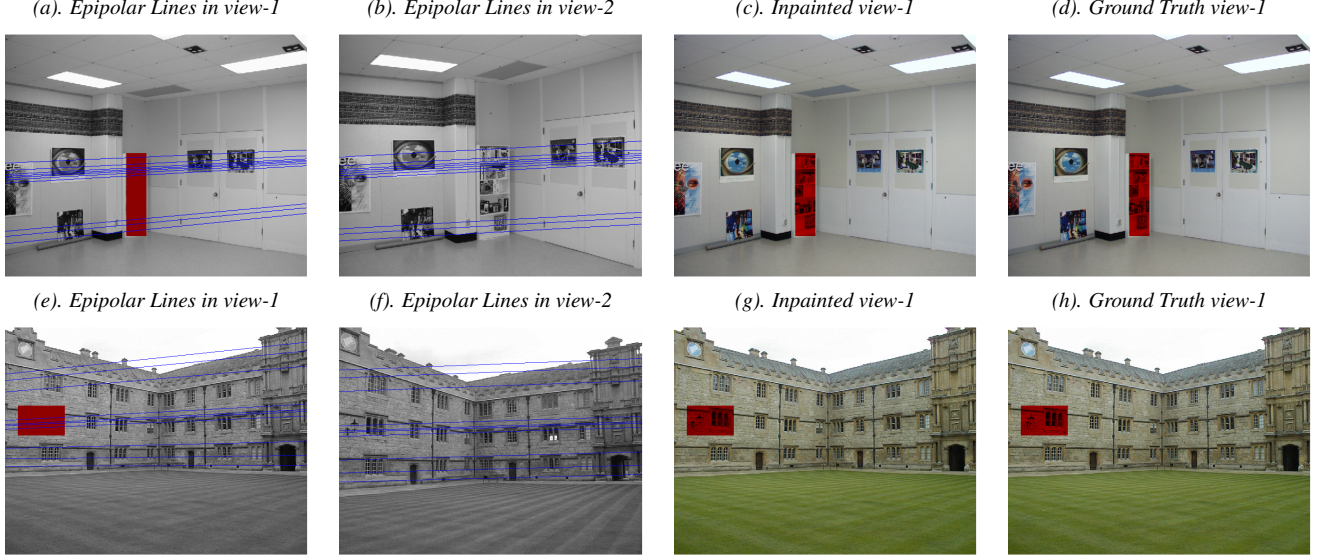


Figure 2. Image Inpainting for an indoor scene (courtesy NRC-IIT VIT Group), and an outdoor scene (courtesy Oxford Univ. VGG)

$$N_p \mathbf{h}_p = \begin{bmatrix} \left\{ \begin{bmatrix} 0 & 0 & -\varepsilon'_{y1} \\ 0 & 0 & -\varepsilon'_{y1} \end{bmatrix} \otimes \mathbf{x}_1^T \right\} \text{vec}(H_p^T) \\ \dots \\ \left\{ \begin{bmatrix} 0 & 0 & -\varepsilon'_{ym} \\ 0 & 0 & -\varepsilon'_{ym} \end{bmatrix} \otimes \mathbf{x}_M^T \right\} \text{vec}(H_p^T) \end{bmatrix} \quad (\text{A3})$$

Using the property $\text{vec}(A B C) = (C^T \otimes A) \text{vec}(B)$, we find that

$$N_p \mathbf{h}_p = \begin{bmatrix} \text{vec} \left(\mathbf{x}_1^T H_p^T \begin{bmatrix} 0 & 0 & -\varepsilon'_{y1} \\ 0 & 0 & -\varepsilon'_{y1} \end{bmatrix}^T \right) \\ \dots \\ \text{vec} \left(\mathbf{x}_M^T H_p^T \begin{bmatrix} 0 & 0 & -\varepsilon'_{ym} \\ 0 & 0 & -\varepsilon'_{ym} \end{bmatrix}^T \right) \end{bmatrix} = (W \otimes I_2) \Xi \quad (\text{A4})$$

$$W = \text{diag}([H_p(3,:)\mathbf{x}_1 \dots H_p(3,:)\mathbf{x}_M]), \Xi = [\varepsilon'_{x1} \ \varepsilon'_{y1} \ \dots \ \varepsilon'_{xm} \ \varepsilon'_{ym}]^T$$

Using eq.(A4) in eq.(A2), we can express δ_{h_p} in compact form as

$$\delta_{h_p} = A_p^\dagger (W \otimes I_2) \Xi \quad (\text{A5})$$

The uncertainty in the DLT homography is estimated as

$$\Lambda_{h_p} = E\{\delta_{h_p} \delta_{h_p}^T\} = A_p^\dagger (W \otimes I_2) E\{\Xi \Xi^T\} (W^T \otimes I_2) A_p^{\dagger T} \quad (\text{A6})$$

Using the noise model of 3.2.2, we can show that

$$E\{\Xi \Xi^T\} = I_M \otimes \text{diag}([\sigma_x^2, \sigma_y^2]) \quad (\text{A7})$$

Substituting eq.(A7) in eq.(A6), and using the property $(A \otimes B)(C \otimes D) = (A \otimes C) \otimes (B \otimes D)$, we obtain the expression for Λ_{h_p} (eq.(7)) in compact form.

Appendix-2. Uncertainty in Normalized DLT

From eq.(13), the perturbation in the normalized homography is $\delta_{h_p^n} = A_p^{\dagger n} (W^n \otimes I_2) \Xi^n$. The corresponding uncertainty $\Lambda_{h_p^n}$ is

$$\Lambda_{h_p^n} = E\{\delta_{h_p^n} \delta_{h_p^n}^T\} = A_p^{\dagger n} (W^n \otimes I_2) E\{\Xi^n \Xi^{nT}\} (W^{nT} \otimes I_2) A_p^{\dagger n T} \quad (\text{A8})$$

To model the statistics of the noise vector Ξ^n , we need to examine the effect of normalizing the noise terms $\{(\varepsilon'_{xi}, \varepsilon'_{yi})\}$. From eq.(10),

$$\begin{bmatrix} \varepsilon'_{xi} \\ \varepsilon'_{yi} \end{bmatrix} = \begin{bmatrix} N_2(1,1) & N_2(1,2) \\ N_2(2,1) & N_2(2,2) \end{bmatrix} \begin{bmatrix} \varepsilon'_{xi} \\ \varepsilon'_{yi} \end{bmatrix} = N_{2S} \begin{bmatrix} \varepsilon'_{xi} \\ \varepsilon'_{yi} \end{bmatrix} \quad (\text{A9})$$

$$\Rightarrow \Xi^n = [\varepsilon'_{x1} \ \varepsilon'_{y1} \ \dots \ \varepsilon'_{xm} \ \varepsilon'_{ym}]^T = (I_M \otimes N_{2S}) \Xi \quad (\text{A10})$$

Using eqs.(A7,A10) in eq.(A8) and simplifying, we obtain the expression (eq.(14)) for $\Lambda_{h_p^n}$ in compact form.

References

- [1] M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester. Image inpainting. *Siggraph 2000, Computer Graphics Proceedings*, pages 417–424, 2000. 1
- [2] T. Chan and J. Shen. Non-texture inpaintings by curvature-driven diffusions. *J. Visual Communication and Image Representation*, 12:436–449, 2001. 1
- [3] P. Chen and D. Suter. Homography estimation and heteroscedastic noise - a first order perturbation analysis, Technical Report MECSE-32-2005. 2
- [4] A. Criminisi. Modelling and using uncertainties in video metrology, Technical Report, 1997. 2, 3
- [5] A. Criminisi. Object removal by exemplar-based inpainting. In *IEEE CVPR*, volume 2, pages 721–728, 2003. 1
- [6] G. H. Golub and C. F. V. Loan. *Matrix Computations*. Third edition, 1996. 3, 4, 7
- [7] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2, 3, 4, 5, 6, 7
- [8] H. V. Henderson and S. R. Searle. The vec-permutation matrix, the vec operator and Kronecker products: A review. *J. Linear Multilinear Algebra*, 9:271–288, 1981. 3
- [9] V. Jain and P. Narayanan. Video Completion for Indoor Scenes. In *ICCVGIP*, pages 409–420, 2006. 1
- [10] K. Kanatani and N. Ohta. Accuracy bounds and optimal computation of homography for image mosaicing applications. In *IEEE ICCV*, volume 1, pages 73–78, 1999. 2, 4
- [11] S. Kang, T. Chan, and S. Soatto. Inpainting from Multiple View, UCLA CAM Report 02-31, 2002. 1
- [12] S. Osher, G. Sapiro, L. Vese, and M. Bertalmio. Simultaneous structure and texture image inpainting. In *IEEE CVPR*, volume 2, pages 707–712, 2003. 1
- [13] S. Periaswamy and H. Farid. Elastic registration with partial data. In *International Workshop on Biomedical Image Registration*, 2003. 7
- [14] E. Rosten and T. Drummond. Machine learning for high-speed corner detection. In *European Conf. on Computer Vision*, volume 1, pages 430–443, May 2006. 1
- [15] Y. Wexler, E. Shechtman, and M. Irani. Space-time completion of video. *IEEE Trans. PAMI*, 29(3):463–476, March 2007. 1
- [16] S. Zokai and G. Wolberg. Image registration using log-polar mappings for recovery of large-scale similarity and projective transformations. *IEEE Trans. on Image Processing*, 14(10):1422–1434, 2005. 1