

# An Algorithm for the Contextual Adaption of SURF Octave Selection with Good Matching Performance: Best Octaves

\*Shoaib Ehsan, Nadia Kanwal, Adrian F. Clark and Klaus D. McDonald-Maier

School of Computer Science & Electronic Engineering,

University of Essex, Wivenhoe Park, Colchester CO4 3SQ, UK

[sehsan@essex.ac.uk](mailto:sehsan@essex.ac.uk), [nkanwa@essex.ac.uk](mailto:nkanwa@essex.ac.uk), [alien@essex.ac.uk](mailto:alien@essex.ac.uk), [kdm@essex.ac.uk](mailto:kdm@essex.ac.uk)

Tel: +441206872432; Fax: +441206872900

**EDICS Category:** ARS-IIU Image & Video Interpretation and Understanding

## Abstract

Speeded-Up Robust Features (SURF) is a feature extraction algorithm designed for real-time execution, though this is rarely achievable on low-power hardware such as that in mobile robots. One way to reduce the computation is to discard some of the scale-space octaves, and previous research has simply discarded the higher octaves. This paper shows that this approach is not always the most sensible and presents an algorithm for choosing which octaves to discard based on properties of the imagery. Results obtained with this *best octaves* algorithm show that it is able to achieve a significant reduction in computation without compromising matching performance.

## I. INTRODUCTION

Recent years have seen much effort expended within the research community towards techniques that are able to detect and describe image features in a way that makes them reasonably independent of scale and orientation changes between the images being matched [1-10]. The technique known as SURF (Speeded-Up Robust Features) has a number of adaptations over earlier techniques such as SIFT [1] and the Harris-Laplace feature detector [6] that are intended to improve execution speed without compromising the effectiveness of feature detection [2, 11-13]. The criticality of real-time performance for some applications provides a stimulus to investigate efficient software and/or hardware solutions, not only in terms of execution speed but also for computational resources, chip area, weight and power consumption.

Despite being faster than contemporary methods, software implementations of SURF do not necessarily achieve real-time performance on desktop computers [2]. To overcome this shortcoming, recent research has targeted software-based optimization and/or hardware acceleration of SURF [8, 14-15].

This paper explores the reduction of scale-space 'octaves' as a key speed improvement of SURF. Rather than simply discarding the higher octaves, which may reduce accuracy, this paper develops and assesses a more sophisticated approach, termed *best octaves*, that examines all octaves and then chooses two octaves that provide the best matching performance according to criteria expounded below. It will be demonstrated that *best octaves* out-performs SURF variants obtained using the conventional approach. To our knowledge, this is the first systematic approach to SURF octave selection.

The remainder of this paper is structured as follows. Following a brief overview of the SURF algorithm in the next section, Section III examines the reduction of SURF octaves as a key method to improve execution speed and presents an approach to octave selection. An assessment of *best octaves* in terms of matching performance and reduction in computation is presented in Sections IV and V. Finally, conclusions are drawn in Section VI.

## II. AN OVERVIEW OF THE SURF ALGORITHM

This section provides a brief overview of the SURF algorithm; see [2] for an in-depth exposition. There are two main (and distinct) stages, detection and description, followed by feature matching. SURF constructs a scale space by convolving rectangular masks of increasing size, corresponding to different scales, with the input image, using an integral image representation [16] for speed. This results in a series of blob response maps at different scales. The scale space is divided into a number of octaves, formed by grouping blob response maps for adjacent scales. Normally, four scales per octave are used as this is considered sufficient for scale space analysis [2]. The algorithm also doubles the spatial sampling interval with increasing octave to reduce computation. Once the scale-space is constructed, 3-D non-maximum suppression is performed [17], followed by 3-D quadratic interpolation [18], to achieve sub-pixel, sub-scale accuracy. A blob response threshold is normally applied to select high-contrast interest points. The descriptors for the detected interest points, based on sums of Haar wavelet responses, are calculated after

orientation assignment, to achieve rotation invariance. The final stage is image feature matching on the basis of computed descriptors by employing a nearest neighbor strategy [1].

### III. REDUCING THE NUMBER OF SURF OCTAVES

Table I summarizes how the various stages of SURF algorithm scale, where  $n$  is the image resolution,  $m$  is the number of detected local maxima,  $i$  is the number of detected interest points and  $k$  is the product of the number of feature descriptors for a test image and the number of feature descriptors for a reference image.

It is common to reduce the number of SURF octaves by discarding higher octaves in favor of lower ones; but the authors have observed that, with real-world images, this can actually degrade image matching performance in cases where higher octaves yield more interest points than lower ones after the blob response threshold is applied. To elucidate this point more, three sample cases are discussed here: the first and the second image of Boat dataset [19], the fifth and the sixth image of the Bike dataset [19] and two aerial images (see Fig. 1). The image matching results for these cases using OpenSURF [20] with four and two octaves are detailed in Table II. Although the matching performance is not affected much by decreasing the number of octaves for the first and second image of Boat dataset [19], it is evident that there is a considerable decrease in the number of matched interest points (58.3 % and 64.4%) for the other two cases when the number of octaves is reduced to two. It is to overcome these limitations that *best octaves* has been devised.

TABLE I. COMPUTATIONAL COMPLEXITY OF SURF-BASED IMAGE MATCHING

S.No.	Stage	Computational Complexity
<b>Detection</b>		
S1.	Integral Image Calculation	$O(n)$
S2.	Blob Response Calculation	$O(n)$
S3.	3-D Non-Maximum Suppression	$O(n)$
S4.	3-D Quadratic Interpolation	$O(m)$
<b>Description</b>		
S5.	Orientation Assignment	$O(i)$
S6.	Descriptor Calculation	$O(i)$
<b>Matching</b>		
S7.	Nearest Neighbor Algorithm	$O(k)$



Fig. 1. The 47<sup>th</sup> (left) and the 48<sup>th</sup> image (right) of an aerial sequence

TABLE II. RESULTS FOR THE FIRST AND THE SECOND IMAGE OF BOAT DATA SET [19], THE FIFTH AND THE SIXTH IMAGE OF BIKE DATA SET[19], THE 47<sup>TH</sup> AND 48<sup>TH</sup> IMAGE OF AN AERIAL SEQUENCE

Image	Octaves	Scales Per Octave	Threshold	Interest Points	Matches	Performance Decrease
Boat 1	1 2 3 4	4	0.002	2360		--
Boat 2	1 2 3 4	4	0.002	2500	210	
Boat 1	1 2	4	0.002	2050		15.7%
Boat 2	1 2	4	0.002	2227	177	
Bike 5	1 2 3 4	4	0.0022	101		--
Bike 6	1 2 3 4	4	0.0022	74	36	
Bike 5	1 2	4	0.0022	59		58.3%
Bike 6	1 2	4	0.0022	30	15	
Aerial 47	1 2 3 4	4	0.0022	349		--
Aerial 48	1 2 3 4	4	0.0022	278	59	
Aerial 47	1 2	4	0.0022	202		64.4%
Aerial 48	1 2	4	0.0022	144	21	

### A. UNDERLINING PRINCIPLES

Since the accuracy of any SURF-based vision system relies heavily on the number of matched points, it is essential to keep the number of *matched* points high; the detection of a large number of interest points does not guarantee a large number of matches while, conversely, one may obtain a significant number of matches from a comparatively small number of interest points. Rather than preferring lower octaves based on the assumption that they detect more interest points, the focus here is on obtaining the maximum number of interest point matches without preferring any particular octave, while keeping the computational cost as low as possible.

To reduce computation, the developers of SURF recommended doubling the spatial sampling interval when moving from lower octaves to higher octaves during detection stage [2]. Since detected interest points can be arbitrarily close together in the image, this non-uniform sampling inevitably incurs loss of

accuracy [1-2], which can be overcome by sampling all four octaves with a sampling rate of unity. This provides the maximum performance configuration (MPC) in terms of detected and matched interest points at a particular blob response threshold. Table III demonstrates the significance of sampling rate on the images of Fig. 1, where the MPC results are obtained with a unit-sampling modification to OpenSURF. MPC provides 135 interest point matches (Table III) as opposed to the 59 interest point matches (Table II) provided by the non-uniformly sampled configuration with four octaves at a threshold of 0.0022. Thus, in this particular case, the performance of the non-uniformly sampled approach with four octaves is only 43.7% of the performance of MPC processed at uniform sampling rate of unity, a significant reduction in matching performance. Reducing it further by discarding octaves compromises accuracy. Since the number of matches for the non-uniformly sampled SURF with four octaves cannot be increased beyond 59 at threshold of 0.0022, the threshold can be lowered to 0.0009 (Table III) to find an equal number of matched points (135), at the cost of processing 44.4% more interest points than MPC.

TABLE III. RESULTS FOR THE 47<sup>TH</sup> AND THE 48<sup>TH</sup> IMAGE OF AN AERIAL SEQUENCE WITH SAMPLING INTERVAL = 1 AND SAMPLING INTERVAL = 1, 2, 4 AND 8

Image	Octaves	Scales Per Octave	Threshold	Sampling Interval	Interest Points	Matches	Interest Points Processed Per Match
47	1 2 3 4	4	0.0022	1	767	135	10.4
48	1 2 3 4	4	0.0022	1	648		
47	1 2 3 4	4	0.0009	1 2 4 8	1317	134	19.0
48	1 2 3 4	4	0.0009	1 2 4 8	1231		

## B. THE BEST OCTAVES APPROACH

In *best octaves*, all four octaves are uniformly sampled to provide a fair opportunity for all octaves to show their maximum performance and to allow better evaluation of their relative performance. The proposed method then finds the two octaves that provide the best matching performance. For any given pair of images, the main steps of the proposed method are outlined in the following paragraphs.

**Step 1.** The matched interest points are calculated for the two given images for MPC (unity sampling rate and four octaves) and the number of them is considered as a reference for the later steps.

**Step 2.** The matched interest points are calculated for the two given images with unity sampling rate and the first two octaves. The ratio of the number of matched points,  $R$ , for octaves 1 and 2 against the

reference is computed to assess the effect of octave reduction. If  $R \geq 0.5$ , more than half the matches lie in octaves 1 and 2, so they are selected. Otherwise, we proceed directly to step 4, omitting step 3.

**Step 3.** Since the first two octaves are sampled at rates of 1 and 2 respectively in the original SURF algorithm, sampling every pixel of octave 2 for the selected best octaves in step 2 above may not make a significant difference to matching performance. To ensure that no extra computation is done, the matched interest points for the two images are calculated for the first two octaves with sampling rates of 1 and 2. The number of matched points for this non-uniform sampling case is compared with the reference; if greater than 0.4, then a non-uniform sampling rate of 1 and 2 is chosen for the selected best octaves in step 2 above. Otherwise, unity sampling rate is selected and the next three steps are skipped.

**Step 4.** The matched interest points are calculated for the two images with unity sampling rate and octaves 2 and 3 only. The Gaussian filters applied at different scales of octaves 2 and 3 are the same as in the original SURF algorithm. Similarly, after discarding octaves 1 and 2, the number of matched points is computed for octaves 3 and 4 by applying Gaussian filters to different scales, as in the original algorithm. The ratio of the number of matched points to the reference is then computed for the two cases and compared with each other to determine which is greater. If the maximum ratio is  $\geq 0.5$ , then octaves corresponding to that ratio are selected as the best octaves. A sampling rate of unity is chosen and the next two steps are skipped.

**Step 5.** For the two images, the matched interest points are computed with unity sampling rate for each of octaves 1 and 3, octaves 1 and 4, and octaves 2 and 4. The ratio of the numbers of matched points to the reference is then computed for the three cases and compared to determine which is greatest. If the maximum ratio is  $\geq 0.5$ , then octaves corresponding to that ratio are selected as the best octaves. A sampling rate of unity is chosen and the next step is skipped.

**Step 6.** Finally, the maximum of the six ratios calculated for octaves 1 and 2, octaves 2 and 3, octaves 3 and 4, octaves 1 and 3, octaves 1 and 4, octaves 2 and 4 with unity sampling rate is determined and the octaves corresponding to the maximum ratio are selected as the best octaves. A sampling rate of unity is

chosen for the best octaves except for the case when octaves 1 and 2 are selected, in which case the sampling rate is determined as described in step 3.

It should be noted that due to low frame-to-frame motion in image sequences with medium to high frame rate, the number of matches in Step 1 is computed for the first pair of images, then utilized as reference for the next few images based on frame rate, and then updated.

#### IV. PERFORMANCE

Fig. 2 shows the results for *best octaves* on the aerial images of Fig. 1 and the widely used UBC dataset [19]; the bars represent the number of points (read values from the left ordinate axis) while the line shows the matching ratio (read values from the right ordinate axis). Octaves 3 and 4 are the best for the aerial images; with these, the decrease in matching performance with respect to MPC is less than 30%, a sharp contrast to the results achieved with non-uniformly sampled SURF (Table II). Octaves 2 and 3 are the best for UBC, with a matching performance 32.3% less than MPC. To demonstrate the quality of interest point matches, Receiver Operating Characteristic (ROC) and Sensitivity-Specificity curves (Fig. 3) show that *best octaves* out-performs the other approaches comprehensively for the two sample cases, even including the reference (MPC) for the second one. Fig. 4 shows the interest point matches obtained using *best octaves* for image 1 and 6 of the UBC dataset.




Fig. 2. Results of *best octaves*: octaves 3 and 4 are selected as the best octaves for the 47<sup>th</sup> and the 48<sup>th</sup> image of the aerial sequence (left); octaves 2 and 3 are selected as the best octaves for image 1 and 6 of the UBC data set [19] (right)

McNemar’s test [21] was used to confirm that these performance differences are statistically significant, using the criterion that an algorithm must achieve at least 45% of the number of matched points obtained by MPC for it to be considered to have succeeded; although this figure is arbitrary, its value was found to have little effect on the conclusions drawn. To avoid inadvertent dataset dependencies, a total of 776 image pairs were employed from the Oxford [19], Copydays [22] and Blur [23] datasets used in [24-26]. The values tabulated in Table IV are the so-called *Z-scores* obtained from McNemar's test using:

$$Z = \frac{|N_{sf} - N_{fs}| - 1}{\sqrt{N_{sf} + N_{fs}}} \quad (1)$$

where  $N_{sf}$  is the number of occurrences where the first algorithm succeeded and the other algorithm failed, and so on. It should be noted that larger *Z-score* values indicate a more significant result. *Z-scores* of about 3 are equivalent to a confidence of about 99.5%, so the values in the table, which are substantially larger than this, provide incontrovertible evidence that *best octaves* out-performs the alternatives.

TABLE IV. RESULTS OF MCNEMAR’S TEST FOR BEST OCTAVES AND NON-UNIFORMLY SAMPLED SURF CONFIGURATIONS WITH 2 OCTAVES, 3 OCTAVES AND 4 OCTAVES

	2-octaves SURF (Sampling = 1, 2) PASS	2-octaves SURF (Sampling = 1, 2) FAIL	3-octaves SURF (Sampling = 1, 2, 4) PASS	3-octaves SURF (Sampling = 1, 2, 4) FAIL	4-octaves SURF (Sampling = 1, 2, 4, 8) PASS	4-octaves SURF (Sampling = 1, 2, 4, 8) FAIL
Best Octaves PASS	402	358	592	168	744	16
Best Octaves FAIL	0	16	16	0	16	0
Computed Z-Score	18.8		11.1		0	

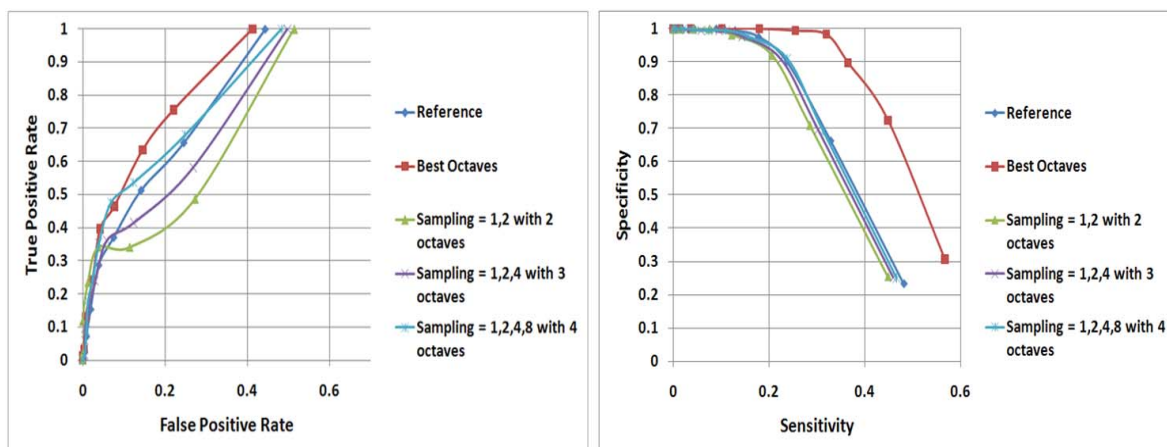


Fig. 3. ROC curve for the 47<sup>th</sup> and the 48<sup>th</sup> image of the aerial sequence (left); Sensitivity-Specificity curve for the first and the sixth image of the UBC data set [19](right)






Fig. 4. Interest point matches obtained using *best octaves* for the first and the sixth image of the UBC data set [19]

## V. REDUCTION IN COMPUTATION

For stage S1, the computation is equal for the selected *best octaves* and non-uniformly sampled SURF configurations (with 2, 3 and 4 octaves). The reductions in computation achieved by the non-uniformly sampled SURF alternatives for stages S2 and S3 with respect to the selected *best octaves* are not significant (less than 20% and 26% respectively). For the particular case when octaves 1 and 2 are selected as the best octaves with non-uniform sampling rate, the computation of *best octaves* is equal to non-uniformly sampled SURF with 2 octaves and achieves similar matching performance. As a result, the selected *best octaves* achieve small reduction in computation for S2 and S3 with respect to non-uniformly sampled SURF variants with 3 and 4 octaves (less than 15% and 21% for S2 and S3 respectively).

For the two image sets (aerial images and UBC), the reduction in computation for the last four stages of SURF achieved by the *best octaves* with respect to non-uniformly sampled SURF configurations having matching performance equal to *best octaves* is presented in Fig. 5. It is evident that the *best octaves* approach compensates for the extra computation done in S1-S3 by achieving a major reduction in computation in S4-S7 and hence comprehensively out-performs the non-uniformly sampled variants. As an example, we present here the timings obtained on a mobile robotic platform based on the Intel Atom (CPU N450 running at 1.66 GHz) for extraction of image features and nearest neighbor matching for

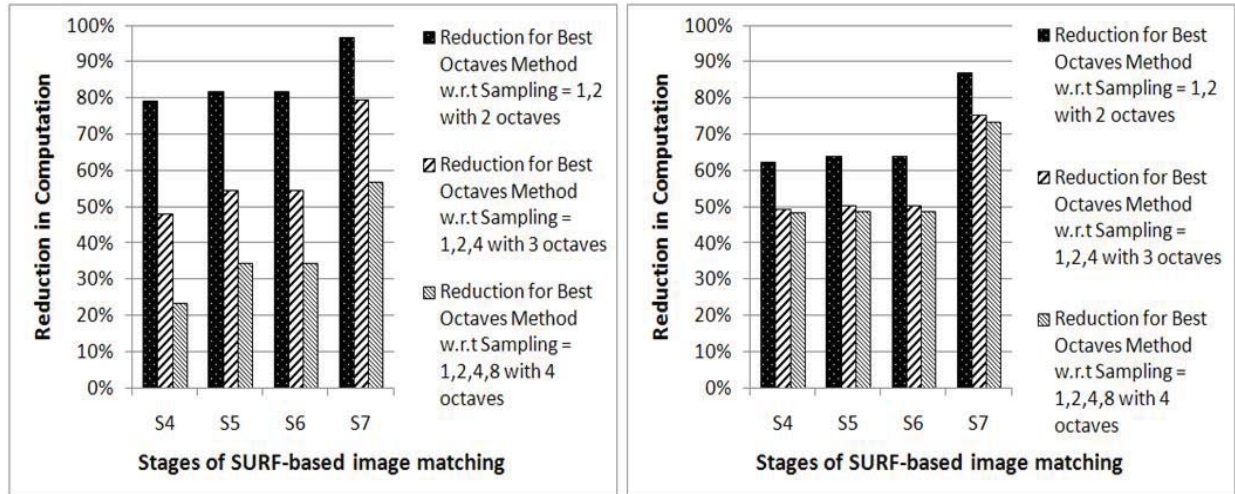


Fig. 5. Reduction in computation for stages S4, S5, S6 and S7 for *best octaves* with respect to the non-uniformly sampled SURF configurations having equal matching performance for 47<sup>th</sup> and 48<sup>th</sup> image of aerial sequence (left), and for image 1 and 6 of UBC dataset (right)

UBC (image 1 and 6): 30.58, 26.79 and 26.55 sec for non-uniformly sampled SURF variants with 2, 3 and 4 octaves respectively; 15.81 sec for *best octaves*. Clearly, the *best octaves* approach out-performs the others.

Since the reduction in computation for the first three stages (S1-S3) of the algorithm is independent of the images being analyzed, it is interesting to examine the performance of *best octaves* in terms of computation for the last four stages (S4-S7) of the algorithm. To make this analysis thorough, the amount of computation required by every algorithm in the last four stages for the 776 image pairs used in McNemar's test above was measured. Since it is tedious to equate the matching performance of all the algorithms by varying their threshold values (as done in Table III) for such a large number of image pairs, the number of interest points processed per match is used for comparing the computation of algorithms. As the amount of computation in stages S4, S5 and S6 of SURF is a function of the number of detected local maxima and interest points (Table I), this allows a fair comparison between *best octaves* and non-uniformly sampled SURF variants for these stages. For a comparative analysis of computation for the last stage (S7), interest point matching, the number of descriptor comparisons per match is used.

For every image pair, the number of interest points processed per match is calculated for *best octaves* and the 2-octaves SURF (sampling = 1, 2). The values obtained for *best octaves* are subtracted from the number of interest points processed per match for the 2-octaves SURF configuration so that a positive

value indicates that *best octaves* does less computation per match. To gauge the significance and magnitude of any computation reduction achieved by *best octaves*, Fig. 6 shows the histogram of the difference in number of interest points processed per match (*i.e.*, for stages S4–S6). The mean difference in the number of interest points processed per match in this particular case is 44.7, showing that *best octaves* computes and processes *nearly 45 times fewer* interest points than non-uniformly sampled SURF with 2 octaves, a huge reduction in computation. The histogram of the difference in number of descriptor comparisons per match (Stage S7) in Fig. 6 also demonstrates that the 2-octaves SURF (sampling = 1, 2) performs 14,940 more descriptor comparisons per match on average than *best octaves*.

Although the performance of 3-octaves SURF is better than 2-octaves (see Fig. 6), it still computes *12 times more* interest points per match than *best octaves* on average. The non-uniformly sampled 3-octaves SURF also requires, on average, 6,159.5 more descriptor comparisons per match. Finally, the two histograms for the 4-octaves (sampling = 1, 2, 4, 8) case again demonstrate the dominance of *best octaves*: on average, it processes *5 times fewer* interest points per match than 4-octaves SURF. Similarly, for the matching stage, *best octaves* achieves a significant reduction in computation as 4-octaves SURF requires 3,991.4 more descriptor comparisons per match on average.

## VI. CONCLUSIONS

This paper has proposed a novel method for reducing the computational complexity of SURF, namely an intelligent reduction in the number of SURF octaves. The approach focuses on the description and matching stages of SURF, yielding a significant reduction in computation at the cost of little extra calculation in the detection stages, and the reduction in computation can benefit both software and hardware implementations. It was found that the impact on matching performance is slight. It is hoped that this work may pave the way for the use of techniques like SURF in battery-operated robots, for which low power consumption is critical.

## ACKNOWLEDGEMENTS

The authors would like to thank the anonymous reviewers for their valuable comments and suggestions.

This work was supported in part by the UK EPSRC under grant EP/I500952/1.




Fig. 6. Histogram of difference in number of interest points processed per match for *best octaves* and non-uniformly sampled SURF with 2-octaves, 3 octaves and 4-octaves (top left to bottom left); Histogram of difference in number of descriptor comparisons per match for *best octaves* and non-uniformly sampled SURF with 2-octaves, 3 octaves and 4-octaves (top right to bottom right)

## References

- [1] D. Lowe, "Distinctive Image Features from Scale-Invariant Keypoints," *International Journal of Computer Vision*, vol. 60, no. 2, pp. 91-110, 2004.
- [2] H. Bay, A. Ess, T. Tuytelaars *et al.*, "Speeded-Up Robust Features (SURF)," *Computer Vision and Image Understanding*, vol. 110, no. 3, pp. 346-359, 2008.
- [3] Y. Ke, and R. Sukthankar, "PCA-SIFT: A More Distinctive Representation for Local Image Descriptors," *Proc. IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 2, pp. 506-513, 2004.
- [4] W. Cheung, and G. Hamarneh, "n-SIFT: n-Dimensional Scale Invariant Feature Transform," *IEEE Transactions on Image Processing*, vol. 18, no. 9, pp. 2012-2021, 2009.
- [5] A. Abdel-Hakim, and A. Farag, "CSIFT: A SIFT Descriptor with Color Invariant Characteristics," *Proc. IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 2, pp. 1978-1983, 2006.
- [6] K. Mikolajczyk, "Detection of Local Features Invariant to Affine Transformations Application to Matching and Recognition," PhD Thesis, Ecole Doctorale Mathématiques, Sciences et Technologies de l'Information, Informatique (MSTII), Institut National Polytechnique de Grenoble (INPG), France, 2002.
- [7] K. Mikolajczyk, and C. Schmid, "A Performance Evaluation of Local Descriptors," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 27, no. 10, pp. 1615-1630, 2005.

- [8] N. Zhang, "Computing Parallel Speeded-Up Robust Features (P-SURF) via POSIX Threads," *Lecture Notes in Computer Science*, vol. 5754, pp. 287-296, Sep. 2009.
- [9] S. Smith, and J. Brady, "SUSAN—A New Approach to Low Level Image Processing," *International Journal of Computer Vision*, vol. 23, no. 1, pp. 45-78, 1997.
- [10] E. Rosten, R. Porter, and T. Drummond, "Faster and Better: A Machine Learning Approach to Corner Detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, no. 1, pp. 105-119, 2010.
- [11] T. Tuytelaars, and K. Mikolajczyk, "Local Invariant Feature Detectors: A Survey," *Foundations and Trends in Computer Graphics and Vision*, vol. 3, no. 3, pp. 177-280, 2008.
- [12] H. Bay, "From Wide-baseline Point and Line Correspondences to 3D," Doctor of Sciences Thesis, Swiss Federal Institute of Technology, ETH Zurich, 2006.
- [13] H. Bay, T. Tuytelaars, and L. Van Gool, "Surf: Speeded Up Robust Features," *Proc. 9th European Conference on Computer Vision*, pp. 404-417, 2006.
- [14] S. Lee, Y. Zhang, Z. Fang *et al.*, "Accelerating Mobile Augmented Reality on a Handheld Platform," *Proc. 27th IEEE International Conference on Computer Design*, pp. 419-426, October 2009.
- [15] N. Cornelis, and L. Van Gool, "Fast Scale Invariant Feature Detection and Matching on Programmable Graphics Hardware," *Computer Vision and Pattern Recognition (CVPR) Workshop*, June 2008.
- [16] P. Viola, and M. Jones, "Rapid Object Detection using a Boosted Cascade of Simple Features," *Proc. IEEE Computer Society Conference on Computer Vision and Pattern Recognition* vol. 1, pp. 511-518, 2001.
- [17] A. Neubeck, and L. Van Gool, "Efficient Non-Maximum Suppression," *Proc. 18th International Conference on Pattern Recognition (ICPR)*, vol. 3, pp. 850-855, 2006.
- [18] M. Brown, and D. Lowe, "Invariant Features from Interest Point Groups," *Proc. BMVC*, 2002.
- [19] K. Mikolajczyk. "Oxford Data Sets," <http://www.robots.ox.ac.uk/~vgg/research/affine/>, 2007.
- [20] C. Evans. "Open SURF," <http://code.google.com/p/opensurf/>.
- [21] A. Clark, and C. Clark. "Performance Characterization in Computer Vision A Tutorial," <http://peipa.essex.ac.uk/benchmark/tutorials/essex/tutorial.pdf>, 1999.
- [22] H. Jegou. "Copydays Data Set," <http://lear.inrialpes.fr/~jegou/data.php>.
- [23] J. van de Weijer. "Blur Data Set," <http://lear.inrialpes.fr/people/vandeweijer/data>.
- [24] K. Mikolajczyk, T. Tuytelaars, C. Schmid *et al.*, "A Comparison of Affine Region Detectors," *International Journal of Computer Vision*, vol. 65, no. 1, pp. 43-72, 2005.
- [25] M. Douze, H. Jégou, H. Singh *et al.*, "Evaluation of GIST descriptors for Web-Scale Image Search," *Proc. CIVR*, 2009.
- [26] J. van de Weijer, and C. Schmid, "Blur Robust and Color Constant Image Description," *Proc. IEEE ICIP*, pp. 993-996, 2006.



Shoaib Ehsan obtained a BSc in Electrical Engineering from University of Engineering and Technology, Taxila, Pakistan in 2003. He then worked in the industry for four years and developed expertise in embedded systems design. Currently, he is doing PhD in Computer Vision in the School of Computer Science and Electronic Engineering at the University of Essex, U.K. He is a recipient of the University of Essex Post Graduate Research Scholarship. He was also awarded the prestigious Overseas Research Student Scholarship. His current research interests include image feature detection, description and matching, performance characterization of vision algorithms and real-time embedded vision applications. He is a member of the IEEE and of the British Machine Vision Association.



Nadia Kanwal received her MSc degree in Computer Science from University of Essex, United Kingdom in 2009. Currently, she is a PhD scholar at School of Computer Science and Electronic Engineering, University of Essex. She is a member of the IEEE, the IET and the British Machine Vision Association. Her research interests include local feature detection, description and matching techniques for tracking and navigation applications.



Adrian Clark obtained a BSc in Physics at the University of Newcastle upon Tyne in 1979 and a PhD in image processing from Queen Elizabeth College, University of London in 1983. After postdoctorate research into parallel image processing, he worked for British Aerospace, researching in the general area of object recognition. He joined the University of Essex in 1988, where he is now a Reader in the School of Computer Science and Electronic Engineering. During 1991-93, he headed UK involvement in the development of the International Standard for Image Processing and Interchange and was editor of the first part of that Standard. He is a former secretary of the British Machine Vision Association and has chaired Technical Committee 5 of the International Association for Pattern Recognition.

He is a fervent evangelist for quantitative evaluation and comparison in vision, leading him to set up PEIPA, one of the first imaging resources on the web. His current research explores the interaction of computer vision and computer graphics, remote sensing, and building vision systems using machine learning. He is a member of the IEEE and of the Institute of Physics.



Professor Klaus D. McDonald-Maier received Dipl.-Ing. and MS degrees in electrical engineering from the University of Ulm, Germany, and the École Supérieur de Chimie Physique Électronique de Lyon, France, in 1995, respectively. In 1999, he received a doctorate in computer science from the Friedrich Schiller University, Jena, Germany. Prof. McDonald-Maier worked as a systems architect on reusable microcontroller cores and modules at Infineon Technologies AG's Cores and Modules Division in Munich, Germany and as a lecturer in electronic engineering at the University of Kent, Canterbury, United Kingdom. In 2005, he joined the University of Essex, Colchester, United Kingdom, where he is a Professor in the School of Computer Science and Electronic Engineering. His current

research interests include embedded systems and system-on-chip (SoC) design, development support and technology, parallel and energy efficient architectures, the application of soft computing and image processing techniques for real world problems. He is a senior member of the IEEE and a Fellow of the IET.