Vision, Video, and Graphics (2003)
P. Hall, P. Willis (Editors)

Models from Image Triplets using Epipolar Gradient Features

Étienne Vincent and Robert Laganière

School of Information Technology and Engineering, University of Ottawa, Ottawa, Canada, K1N 6N5
evincent,laganier@site.uottawa.ca

Abstract

In an application where sparse matching of feature points is used towards fast scene reconstruction, the choice of the type of features to be matched has an important impact on the quality of the resulting model. In this work, a method is presented for quickly and reliably selecting and matching points from three views of a scene. The selected points are based on epipolar gradients, and consist in stable image features relevant to reconstruction. Then, the selected points are matched using edge transfer, a measure of geometric consistency for point triplets and the edges on which they lie. This matching scheme is invariant to image deformations due to changes in viewpoint. Models drawn from matches obtained by the proposed technique are shown to demonstrate its usefulness.

1. Introduction

Several applications require scene reconstruction from a few fixed cameras. To performed this task, the camera system must first be calibrated, and then, some features must be matched between the different views. One approach is to use feature points, i.e. some interest points are selected in each view, and matching is then restricted to these points only. In applications such as telerobotics, this feature point selection and matching process must also be rather efficient to allow a 3D model to be continuously updated. Also, the reliability of the matches is crucial, as the integrity of the task would be compromised by wrong information that mismatches provide.

In the context of applications where a three dimensional model of a changing environment must be continuously updated, matches between points in images produced by at least two cameras are needed. From these matches, the position in space of the points in the scene can be estimated by backprojection. In order to minimize the refreshment rate of the model, only a limited number of points can be matched. This would be done by first selecting special feature points in the images, then attempting matching only those points using some similarity measure between their neighborhoods.

When it comes to fast sparse matching in general, the most common approach is based on matching Harris feature points (see Section 3.1), using some correlation-based measure[11, 16, 20]. We have previously presented a system for fast calibrated matching based on this approach[18]. However, although Harris feature points are relatively stable and fast to compute, it was found that on arbitrary scenes, not chosen to contain suitable textural content, the detected feature points might not be well distributed or stable enough to allow accurate reconstruction. Harris features have proven very useful in tasks such as calibration. However, when distribution is more important, they can be insufficient. Often, few matching Harris corners will lie on the borders or on portions of many significant scene elements.

This problem is addressed here. An alternative feature detector is presented which is fast and often results in more and better distributed matches in a calibrated system of cameras. This detector relies on the concept of epipolar gradients, introduced here. A simple way to match feature points is also presented which works on a calibrated triplet of images and is invariant to deformations due to viewpoint variations.

The next section quickly reviews some features of trinocular geometry. Then, Section 3 discusses the problem of feature points detection. Next, Section 4 describes how feature points are matched. Section 5 describes a constraint which weeds out the possible remaining false matches. Section 6 shows some experimental results of the proposed matching scheme, and Section 7 presents models constructed from the matched feature points. Section 8 is a conclusion.

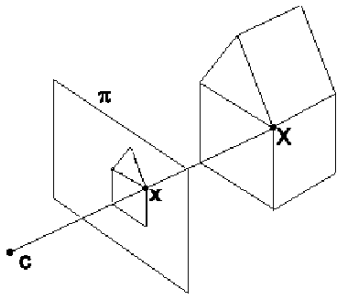

Figure 1: *Pinhole camera model.*

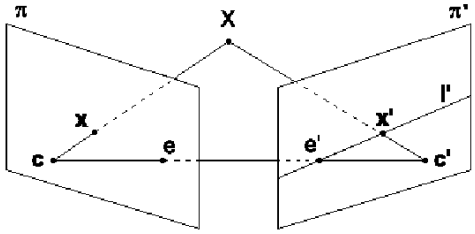

Figure 2: *Two-view geometry.*

2. Trinocular Geometry

A simple *pinhole model* can be used to represent cameras[5] (see Figure 1). A point in space $\mathbf{X}$ is projected onto the image plane π, to a point $\mathbf{x}$, which is at the intersection with the ray joining $\mathbf{X}$ and the camera's focal point $\mathbf{c}$.

When two cameras look at the same scene, the projection $\mathbf{x}$, on one camera plane π, of an unknown point in space $\mathbf{X}$, can tell us something about where the point will land on the other camera plane (see Figure 2). More precisely, it is known that a matching point in the second image, must be on the *epipolar line*.

Since $\mathbf{X}$ must be somewhere along the ray defined by $\overline{\mathbf{c}x}$, its projection on π' must be on the projection of that ray. This projection, $\mathbf{l}'$ is the epipolar line of $\mathbf{x}$, and thus points in one image are related to lines in the other through the image pair's epipolar geometry.

For a pair of cameras, the relationship between points in one view, and their epipolar line in the other, is called the *epipolar geometry*. It can be represented by a 3×3 matrix F, of rank 2: its *fundamental matrix*. Points $\mathbf{x}$ in the first image, are related to their epipolar lines $\mathbf{l}'$ in the second image, by:

$$F\mathbf{x} = \mathbf{l}' \tag{1}$$

where the point $\mathbf{x}$ is represented with homogeneous coordinates as $(x,y,1)^\top$, and the epipolar line is the set of points $\mathbf{x}'$, represented in homogeneous coordinates, such that $\mathbf{l}'^\top\mathbf{x}' = 0$. This matrix can be estimated from known pairs of

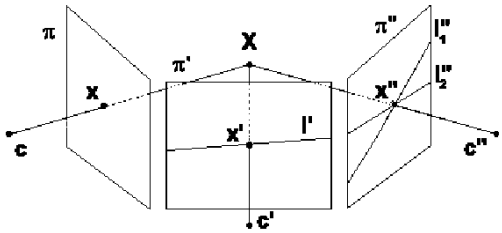

Figure 3: *Three-view geometry.*

points between images produced by the two cameras using the fact that:

$$\mathbf{x}'^\top F\mathbf{x} = 0 \tag{2}$$

The 8-point algorithm[4], which was used to estimate fundamental matrices is based on this equation, and on a normalization of point coordinates. Alternatively, more accurate estimates can be obtained through the non-linear minimization of more geometrically significant quantities such as reprojection error[5].

If three cameras are used, and a match $(\mathbf{x},\mathbf{x}')$ is already known between the first two images, the position of the matching point $\mathbf{x}''$ in the third image can be determined exactly. Indeed, the projection of $\mathbf{X}$ on π'' should be at the intersection of the epipolar lines $\mathbf{l}_1{}''$ and $\mathbf{l}_2{}''$ of $\mathbf{x}$ and $\mathbf{x}'$ respectively (see Figure 3).

This relationship between points in three images is called *trinocular geometry*[12]. In fact, it can be shown that the position of $\mathbf{x}''$ can be determined even if the lines $\mathbf{l}_1{}''$ and $\mathbf{l}_2{}''$ are the same and thus have no single point of intersection. Trinocular geometry can be represented by the *trifocal tensor*, a $3\times 3\times 3$ tensor T for which:

$$\mathbf{x}''_k = \sum_{i,j\in\{1,2,3\}} \mathbf{x}_i \mathbf{l}'^{\perp}_j T_{ijk} \tag{3}$$

where $\mathbf{l}'^{\perp}$ is the line going through $\mathbf{x}'$, and perpendicular to $\mathbf{l}'$, the epipolar line of $\mathbf{x}$, and where i,j,k, and l are indices of the vectors and tensor. This formula can be used to compute the position of $\mathbf{x}''$, but our experiments have shown that more stable results can be obtained using:

$$\mathbf{x}''_l = \mathbf{x}'_i \sum_{k=1}^{3} \mathbf{x}_k T_{kjl} - \mathbf{x}'_j \sum_{k=1}^{3} \mathbf{x}_k T_{kil} \tag{4}$$

which defines 9 trilinearities for $i,j\in\{1,2,3\}$, 4 of which are linearly independent. Image point $\mathbf{x}''$ can then be estimated by solving the over-constrained system of equations.

Trifocal geometry is a powerful tool for matching. Indeed, for a point $\mathbf{x}$ in a first image, the search for correspondence is restricted to the line $\mathbf{l}'$ in the second image. And once the correspondence $(\mathbf{x},\mathbf{x}')$ is known, the matching point in a third

image is directly determined by the trifocal tensor. In practice, a third image can therefore be used to validate matches between the first two. Once a match is found, it is verified by *transferring* it to the third image using equation 4, and evaluating a similarity measure between the transferred point and the first two[11, 16, 18]. The fact that a transferred image point is visually very similar to the putatively matched points in the other two images is indeed a good indication of the validity of that match.

Trifocal tensors can also relate lines between images. Thus, when the equation of an image line is known in two views, it can be *transferred* to another one using:

$$\mathbf{l}_i = \sum_{j,k \in \{1,2,3\}} \mathbf{l}'_j \mathbf{l}''_k T_{ijk} \tag{5}$$

This result will be used in Section 4 to compute the expected slope of the tangent to an edge in the third image.

3. Feature Point Selection

In many applications, it would be too costly to compare points in the first image with all points along their epipolar lines in the search for matches. Thus, matching is often limited to selected feature points. Feature points in the first image are only compared to feature points along their epipolar line in the second image. Note, however, that there is no need for selecting feature points in the third image, as the trinocular geometry limits the search there to a single point neighborhood.

Good feature points are those that are likely to be easily distinguishable from each other, and which can be identified robustly, with respect to changes in viewpoint. These points should also, as much as possible, represent significant scene features, such as points on the borders of scene objects, as they will be the points used in the constructed model.

3.1. Harris Feature Points

The most commonly used feature points are Harris corners[3]. These correspond to high curvature points on image edges. This feature point detector finds points where the image intensity gradient has a high magnitude in more than one direction using the gradient's autocorrelation matrix:

$$C(x,y) = S * (\bigtriangledown I \cdot \bigtriangledown I^T) = S * \begin{bmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{bmatrix} \tag{6}$$

where S is a smoothing operator. At the point where it is computed, this matrix's greatest eigenvalue corresponds to the image's rate of change in the direction of highest variation, while its smallest eigenvalue corresponds to the rate of change in the perpendicular direction. If the smallest eigenvalue has a high magnitude, it means that, at the considered point, the image has a high rate of variation in at least two directions, and thus, that the point is in a high curvature region.

These features have been shown to be relatively stable[15]. However, it was found that in practice, it can be dependent on the presence of appropriate textural content. When this is not the case, few corresponding points might be detected in both images. Furthermore, the corners are often not distributed well enough to permit a reasonable reconstruction of the complete scene from matched points. Epipolar gradient features overcome these problems.

3.2. Epipolar Gradient Feature Points

The idea behind epipolar gradient features is to select points where the image's intensity gradient is locally perpendicular to epipolar lines. For most regular camera configurations, these feature's corresponding point in other views would also lie on edges which are nearly perpendicular to epipolar lines. Thus epipolar gradient features in one image should have corresponding points in other images which are also epipolar gradient features. This stability makes them good candidates for matching. Additionally, being points on epipolar lines which cross strong edges, these features can be accurately localized and should be easily discernable from other points on the same epipolar line, making matching less ambiguous. Finally, since they lie on important image edges, such points are often found on the border of significant scene features and are thus important for scene reconstruction.

More formally, let $\mathbf{I}$ and $\mathbf{I}'$ be two images, with $\mathbf{x}$ a point on $\mathbf{I}$, and $\mathbf{x}'$ its corresponding point on $\mathbf{I}'$. Then, it is clear from Figure 2 that $\mathbf{x}'$ will lie on $\mathbf{l}'$, the epipolar line of $\mathbf{x}$ in $\mathbf{I}'$. Similarly, $\mathbf{x}$ will lie on $\mathbf{l}$, the epipolar line of $\mathbf{x}'$ in $\mathbf{I}$. Now $\mathbf{l}$ and $\mathbf{l}'$ should also correspond so all points on $\mathbf{l}$ will have their corresponding point lying on $\mathbf{l}'$. Thus if $\mathbf{X}$, the world point projected onto $\mathbf{x}$ and $\mathbf{x}'$, is centered on a locally planar surface, the points on $\mathbf{l}$ that are immediately next to $\mathbf{x}$ should correspond to points in $\mathbf{I}'$ that lie on $\mathbf{l}'$ and are immediately next to $\mathbf{x}'$. Thus, the intensity gradient of $\mathbf{I}$, at $\mathbf{x}$, in the direction of $\mathbf{l}$, should be similar to the intensity gradient of $\mathbf{I}'$, at $\mathbf{x}'$, in the direction of $\mathbf{l}'$.

The intensity gradient in the direction of the epipolar line will be referred to as the *epipolar gradient*. It can be computed by projecting $\bigtriangledown I(\mathbf{x})$ onto the epipolar line $\mathbf{l} = (l_1, l_2, l_3)$, giving the explicit formula:

$$\bigtriangledown_{ep}(\mathbf{x}) = \frac{\bigtriangledown I(\mathbf{x}) \cdot (\frac{-l_3}{l_1}, \frac{l_3}{l_2})}{\|(\frac{-l_3}{l_1}, \frac{l_3}{l_2})\|} \tag{7}$$

where $\mathbf{l}$ can be obtained from $\mathbf{x}$ and F using an arbitrary line $\mathbf{k}'$ not going through the second image's epipole as:

$$\mathbf{l} = F^T \mathbf{k}' \times F\mathbf{x} \tag{8}$$

Thus, in a pair of images for which the epipolar geometry is known, a point having a high epipolar gradient in one image should have a high epipolar gradient in the other as well. Of

Figure 4: *Detected Harris feature points.*

course, the stability can be limited by the importance of the changes in angles between epipolar lines and image edges caused by changes in viewpoint. Nevertheless, a moderate change in this angle will not significantly reduce epipolar gradients. Imposing a threshold on the magnitude of these epipolar gradients will lead to a set features that are reasonably robust to change in viewpoint. Furthermore, these points will be found on strong image edges with orientations perpendicular to epipolar lines. This is a desirable property, as there is an ambiguity in attempting to match other points, such as those that are in low contrast areas, or on contours which are oriented along epipolar lines.

Figure 5 shows the detected epipolar gradient feature points on one image of a test pair (obtained from the model house image sequence available at http://www.robots.ox.ac.uk/ vgg/data/). In one image, points may be detected only on every few lines to limit their number. Figure 4 presents the detected Harris corners on that same image.

It can be seen, when epipolar gradient feature points are compared to the same number of Harris feature points, that the former are more evenly distributed among the different scene surfaces. Harris features are mostly concentrated on the house's front wall, while epipolar gradient features are often found on the boundaries between different surfaces which should allow them to quickly yield a more complete reconstruction of scene objects.

4. Matching based on Edge Transfer

Now that feature points suitable for matching have been selected, these points must be matched. A common way of comparing potentially matching points is variance normalized correlation. Such a correlation based approach can give good results when the difference between viewpoints is limited, but will not be an accurate measure of similarity in the case of more widely separated views. Then, a measure which is invariant to the reprojection deformation of the area around feature points is needed.

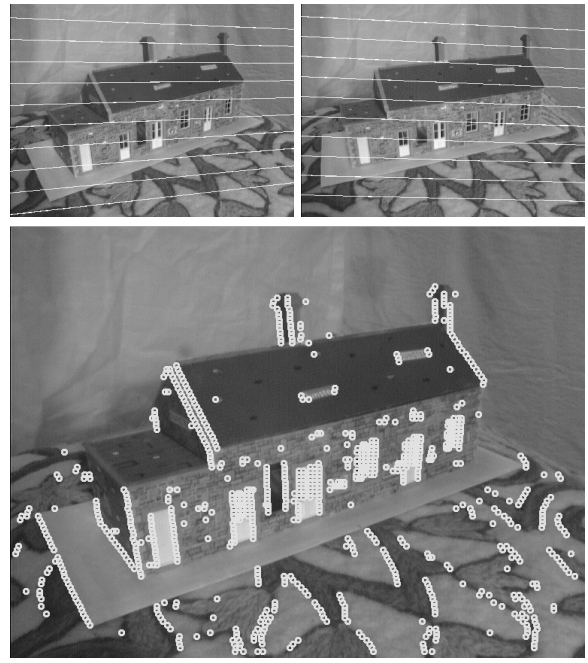

Figure 5: *Detected epipolar gradient feature points, and the image pair's epipolar geometry.*

Many such measures, invariant to rotation or affine transformations of point neighborhoods have been proposed[1, 10, 13, 14, 17]. These are usually computationally expensive and inappropriate for calibrated matching as they do not exploit the camera system's geometry. Since here, matching is guided by the system's trifocal geometry, points only have a few candidate matches, so a more general, but less discriminating comparison measure can be used.

The similarity measure presented here is based on consistency of edge orientations between the views. Several authors have proposed to impose a bound on edge orientations between views to constrain matching[2, 7]. However, such a constraint is only satisfied in cases of small changes in viewpoint. Horaud and Monga[6] have also presented an orientation constraint which measures the consistency of the change in angle with the change in viewpoint, but requires the camera projection matrices.

Two simple descriptors are used, together with a similarity measure defined between them. The most important descriptor is based on the transfer of lines perpendicular to intensity gradients from the first two images to the third one. Similarity measures which use gradient directions have been used

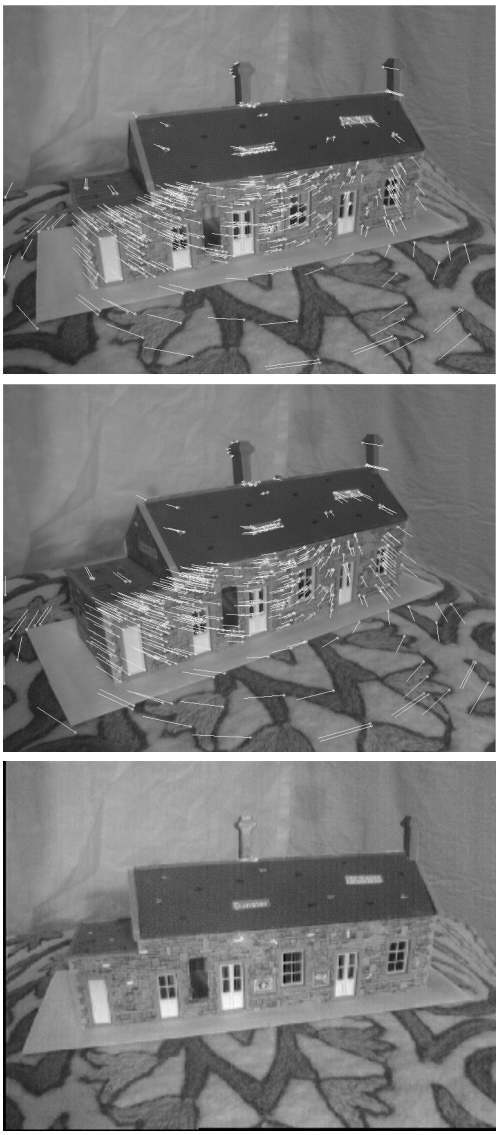

Figure 6: *Matches from Harris features and correlation, with disparities between images 1-2 (top), and 2-3 (middle).*

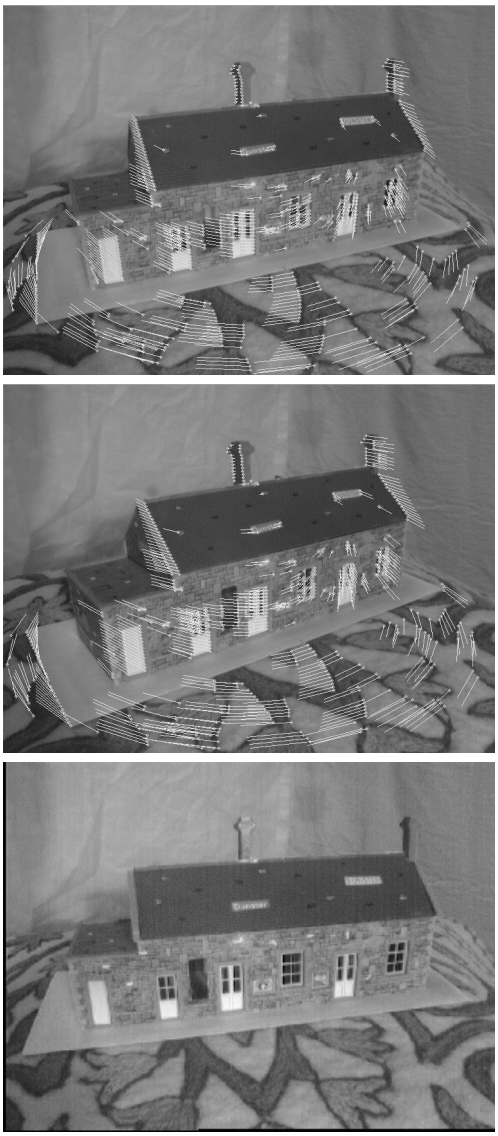

Figure 7: *Matches from proposed method, with disparities between images 1-2 (top), and 2-3 (middle).*

by many authors[8, 9]. These methods will usually align gradients before further comparing feature point neighborhoods, thus making the process invariant to rotation. Using trifocal geometry together with the intensity gradients, a higher degree of invariance can be achieved.

The edge transfer similarity measure is based on the fact that, using the edge orientation computed at corresponding points in two images, the orientation of the edge at the corresponding point in a third image can be computed using equation (5). This is because the lines going through the points and tangent to the edges going through these points (the lines perpendicular to the intensity gradients at the points) should correspond. Thus, a measure of similarity between three points would be the difference between the orientation of the tangent to the edge of one of the points, and the orientation of the line obtained by transferring the tangent to the edges of the two other points.

The other descriptor is based on the image intensity values in the area around the point. The average intensities on each side of the tangent to the edge going through the point are considered. These values should be preserved in different views of the same point taken simultaneously. First, it

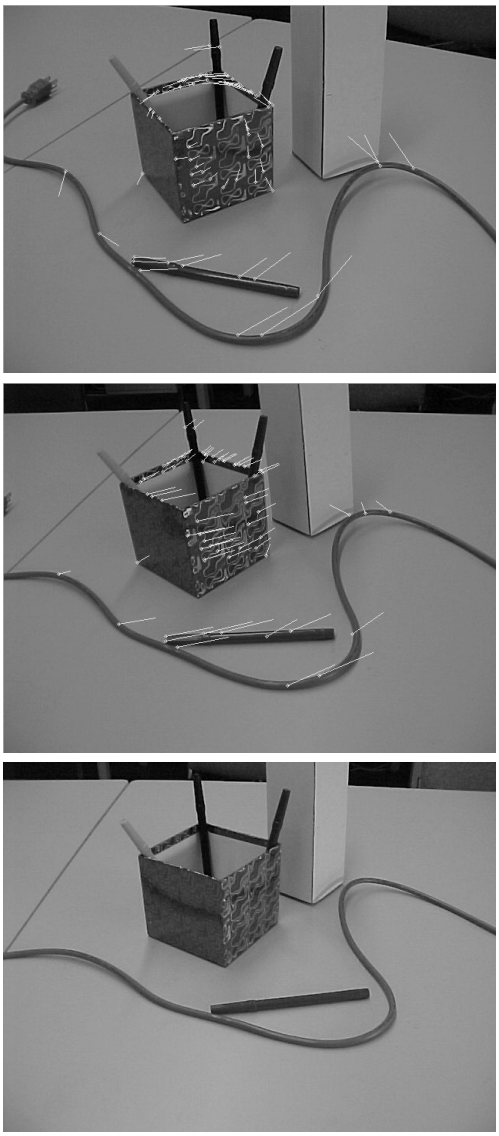

Figure 8: *Matches from Harris features and correlation, with disparities between images 1-2 (top), and 2-3 (middle).*

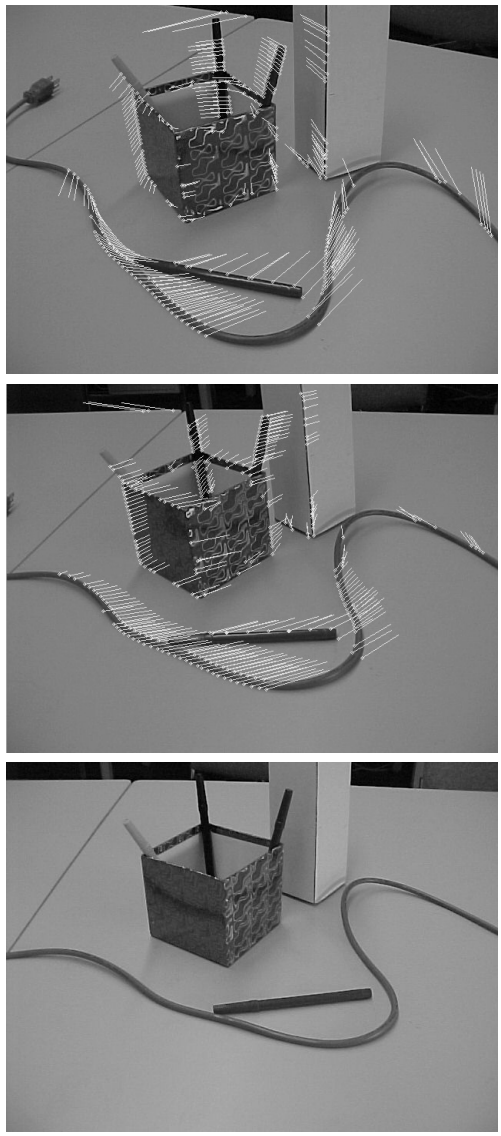

Figure 9: *Matches from proposed method, with disparities between images 1-2 (top), and 2-3 (middle).*

is determined which side of the edge corresponds to which in the other image. Then, the measure of similarity is taken as the difference between the two average intensities of the most different corresponding sides.

Let $\Delta I(\mathbf{x},\mathbf{x}',\mathbf{x}'')$ be the maximum difference between the intensities of $\mathbf{x}$,$\mathbf{x}'$ or $\mathbf{x}''$, and $\Delta\theta(\mathbf{x},\mathbf{x}',\mathbf{x}'')$ be the difference between the gradient orientation at $\mathbf{x}''$ measured in $\mathbf{I}''$ and computed from the gradients at $\mathbf{x}$ and $\mathbf{x}'$. Then the chosen similarity measure between $\mathbf{x}$, $\mathbf{x}'$ and $\mathbf{x}''$ will be:

$$s(\mathbf{x},\mathbf{x}',\mathbf{x}'') = \max\left(\frac{\Delta I(\mathbf{x},\mathbf{x}',\mathbf{x}'')}{\sigma_{\Delta\theta}}, \frac{\Delta\theta(\mathbf{x},\mathbf{x}',\mathbf{x}'')}{\sigma_{\Delta I}}\right) \quad (9)$$

where $\sigma_{\Delta I}$ and $\sigma_{\Delta\theta}$, the standard deviations of the descriptors, are used to normalize the descriptors to a similar range. This measure will have a low value when the points correspond.

5. Disparity Consistency Constraint

Sometimes, the similarity measure presented in the previous section might not be discriminating enough. Consequently, even when the search for matches is guided by the trinocular geometry, mismatches can be expected. However, mismatches are very undesirable when the goal is re-

construction. Fortunately, when many matches are identified throughout the images, and mismatches are relatively few, they can generally be eliminated by enforcing that close matches have similar disparities.

To this end, a disparity gradient constraint is applied[19]. The disparity gradient is a measure of the compatibility of two matches. It is essentially the norm of the difference of the disparities, normalized by the distance between the matches. For two pairs $(\mathbf{x},\mathbf{x}')$ and $(\mathbf{y},\mathbf{y}')$, having disparities $d(\mathbf{x},\mathbf{x}')$ and $d(\mathbf{y},\mathbf{y}')$ respectively, the cyclopean separation, $d_{cs}(\mathbf{x},\mathbf{x}';\mathbf{y},\mathbf{y}')$ is the distance between the midpoint of the disparity vectors, and their disparity gradient is defined as:

$$\Delta d(\mathbf{x},\mathbf{x}';\mathbf{y},\mathbf{y}') = \frac{|d(\mathbf{x},\mathbf{x}') - d(\mathbf{y},\mathbf{y}')|}{|d_{cs}(\mathbf{x},\mathbf{x}';\mathbf{y},\mathbf{y}')|} \qquad (10)$$

A pair is considered a mismatch when its disparity gradients with many of its closest neighbors are too high (It was usually required that a correspondence have a low disparity gradient with two out of its three nearest neighbors). Since there are three images, the disparities between two pairs of images are actually compared. This procedure eliminates false matches as long as they are not surrounded only by similar false matches, an unlikely situation.

6. Experimental Results

Figure 7 shows the result of applying the proposed matching scheme to an image triplet. The disparities between the first and second images are shown (drawn on the first image), as well as the disparities between the second and third (drawn on the second image). Thus, the lines in the first image join the coordinates of feature points there, to their corresponding coordinates in the second image, and similarly for the lines in the second image with respect to the third one.

Figure 6 shows the disparities obtained when a Harris detector and correlation are used instead. The same number of feature points were used in both experiments, and the thresholds relevant to the matching process were chosen empirically to maximize the resulting number of matches. It can be seen that the method based on epipolar gradients obtained more matches (601 versus 423), and provides scene features which are more relevant to scene reconstruction. The matches obtained through the Harris detector are mostly located on the front wall of the house, while the matches obtained using the epipolar gradient features are distributed more evenly among the different surfaces, and often lie on the borders between them.

Figures 9 and 8 also show matches found using the proposed approach and the Harris/correlation approach respectively for simple images of a few objects. Disparities between the images are also shown. Here, 273 matches were found using the proposed method. With the same number of feature points, the Harris detector with correlation only found 64, and it was not possible to modify the thresholds to accept more matches without introducing a significant number of mismatches. The success of the proposed method, in contrast to the Harris/correlation approach can be attributed to the fact that the scene objects contain few clear corners, and little textural information, but still enough significant edges to permit their detection as epipolar gradient features.

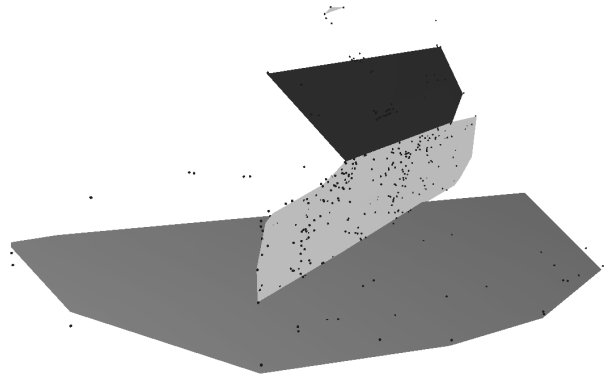

Figure 10: *A model constructed from the matches shown in Figure 6, found using Harris features.*

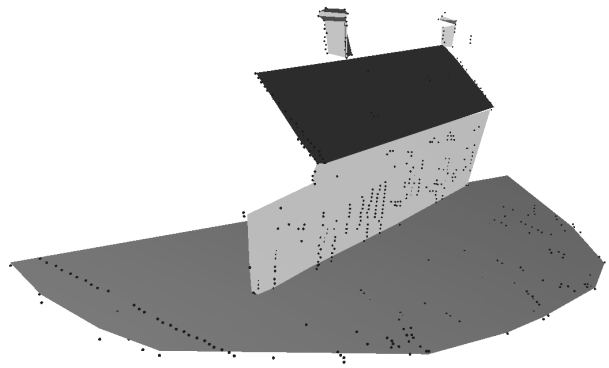

Figure 11: *A model constructed from the matches shown in Figure 7, found using epipolar gradient features.*

7. Reconstruction

To demonstrate the usefulness of the proposed matching approach in fast model building, the correspondences shown in Figures 6 and 7 were used, together with the known camera calibrations, to construct models of the scene. The position of scene points in space were computed as the intersection of backprojected rays from the image points. Then, points from common planar surfaces were used to estimate the position of these points in space, and the section of these planes defined by the points were drawn in Figure 10 and 11. These figures also show, as black dots, the point positions computed from the matches.

It can be seen that the model generated from the proposed

matching method is far more representative of the scene, mainly as the points it is generated from contain more relevant information. These are distributed more evenly among scene objects, thus allowing for instance the drawing of some parts of the chimneys. They also define more precisely the border of objects, as we can see, for example, that the areas of the front wall and roof do not cover their entire true area in the other model.

8. Conclusion

In summary, two new techniques were introduced for fast and reliable calibrated sparse matching. A new feature was used based on epipolar gradients, and a new correspondence measure was introduced which relies on transferring edges.

These new techniques are improvements over other approaches. The features based on epipolar gradients are more stable, constitute features which are more relevant to the structure of scenes, and are usually well distributed over images. Matching based on edge direction is fast, and viewpoint independent. Beyond calibrated sparse matching, we believe that epipolar gradients and edge transfer are interesting concepts susceptible of finding other applications.

References

1. A. Baumberg, Reliable Feature Matching Across Widely Seperated Views, *Proc. of Int. Conf. on Computer Vision and Pattern Recognition*, **1**:774–781, 2000.
2. R. Deriche, Z. Zhang, Q.-T. Luong, O. Faugeras, Robust Recovery of the Epipolar Geometry for an Uncalibrated Stereo Rig, *Proc. of European Conf. on Computer Vision*, 567–576, 1994.
3. C. Harris, M. Stephens, A Combined Corner and Edge Detector, *Proc. of Alvey Vision Conf.*, 147–151, 1988.
4. R. Hartley, In Defense of the Eight-Point Algorithm, *IEEE trans. on Pattern Analysis and Machine Intelligence*, textbf19(6):580–593, 1997.
5. R. Hartley, A. Zisserman, Multiple View Geometry, *Cambridge University Press*, 2000.
6. R. Horaud, O. Monga, Vision par ordinateur, *Hermès*, 1995.
7. A. Lacey, N. Thacker, P. Courtney, S. Pollard, TINA 2001: The Closed Loop 3D Model Matcher, *Proc. of British Machine Vision Conf.*, 2001.
8. D. Lowe, Object Recognition from Local Scale-Invarient Features, *Proc. of Int. Conf. on Computer Vision*, 1150–1157, 1999.
9. K. Mikolajczyk, C. Schmid, Indexing based on Scale Invariant Interest Points, *Proc. of Int. Conf. on Computer Vision*, 525–531, 2001.
10. P. Montesinos, V. Gouet, R. Deriche, D. Pelé, Matching Color Uncalibrated Images Using Differential Invariants, *Image and Vision Computing*, **18**(9):659–671, 2000.
11. G. Roth, A. Whitehead, Using Projective Vision to find Camera Positions in an Image Sequence, *Proc. of Vision Interface*, 225–232, 2000.
12. A. Shashua, Trilinearity in Visual Recognition by Alignment, *Proc. of European Conf. on Computer Vision*, 479–484, 1994.
13. C. Schmid, Constructing Models for Content-Based Image Retrieval, *Proc. of Int. Conf. on Computer Vision and Paterrn Recognition*, **2**:39–45, 2001.
14. C. Schmid, R. Mohr, Local Grayvalue Invariants for Image Retrieval, *IEEE trans. on Pattern Analysis and Machine Intelligence*, **19**(5):530–535, 1997.
15. C. Schmid, R. Mohr, C. Bauckhage, Comparing and Evaluating Interest Points, *Proc. of Int. Conf. on Computer Vision*, 230–235, 1998.
16. P. Torr, A. Zisserman, Robust Computation and Parameterization of Multiple View Relations, *Proc. of Int. Conf. on Computer Vision*, 727–732, 1998.
17. T. Tuytelaars, L. Van Gool, L. D'haene, R. Koch, Matching of Affinely Invariant Regions for Visual Servoing, *Proc. of Int. Conf. on Robotics and Automation*, 1601–1606, 1999.
18. E. Vincent, R. Laganière, Matching Feature Points for Telerobotics, *Proc. of Int. Workshop on Haptic Virtual Environments and their Applications*, 13–18, 2002.
19. E. Vincent, R. Laganière, Matching Feature Points in Stereo Pairs: A Comparative Study of Some Matching Strategies, *Machine Graphics & Vision*, **10**(3):237–259, 2001.
20. Z. Zhang, R. Deriche, O. Faugeras, Q.-T. Luong, A Robust Technique for Matching Two Uncalibrated Images Through the Recovery of the Unknown Epipolar Geometry, *Artificial Intelligence*, **78**:87–119, 1995.