**QUEEN'S
UNIVERSITY
BELFAST**

Following the Footsteps of Others: Techniques for Automatic Shoeprint Classification

Crookes, D., Gueham, M., Bouridane, A., & Su, H. (2007). *Following the Footsteps of Others: Techniques for Automatic Shoeprint Classification*. 67-74. Paper presented at NASA/ESA Conference on Adaptive Hardware and Systems, AHS2007, Edinburgh, United Kingdom. <https://doi.org/10.1109/AHS.2007.56>

Queen's University Belfast - Research Portal:

[Link to publication record in Queen's University Belfast Research Portal](#)

General rights

Copyright for the publications made accessible via the Queen's University Belfast Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The Research Portal is Queen's institutional repository that provides access to Queen's research output. Every effort has been made to ensure that content in the Research Portal does not infringe any person's rights, or applicable UK laws. If you discover content in the Research Portal that you believe breaches copyright or violates any law, please contact openaccess@qub.ac.uk.

Following the Footsteps of Others: Techniques for Automatic Shoeprint Classification

D Crookes, A Bouridane, H Su and M Gueham

Institute for Electronics, Communications and Information Technology,

Queen's University Belfast, Belfast BT3 9DT, UK

d.crookes@qub.ac.uk, a.bouridane@qub.ac.uk, h.su@qub.ac.uk, mgueham02@qub.ac.uk

Abstract

Shoeprint evidence is often left at crime scenes, but is not always exploited. There is an increasing amount of research in developing systems to provide more rapid identification of footwear tread patterns. The main problem is that scene of crime shoeprint images can be very significantly degraded. In this paper we identify some of the challenges of this emerging research area. We then review current approaches to this problem, and we present some novel methods and results for two different ways of addressing the problem - namely in the spatial domain and in the transform domain. In the spatial domain, improvements to existing techniques lead to two novel variations which we call the Modified Harris-Laplace (MHL) detector, and the enhanced SIFT descriptor. In the transform domain, we present results of a technique based on Phase-Only Correlation.

1. Introduction

Shoeprints are not normally regarded as biometric information, but the processing of shoeprints from scenes of crime shares some of the features and challenges associated with biometric identification. Shoeprint evidence is often found at crimes scenes – it has been estimated that 15–30% of all burglaries provide usable shoeprints that can be recovered from the crime scene [1]. Because of the pattern of repeated offences, rapid classification of such shoeprints would enable investigating officers not only to link different crimes, but to identify potential suspects while the crime is still ‘hot’.

The objectives of the forensic investigator are usually: (i) to determine the make and model of a shoe; (ii) to determine if a particular shoeprint was made by a specific shoe; (iii) to match the shoeprint with other

shoeprints, possibly from other crime scenes. In this paper we will be focusing on the first of these three objectives.

An image of a shoeprint is first obtained using a technique such as photography, gel or electrostatic lifting, or by making a cast when the impression is in soil, snow or sand. A matching of the shoeprint is then made against specific image databases containing current and previous shoeprint images and/or images of shoeprints found at other crime scenes, leading to a few candidate images, which is then used in the second stage. This matching process is usually in two stages: an initial classification which results in a hitlist of likely candidates, followed by a visual identification from this reduced set. For this paper, we are concerned with automating the first classification process.

2. Challenges for Shoeprint Classification

The identification of the classification of a shoeprint (typically the manufacturer and model) recovered from a scene of crime can be problematic, even when done manually by experts. This section outlines some of these. In the area of automatic classification, there are also certain *research* challenges facing researchers in this emerging discipline. This section also aims to identify some of these.

Fig. 1 shows a typical reference shoeprint image, and a scene of crime shoeprint image. The scene image shows one shoeprint superimposed on a less distinct one. The database of reference prints normally contains just one image of each class, captured under good conditions with a new shoe. Scene images are often very significantly degraded because of problems either in making the impression or in capturing it. Factors which give rise to degradation include: partial prints (only part of the shoeprint leaves an impression),

Fig. 1 A reference image (L) and scene image (R)

artefacts (visible pattern features on the surface, or occluding objects), superimposed shoeprints, pattern distortion due to uneven pressure or shoe wear, and blurring due to foot slippage or a fluid surface. There are also the standard problems of invariance with respect to rotation, translation, scale and noise. There are then further degradations arising from the capture process, which depend on the method used (e.g. ripples when gel is lifted).

As a research discipline, automatic classification of shoeprints is still relatively new and immature. For the discipline to mature, our experience in this area has identified a number of challenges which research groups and the research community needs to address:

(i) *Good data sets.* Building up a set of reference prints and scene prints to evaluate algorithm performance can be expensive and time-consuming. Commercial companies in the forensic sector have significant data sets and expertise, but these constitute valuable IP, and companies are understandably careful to protect their IP. Obtaining suitable scene images is particularly problematic, since there are also legal issues where these constitute evidence of real crimes. An additional difficulty is to make sure that the scene images have their counterparts in the reference database (to enable them to be identified!)

(ii) *Standard data sets.* As a discipline matures, it is vital to evaluate the performance of different approaches from different groups by comparing various methods using equivalent data sets. At present, it is common to synthesize scene images by degrading reference images in ways which mimic real world degradations. However, because shoeprint data sets are not (yet) public domain, it is difficult for researchers to compare their methods with others. One of the challenges for the emerging community is to

develop a public domain data set of reference shoeprint images, and scene images (even if the latter are synthesized rather than real scene images).

(iii) *Standard comparison measures.* There is a range of statistical measures used to evaluate classification performance and image retrieval. Some use precision vs recall; others use Receiver Operating Characteristic (ROC) curves, as is common in fingerprint analysis. Another is Cumulative Matching Characteristics. Other proprietary measures are sometimes used. The lack of a standard for presenting and comparing performance results makes progress harder to identify.

One of the challenges for the maturing research community in this area is to address these 'infrastructure' issues in a collaborative manner.

In the following sections we will review the main approaches to the automatic classification of shoeprint images, and in particular, we will present two novel approaches developed by the team at Queen's University Belfast – one in the spatial domain, and one in the transform domain.

3. Main Current Approaches

Research in automatic shoeprint classification has been reported for a little over a decade. Some of the techniques reported have focused on representing the small shape components which make up a pattern, whereas others extract features from the shoeprint without any subdivision. Similarly, some techniques are based in the spatial domain, while others operate in the transform domain, or a combination of the two.

Early research used techniques such as Fourier descriptors to model the pattern components, or fractals to model the complete shoe pattern [1,2,3].

More recent work by the group at UCD, Dublin [4] has used the Power Spectral Density (PSD) coefficients of the image, which are calculated using the Fourier Transform and used as features. Rotation invariance is achieved using a 'brute force' approach in which the query image is rotated in 1 degree steps.

The group at Sheffield use the technique of matching Edge Directional Histograms (EDH) [5]. The authors describe the interior shape information of a shoeprint image in terms of its significant edges, and use a histogram of the edge direction as the signature of a shoeprint image. This method first extracts the edges using a Canny edge detector. To obtain rotation invariance, they compute the 1-D FFT of the normalized edge direction histogram and take it as the final signature of the shoeprint image.

The team at Queen's University Belfast has explored a number of approaches. In the spatial

domain, one initial approach used an enhancement of the pattern spectrum method proposed by Maragos. This was extended by defining the topological spectrum and combining it with the pattern spectrum to give a more effective hybrid approach. The method was evaluated using a database of 500 clean shoeprints (supplied by Foster and Freeman, and by Forensic Science Services), and several synthesized degraded versions of this data set including rotation, noise, scale, partial images and background addition. Results under these degraded conditions were an improvement on the available results from other methods.

More recently, two further methods have been developed which perform significantly better than the topological/pattern spectrum method. These two methods are presented in more detail below, together with some indicative results. The methods are (i) an enhancement of the Local Image Features approach (SIFT) in the spatial domain, and (ii) the use of Phase Only Correlation (POC) in the transform domain.

4. Using Local Image Features

A local feature here refers to any point and its neighbourhood in an image where the signal changes significantly in 2D space. Conventional “corners”, such as L-corners, T-junctions and Y-junctions satisfy this, but so do isolated points, the endings of branches and any location with significant 2D texture. Local Feature Descriptors have proved to be very successful in applications such as image retrieval and matching [6,7], object recognition and classification [8,9], and wide baseline matching [10]. There is a range of different scale-invariant and affine-invariant local feature detectors and robust local feature descriptors, which have been investigated [11-16].

Mikolajczyk et al in [14] have given a detailed comparison of six state-of-the-art local feature detectors. They conclude firstly that all detectors involved were complementary, and could be used together to achieve the best performance. Secondly, the Harris-Affine [11] and Hessian-Affine [14] detectors provide more features than the other detectors, and are useful in matching scenes with occlusion and clutter, although the MSER (Maximal Stable Extremal Regions) detector [13] often obtains the highest score for repeatability. In our case, an invariant local feature refers only to the translation-, rotation- and scale-invariant local region – perspective invariance is not normally relevant. Mikolajczyk et al [15,16] have evaluated ten state-of-the-art local feature descriptors in the presence of real geometric and photometric transformations. They conclude that an

extension of the SIFT descriptor [8, 12], called GLOH [14], performs slightly better than SIFT itself, and both beat the other descriptors involved. They have also suggested that the local feature detectors, such as Hessian-Affine and Hessian-Laplace, which mainly detect blob-like structures, only work well with a larger neighbourhood. But this conflicts with the locality property of local image features.

4.1 Outline of the enhanced LIF method

Typically, a local image feature should have four properties: locality, repeatability, distinctiveness, and robustness to different degradations. The results in this section are based firstly on an enhanced feature detector, which combines a scale adaptive Harris corner detector with a Laplace-based scale selection. Here, the location of the features is determined by the scale adaptive Harris corner detector, and the characteristic size of a local feature depends on the scale of the blob-like structure around this corner. This is determined by the automatic Laplace-based scale selection.

Secondly, for each local feature, an enhanced SIFT descriptor is computed to represent the feature. This descriptor actually further enhances the GLOH by using a circle binary template for rotation invariance, and binning the GLOH histogram with a range of 180° rather than 360° for inverse lighting robustness.

Finally, the matching of descriptors in two images has been conducted by combining the nearest neighbour with the threshold-based screening, i.e. two descriptors are matched only if one is the nearest neighbour to the other one, and at the same time the distance between them is less than a threshold. The distance between two images is computed from only the matched pairs.

Thus the novel aspects of this approach are: first, a new robust and distinctive local feature detector, Modified Harris-Laplace; and second, an improved SIFT descriptor for the local features detected at the first stage. We have investigated the application of these new local image features to automatic shoeprint image classification, and compared this technique with three other methods mentioned above, using the same data sets of clean and degraded shoeprints.

4.2 Applying the enhanced LIF approach

The similarity between two images depends on the matching strategy of the local features. For the sake of retrieval speed, we apply the nearest neighbour and thresholding jointly to compute the distance between

Fig. 2 Matching result of two images: here 23 out of 32 matches are correct.

two images, i.e. for each descriptor in one image, the nearest neighbour in another image is found as a potential match, then only those matches whose distance is below a threshold are selected as the final matches. The similarity of two images is computed from the summation of $\exp(-d)$, where d denotes the distances of the final matches. Of course, there are many other strategies for computing the similarity or matching score between two images. The example of image matching in Fig. 2 applies the nearest neighbour to obtain the initial matches, and then the RANSAC (Random Sample Consensus) method is used to fit the geometric transformation and then reject the mismatches. The correspondences are shown in Fig. 2.

4.3 Test shoeprint image databases

Our experiments and evaluations for shoeprint image retrieval are conducted on a database of 500 ‘clean’ reference images, called *dClean*. To simulate the scene images, six degraded image sets have been synthesized from this base dataset, as follows:

(a) *dNoisy* – Five different levels of Gaussian noise are added to each shoeprint in the base data set, giving 2500 noisy prints. The noise level (σ) varies from 10 to 50 in steps of 10.

(b) *dRotate* – Each clean shoeprint is rotated in five random orientations in the range of 0° - 90° . The selection of this range, rather than 0° - 360° , is based on the fact that the algorithms developed in this paper are flip invariant both in horizontal and vertical directions.

(c) *dRescale* – Each clean print is rescaled with five random scale ratios in the range of 0.35 to 0.65.

(d) *dPartial* – Five partial shoeprint images are generated for each shoeprint in the first dataset. The percentage of the partial shoeprint which remains

varies from 40% to 95%. To create a partial shoeprint image, the silhouette of a complete shoeprint is extracted first, and then two appropriate points on the silhouette perimeter are chosen to generate a straight line across the shoeprint.

(e) *dScene* – To simulate the effect of surface patterns and background interfering with the scene image, for each clean shoeprint we randomly select one of five genuine scene images as the background, and superimpose the clean shoeprint on the selected background. This again gives a data set of 2500 synthesized scene images.

(f) *dComplex* – This dataset contains 50 shoeprint images with a combination of synthesized degradations – complex backgrounds, geometric transformations, significant ‘cut-out’, or their combinations.

To make computation time feasible, in the results below we then made a selection of 50 images from each of the first five degraded datasets.

4.4 Results and performance evaluation

Our evaluation compares several of the methods described above, runs them using the test databases described in 4.3, and presents the results using *Cumulative Matching Characteristics* (CMC). This measure is suitable for the case where the reference database contains only one relevant record for the query record, and it is described in [4] as being able to answer the question ‘‘What is the probability of a match if one looks at the first n percent of the database records?’’. It can be estimated by computing the proportion of times during trails of the searching when the relevant record appears in the first n percent of the sorted records in the database. The performance of a retrieval system is reflected as a graph which plots

the probability of a match (also called cumulative matching score, vertical axis) against the percentage of the records reviewed (horizontal axis).

We use CMC to compare four methods: Edge Directional Histogram (EDH), Power Spectral Distribution (PSD), Pattern & Topological Spectra (PTS), and our enhanced Local Image Feature (LIF). The six experiments are conducted as follows: we take each of the 50 shoeprint images from the degraded data sets as a query image, and search against the set *dClean* of reference prints. For each data set, we compute the CMC curve. For brevity, only three of the six sets of results are shown in Fig. 3.

Fig. 3 Performance evaluation of four methods (EDH, PTS, PSD, and LIF) using CMC on three degraded datasets. The main diagonal is the worst case.

The results of the six tests suggest that:

- (i) For noise, rescaling and partial image degradations, PSD and LIF can achieve almost perfect results. For rotation and scene background addition, only LIF can achieve almost perfect results.
- (ii) EDH and PTS are not quite as good as PSD and LIF for noise, rescaling, partial image and rotation degradations, but are still efficient approaches considering that the cost of the two methods is significantly less than the other two. PTS beats EDH in most cases. The only exception is on the rescaled database. Furthermore, it is a surprise that both PTS and EDH beat PSD on the dScene dataset, even though they both have a much lower cost than PSD.
- (iii) LIF works very well for all kinds of degradations. It clearly beats other methods on the dataset with combined degradations. However, it is more time-consuming than EDH and PTS.

Of course, we should bear in mind the limitations of these particular tests, and the synthetic nature of the ‘scene’ images. Nevertheless, the experiments indicate that promising results can be achieved by operating in the spatial domain.

5. Using Phase Only Correlation

An alternative approach to spatial feature selection is to operate in the transform domain. Options include Fourier space, and various multi-resolution wavelet spaces. This section describes one promising approach which has been developed by the team at Queen’s University Belfast. This uses the phase component of the Fourier transform.

In the Fourier domain, the phase information is much more important than the magnitude in preserving the features of image patterns, as proved by Oppenheim et al [17]. A simple illustration for shoeprint images is given in fig. 4.

Fig. 4. (a) Original shoeprint image A. (b) Original shoeprint B. (c) Image synthesized from the FFT phase of image B and the magnitude of image A.

If any two images g_1 and g_2 are identical, their POC function will be a Dirac δ -function centered at the origin and having the peak value 1. When matching similar images, the POC approach produces

Fig. 5. (a) Original shoeprint image A. (b) Noisy partial shoeprint B from A. (c) POC of A and B. (d) Conventional correlation of A and B

a sharper correlation peak compared to conventional correlation, as shown in Fig. 5.

5.1 The proposed POC-based method

The proposed method uses the POC approach combined with a spectral weighting function. Spectral weighting functions have been already used with the POC technique in image registration in order to enhance the registration accuracy [18]. In this work, we propose to use a band-pass-type spectral weighting function to improve the recognition rate by eliminating high frequency components which have low reliability and which preserve the sharpness of the correlation peak. The proposed weighting function $W(u,v)$ has the same shape as the spectrum of a Laplacian of Gaussian (LoG) function and is given by:

$$W(u,v) = \left(\frac{u^2 + v^2}{\alpha} \right) e^{-\frac{u^2 + v^2}{2\beta^2}} \quad (1)$$

where β is a parameter that controls the function width and α is used for normalization only. Thus, the modified phase-only correlation (MPOC) function $\tilde{q}_{g_1 g_2}(x, y)$ of images g_1 and g_2 is given by:

$$\tilde{q}_{g_1 g_2}(x, y) = F^{-1} \left\{ W(u, v) \frac{G_1(u, v) G_2^*(u, v)}{|G_1(u, v) G_2^*(u, v)|} \right\} \quad (2)$$

The peak value of the MPOC function $\tilde{q}_{g_1 g_2}(x, y)$ is also invariant to translation and brightness changes.

5.2 POC-based shoeprint matching algorithm

A schematic of the POC-based shoeprint matching algorithm is shown in Fig. 6. In response to an unknown shoeprint image g_i , the algorithm matches g_i to each data base image g_n ($n=1 \dots M$ where M is the size of the database) and determines the corresponding matching score. The matching algorithm consists of the following steps:

- i) Calculate the Fourier transform of g_i and g_n using FFT to obtain G_i and G_n .
- ii) Extract the phases of G_i and G_n and calculate the cross-phase spectrum $Q_{g_n g_i}$.
- iii) Calculate the modified cross-phase spectrum $\tilde{Q}_{g_n g_i}$ by modifying $Q_{g_n g_i}$ using the spectral weighting function W .
- iv) Calculate the inverse Fourier transform of $\tilde{Q}_{g_n g_i}$ to obtain the MPOC function $\tilde{q}_{g_n g_i}$.
- v) Determine the maximum value of $\tilde{q}_{g_n g_i}$. This value will be considered as the matching score between images g_i and g_n .

Fig. 6. Schematic of the MPOC matching algorithm

The use of the band-pass-type weighting function W (equation (1)) will eliminate meaningless high frequency components without significantly affecting the sharpness of the correlation peak (since very low frequency components will also be attenuated).

In this work, we take the peak value of the MPOC function to be the similarity measure for image matching: if two images are similar, their MPOC function will give a distinct sharp peak; if they are dissimilar, then the peak drops significantly.

The result of matching an input image g_i with all database images, using the algorithm described above, is a ranked list of m shoeprints from the database ($m \ll$ database size), ranked from best match to the worst.

5.3 Experimental results

The tests for this algorithm were performed using a database of 100 clean shoeprint images as above. To evaluate the robustness of the method to different degradations, a total of 64 synthetic scene images were generated from each original shoeprint image (giving a total of 6400 test images). The test images were grouped into 4 data sets:

Set 1: four clean *partial* shoeprint images are synthesized from each original image by dividing it into four quarters: i) left toes and midsole, ii) right toes and midsole, iii) left heel and iv) right heel.

Set 2: five noisy versions of each partial image in Set 1 were generated by adding white Gaussian noise (with zero mean and standard deviations $\sigma = 20, 40, 60, 80$ and 100). Set 2 has 2000 images.

Set 3: five blurred versions of each partial image in Set 1 were generated by blurring, with motion blur of length L ($L=10, 20, 30, 40, 50$ pixels) and angle $\theta=90$ degrees (vertical blur) to simulate shoeprint blurring caused by foot slippage in the real world.

Set 4: five textured background versions of each partial image in Set 1 were generated by pasting each partial shoeprint image into five texture images selected from the Brodatz album [19]: D16 (Weave), D19 (Wool), D24 (Leather), D68 (Wood) and D94 (Brick). See three samples in fig. 7.

Fig. 7. Some texture images used for generating Set 4

During the evaluation process, each of the 6400 test images was used as input, and matched against all 100 original images. The rank of the correct match was determined. For each type of degradation, the proportion of times a correct match appeared first in the hitlist was determined

Three algorithms were tested in this way: our modified MPOC method, the standard POC method, and the previously-mentioned PSD-based algorithm. MPOC and POC denote the phase only correlation algorithms with and without the spectral weighting function, respectively. The parameters of the weighting function used during the tests were $\beta=10, 20, 30, 40, 50$ and 60 , with $\alpha = 4\pi\beta^4$ (to normalize the maximum of the MPOC function to 1: when matching two identical images). The results obtained are shown in table I. Only results corresponding to $\beta = 50$ (the best value) are shown in table I.

From these results, it can be seen that the phase based algorithms (POC and MPOC) outperform the PSD based one even without the use of the spectral weighting function. It can be also observed that the PSD based algorithm is very sensitive to blur and textured background. For the phase based approaches, the use of the weighting function (in MPOC) results in an improvement for blurred images without affecting the performance of the method when processing clean, noisy or textured background images. The best results were obtained for a weighting function with $\beta=50$, where 100% identification was obtained.

TABLE I. First rank recognition rate (%) using PSD [4], POC and MPOC based algorithms

Algorithms		PSD	POC	MPOC $\beta=50$
Test images				
1. Clean partial prints		96.25	100	100
2. Noisy partial prints	$\sigma=20$	95.75	100	100
	$\sigma=40$	93.5	100	100
	$\sigma=60$	88.5	100	100
	$\sigma=80$	76.5	100	100
	$\sigma=100$	60.75	100	100
3. Blurred partial prints	$L=10$	28.5	100	100
	$L=20$	13.25	100	100
	$L=30$	13	100	100
	$L=40$	13	97.75	100
	$L=50$	9.75	95.5	100
4. Partial prints with textured background	D16	13	100	100
	D19	4	100	100
	D24	1.25	100	100
	D68	5.5	100	100
	D94	64.5	100	100

The MPOC method appears to offer great potential for robustness to several kinds of degradation. But because we are still in the process of standardising data sets and performance measures, we are not yet in a position to compare the results in detail with the enhanced LIF method discussed in section 4. The main disadvantage encountered so far with the POC-based method is that it is not rotation invariant. Methods of addressing this are under investigation, including variants of the 'brute force' approach.

6. Conclusions

This paper has presented some of the challenges of automatic shoeprint classification using scene of crime images. While the discipline is still emerging, there are several approaches which demonstrate promising robustness in the face of several simulated degradations which are characteristic of real scene images. In particular, the paper has presented results which show that an enhanced Local Image Feature-based method, and a modified Phase Only Correlation approach consistently out-perform other methods.

In order for research in this area to mature, the paper has also proposed a number of steps which the research community could usefully take to enable more rigorous performance evaluation and comparison. These recommendations include the establishment of a set of standard data sets which are available to all researchers. This would include not only a large set of reference shoeprint images, but also sets of scene images (probably synthesized from the reference images) which have been degraded in various appropriate ways. Another important step would be to standardize on a set of evaluation measures, so that new methods can be compared rigorously with existing methods on a level playing field.

7. Acknowledgements

This work has been supported by EPSRC Grant No. EP/C008057/1. The authors would also like to acknowledge with gratitude the support of Foster & Freeman Ltd, and Forensic Science Services.

8. References

- [1] G. Alexandre, "Computerized classification of the shoeprints of burglar's soles", *Forensic Science International*, Vol. 82, pp. 59-65, 1996.
- [2] Z. Geradts and J. Keijzer, "The image data REBEZO for shoeprint with developments for automatic classification of shoe outsole designs", *Forensic Science International*, Vol. 82, pp. 21-31, 1996.
- [3] A.G. Alexander, A. Bouradine, D. Crookes, "Automatic classification and recognition of shoeprints," Special issue of the information bulletin for shoeprint/toolmark examiners, Vol. 6. No. 1, pp. 91-104, Mar. 2000.
- [4] P. D. Chazal, J. Flynn, and R.B. Reilly, "Automated processing of shoeprint images based on the Fourier transform for use in forensic science," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 27, No. 3, pp. 341-350, Mar. 2005.
- [5] L. Zhang, and N.M. Allinson, "Automatic Shoeprint Retrieval System for use in Forensic Investigations," 5th Annual UK Workshop on Computational Intelligence, 2005
- [6] C. Schmid and R. Mohr, "Local greyvalue invariants for image retrieval," *IEEE Trans. on PAMI*, Vol 19-5, pp. 530-535, 1997.
- [7] L. Ledwich, and S. Williams, "Reduced SIFT Features for Image Retrieval and Indoor Localisation," *Australasian Conf. on Robotics and Automation*, Canberra, 2004.
- [8] D. Lowe, "Object Recognition from Local Scale-Invariant Features," *ICCV'99*, pp.1150-1157. 1999.
- [9] K. Mikolajczyk, B. Leibe, and B. Schiele, "Local Features for Object Class Recognition," *ICCV'05*, pp.1792-1799, 2005.
- [10] T. Tuytelaars and L. Van Gool. "Matching widely separated views based on affine invariant regions," *IJCV*, Vol. 59, No. 1, pp. 61-85, 2004.
- [11] K. Mikolajczyk and C. Schmid, "Scale and Affine invariant interest point detectors," *IJCV*, Vol. 60, No. 1, pp.63-86, 2004.
- [12] D. Lowe, "Distinctive image features from scale-invariant keypoints", *IJCV*, Vol.60, No. 2, pp. 91-110, 2004.
- [13] J. Matas, O. Chum, M. Urban, and T. Pajdla, "Robust wide baseline stereo from maximally stable extremal regions," *BMVC*, Vol. 1, pp. 384--393, 2002.
- [14] K. Mikolajczyk, T. Tuytelaars, C. Schmid, A. Zisserman, J. Matas, F. Schaffalitzky, T. Kadir, L. Gool, "a comparison of affine region detectors", *IJCV*, Vol. 65, No.1-2, pp.43-72, 2005.
- [15] K. Mikolajczyk, C. Schmid, "A performance evaluation of local descriptors," *CVPR'03*, Vol. 2, pp. 257-264, 2003.
- [16] K. Mikolajczyk, C. Schmid, "A performance evaluation of local descriptors," *IEEE Trans. on PAMI* 27(10):1615-1630. 2005.
- [17] A.V. Oppenheim and J.S. Lim, "The importance of phase in signals," *IEEE Proc.*, vol. 69, no. 5, pp. 529-541, 1981.
- [18] K. Takita, T. Aoki, Y. Sasaki, T. Higuchi, and K. Kobayashi, "High-accuracy subpixel image registration based on phase-only correlation," *IEICE Trans. Fundamentals*, vol. E86-A, no. 8, pp. 1925-1934, Aug. 2003.
- [19] P. Brodatz, *Textures: A Photographic Album for Artists and Designers*, Dover, New York, 1966.