Defence Research and Development Canada
Recherche et développement pour la défense Canada

Depth Maps from Simulated Stereo Images Using Kernel Matching

S.A. Barton and D.J. Mackay
DRDC Suffield

Technical Memorandum
DRDC Suffield TM 2006-198
November 2006

Canada

Depth Maps from Simulated Stereo Images Using Kernel Matching

S.A. Barton and D.J. Mackay
DRDC Suffield

Defence R&D Canada – Suffield

Technical Memorandum

DRDC Suffield TM 2006-198

November 2006

Author

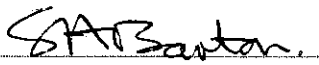

S.A. Barton

Approved by

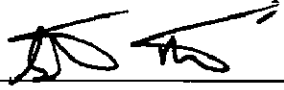

Per D.M. Hanna
A/Head/AISS

Approved for release by

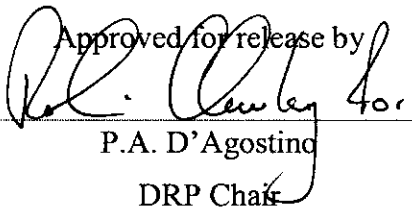

P.A. D'Agostino
DRP Chair

Abstract

Simulated stereo images and depth maps were generated using the open-source ray tracing program POV-Ray, operating in an environment containing obstacles. The image pairs were used to generate depth maps by finding the relative positions of matching groups (kernels) of image pixels. Kernel matching generated an average depth accuracy of better than 90% over all images and pixels tested. When neighbourhood average and median filter smoothing techniques were applied to the resulting depth maps, the average accuracy exceeded 95%. Occlusion and lighting differences between the images were attenuated by smoothing and kernel normalization. Computation times for the simulated depth maps indicate that kernel matching could be used in real video imagery on a vehicle moving at 30 km/hr.

Résumé

Des images stéréoscopiques simulées et des cartes de profondeur ont été générées au moyen du programme de lancer de rayons de source non scellée, POV-Ray, opérant dans un milieu contenant des obstacles. On a utilisé des paires d'images pour générer des cartes de profondeur en trouvant les positions relatives des groupes correspondants (noyaux de convolution) des pixels des images. La correspondance des noyaux de convolution a généré une moyenne d'exactitude de profondeur supérieure à 90% pour toutes les images et tous les pixels testés. Quand les techniques de lissage par filtres des moyennes et des médianes de voisinage ont été appliquées aux cartes de profondeur résultantes, la moyenne d'exactitude dépassait les 95%. Les différences d'occlusion et de luminosité entre les images ont été atténuées par le lissage et la normalisation des noyaux de convolution. Les durées de calcul des cartes de profondeur simulées indiquent que la correspondance des noyaux de convolution pourrait être utilisée en imagerie vidéo réelle sur un véhicule se déplaçant à une vitesse de 30 km/h.

This page intentionally left blank.

Executive summary

Depth Maps from Simulated Stereo Images Using Kernel Matching

S.A. Barton and D.J. Mackay; DRDC Suffield TM 2006-198; Defence R&D Canada – Suffield; November 2006.

Navigation of autonomous unmanned ground vehicles is part of the R&D program at DRDC Suffield. This requires distance measurements in the form of depth maps. This report shows that simulated stereo image pairs may be used to generate depth maps by finding the relative positions of matching groups (kernels) of image pixels in the two images. The images and depth maps were generated using the open-source ray tracing program POV-Ray operating in a simulated environment containing obstacles.

Kernel matching does not require explicit feature identification in the images, it is easy to implement and performs well in generating depth maps for the simulated environment. The average depth accuracy over all images and pixels tested is better than 90%. When neighbourhood average and median filter smoothing techniques are applied to the resulting depth maps, the average accuracy exceeds 95%. Occlusion, where a region may be visible in one image but not the other, is unavoidable and will introduce local error in a depth map produced by kernel matching. Lighting differences between the images may be minimized by kernel normalization. Both occlusion errors and lighting differences are attenuated by the application of the smoothing techniques.

Computation times for the simulated depth maps indicate that kernel matching could be used in real video imagery on a moving vehicle.

Sommaire

Depth Maps from Simulated Stereo Images Using Kernel Matching

S.A. Barton and D.J. Mackay; DRDC Suffield TM 2006-198; R & D pour la défense Canada – Suffield; novembre 2006.

La navigation des véhicules terrestres autonomes sans équipage fait partie d'un programme de R et D à RDDC Suffield. Elle requiert de mesurer des distances sous forme de cartes de profondeur. Ce rapport montre que les paires d'images stéréoscopiques simulées pourraient être utilisées pour générer des cartes de profondeur en trouvant les positions relatives des groupes correspondants (noyaux de convolution) des pixels d'images dans les deux images. On a généré les images et les cartes de profondeur au moyen d'un programme de lancer de rayons de source non scellée, POV-Ray, opérant dans un milieu contenant des obstacles.

La correspondance des noyaux de convolution n'exige pas l'identification explicite des caractéristiques dans les images; elle est facile à implémenter et fonctionne bien pour générer des cartes de profondeur d'un environnement simulé. L'exactitude moyenne de profondeur parmi toutes les images et tous les pixels testés est supérieure à 90%. Quand les techniques de lissage par filtres des moyennes et des médianes de voisinage sont appliquées aux cartes de profondeur résultantes, la moyenne d'exactitude dépasse les 95%. L'occlusion, quand une région est visible dans une image mais pas dans l'autre, est inévitable et introduira des erreurs locales sur une carte de profondeur produite par la correspondance des noyaux de convolution. Les différences de luminosité entre les images peuvent être atténuées par la normalisation des noyaux de convolution. Les erreurs d'occlusion et les différences de luminosité sont toutes deux atténuées par l'application de techniques de lissage.

Les durées de calcul pour les cartes de profondeur simulées indiquent que la correspondance des noyaux de convolution pourrait être utilisée, dans le domaine de l'imagerie vidéo réelle, sur un véhicule en déplacement.

Table of contents

List of figures

1 Introduction

Autonomous navigation of an unmanned ground vehicle (UGV) is an essential component of the Autonomous Land Systems project at DRDC Suffield. Such navigation requires a suite of sensors that define the UGV state, position and the distances to nearby obstacles. Ideally, distances would be continuously available as a depth map for an extended field of view around the direction of motion. Depth maps may be obtained from sensors such as laser scanners and sonar transceivers, but the depth map construction is quite slow using lasers and difficult to interpret accurately using sonar. The speed at which a UGV could move through a complex real-world environment would therefore be quite limited using those techniques.

Stereo images in the visible wavelengths offer an alternative basis for the rapid generation of extended depth maps. The difference in the positions of an image point in the separate views of a stereo system, together with the optical parameters of the system, can be used to derive the depth of the point; the basic equations are given in several elementary texts, e.g. [1], and modern techniques are described in a recent review [2]. Computational stereopsis is the term used to describe more general calculations of depth from stereo images. A recent discussion, including real-world issues, is given by Brown et al [3], and Hirschnuller [4] details a basic algorithm for performing stereopsis.

Any attempt to relate the position of a single pixel in the left image to a corresponding pixel in the right image relies on a computation involving the surrounding regions of both pixels. We use the term *kernel* to describe a precisely-defined region around a pixel. The relative location of a kernel in the right image that matches one in the left image may be used to estimate depth at the kernel centre. Some pixels may be partially occluded; i.e. visible in only one view. This is unavoidable, and may result in unmatched kernels and undefined depth points. Smoothing techniques, such as neighbourhood averaging and median filtering, may be applied to the resulting depth map to minimize this problem.

Lighting differences, shadows and reflections also introduce differences between the left and right images, which may be attenuated by image processing techniques; kernel normalization and cosine matching are two such methods that we have employed (Chapter 3). In the case where there are regions with no distinguishable features in both images, such as a smooth surface of a uniform colour, structured lighting has been used by others [5],[6] to greatly improve depth map accuracy.

This work describes the application of a kernel matching technique to a set of simulated stereo images and their corresponding depth maps. The images were generated from a simulated environment containing objects using the Persistence 0f Vision Raytracer (POV-Ray), an open-source software package; details are given in Chapter 2.

Chapter 3 describes the kernel matching process and how a set of calibration images (stereo image pairs and their associated depth map) may be used to assign depth values to the shift required to produce a match for each pixel. The matching process is then applied to a different set of left and right images (the test set) and a depth map is generated. The

application of smoothing techniques to the crude depth map is shown to greatly improve the accuracy of the final depth image.

There is more information in stereo images of a view than that available by only finding the shifts required for kernel or feature matching. The perspective seen when several objects or features are present in a view gives an indication of relative depth, and this coupled with the knowledge of the true size of one or more objects is used in human vision to create indications of real distance. This work was intended to create a baseline of performance against which other techniques could be judged. In particular it was thought that it might be possible to capture more information from the stereo images using recurrent neural networks, and that approach is an ongoing area of research. However, even with this simple kernel matching process, depth maps of practical use may be obtained, and Chapter 4 discusses these results and suggests how the technique may be applied to a real mobile vehicle.

2 Simulated Image Generation

The generation of stereo and depth map images using POV-Ray has been described in detail by Mackay [7].

The simulated environment that was constructed is shown from above in Figure 1. It can be seen to contain 15 three-dimensional rectangular objects. User-defined patterns and textures may be superimposed on the surface of each obstacle, as well as on the floor and wall surfaces.

A pair of cameras is moved along a path through the objects, and at specified points stereo images and depth maps are generated. Two paths for are shown in Figure 1, providing two slightly different data sets for training and testing a neural network. However, for the kernel shifting technique described here, only the data from the lower path at the bottom of Figure 1 was used as a test set.

Figure 2 shows a 160 x 120 pixel stereo image pair seen from one of the positions along the test path. The associated depth map calculated from the centre position of the two cameras is shown in Figure 3. The patterns that were applied to the floor, walls and objects are apparent.

A *calibration set* of left and right images and their associated depth maps was obtained by moving the optical system in a direction perpendicular to a flat wall whose surface was covered with a randomly patterned texture. An example of the left and right stereo images (160 x 120 pixels) is shown in Figure 4. The depth map is not shown since it is essentially a simple grey scale, because all distances to the wall are almost equal. The slight differences in distance between the central position and the edge positions are normally only one or two units in the distance scale.

A second set of wall images was also generated, with a different superimposed pattern. This was used to verify the distance correlation obtained with the calibration set, as described in Chapter 3. A pair of left and right images is shown in Figure 5, for the same distance as those shown in Figure 4.

The distances generated by POV-Ray are relative grey-scale values that would be between 0 and 255 (one byte per image pixel) if distances could be measured up to the camera plane of the optics. In practice, only distances greater than the optical focal length were measured, resulting in an effective distance range of about [140, 255]. The stereo images have intensity values in the full range, [0, 255] per pixel.

The fog technique used to generate the depth maps [7] produces values that are quite linear with distance from the camera plane up to about 10 POV-Ray distance units, where the fog technique saturates. Thus any distance above that is assigned a byte value of 255. POV-Ray distance units may be related to real world values by the focal length, field of view and horizontal aperture (in pixels) of the optical system. The values that were chosen

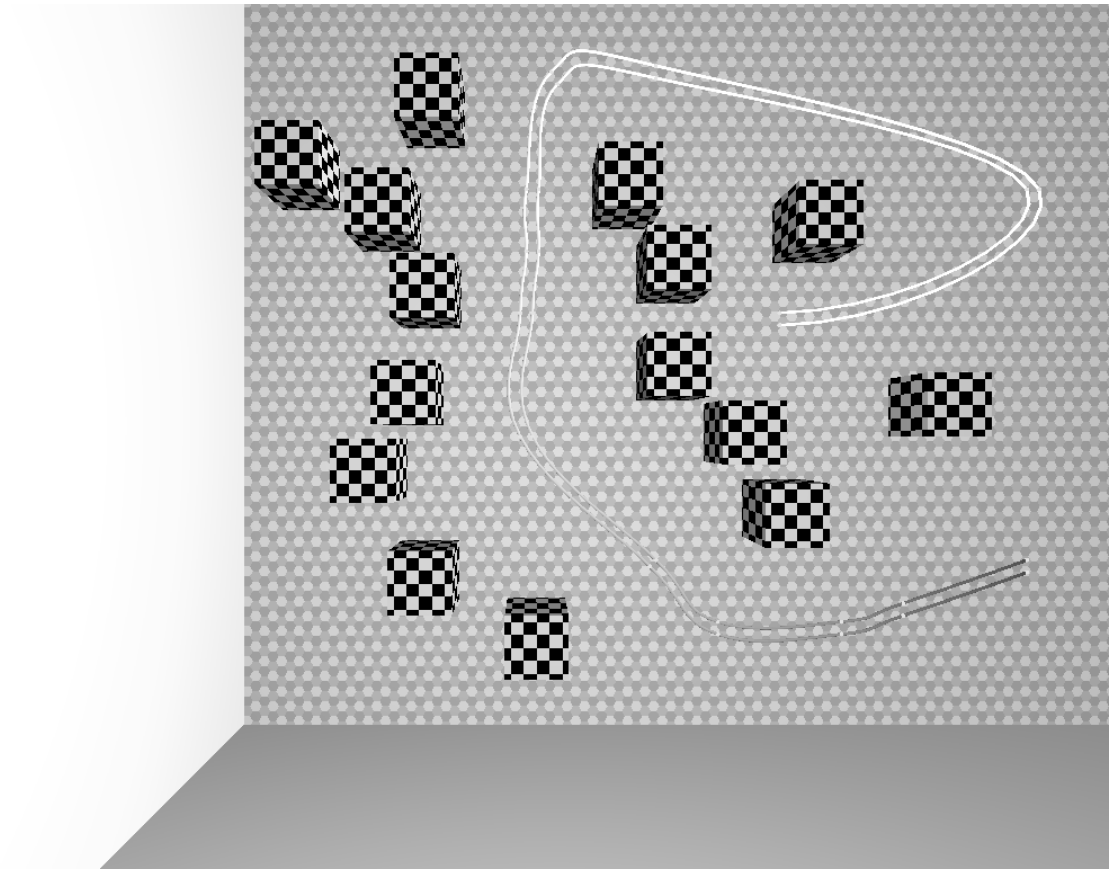

Figure 1: *Simulated environment showing objects and camera paths*

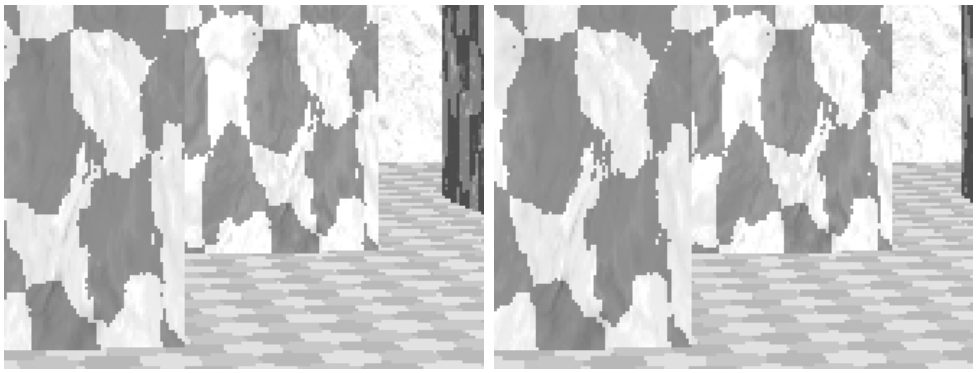

Figure 2: *Left and right views at one point in the lower path*

Figure 3: *Depth map for Figure 2*

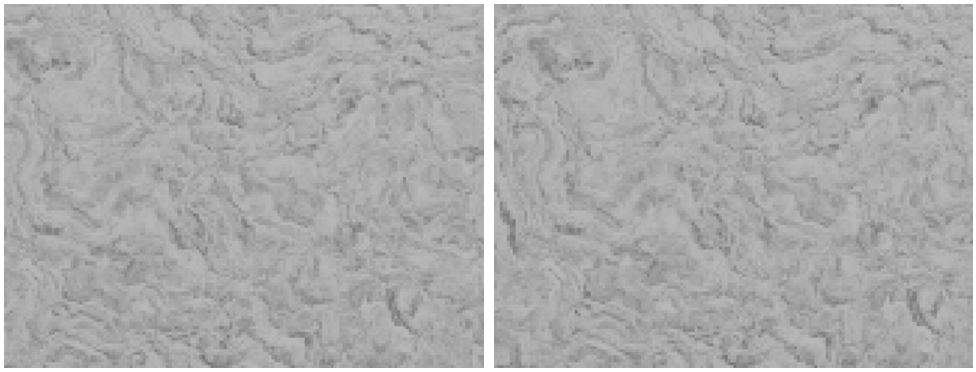

Figure 4: *Left and right views of wall (calibration pattern)*

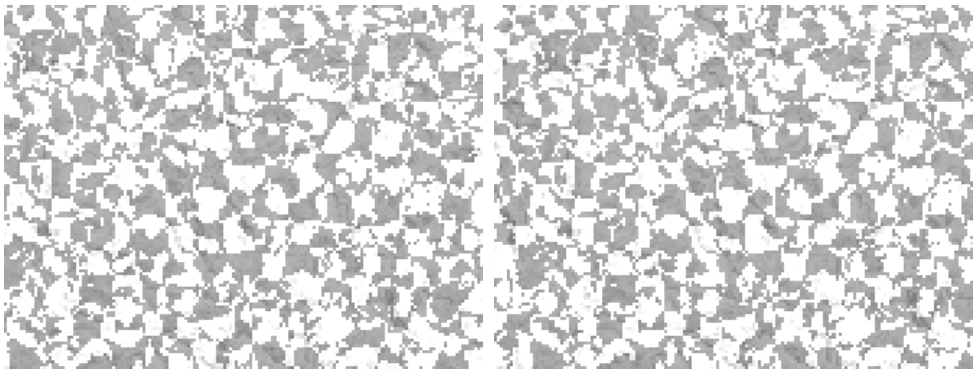

Figure 5: *Left and right views of wall (verification pattern)*

in the POV-Ray calculation approximate the setup of a realistic array detector and optical system.

Only those parts of the stereo and depth images that have a common field of view are retained for this study since it would be futile to search for a matching kernel if it were only visible in one of the views. This requires trimming the images in such a way that the coordinate origin for each image is altered, as described in detail in [7]. After this transformation, the difference between the positions of related points in the left and right stereo images is greatest when the distance is large, and goes to zero as the distance approaches the focal length of the optics. The images of Figures 4 and 5 contain features that are easily seen to be at different locations, the distance being close to the maximum (255 units).

We define **disparity** to be the shift in pixels that is required to match a kernel in the left image with one in the right. A positive disparity means that the shift is to the right, the conventional positive x-direction. The maximum disparity depends on the separation of the camera axes, and as we show in Chapter 3, the maximum is about 14 pixels for this optical setup.

3 Kernel Matching

We use the term *kernel* to describe a group of pixels surrounding a point in one of the stereo images that will be related to an identically shaped group in the other stereo image in order to assign a depth value to a related point in the depth image. The size and shape of the kernel affects the value obtained in the matching process, as does the criterion used to define the closeness of the match. A square kernel of variable size was used here, as only marginal differences were found with rectangular and circular shapes.

All distance values and image intensities are converted to double precision values in the range 0.0 to 1.0 for use in our computations; i.e. the byte values are divided by 255.0.

The simplest matching criterion used was the kernel position that gives the minimum of the sum of the absolute values of the differences of corresponding kernel elements; the algorithmic steps for this are:

1. construct a kernel (as a vector) centred on the point in the left stereo image that corresponds to a point in the depth image (the kernel contains the stereo image values);
2. construct the same sized kernel at the same point in the right stereo image;
3. calculate the matching function:

$$S = \sum_i |KL_i - KR_i| \tag{1}$$

   where KL and KR are the left and right kernel vectors;
4. move the right kernel centre by one pixel at time, recalculating S at each point, and find the shift that minimizes S.

The maximum shift that needs to be tested depends on the separation of the camera centres in pixel units. In practice it is easily found over the calibration range of distances (Section 3.1).

If the lighting is significantly different when a scene is viewed from the left and right, then the simple matching function may be inadequate. To minimize this potential problem, our program includes the options of kernel normalization and cosine maximization. In kernel normalization, the maximum value in each kernel is first set to 1.0, and all the other values are scaled accordingly. The matching function given above is then used. If cosine maximization is chosen, the cosine of the angle between the left and right kernel vectors is used as the matching function. The optimal shift is that which gives the smallest angle between the vectors (maximum cosine). This angle is less susceptible to lighting differences; e.g. when all the pixels in one scene are uniformly brighter, the kernel vectors have the same direction. The cosine of the angle between the kernel vectors is given by:

$$C = \sum_i (KL_i.KR_i)/|KL|.|KR| \tag{2}$$

3.1 Calibration

During calibration of the optical system, the depth values at each point in each of the depth images must be known. In this work, the depth values are obtained by calculation in the simulation. In a real system they would be supplied by some form of physical measurement.

Figure 4 shows the first of 60 pairs of stereo images from the calibration set. This set is simply a series of views of a flat patterned wall at decreasing distances from the view shown. The byte values for relative distances range from 255 to 142 in the calibration set. Thus the maximum normalized distance is 1.0 and the minimum used in the calibration is 0.56. The minimum distance viewable would be at the focal length of the optics, which in this system is about 0.4.

At each of the calibration distances, kernel shifts for many points in the images must be found. It is necessary to use a large number of points for the following reasons: some kernels may not be correctly matched due to similar or ill-defined patterns being present; each kernel shift is an integer number of pixels so that even when good matches are found they may differ by at least one pixel; and there are a limited number of possible shift values (disparities). This last reason arises because the entire range of distances must be represented by a fixed number of pixels (related to the camera separation). In this optical setup it means that the 114 possible calibration distances (255-142) must be compressed into about 14 shifts. Thus a range of distances will generate each kernel shift. Conversely, when a disparity is found it can only be associated with a range of distances.

We used 100 points (10x10 square) in each calibration image to define an average distance to be associated with each disparity. The result is shown in Figure 6. In obtaining this curve, a square kernel of 11x11 pixels was used. This size was arrived at by iteratively calibrating then using the curve to check distances using the verification set of wall images, an example of which is shown in Figure 5. The kernel size can be optimized by finding the value that minimizes the overall error in the distances for all the images.

3.2 Testing

First, the calibration and verification sets were tested to demonstrate that the correct distances could be reproduced. Larger regions of each image (20x20 pixels) were scanned, and the disparities that were found generated distances that have an average error of 0.020 (standard deviation, $\sigma = 0.014$). As the average distance being measured is about 0.8, these sets are reproduced with about a 3% average error.

Using disparities that are found in views containing three dimensional objects is quite different from the calibration and verification scenario that used a flat, patterned wall. Sixty images were taken from the lower path shown at the bottom of Figure 1, and a region of 80x80 pixels was tested from each of the image pairs. The shifts that produced matched kernels for each pixel again used the relationship shown in Figure 6 to generate depth maps. A typical example of the raw depth map generated is shown on the right of Figure 7. This can be compared with the rectangle marked in the true depth map shown between

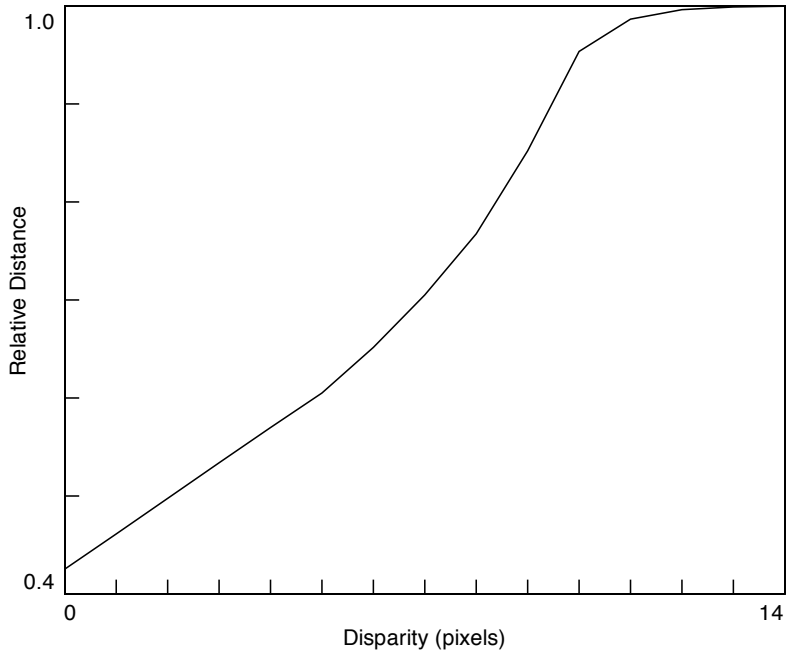

Figure 6: *Average distance as a function of disparity*

the left (top) and right (bottom) stereo images. The ragged nature of the output depth image demonstrates the intrinsic quantization of the kernel matching technique. There are also several points that have clearly not found well matched kernels. However, the overall indicator of depth is quite good, the average error for all the pixels over all the images being 0.026 ($\sigma = 0.007$).

We have also applied two smoothing techniques to the output depth maps, using a neighbourhood average (NA) and a median filter (MF). As the name implies, NA simply calculates the value of each pixel averaged with its neighbours. MF replaces each pixel value with the median of the neighbourhood values. NA produces a smoothed effect but can leave anomalous values (outliers), as well as blurring sharp edges that may be real. Edges are better retained with MF and outliers are generally removed. The range over which NA and/or MF are applied is empirical. We found that the error over all the images is minimized with a range of 5 pixels; i.e. a square extending 5 pixels from the central pixel in each direction. The range is not critical; changing it by plus or minus two pixels only affects the average error by about 3 percent.

Figure 8 shows the result of applying NA to the images shown in Figure 7. Using NA the average error over all images and pixels is reduced from 0.026 to 0.019 ($\sigma = 0.006$), almost a 27% error reduction.

The effect of MF is shown in Figure 9. The average error with MF is 0.023 ($\sigma = 0.007$). Sharper discrimination between the depth levels is retained, real edges are not blurred and the outliers of Figure 6 are removed, but the overall error reduction is not as good as NA. However, MF may be better for indicating a clear path to bigger distances.

Combinations of NA and MF have also been investigated. The optimal result is obtained with NA followed by MF, which gives a slightly lower average error (0.0186 vs 0.0188).

3.3 Computation Times

The program that performed these calculations was written in C, with OpenGL used for graphical display. It was compiled using gcc (version 4.0.3) with O3 optimization, and run as a single thread on a 3GHz pentium 4 system.

With this configuration 80x80 pixel depth map regions are generated for 60 images in about 2.5 seconds, with NA smoothing. The C code has not been rigorously optimized for execution speed, and does not employ fast vector processing routines.

The calculations are well suited to multithreaded execution, so running on a multiple CPU system with multicore processors could easily reduce the execution time by an order of magnitude.

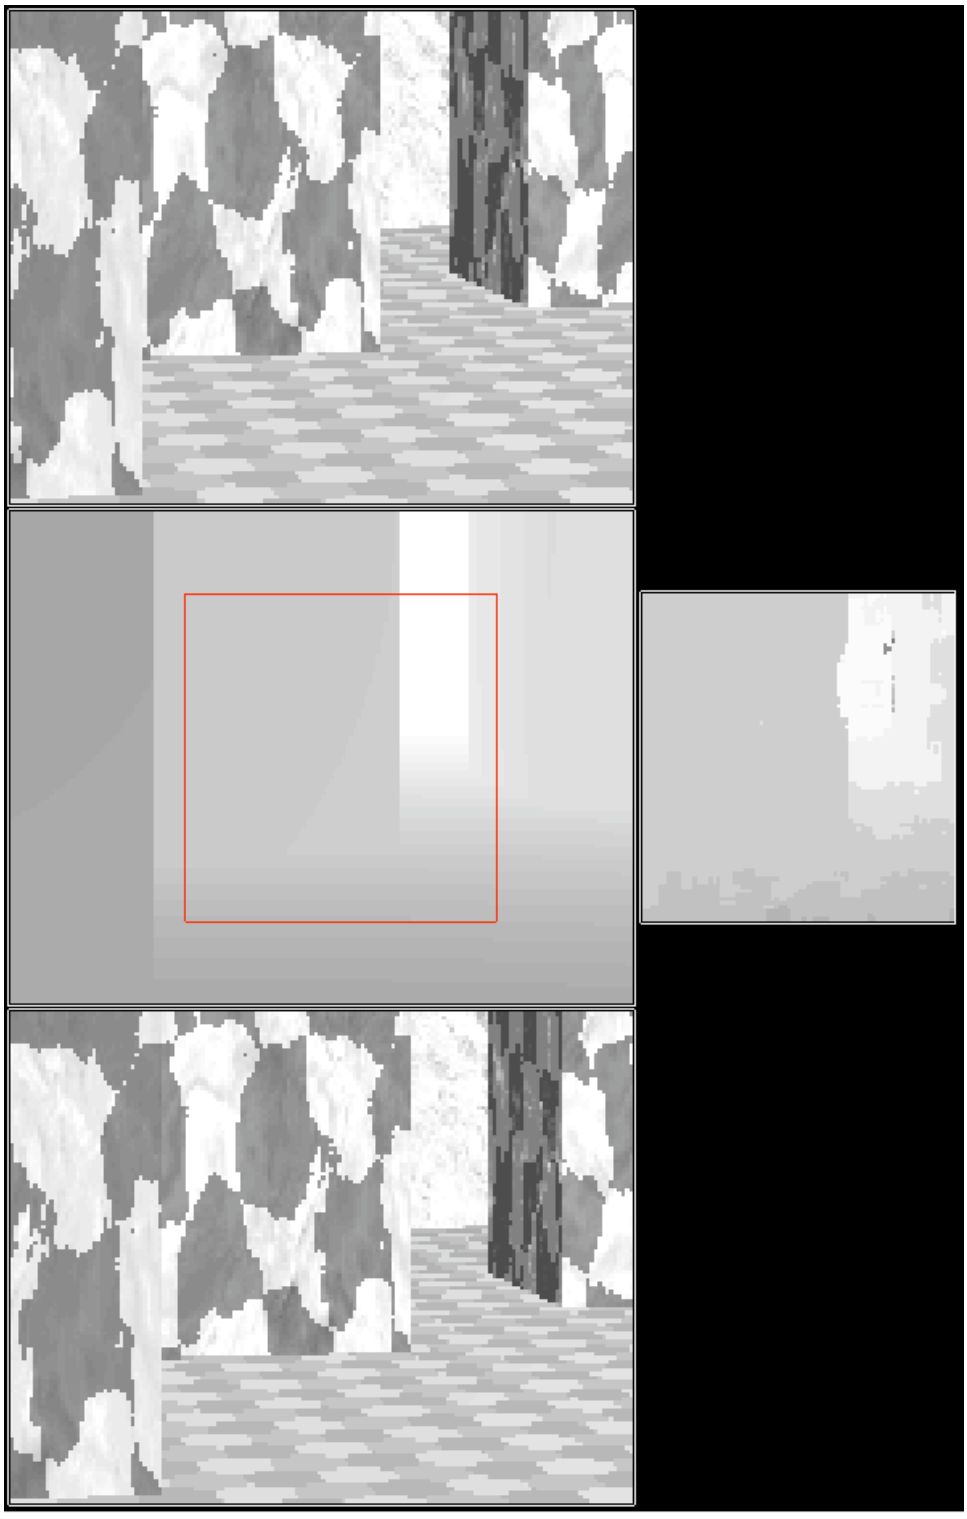

Figure 7: *Stereo and depth images with no smoothing*

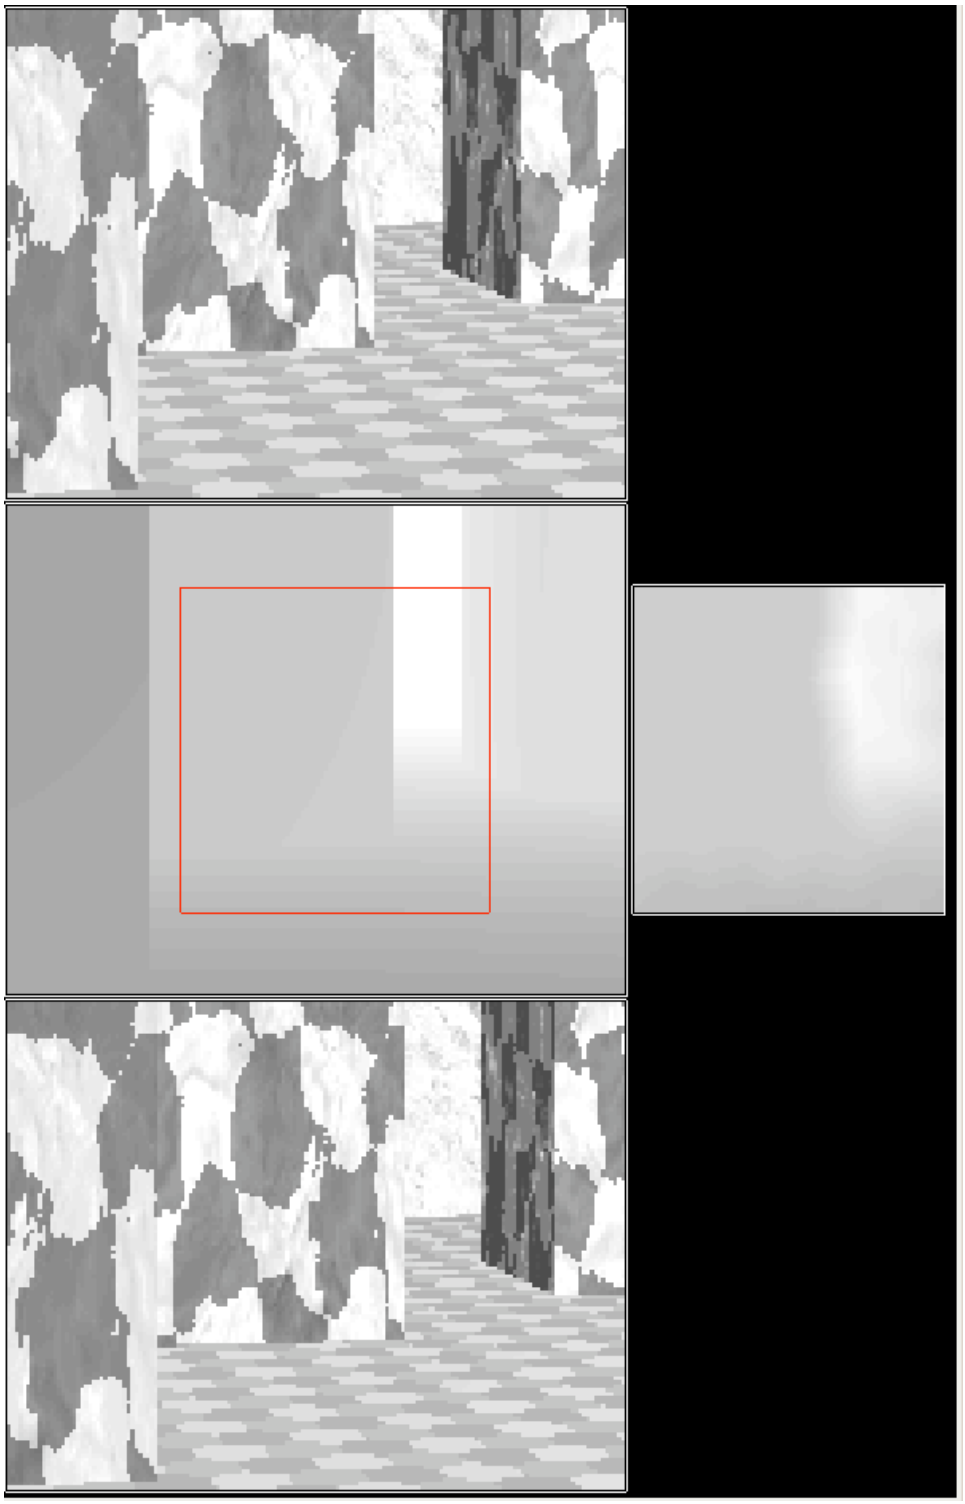

Figure 8: *Stereo and depth images with neighbourhood averaging*

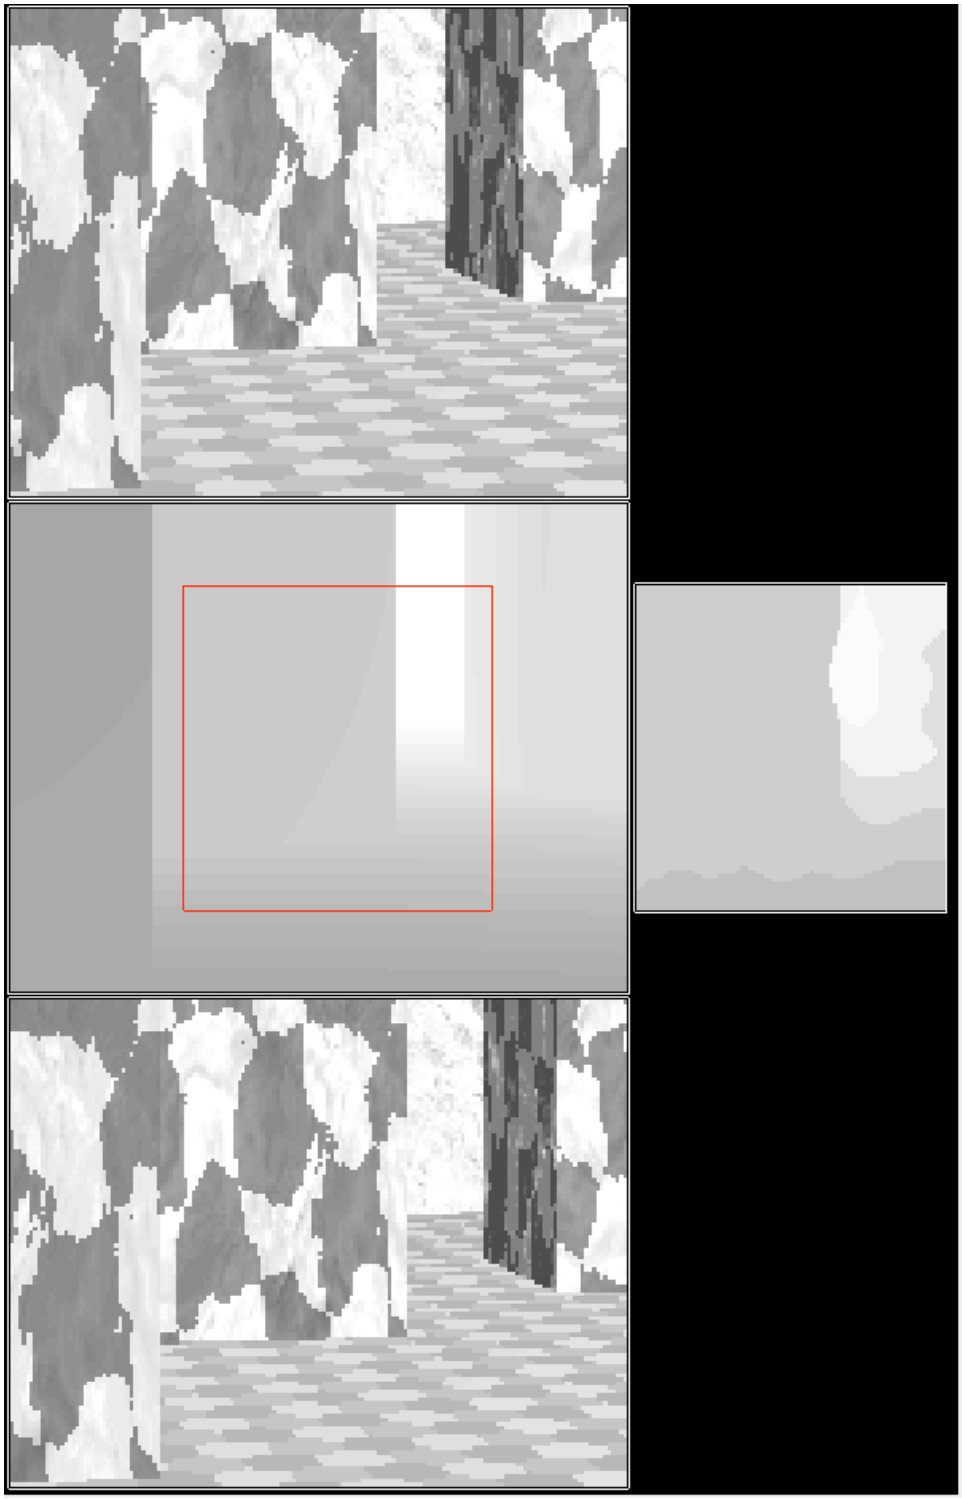

Figure 9: *Stereo and depth images with median filtering*

4 Conclusions

The kernel matching technique, which does not require explicit feature identification within the stereo images, is easy to implement and performs well in generating depth maps in a simulated environment. If the accuracy achieved in simulation can be reproduced in a real environment, the technique could provide maps that would be usable for navigation.

In a real camera system, the size of the detector array, number of elements, focal length and aperture of the optical system would be needed to estimate the granularity caused by the compression of the depth range into a limited number of kernel shifts (in pixels). The simulation generates a granularity that is belived to be approximately realistic. High resolution systems may improve this intrinsic limitation.

To calibrate a real system, it would only be necessary to move it toward a suitably patterned flat surface. The kernel size and smoothing ranges are empirical parameters and would have to be determined by experiment in a real environment. The use of projected light, containing random patterns, would improve the depth accuracy when uniform surfaces are present.

The technique has real-time capability for a moving vehicle. With well optimized code and a mutiple multi-core CPU computer, it would be quite feasible to generate 320x240 depth map images at about 30 per second, which is a reasonable rate for operation on a vehicle moving at about 30 km/hr.

This study also provides useful data for comparison with other techniques that may extract additional information from stereo images, e.g. using perspective and object sizes.

References

[1] Gonzalez, R.C. and Wintz, P. (1987). Digital image processing. Addison-Wesley Publishing Company Inc., Reading, Massachusetts, pp. 52-54.

[2] Scharstein, D. and Szeliski, R. (2002). A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *Int. J. Comp. Vision*, 47(1), 7-42.

[3] Brown, M.Z., Burschka, D. and Hager, G.D. (2003). Advances in computational stereo. *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, 25(8), 993-1008.

[4] Hirschnuller, H. (2001). Improvements in real-time correlation-based stereo vision. In *Proc. IEEE Workshop on Stereo and Multi-Baseline Vision*, 141-148. Kauai, Hawaii.

[5] Chen, C., Hang, Y. and Wu, J. (1997). Range data aquisition using color structured lighting and stereo vision. *Image and Vision Computing*, 15(6), 445-456.

[6] Scharstein, D. and Szeliski, R. (2003). High-accuracy stereo depth maps using structured light. *IEEE Comp. Soc.*. 1, 195-202.

[7] Mackay, D.J. (2006). Generating synthetic stereo pairs and a depth map with POVRay. (DRDC Suffield TM 2006-197). Defence R&D Canada - Suffield.

DOCUMENT CONTROL DATA

(Security classification of title, body of abstract and indexing annotation must be entered when the overall document is classified)

1. ORIGINATOR (the name and address of the organization preparing the document. Organizations for who the document was prepared, e.g. Establishment sponsoring a contractor's report, or tasking agency, are entered in Section 8.)

Defence R&D Canada – Suffield
PO Box 4000, Station Main
Medicine Hat, AB T1A 8K6

2. SECURITY CLASSIFICATION
(overall security classification of the document, including special warning terms if applicable)

Unclassified

3. TITLE (the complete document title as indicated on the title page. Its classification should be indicated by the appropriate abbreviation (S, C or U) in parentheses after the title).

Depth Maps from Simulated Stereo Images Using Kernel Matching

4. AUTHORS (Last name, first name, middle initial. If military, show rank, e.g. Doe, Maj. John E.)

Barton, Simon A. and Mackay, David J.

5. DATE OF PUBLICATION (month and year of publication of document)

November 2006

6a. NO. OF PAGES (total containing information, include Annexes, Appendices, etc) 23

6b. NO. OF REFS (total cited in document) 7

7. DESCRIPTIVE NOTES (the category of the document, e.g. technical report, technical note or memorandum. If appropriate, enter the type of report, e.g. interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered.)

Technical Memorandum

8. SPONSORING ACTIVITY (the name of the department project office or laboratory sponsoring the research and development. Include the address.)

9a. PROJECT OR GRANT NO. (If appropriate, the applicable research and development project or grant number under which the document was written. Please specify whether project or grant.)

9b. CONTRACT NO. (If appropriate, the applicable number under which the document was written.)

10a. ORIGINATOR'S DOCUMENT NUMBER (the official document number by which the document is identified by the originating activity. This number must be unique to this document.)

DRDC Suffield TM 2006-198

10b. OTHER DOCUMENT NOs. (Any other numbers which may be assigned this document either by the originator or by the sponsor.)

11. DOCUMENT AVAILABILITY (any limitations on further dissemination of the document, other than those imposed by security classification)

(x) Unlimited distribution
( ) Distribution limited to defence departments and defence contractors; further distribution only as approved
( ) Distribution limited to defence departments and Canadian defence contractors; further distribution only as approved
( ) Distribution limited to government departments and agencies; further distribution only as approved
( ) Distribution limited to defence departments; further distribution only as approved
( ) Other (please specify):

12. DOCUMENT ANNOUNCEMENT (any limitation to the bibliographic announcement of this document. This will normally corresponded to the Document Availability (11). However, where further distribution (beyond the audience specified in 11) is possible, a wider announcement audience may be selected).

Unlimited

13. ABSTRACT (a brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C) or (U). It is not necessary to include here abstracts in both official languages unless the text is bilingual).

Simulated stereo images and depth maps were generated using the open-source ray tracing program POV-Ray, operating in an environment containing obstacles. The image pairs were used to generate depth maps by finding the relative positions of matching groups (kernels) of image pixels. Kernel matching generated an average depth accuracy of better than 90% over all images and pixels tested. When neighbourhood average and median filter smoothing techniques were applied to the resulting depth maps, the average accuracy exceeded 95%. Occlusion and lighting differences between the images were attenuated by smoothing and kernel normalization. Computation times for the simulated depth maps indicate that kernel matching could be used in real video imagery on a vehicle moving at 30 km/hr.

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifies, such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus, e.g. Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus-identified. If it is not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title.)

Autonomous vehicles, depth maps, robotics, stereopsis