

INTEREST OPERATORS IN CLOSE-RANGE OBJECT RECONSTRUCTION

I. Jazayeri, C.S. Fraser

Department of Geomatics, The University of Melbourne, Melbourne, VIC 3010 Australia –
i.jazayeri@pgrad.unimelb.edu.au &- c.fraser@unimelb.edu.au

Commission V, WG V/1

KEY WORDS: Interest Operators, Object Reconstruction, 3D modelling, Close-Range Photogrammetry, Industrial Measurement, Feature Detection

ABSTRACT:

This paper investigates the use of interest operators for feature detection and the generation of high-accuracy object measurement, modelling, texturing and visualization. Interest operators are required as a preliminary step in the surface measurement process, to find features points for matching across multiple images. Three interest operators, namely the Förstner, SUSAN and FAST algorithms are assessed based on speed, feature detection rate and localization. The FAST operator is found to be the optimal interest operator for close-range object reconstruction. An image pre-processing algorithm, namely the Wallis filter, is also investigated and results indicate that it greatly enhances interest operator performance.

1. INTRODUCTION

Image-based 3D modelling can be classified into two categories, based on the intended applications. The first category involves creating high-accuracy 3D models. Automated reconstruction methods currently in existence generally do not yield high accuracy and are limited in their application. For high-accuracy solutions, image point correspondence determination and photogrammetric orientation of complex network geometries must be fully automated and issues such as camera calibration should also be considered. Some applications that demand the accurate and complete modelling of object detail include industrial vision metrology and accident scene reconstruction. The second category involves creating visually-pleasing 3D models that focus on the photo-realism of results, rather than the accuracy achieved. These models are used for applications where accuracy is not critical, such as animation and some types of cultural heritage recording. The research reported in this paper investigates the use of interest operators in the creation of high-accuracy 3D models and therefore falls into the first category.

The image-based modelling of an object has been defined by Remondino (2006) as a complete process that starts with image acquisition and ends with an interactive 3D virtual model. The photogrammetric approach to creating 3D models involves the following phases: image pre-processing, camera calibration and network orientation, image scanning for point detection, surface measurement and point triangulation, blunder detection and statistical filtering, mesh generation and texturing, and visualization and analysis. Currently there is no single software package available that allows for each of these steps to be executed within the same environment. Instead, each phase is typically completed separately using different software systems that require the 3D data to be translated and interchanged between each of the various measurement, modelling and visualization packages.

This paper reports on the first phase of research into the generation of high-accuracy 3D object measurement, modelling, texturing and visualization, focussing upon the application of

three different interest operators for accurate feature point extraction. The implementation of interest operators is discussed and the performance of the different operators is analysed. The optimal operator for high-accuracy close-range object reconstruction is then highlighted. This research has facilitated the development of a feature extraction and image measurement process that will be central to the development of an automatic procedure for high-accuracy point cloud generation in multi-image networks, where robust orientation and 3D point determination will facilitate surface measurement and visualization within a single software system.

2. INTEREST OPERATORS

Interest operators are algorithms which detect features of interest in an image such as corners, edges or regions. For high-accuracy 3D object reconstruction, interest operators are required as a preliminary step in the surface measurement process, to find features that serve as suitable points when matching across multiple images. Interest operators were first developed in the 1970's and since then a number of different algorithms have been developed and presented. No single algorithm, however, has been universally accepted as optimal for all applications. This paper will report on the following interest operators and identify which is the most suitable for close-range high-accuracy 3D object reconstruction: Förstner (Förstner and Gülch, 1987), SUSAN (Smith & Brady, 1997) and the recently developed FAST interest operator (Rosten & Drummond, 2006). The suitability of each algorithm is assessed based on speed, feature detection rate and localization.

2.1 Förstner Operator

The Förstner Operator (Förstner and Gülch, 1987) has been widely adopted in photogrammetry and computer vision over the last two decades. It was developed with the aim of creating a fast operator for the detection and precise location of distinct points, corners and centres of circular image features, with photogrammetric image matching applications in mind.

The algorithm identifies interest points, edges and regions using the autocorrelation function \mathbf{A} . The derivatives of \mathbf{A} are computed and summed over a Gaussian window. Error ellipses are computed and based on the size and shape properties of each ellipse, the interest points found are classified as points, edges or regions. Förstner calculates the size and shape of the error ellipses using two eigenvalues λ_1 and λ_2 as well as the inversion of \mathbf{A} .

The error ellipse size is determined by:

$$w = \frac{1}{\lambda_1 + \lambda_2} = \frac{\det(\mathbf{A})}{\text{trace}(\mathbf{A})}, \quad w > 0 \quad (1)$$

The error ellipse shape (roundness of the ellipse) is determined by:

$$q = 1 - \left(\frac{\lambda_1 - \lambda_2}{\lambda_1 + \lambda_2} \right)^2 = \frac{4 \cdot \det(\mathbf{A})}{\text{trace}(\mathbf{A})^2}, \quad 0 \leq q \leq 1 \quad (2)$$

Based on the values of w and q , the algorithm classifies each area as follows (Rodehorst and Koshan 2006):

- Small circular ellipses define a salient point
- Elongated error ellipses suggest a straight edge
- Large ellipses mark a homogeneous area

An interest point is only precisely located when the given threshold values w_{min} and q_{min} are exceeded. Practical experience with the Förstner operator has found it to be limited by high computational cost, since it is relatively slow and often impractical for high-level data analysis.

2.2 SUSAN Operator

The SUSAN operator was developed in the mid nineties by Smith and Brady (1997) in an attempt to create an entirely new approach to low level image processing. It is an accurate, noise resistant and fast edge and corner detector that addresses some limitations of other operators, such as high computation time.

The SUSAN operator is based on the concept that each point of interest in the image will have associated with it a local area of similar brightness values and that these areas can be used as a guide to help find features of interest such as corners and edges in the image. The SUSAN operator searches for areas of similar brightness, and consequently for interest points within a weighted circular window. The algorithm denotes the central pixel in the search window as the nucleus. The area within the window that has similar intensity values to the nucleus is computed and referred to as the USAN (Univalued Segment Assimilating Nucleus). A low value for the USAN indicates a corner since the central pixel would be very different from its surroundings. After assessing results and eliminating outliers, the local minima of the SUSANs (smallest USAN) remain as valid interest points. The comparison between pixel brightness values is computed using the following equation:

$$c(\vec{r}, \vec{r}_0) = \begin{cases} 1 & \text{if } |I(\vec{r}) - I(\vec{r}_0)| \leq t \\ 0 & \text{if } |I(\vec{r}) - I(\vec{r}_0)| > t \end{cases} \quad (3)$$

Where \vec{r}_0 is the position of the nucleus in the two-dimensional image, \vec{r} is the position of any other point within the circular window, $I(\vec{r})$ is the brightness value of any pixel, t is the brightness value threshold and c is the output of the comparison. The comparison is calculated for each pixel in the circular window and the total number of pixels with similar brightness values as the nucleus is summed by:

$$n(\vec{r}_0) = \sum_{\vec{r}} c(\vec{r}, \vec{r}_0) \quad (4)$$

In the next step, the $n(\vec{r}_0)$ value is compared with a geometric threshold, g . The algorithm uses a threshold value in order to distinguish between features that make suitable interest points and non-suitable features. To find a corner in the image, the threshold value g is set to half of the maximum value of $n(\vec{r}_0)$, n_{max} . If n_{max} is less than the threshold g then a corner exists.

2.3 FAST Operator

The FAST interest operator, which was only recently developed by Rosten and Drummond (2006), was designed as a high speed feature detector for application in real-time frame-rate digital photogrammetry and computer vision. It employs a high-speed algorithm that is computationally simple and effective for real-time applications. The FAST operator has been found to generally outperform other interest operators in speed, being faster than the Harris, SUSAN and SIFT interest operators commonly used in imaging applications (Rosten and Drummond, 2006). Beyond its speed, a second advantage of the FAST operator is its invariance to rotation and changes in scale.

The FAST algorithm functions in a similar way to the SUSAN algorithm detailed in Smith and Brady (1997) in that the detectors all examine a small patch in an image and assess whether or not it 'looks' like a corner. A circular window is scanned across the image and the intensity values of the pixels within or around the window are compared to that of the central pixel. The SUSAN algorithm is efficient and fast, and in contrast to other interest operator algorithms, there is no need to calculate a second derivative in the computational process as only a small number of pixels are examined for each detected corner. As a result, the SUSAN operator is not computationally expensive and it offers an effective option for applications which require speed.

The development of the FAST operator by Rosten and Drummond (2006) grew out research that aimed to improve the speed and invariance to rotation, transformation and changes in scale of the SUSAN operator. In comparison to the SUSAN detector and other operators that follow a similar computation process, the FAST operator is indeed proving to be faster and more reliable and it is thus being adopted in a variety of applications including real-time frame-rate imaging.

The algorithm operates by considering a circle of sixteen pixels around the corner candidate p . A corner is deemed to exist when a set of n contiguous pixels in the circle are all brighter than the intensity of the candidate pixel I_p plus a threshold t , or all darker than $I_p \leq t$. The operator uses machine learning and decision trees to classify the interest points. A detailed description of the FAST algorithm can be found in Rosten & Drummond (2006).

3. IMAGE PRE-PROCESSING

An image pre-processing algorithm can be applied to enhance the images for subsequent image processing. Research conducted within the reported project has found that a filter is warranted to provide greater detail in shadowed areas and saturated areas simultaneously, and thus to allow a greater number of interest points to be detected.

The Wallis filter (Wallis, 1976), is a digital processing function that enhances the contrast levels of an image and it can be applied in order to optimise images for subsequent image-matching (Baltsavias, 1991). Studies have found that interest operators typically find more suitable points on imagery that has been pre-processed with the Wallis filter (eg Remondino, 2006). The algorithm is adaptive and adjusts pixel brightness values in local areas only, as opposed to a global contrast filter which applies the same level of contrast throughout an entire image. The resulting image contains greater detail in both low and high level contrast regions concurrently, ensuring that good local enhancement is achieved throughout the entire image. After testing a number of smoothing filters it was found that the Wallis filter is the optimal choice for this work.

The Wallis filter requires the user to input a target mean and standard deviation, and adjusts local areas to match the target values accordingly. A weight value is used to control to what extent the target values are used and how much of the original image is kept.

The general form of the Wallis filter is given by:

$$i_{wallis}(x, y) = i_{original}(x, y)r_1 + r_0 \quad (5)$$

where

$$r_1 = \frac{cS_{original}}{(cS_{original} + \frac{S_{target}}{c})} \quad (6)$$

and

$$r_0 = bm_{target} + (1 - b - r_1)m_{original} \quad (7)$$

In Eqns. 5-7, i_{wallis} and $i_{original}$ are the filtered and the original images, respectively; r_0 and r_1 the additive and multiplicative parameters, respectively; $m_{original}$ and $S_{original}$ the mean and standard deviation of the original images; m_{target} and S_{target} the user-specified target mean and standard deviation values for the filtered output images; c the constant expansion constant; and b , a weight parameter, the brightness forcing constant.

The Wallis filter first divides the input image into neighbouring square blocks of user-defined size in order to calculate local statistics. The choice of block size is governed by the amount of detail and the level of contrast that the user requires. Small blocks, for example 5x5 pixels in size, result in the strong enhancement of small and unimportant features which are problematic in subsequent feature extraction and image matching procedures. Conversely, large blocks, for example 120x120 pixels in size, result in a significant loss of detail. A medium size block of 31x31 pixels is small enough to recover the required detail in the input image and large enough to appropriately filter out and ignore unimportant features.

The original mean and standard deviation of the unfiltered image, $m_{original}$ and $S_{original}$, are calculated for each individual block based on the grey values of the pixels. The resulting value is assigned to the central cell in each block. The mean and standard deviation values of all other cells in the block are calculated from this central cell by bilinear interpolation. In this way, each individual pixel is assigned its own initial local mean and standard deviation based on surrounding pixel values.

The target mean and standard deviation values, specified by the user, are used to adjust the brightness and the contrast of the input cells respectively. The target mean value is set at 127 by default and the user may input any value between 0-255. Higher local mean values than 127 will brighten the image, while an input of less than 127 will create a darker output. The target standard deviation value is set at 60 by default and may be assigned any value within the data range of the image. Higher values for this parameter result in a greater contrast stretch, producing higher local contrast and greater detail throughout the output.

The resulting Wallis filtered image is a weighted combination of the mean and standard deviation of the original image and the target mean and standard deviation values specified by the user's input. The weight is determined by the brightness forcing constant, b , which can take on any value between 0 and 1. A weight value of 1 will generate an output image equal to the Wallis filter target values specified, while a weight of 0 will keep the original pixel values. A weight value of $b = 0.6$ for example will calculate an output image with 60% of the user-specified target values and 40% of the original unaltered pixel values.

The contrast expansion constant, c , is a general contrast function that may be any value between 0 and 1. Lower values such as 0.1 produce an overall grey image where only the outlines of objects are visible, with very little contrast and detail. Values closer to 1, produce a highly contrasted image with greater detail and distinct black and white areas.

4. EXPERIMENTAL TESTING PROGRAM

The experimental testing program for this work has two phases. The first assessed both the Wallis filter and the degree to which it enhanced interest operator performance. Different values for the parameters of the filter (mean; standard deviation; block size; brightness forcing constant; contrast expansion constant) were tested and based on the trends found, a range of optimum values is suggested.

In the second testing phase, the Förstner, SUSAN and FAST interest operators were evaluated to ascertain which is optimal

for high-accuracy close-range photogrammetric object reconstruction. The first test object comprised the alphabet printed on an A3 piece of paper and imaged in a strong, convergent network with optimal camera exposure settings. Colour coded targets (Cronk et al., 2006) were also placed around the object for automatic camera calibration and network orientation, as indicated in Figure 1. A quantitative assessment of each interest operator was then performed, based on the following criteria:

- Speed: A fast and efficient algorithm is required for high-accuracy object reconstruction. It is important that the interest operator has minimal computational time without compromising the accuracy of the results.
- Feature detection rate: This is a measure of the number of true interest points found relative to the number of points missed and points wrongly detected.
- Localization: This is a measure of how well the operator positions the interest point found relative to its true position and can be assessed visually at pixel level in the image.

Through the use of coded targets, the 3D positions of the points found with the interest operators were automatically computed and shown graphically in 3D space. A 3D point cloud of the alphabet was generated, showing all the interest points found in the multi-image network.



Figure 1: Original alphabet image

The Förstner, SUSAN and FAST operators were subsequently further tested on a number of other objects to verify the alphabet test results and to confirm the optimal operator for this work.

5. RESULTS

5.1 Wallis Filter

Results indicated that the Wallis filter is a necessary pre-processing function that will enable the three interest operators considered to find more suitable interest points. By applying the Wallis filter, issues arising from changes in contrast and illumination are overcome, leading to more repeatable and reliable results. The building shown in Figure 2a is an example of an object with largely varying pixel brightness values. By applying the Wallis filter to the image, the shadowed areas are brightened and good local enhancement is achieved throughout the entire image, as illustrated in Figure 2c. This allows the

interest operator to detect suitable corresponding points in all areas of the image. Figures 2b and 2d show the results of the FAST interest operator. The algorithm is able to efficiently detect many more interest points in the Wallis filtered image, as illustrated in Figures 1e and 1f.

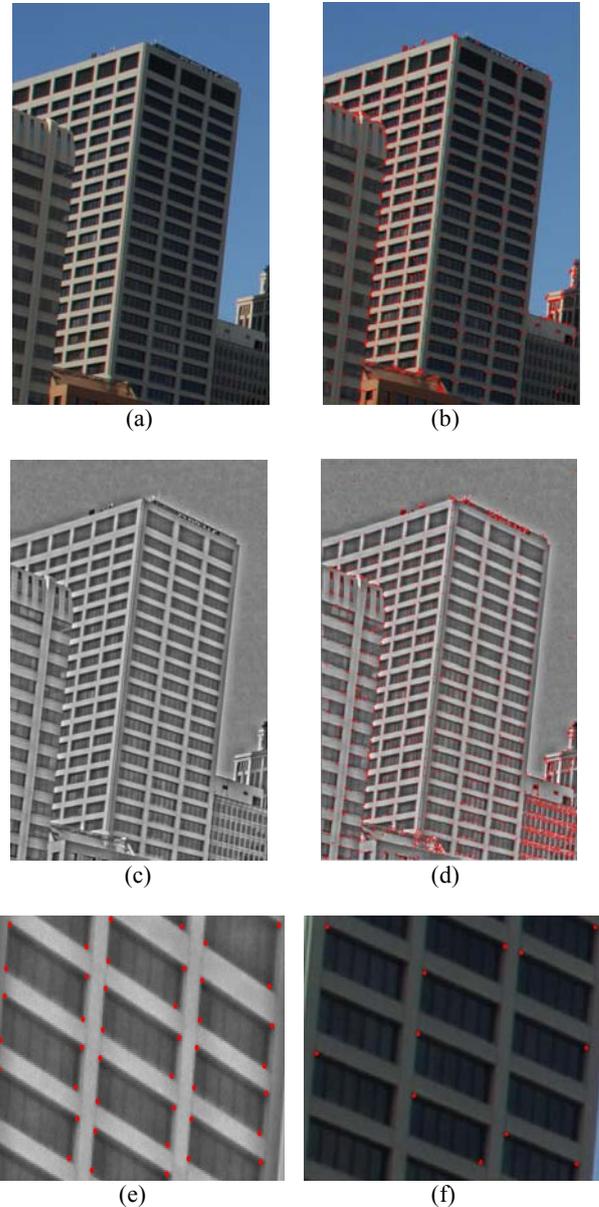


Figure 2: (a) Original image; (b) Results of the FAST interest operator; (c) Wallis filtered image; (d) Results of the FAST interest operator superimposed on Wallis filtered image; (e) Enlarged area showing FAST interest operator results in detail on Wallis filtered image; (f) Enlarged area showing FAST interest operator results in detail on original image.

As will be discussed in the following section, the Wallis filter was further assessed in the interest operator testing phase of this work. The results showed that while the SUSAN operator could not be applied to the Wallis filtered images, the FAST and the Förstner algorithms detected up to 25 times more interest points on pre-processed images.

The results from testing the parameters in the Wallis filter demonstrated that an optimum range of values exists and depending on the requirements of the user, i.e. upon the level of detail and contrast required, an appropriate value can be selected. A representative summary of the results is provided in Table 1.

Parameter	De-fault	Optimum Range	Trend Explanation (in pixel brightness values)
Mean	127	107-147	67 87 107 127 147 167 187 Dark → Medium → Light
Standard Deviation	60	60-90	10 30 60 90 120 D: Low → Med → High C: Low → Med → High
Block Size (in pixels)	31	21-41	11 21 31 41 51 61 71 81 91 D: High → Med → Low C: Low → Med → High
Brightness Forcing Constant	0.95	0.95-0.99	0 0.1 0.3 0.6 0.9 0.95 0.99 D: Low → Med → High C: High → Med → Low
Contrast Expansion Constant	0.9	0.9-0.99	0 0.1 0.3 0.6 0.9 0.95 0.99 D: Low → Med → High C: Low → Med → High

D = Detail; C = Contrast

Table 1: Optimum Range of Wallis filter parameters for the detection of interest points.

5.2 Interest Operator Performance

The results from the alphabet test, illustrated in Figure 1, clearly indicated that the FAST operator (figure 3c) is the fastest and most robust algorithm, repeatedly yielding reliable results with good localization, as illustrated in Table 2. The FAST algorithm displayed speed advantages over both the SUSAN and Förstner operators and detected nearly all the corners in the image (95 out of a total of 110 points), with very few erroneous points being detected. The SUSAN operator yields reasonable results, as shown in Figure 3b, finding most of the interest points in the image at a good speed. It however wrongly detected 67 points. This is evident in the letter 'O', where no corners exist but seven interest points have been wrongly detected. As illustrated in Figure 3a, the Förstner operator found few interest points, in fact less than half the points found by FAST, but with good localization and with no outliers. The algorithm operated at a very low speed with the original image but was much faster when applied to the Wallis filtered image.

	Förstner	SUSAN	FAST
Alphabet: 110 points Total			
Speed (sec)	70	0.4	0.1
Detection Rate (Total pts)	45	157	95
Points missed	65	20	3
Incorrect points detected	0	67	12
Speed (sec) with Wallis	3	N/A	0.4
Detection Rate with Wallis	54	N/A	240
Localization (in pixel)	0.5	0.6	0.3

Table 2: Alphabet test results

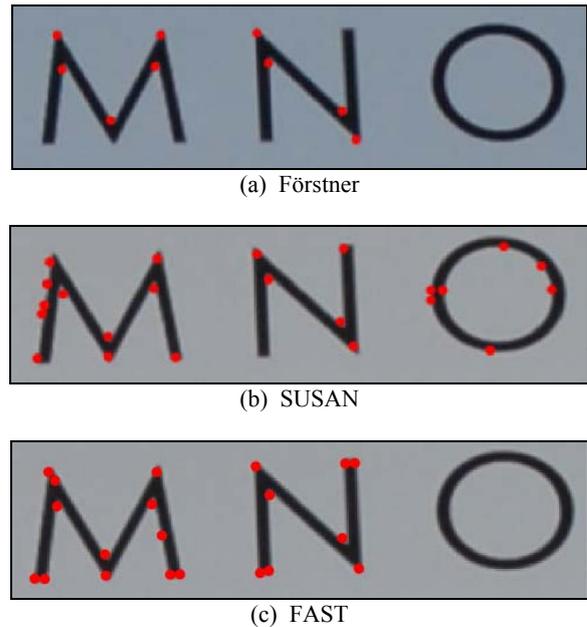


Figure 3: Enlarged area showing interest operator results.

Figure 4 shows the final 3D alphabet test network, where the coordinates of all the interest points found are known. This strong network forms a starting platform for subsequent image matching and other image processing steps.

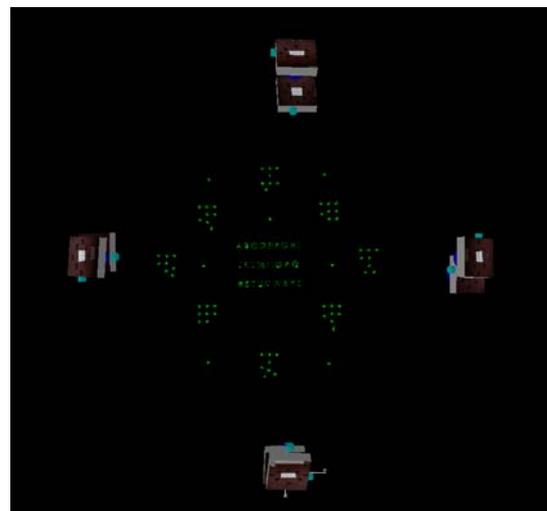


Figure 4: 3D point cloud of alphabet test.

Further testing of a number of different objects showed that the implementation of the FAST operator repeatedly produced the best results during the experiments, as indicated in Table 3. In these tests, the FAST operator was able to detect close to 300,000 points in less than a second. It found more than double the number of interest points found by the SUSAN and Förstner operators, with good localization.

Figure 5a shows the results of applying the FAST operator to one of the further test objects, namely a section of a brick wall, whereas Figure 5b shows the resulting 3D point cloud.

6. CONCLUDING REMARKS

The main findings from the experimental testing reported in this paper are, firstly, that the FAST operator is the optimal algorithm of the three evaluated for finding feature points in images utilised for high-accuracy object reconstruction. Secondly, application of the Wallis filter is a necessary image pre-processing step that will facilitate a greater number of interest points being detected. The reported research has resulted in the development of a feature extraction and image measurement process that will be central to the development of an automatic procedure for high-accuracy point cloud generation in multi-image networks, where robust orientation and 3D point determination will enable surface measurement and visualization to be implemented within a single software system.

	Förstner	SUSAN	FAST
Brick Wall			
Speed (sec)	210	1.4	0.1
Detection Rate (No. of pts)	18938	4572	26130
Speed (sec) with Wallis	16.2	N/A	0.4
Detection Rate with Wallis	48944	N/A	247294
Localization (in pixel)	0.5	0.6	0.4

(a) Brick wall.

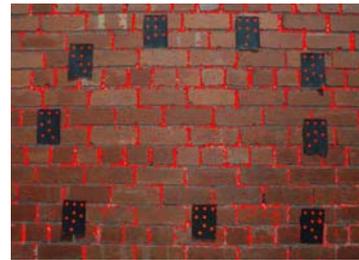
	Förstner	SUSAN	FAST
Stained Glass			
Speed (sec)	194.1	1.5	0.1
Detection Rate (No. of pts)	18938	4572	10324
Speed (sec) with Wallis	16.2	N/A	0.2
Detection Rate with Wallis	48944	N/A	71148
Localization (in pixel)	0.5	0.9	0.4

(b) Stained Glass Windows

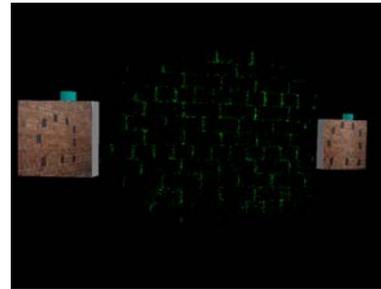
	Förstner	SUSAN	FAST
Archway			
Speed (sec)	12min	1.4	0.1
Detection Rate (No. of pts)	2132	1847	4816
Speed (sec) with Wallis	16.5	N/A	0.4
Detection Rate with Wallis	56988	N/A	125078
Localization (in pixel)	0.5	0.9	0.4

(c) Building Archway

Table 3: Performance of interest operators as assessed by speed, detection rate and localization.



(a) Results of FAST interest operator



(b) Resulting Point Cloud

Figure 5: Brick wall test.

REFERENCES

- Baltsavias, E. P. 1991, Multiphoto geometrically constrained matching, *PhD dissertation*, Institute of Geodesy and Photogrammetry, ETH Zurich, Mitteilungen No. 49, 221 pp.
- Cronk, S., Fraser, C.S. and Hanley, H.B. 2006, Automatic calibration of colour digital cameras. *Photogrammetric Record*, 21(116) pp. 355-372.
- Förstner, W. and Gülch, E. 1987, A fast operator for detection and precise location of distinct points, corners and centres of circular features. *ISPRS Conference on Fast Processing of Photogrammetric Data*, Interlaken, Switzerland, pp. 281-305
- Remondino, F. 2006, Image-based modelling for object and human reconstruction. *PhD dissertation*, Institute of Geodesy and Photogrammetry, ETH Zurich, Mitteilungen No. 91, 174pp
- Rodehorst, V. and Koshan, A, 2006, Comparison and evaluation of feature point detections. Berlin, Germany.
- Rosten, E. and Drummond, T. 2006, Machine learning for high-speed corner detection, *European Conference on Computer Vision*, Graz, Austria, pp 430-443
- Schmid, C, Mohr, R and Bauckhage, C. 2000, Evaluation of interest point detectors. *International Journal of Computer Vision*, 37(2) pp. 151-172.
- Smith, S.M. and Brady, J.M. (1997) SUSAN – A new approach to low level image processing. *International Journal of Computer Vision*, 23(1) pp. 45-78.
- Wallis, K.F. (1976) Seasonal adjustment and relations between variables *Journal of the American Statistical Association*, 69(345) pp. 18-31.