Australian Government

Department of Defence

Defence Science and
Technology Organisation

Video Metrology

Tony Scoleri

Intelligence, Surveillance and Reconnaissance Division

Defence Science and Technology Organisation

DSTO–TN–0823

ABSTRACT

This report describes a range of techniques to carry out real-world measurements from a single image of a scene. The methods have been incorporated in an interactive system where they were tested and compared. The performance of the system has been evaluated on various images for which the aim was to measure a person's height. Experiments revealed that the methods yielded accurate and consistent results, which matched the true person's height within a small percentage of error.

Video Metrology

Executive Summary

This report gives an overview of a range of techniques to carry out real-world measurements from a single image of a scene. The methods have been incorporated in an interactive system where they were tested and compared. The motivation for development of such a system was its great practical benefit in the context of video surveillance and complement to existing methods in the Analysts' Detection Support System (ADSS).

Given an ordinary image of a scene and no further details about the camera configuration, the methods developed are able to retrieve a minimum of critical information to interpret the underlying geometry of the scene. Knowledge of this geometry is subsequently used to measure the dimensions of objects of interest in the image.

The performance of the measuring system has been evaluated on various images for which the aim was to estimate a person's height. Experiments revealed that the methods executed rapidly and yielded accurate results, which matched the person's true height within a small percentage of error.

Author

Tony Scoleri
Intelligence, Surveillance and Reconnaissance Division

Tony Scoleri graduated with first class honours in Pure Mathematics and Computer Science at the University of Adelaide in 2001. He then started a PhD in the field of computer vision also at the University of Adelaide. His thesis presents an elaboration of specifically designed algorithms for solving a class of general constrained parameter estimation problems encountered in geometric computer vision applications and is currently submitted for examination. In 2007, he joined the Intelligence, Surveillance and Reconnaissance Division of the Defence Science and Technology Organisation. His research focuses on structure from motion problems, video security applications and the development of tracking algorithms.

Contents

Figures

1   Introduction

Images carry a significant amount of geometrical information about the scene or objects photographed. A challenge is to retrieve this information in a quantifiably accurate way. This report describes a hierarchy of techniques to obtain real-life measurements of a scene given one image of it. Technically, an image is a two-dimensional snapshot of the world taken by a camera at a given time. Alternatively, a camera can be seen as a genuine geometric device which constructs planar images of the three-dimensional world by a projection through its optical centre. Due to the camera projection mechanism, the input image (or *image plane*) generally suffers from a perspective effect which alters not only the dimensions but also the angles between elements in the real world. Measuring directly on the input image would not give the exact lengths in the scene. The term "metric" is often used to emphasise properties of distances and angles in the world to contrast with their perspective counterparts in the image.

The techniques described in this report aim at removing the perspective effect from the input image before any measurement is carried out. These techniques span situations where no metric information is known about the world (completely uncalibrated camera) through to cases where some reference distances are known but they are not sufficient for a complete camera calibration (partial calibration). This leads to extremely flexible algorithmns which can be applied to a wide range of images such as photographs of buildings and interiors, aerial images, archived images, photographs of crime scenes and even paintings.

The methods are based on minimal geometric information determined from the image and no camera specifications. This minimal information is based on geometric properties of parallel lines. Two preliminary stages are necessary before making actual measurements in the image. These stages rectify the input image by restoring geometric and metric properties such as parallelism, angles, lengths, or area ratios. Broadly speaking, the resulting image then has dimensions proportional to the real world scene. The material presented in the following sections outline different modules to rectify an image in the above sense. The system is interactive and permits ultimately to measure specific image targets.

2   Stratified metric reconstruction

Recovering metric properties of a 3-D structure or object can be stratified in two stages as first an affine and then metric upgrade. Each upgrade is responsible for recovering different aspects of the "metric" world. The length to be measured, or *target segment*, is assumed to belong to a particular plane in the scene (or *world plane*). The task is to compute a transformation from the image plane to the world plane in order to obtain the metric length of the target segment (in the world plane) from its projection in the image.

It turns out that this transformation can be decomposed into a product of two matrices P and A. The first of these matrices, P, depends on an entity called the *vanishing line* and the latter, A, depends on two parameters α and β. A stratified reconstruction of a perspective image consists in an affine upgrade, by calculating the vanishing line, followed by a metric upgrade, by finding α and β. This is detailed further in the next sections.

2.1 Affine upgrade

The image rectification process begins by an affine upgrade. This initial stage modifies the input image to restore properties such as parallelism, ratio of areas, or ratio of lengths on collinear or parallel lines. This requires determining the vanishing line to obtain transformation P. A minimum of two sets of parallel lines are sufficient for the task. Each set enables the computation of a *vanishing point* and two such points define the vanishing line. Figures 1(a) and 1(b) show an example of an image and its corresponding affine rectification.

(a) Input image.

(b) Affine rectified image.

Figure 1: Affine upgrade. Only a minor rectification (green area) was necessary to upgrade the input (perspective) image to an affine image.

Two sets of parallel world lines (Figure 2) are sufficient to compute the vanishing line of the ground plane (Figure 3). Knowledge of this line with respect to the ground plane is used subsequently to measure the man's height.

Figure 2: The two sets of lines selected to compute the vanishing line of the ground plane.

Figure 3: *Vanishing line of the ground plane. This line is a geometric invariant in the image and so can be used to measure the man's height.*

2.2 Metric upgrade

Recovery of metric geometry requires an affine transformation of the image plane, say A, that will restore angles and length of ratios for *non-parallel* line segments. This is perhaps the most difficult and error-sensitive part of the whole stratified reconstruction. Three general methods are available to estimate the affine parameters α and β entering matrix A. These methods again require identifying sets of lines, however, on the *affine* image this time. In each case the lines provide one constraint on α and β in the shape of a circle. A minimum of two distinct constraints are needed to find the intersection of their respective circles and solve for both parameters. A metric upgrade of the affine image can then be performed. A fourth method can also be used specifically for buildings. It is based on a different strategy than the previous three techniques. Details of the various methods are exposed next.

2.2.1 Known angle

One circular constraint can be obtained if the angle θ on the *world* plane is known for two lines seen in the image. It should be noted that the constraint depends on the line orientation. Any parallel line sets which define the same world angle produce the same constraint although they are different sets. An illustration is given in Figure 4. One constraint is obtained from the orthogonality of *any* red line pairs. Another constraint results from orthogonality of the green line pair which is not parallel to any of the three red line sets. Note that the parallelism of the tiles' boundary lines is restored but the relative lengths are visibly not correct.

2.2.2 Equal unknown angles

Another constraint can be computed by identifying two sets of lines for which the angle on the *world* plane between the lines is the same. The true value of the angle need not be known.

(a)   (b)

Figure 4: Known angle constraint. Floor tile (a) input image; (b) affine rectified image with orthogonal line pairs identified.

For instance, referring to Figure 4(b), any set of red lines with the green lines provide enough information to derive such a constraint. Note that in this situation the angle between the lines is known to be orthogonal, however the formula for the equal angles constraint generates a different circle than the known angle one. An additional example is given in Figure 5 where the angle there is unknown.

Figure 5: Equal angle constraint. The repetition of structure in the window allows application of the equal angles constraint to the green line pairs.

2.2.3 Known length ratio

If the length ratio of two non-parallel segments is known on the *world* plane, then it is possible to obtain a third constraint on the affine parameters α and β. The requirement of *unity* length ratio

between the sides of the squares in Figure 4(b) can be used. Once two constraints are known at least, the intersection of the circles provide the values of the sought parameters, see Figure 6(a). The blue circle is generated from the unity ratio constraint whereas the red circle comes from the known angle constraint. Note that the circles intersect in two points differing by the value of the β-parameter. The proof of this result finds explanation in the underlying geometry which is beyond the scope of this report. Nevertheless, the solution chosen is the one where β has positive sign. The case where β is negative corresponds to a rotation of the image which bears no significance on its metric properties. The affine image is then metrically rectified as shown in Figure 6(b). Observe that the square tiles and circles are now proportional.
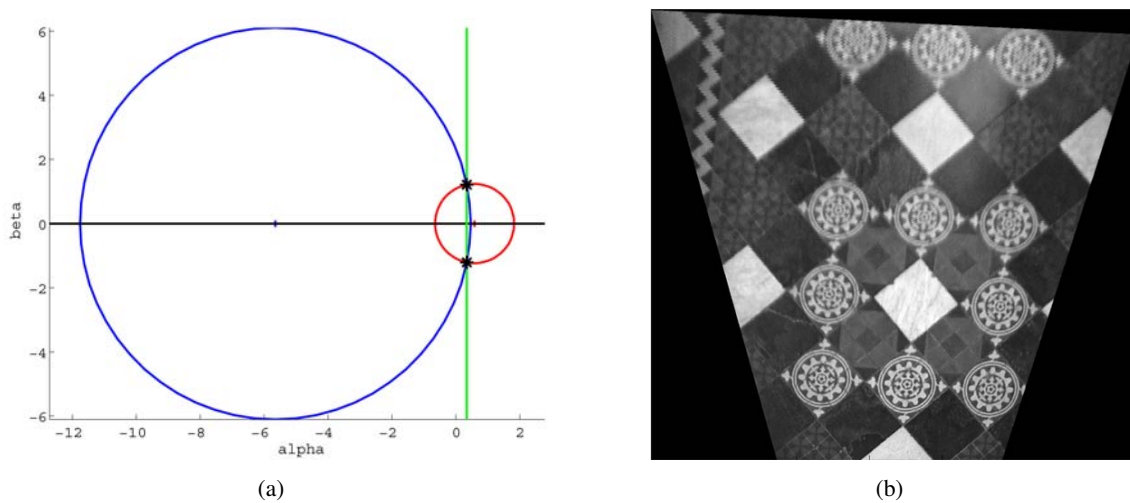

(a)	(b)

Figure 6: *Metric upgrade. (a) Intersection of two circles for the tile floor image giving $\alpha = 0.337798$ and $\beta = \pm 1.21025$. (b) metric rectified image.*

2.2.4 Building rectification

In the case of buildings, a specific method is available to upgrade the input image. The rectangular structure of a building facade presents orthogonal sets of parallel lines in the vertical and horizontal directions which can be used directly for metric reconstruction purposes. A transformation is computed so that the directions of the building are aligned with the horizontal and vertical axes of the image. This transformation ensures that the building is facing the viewer (fronto-parallel view). An example is given in Figure 7.

3 Measuring a target segment from a video image

After the image has been metrically rectified, it is possible to measure an image segment orthogonal to the ground plane. In a video security context, this could be measuring the height of a person as in Figure 8. One reference length in the scene is needed to obtain a measurement of the target segment. In the example of Figure 8, we have used the door's height. The person's height was computed from the image as 184.6cm with an uncertainty of about 2cm depending on the reference measurement. The true person's height is 185cm.

(a) Input image.

(b) Metric rectified image.

Figure 7: Building metric upgrade. The perspective effect of the original image is removed using orthogonal properties of the facade structure.

(a)

(b)

Figure 8: Image measurement from a surveillance camera. (a) Original image; (b) metrically rectified image with target segment.

The person's height in Figure 1(a) was also calculated and found to vary between 177.9cm and 182.3cm depending on the reference length used. The true height is 180cm.

4   Future research

A number of extensions to the current system are possible. Depending on the quality of the input image, a pre-processing stage could calculate and eliminate the radial distortion present in the image. This step would remove any lens distortion effect which complicates the selection of a reliable reference segment in the image. In the case of buildings, the current rectification method does not permit to recover the exact scaling of the structure. This drawback can be remedied using an unstratified calibration technique.

The methods implemented in this work make no assumptions about the camera calibration. If some metadata about the camera system are available, for instance if the focal length is known, this information can be incorporated in the computational mechanism to improve the reliability of the measuring system. Since data and transformations are affected by errors, so is the output mea-

surement. A proper treatment of error propagation and quantification would be of great additional value to the system. An uncertainty estimate could be associated with each output measurement. In this report, the uncertainty was calculated using different reference distances in the image, which is helpful but not ideal. Given a particular reference length, the task would be to quantify the error in the final measurement resulting from its use. Lastly, the current system operates interactively to select the various lines in the image. The line selection process could be fully automated.

5   Acknowledgements

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION DOCUMENT CONTROL DATA		1. CAVEAT/PRIVACY MARKING	
2. TITLE Video Metrology		3. SECURITY CLASSIFICATION Document (U) Title (U) Abstract (U)	
4. AUTHOR Tony Scoleri		5. CORPORATE AUTHOR Defence Science and Technology Organisation PO Box 1500 Edinburgh, South Australia 5111, Australia	
6a. DSTO NUMBER DSTO–TN–0823	6b. AR NUMBER 014-194	6c. TYPE OF REPORT Technical Note	7. DOCUMENT DATE June, 2008

8. FILE NUMBER 2008/1062265/1	9. TASK NUMBER NS 07/114	10. SPONSOR CCT sponsor	11. No OF PAGES 7	12. No OF REFS 0

13. URL OF ELECTRONIC VERSION http://www.dsto.defence.gov.au/corporate/ reports/DSTO–TN–0823.pdf	14. RELEASE AUTHORITY Chief, Intelligence, Surveillance and Reconnaissance Division

15. SECONDARY RELEASE STATEMENT OF THIS DOCUMENT

Approved For Public Release

OVERSEAS ENQUIRIES OUTSIDE STATED LIMITATIONS SHOULD BE REFERRED THROUGH DOCUMENT EXCHANGE, PO BOX 1500, EDINBURGH, SOUTH AUSTRALIA 5111

16. DELIBERATE ANNOUNCEMENT

No Limitations

17. CITATION IN OTHER DOCUMENTS

No Limitations

18. DSTO RESEARCH LIBRARY THESAURUS

video metrology
stratified metric reconstruction
video surveillance

19. ABSTRACT

This report describes a range of techniques to carry out real-world measurements from a single image of a scene. The methods have been incorporated in an interactive system where they were tested and compared. The performance of the system has been evaluated on various images for which the aim was to measure a person's height. Experiments revealed that the methods yielded accurate and consistent results, which matched the true person's height within a small percentage of error.