

Histogram of Weighted Local Directions for Gait Recognition

Sabesan Sivapalan, Daniel Chen, Simon Denman, Sridha Sridharan and Clinton Fookes

Image and Video Research Laboratory,
Queensland University of Technology,
GPO Box 2434, 2 George St.
Brisbane, Queensland 4001.

{sivapalen.sabesan, daniel.chen, s.denman, s.sridharan, c.fookes}@qut.edu.au

Abstract

In this paper, we explore the effectiveness of patch-based gradient feature extraction methods when applied to appearance-based gait recognition. Extending existing popular feature extraction methods such as HOG and LDP, we propose a novel technique which we term the Histogram of Weighted Local Directions (HWLD). These 3 methods are applied to gait recognition using the GEI feature, with classification performed using SRC. Evaluations on the CASIA and OULP datasets show significant improvements using these patch-based methods over existing implementations, with the proposed method achieving the highest recognition rate for the respective datasets. In addition, the HWLD can easily be extended to 3D, which we demonstrate using the GEV feature on the DGD dataset, observing improvements in performance.

1. Introduction

Automatic human recognition is a major research area in computer vision, with gait being an attractive biometric due to its ability to operate on low resolution video, without interrupting or interacting with the subject. Recent algorithms can achieve close to 100% under controlled conditions [10], subject to variations in capture conditions, appearance changes (*e.g.* clothing), and dataset size.

Most recent research has focused on appearance-based approaches to gait recognition, particularly the gait energy image (GEI) [4] and its variants, such as the 3D gait energy volume (GEV) [9], due to greater recognition performance and overall simplicity. The GEI is constructed from silhouettes (or binary voxel volumes in the case of the GEV), encoding the gait features as a single image template. Due to this, the feature can be sensitive to segmentation errors, and like other appearance-based techniques, to changes of appearance in the subject. Extensions to the GEI have been

proposed to address this [6], though their success is limited to certain classes of problems [7].

Though the image domain is different (averaged silhouettes *vs.* gray-scale image), the problems encountered are similar to those found in appearance-based face recognition, where the features extracted need to be robust towards unwanted variations in the image, whether they be due to lighting or pose. Patch-based feature descriptors such as histogram of oriented gradients (HOG) [2], and more commonly, local binary patterns (LBP) are used for this purpose to great effect, offering relatively high recognition performance at lower computational complexity. An extension to the LBP, local directional patterns (LDP) [8] provide even better accuracy under noisy conditions or situations like changes in illumination.

HOG has been used in gait recognition [5], but has yet to be applied directly to GEIs. In this paper, we explore the effectiveness of HOG and LDP when applied to GEIs for gait recognition. We also propose a new feature descriptor, which we call the histogram of weighted local directions (HWLD), borrowing concepts from both HOG and LDP. Like the HOG, the proposed method is a histogram based on local gradient directions. However, the raw directions are discrete, mapped to each pixel's eight neighbours, with the magnitudes to the possible directions determined from directional response kernels commonly used in LDP. Weights are applied based on the relative strengths of the directions at each pixel like the LDP.

This method takes the strengths of the two base systems, in that it keeps the richer local appearance description offered by the LDP, yet keeps a more compact feature size possible with the HOG. The second point allows us to extend the HWLD to 3D volumes, avoiding the prohibitively large number of histogram bins required for a similar LDP implementation.

Performance of the HWLD is compared with HOG and LDP by applying them to the GEI, with experiments carried out on the CASIA [12] and OULP [7] datasets. Classifica-

tion is performed using sparse representation classification (SRC) [11] as it has been shown to perform well for gait recognition [10]. Evaluation of the HWLD when applied to GEVs is also performed on the DGD dataset [10], comparing it with a baseline implementation also using SRC.

The remainder of this paper is organised as follows. Section 2 outlines the applicability of patch-based feature extraction techniques to the GEI. Section 3 discusses the proposed novel feature extraction technique, HWLD and Section 4 details the classification. Experiments and results are discussed in Section 5 followed by the conclusion in Section 6.

2. The GEI and Local Histogram Feature Descriptors

GEIs represent the static and dynamic behaviour of human motion within a gait cycle as a single image template by averaging the normalised binary silhouettes over that cycle. The GEI is computed by,

$$GEI = \frac{1}{T} \sum_{t=1}^T I_t, \quad (1)$$

where T is the number of frames within the gait cycle and I is the silhouette image at frame t .

Due to the nature of the GEI, various forms of errors can be introduced into the feature image. Firstly, errors in the segmentation, as well as actual appearance changes in the subject, can cause changes to the appearance of the GEI. Poor normalisation and alignment of individual silhouettes can also contribute to artifacts, effectively creating a blurring effect on the final GEI.

Finally, the registration between the probe and gallery GEIs can also be an issue when performing classification. This can be addressed using local grid-based histogram methods, as its coarse spatial sampling introduces some tolerance to these alignment issues.

Though more importantly, the use of histograms allows linear comparisons of feature values that are not scalar in nature, such as those produced from LBP/LDP. LBP features were originally used as a texture descriptor, though have since been successfully applied in face recognition for their robustness to appearance changes due to illumination and pose, issues that are similar to those described above for GEIs. The LDP is an extension to the LBP, providing superior performance.

Another commonly used local histogram feature descriptor is the HOG. This section will provide an overview of the LDP and HOG feature descriptors and a way to apply them to GEIs for gait recognition.




Figure 1. 2D local directional kernels.

2.1. Local Direction Pattern

The local directional pattern (LDP) is an extension of the local binary pattern (LBP), and has been shown to be a robust feature for use in various face recognition applications [8]. Like the LBP, it assigns an eight bit binary key to each pixel of an input image (in this case, a GEI), representing the local appearance at that pixel.

To compute the LDP, a set of 8 directional kernels are applied to extract the dynamic response (k_0, k_1, \dots, k_7) in each of the 8 neighbouring pixel directions. The kernels used are based on the Kirsch compass kernels [8], which are shown in Figure 1.

The dynamic response values are sorted in descending magnitude, and the top n values are selected. The 8 possible directions each correspond to a bit in an 8-bit value. The selected n bits are set to 1, and the resulting value is the feature ‘key’ for that pixel. For the purposes of this paper $n = 3$.

The resulting feature map is partitioned into a grid, and a histogram is computed for each local patch. 56 histogram bins are used, corresponding to the 56 unique key values (8C_3). The histogram values from each region are then concatenated to form the final feature vector. Figure 2 shows the process of computing the LDP feature from a GEI.

The LDP however, is not suitable for applications using 3D data, such as the GEV. With 26 neighbours per voxel, the number of unique keys, and therefore required histogram bins, become infeasible. At $n = 3$, the number of bins required is 2600.

2.2. Histogram of Oriented Gradients

The histogram of oriented gradients (HOG) is a very common feature descriptor used for object detection and recognition. Though it has been used in face recognition [2], it is less popular than LBPs due to poorer performance in that context. Recently, it has also been applied to gait recognition [5], showing significant improvements over a GEI/PCA baseline [4].

In [5], the HOG is applied directly to the raw image, and the magnitudes of the gradients are weighted by the silhouette mask. The HOG features are then averaged over the gait cycle to arrive at the final feature vector.

In this paper, the HOG operator will be applied directly

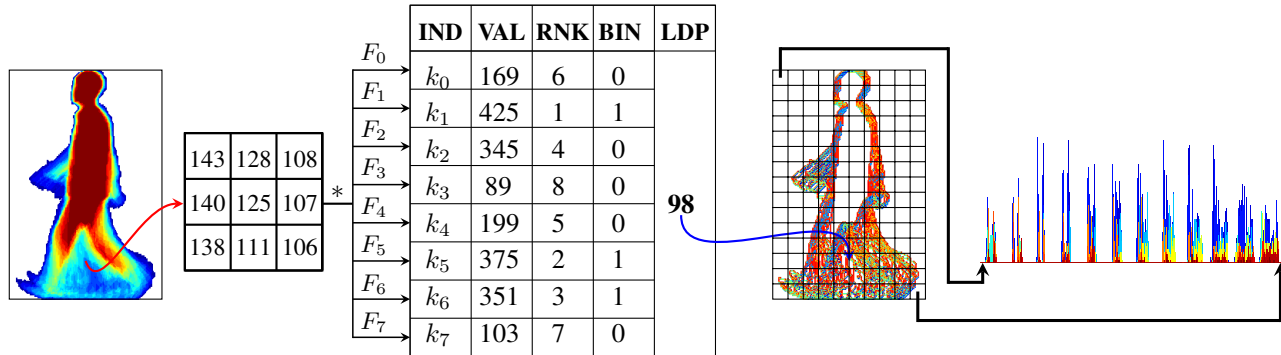


Figure 2. Computation of binary-coded LDP from a GEI image. The GEI is convolved with the directional kernels. The resulting values are sorted by rank (RNK), and the top 3 are assigned a 1 in its corresponding bit in the binary number (BIN). The decimal representation of this number is the LDP value for the pixel. The GEI is partitioned into a grid and a histogram is computed for each patch from the LDP values.

to the GEI image instead. The GEI encodes the temporal features into the image appearance; applying the HOG to individual frames and then averaging will lose this information. Whether this approach is actually preferable or not will not be evaluated in this paper.

To compute the HOG feature, the gradient vector at each pixel in the GEI is first found. The 1st order gradient operator ($[-1, 0, 1]$ and $[-1, 0, 1]^T$) is applied to extract the horizontal (g_h) and vertical gradient (g_v) magnitudes, which is then combined to calculate the final gradient magnitude (g) and orientation (θ) as below,

$$g = \sqrt{g_h^2 + g_v^2} \quad \theta = \text{atan2}(g_h, g_v). \quad (2)$$

The image is then partitioned into a grid, with the gradients placed into histograms at each patch based on the gradient orientation. The histograms contain 9 equal-width bins, with each gradient weighted by the gradient’s magnitude.

Applying the HOG to the GEV is possible, as the number of bins used can be adjusted. Selecting the values required to represent an angle in 3D space can be cumbersome, with the directional kernels used to compute the LDP providing a more elegant solution. This motivates the development of the histogram of weighted local directions.

3. Histogram of Weighted Local Directions

The initial motivation for developing this new feature is to extend the LDP into 3D for use in gait energy volumes. Due to the binary pattern coding used however, the number of unique code combinations (and therefore, the required number of histogram bins) increases to 2600 with the 26 possible directions in 3D voxel space and $n = 3$. This number rises rapidly as n increases, with 65780 unique values at $n = 5$.

The solution is to simply remove the binary pattern coding from the algorithm, moving towards a more HOG-like implementation. However, the proposed algorithm would

retain the 8 discrete directions, with magnitudes computed using the directional kernels. Multiple gradient directions for each pixel are still used, as this should provide a richer description of the local appearance than simply using one, as many pixels would not have a dominant gradient direction. The contributions to the histogram are weighted, as a function of its directional response values. We term this method the histogram of weighted local gradients (HWLD).

The initial computation of the HWLD is identical to that of the LDP; the directional kernels are convolved with the GEI to extract the directional response values, which are used to rank the directions in terms of their magnitude. The GEI is partitioned into a grid, and a histogram is formed for each patch. The first n directions at each pixel are used to populate the corresponding histogram, though their contribution is weighted by their magnitudes. The weighting factor, w , is proportional to the relative dominance of each direction at the pixel, such that,

$$w_j = \|k_j\| \left(\sum_{i=0}^7 \|k_i\| \right)^{-1}, \quad (3)$$

where j is a given direction index, and k is the response value.

3.1. Applying HWLD to 3D Data

The gait energy volume (GEV) [9] is the 3D analogue of the GEI. With the advent of cheap and accurate real-time depth cameras, gait recognition from a frontal perspective may be viable, with many attractive properties [9]. To create the GEV, the depth image is first projected into world coordinates and a binary voxel volume created from the surface reconstruction. The GEV is then constructed from these volumes similar to how a GEI is constructed from a binary silhouette image.

To extend the HWLD feature to 3D, we must first define the directional kernels to be used. There are twenty six




Figure 3. 3D local directional kernels. 3 unique cases are presented, showing the kernel direction and the corresponding weight values. The other 23 kernels can be obtained by applying rotational transforms to these.

$3 \times 3 \times 3$ directional kernels defined for each of the neighbouring voxel directions. Like the Kirsh compass kernels, the direction-facing voxel and its adjacent voxels are given a positive value, while all others, except the centre voxel, are assigned a negative value. The centre voxel is set to 0. Each of the 26 kernels must be of the same magnitude, and sum to 0 (Equation 4).

$$\sum K = 0, \quad \|K\| = 1. \quad (4)$$

Examples of the kernels are shown in Figure 3, depicting the three possible unique cases and their corresponding kernel values. The other 23 kernels can be obtained by applying rotational transformations to these.

The rest of the algorithm follows that of the 2D case. The response values are calculated and weights obtained. The GEV is divided into a series of 3D patches, with histograms constructed for each patch using the top n directions. To capture a greater range of appearances due to the larger number of possible directions, $n = 5$ is used for the 3D HWLD. Note that a similar LDP implementation would require 65780 histogram bins per volume patch.

4. Classification

Recently, sparse representation based classification [3] has gained popularity and has been used as an effective classification method for face [11] and gait [10]. Initial investigation into sparse representation has shown improvement over traditional classification methods. Hence, we use a sparse representation based classifier to evaluate the proposed algorithms and baselines.

Classification is initialised by forming a dictionary that represents all the subjects in the gallery with a total of N extracted feature vectors $(\tau_1, \tau_2, \dots, \tau_N)$ packaged into the columns of a matrix, $A \in \mathbb{R}^{n-1 \times N}$, where n is the number of subjects in the gallery. The objective is to identify the test subject, γ , in terms of the dictionary matrix that satisfies the following linear equation,

$$\gamma = A\alpha, \quad (5)$$

where by determining the coefficient vector $\alpha = [0, \dots, 0, \delta_{i1}, \dots, \delta_{ik}, 0, \dots, 0]$ (where δ are non zero elements and k is the number of feature vectors in the i^{th} subject), we can determine that γ corresponds to the i^{th} subject. The above objective can be achieved by computing the most sparse coefficients for a given subject γ , with respect to the training subjects in the dictionary, using the following optimisation problem,

$$\arg \min \hat{\alpha} = \|\alpha\|_1, \quad s.t. \quad \gamma = A\alpha. \quad (6)$$

Ideally the non-zero elements in the sparse coefficients should represent the test subject. However, due to complicating factors and modelling errors, the sparsest solution can produce non-zero entries that corresponds to multiple subjects. However, we can find the most significant entry by individually analysing the strength of coefficients associated with multiple subjects. To do this, the modified sparse coefficient, $\hat{\alpha}_i = [0, \dots, 0, \delta_{i1}, \dots, \delta_{ik}, 0, \dots, 0]$ is defined, where k is number of feature vectors in the dictionary corresponding to the i^{th} subject. Using α , the normalised distance of the i^{th} subject to the test subject, γ is calculated as follows,

$$D_i(\gamma) = \|\gamma - A\hat{\alpha}_i\|. \quad (7)$$

The smaller the distance, $D_i(\gamma)$, the more likely γ corresponds to the i^{th} subject.

5. Experiments and Results

5.1. Evaluation of 2D approaches

Profile views from CASIA dataset B [12] and OULP-C1V1-A database [7] are used to evaluate the proposed approaches in 2D. Both identification and verification experiments are performed, with CMS and ROC results compared. Evaluations are performed using the framework outlined in the respective datasets.

CASIA dataset B contains multiple gait sequences of 125 subjects for 3 different conditions - 6 sequences of normal walk (nw), 2 sequences of different clothing (cl) and 2 sequences of carrying bag (bg). For the 2D experiments, intra and inter-class experiments follow the evaluation outlined in [12] where intra-class experiments are performed on the nw sequences, with 4 allocated to the gallery and 2 to the probe. In inter-class tests, 4 cycles from nw are used as the gallery, while the 2 sequences in each of the other classes make up the probe in their individual experiments.

In addition to the segmented silhouettes provided with the dataset (labeled as S1 in this experiment), higher quality silhouettes are segmented using a graph-based segmentation algorithm similar to [1] (labeled as S2). The original silhouettes contains significantly more segmentation errors and examples of the differences, as well as its effects on the GEI, can be seen in Figure 4. These two sets of silhouettes

Feature	Seg.	PCAMDA	Rank-1			Rank-3			TAC @ FAR 3%			TAC @ FAR 5%			AUC		
			<i>nwnw</i>	<i>nwbg</i>	<i>nwcl</i>	<i>nwnw</i>	<i>nwbg</i>	<i>nwcl</i>	<i>nwnw</i>	<i>nwbg</i>	<i>nwcl</i>	<i>nwnw</i>	<i>nwbg</i>	<i>nwcl</i>	<i>nwnw</i>	<i>nwbg</i>	<i>nwcl</i>
HWLD	S2	-	100	92.2	96.5	100	96.6	97.6	99.6	87.7	96.8	99.6	91.6	97.3	0.999	0.984	0.991
	S2	✓	100	91.3	96.5	100	94.8	98.2	99.5	86.7	96.3	99.5	90.3	96.3	0.999	0.980	0.993
	S1	-	99.1	90.4	94.8	99.1	93.0	97.4	98.2	87.8	94.8	98.7	92.5	96.0	0.997	0.984	0.986
LDP	S2	-	98.2	85.1	93.8	99.2	92.1	97.3	98.2	84.9	95.0	99.1	89.1	96.3	0.995	0.961	0.984
	S2	✓	100	79.2	92.1	100	87.9	94.7	99.0	81.9	92.6	99.0	85.0	94.0	0.998	0.952	0.978
	S1	-	97.4	83.4	93.6	99.2	90.2	98.8	97.8	82.8	94.2	98.6	87.2	98.1	0.992	0.954	0.991
HOG	S2	-	99.1	80.8	96.3	100	91.3	97.3	98.6	80.2	95.4	99.1	84.6	95.9	0.999	0.962	0.989
	S2	✓	99.2	78.3	93.8	100	88.	97.4	98.6	78.8	94.0	99.5	84.5	94.9	0.999	0.951	0.987
	S1	-	98.2	80.2	96.2	99.1	92.2	98.1	98.2	80.0	94.2	98.7	85.1	96.2	0.998	0.970	0.993
GEI	S2	-	98.2	57.4	80.9	100	65.2	91.3	97.2	58.1	87.6	98.2	62.5	89.3	0.994	0.860	0.954
	S2	✓	100	74.1	91.2	100	75.6	92.0	99.5	66.5	91.0	100	71.5	92.7	0.999	0.920	0.966
	S1	✓	100	64.3	82.6	100	69.6	87.8	99.1	60.2	82.6	99.6	71.3	88.9	0.999	0.864	0.946
GEI [10]	-	✓	100	68.5	80.3	-	-	-	-	-	-	-	-	-	-	-	-
SGEI [6]	-	✓	99.0	72.0	64.0	-	-	-	-	-	-	-	-	-	-	-	-

Table 1. Identification and verification results for CASIA dataset. Seg refers to the silhouette set used (S1 or S2). PCAMDA refers to the use of feature modeling (PCA and MDA). Our experiments are compared with other state-of-the-art results found in the literature for this dataset.




Figure 4. Example silhouettes and computed GEI comparing S1 (top) and S2 (bottom).

are used to compare the performance of the algorithms in the presence of segmentation noise.

GEIs are cropped to the lower half of the image to minimise the appearance changes in the upper body. LDP, HOG and HWLD features are computed from the GEI using a patch size of 5×5 . PCA and MDA is applied to the features, with classification performed using SRC.

From the performance measures in Table 1, we can see that the proposed HWLD is able to perform equal, or better than the other techniques in all test cases and achieves the state-of-the-art performance on the CASIA dataset. We also see good performance of the LDP and HOG, with LDP favouring the *nwbg* case and HOG favouring *nwcl*. This result is slightly unexpected as LBP/LDP approaches generally provide better performance to HOG in face recognition. This could simply be due to the different nature of a GEI compared to a face image, though the performance of the LDP may have been limited due to the small patch size, as we are trying to populate 56 histogram bins with 25 samples, compared to 8 or 9 bins in the HWLD and HOG. [8] finds slightly improved performance in face recognition with a patch size of 10×10 compared to 5×5 .

When comparing the results of the S1 silhouettes to S2, we see an improvement using the cleaner silhouettes,

though the improvement is only slight for many of the test cases using the local patch methods. Contrasting this with the much more significant improvements seen when using just the GEI, demonstrates the greater tolerance to segmentation errors in these algorithms. It is also noted that better performance is achieved in the patch-based methods when PCA and MDA is not applied to the feature vectors in inter-class tests. The reason for this is unknown, though we speculate that this could be due to ‘over-fitting’, with the SRC dictionary templates failing to properly generalise to variations in the subjects’ appearance.

The proposed approach is also evaluated on the OULP-C1V1-A dataset of the OU-ISIR Gait database [7], which consists of silhouettes from a maximum of 3961 subjects which are grouped based on the captured view angle (55, 65, 75, 85, All). Each subject has 2 sequences and from them one centered gait cycle for the particular view is selected. In this paper, we use the distributed cleaned segmented silhouettes from the near-profile view (A-85) and from gait cycles that cover all the angles (A-All). Only intra-class test cases are considered as following the testing protocol in [7]. First sequence in each group are used as gallery and the other one is used as probe.

Table 2 shows the recognition performances based on rank and AUC metrics. Again it shows the effectiveness of the propose HWLD algorithm with the AUC close to 1 for the A-All class. The results shown are computed using PCA-MDA, with experiments without it (not shown) performing slightly less. This provides further support for the over-fitting hypothesis previously mentioned, as these are intra-class tests.

The results for the GEI are lower than the GEI imple-

mentation in [7]. The reason for this is likely due to using only the lower half of the GEI. Since this is an intra-class evaluation, there are no significant appearance changes in the upper body for us to ignore.

Feature	A-85			A-All		
	Rank 1	Rank 5	AUC	Rank 1	Rank 5	AUC
HWLD	87.7	94.7	0.992	95.5-	98.5	0.999
LDP	83.7	91.2	0.992	95.1	98.3	0.999
HOG	56.5	74.5	0.954	75.9	89.9	0.984
GEI+ PCA	81.1	92.0	0.985	90.8	97.2	0.997
[7]	85.7	93.1	-	94.2	97.1	-

Table 2. Recognition Performance on OULP-C1V1-A.

5.2. Evaluation of HWLD in 3D

Experiments in 3D are conducted using the DGD [10]. The DGD contains 3D depth data for 35 subjects in multiple sequences with 6 varying conditions; 5 sequences of normal walk (*nw*) and fast walk (*fw*), 4 sequences of walking without shoes (*ns*) and 2 sequences of front carrying bag (*fc*), side carrying bag (*sc*) and back carrying bag (*bc*). An intra-class test is carried out using the *nw* sequences. 3 of the 5 *nw* cycles for each subject are assigned to the gallery while the remaining cycles are used as the probe. In inter-class tests, all 5 *nw* cycles are used as the gallery, and all available cycles in each of the other classes (*fw*, *sc*, *fc*, *bc*, *ns*) are treated as the probe in their respective experiments. GEVs from frontal depth images in the DGD are extracted using the method described in [9]. Similar to the 2D experiments, HWLD features are extracted as explained in Section 3.1.

The Rank-1 cumulative match scores of this experiment is presented in Table 3. We see a slight, but consistent improvement over the baseline implementation that is in line with the results in [10].

6. Conclusion

In this paper, we have proposed the Histogram of Weighted Local Directions (HWLD) for use in gait recognition. The proposed HWLD features demonstrate state-of-the-art performances, showing significant improvements in

Feature	Rank-1					
	<i>nwnw</i>	<i>nwfw</i>	<i>nwfc</i>	<i>nwsc</i>	<i>nwbc</i>	<i>nwns</i>
HWLD	100	100	100	96.5	97.1	100
GEV+MDA	100	99.6	93.1	94.2	96.7	96.2
[10]	100	98.2	94.1	92.4	95.6	96.8

Table 3. Rank-1 cumulative match scores that compare the proposed and state-of-the-art method on 3D data using DGD database.

inter-class tests over existing implementations for the CASIA dataset. Superior performance of the feature on the high population OULP dataset, which contains more than 3000 subjects, shows that the proposed method is stable over a large population. The HWLD can also be easily extended to 3D, with evaluations using GEVs on the DGD dataset beating all known results.

Furthermore, we demonstrate that local histogram feature extraction techniques are much more stable to minor segmentation errors. They also show improved performance in the absence of feature conditioning processes such as PCA and/or MDA in inter-class tests.

Acknowledgments

Portions of the research in this paper use the CASIA Gait Database collected by Institute of Automation, Chinese Academy of Sciences and the OU-ISIR Gait Database distributed by Osaka University. Computational resources and services used in this work were provided by the HPC and Research Support Group, Queensland University of Technology, Brisbane, Australia.

References

- [1] D. Chen, S. Denman, C. Fookes, and S. Sridharan. Accurate silhouettes for surveillance - improved motion segmentation using graph cuts. In *Proc. IEEE Int. Conf. on DICTA*, pages 369–374, 2010. 4
- [2] O. Déniz, G. Bueno, J. Salido, and F. De la Torre. Face recognition using histograms of oriented gradients. *Pattern Recognition Letters*, 32(12):1598–1603, 2011. 1, 2
- [3] D. L. Donoho. Compressed sensing. *IEEE Trans. on Information Theory*, 52(4):1289–1306, 2006. 4
- [4] J. Han and B. Bhanu. Individual recognition using gait energy image. *IEEE Trans. on PAMI*, 28(2):316–322, 2006. 1, 2
- [5] M. Hofmann and G. Rigoll. Improved gait recognition using radient histogram energy image. In *Proc. IEEE Int. Conf. on ICIP*, pages 1389–1392, 2012. 1, 2
- [6] X. Huang and N. V. Boulgouris. Gait recognition using linear discriminant analysis with artificial walking conditions. In *Proc. IEEE Int. Conf. on ICIP*, pages 2461–2464, 2010. 1, 5
- [7] H. Iwama, M. Okumura, Y. Makihara, and Y. Yagi. The ou-isir gait database comprising the large population dataset and performance evaluation of gait recognition. *IEEE Trans. on IFS*, 7(5):1511–1521, 2012. 1, 4, 5, 6
- [8] T. Jabid, M. Kabir, and O. Chae. Gender classification using ldp. In *Proc. IEEE Int. Conf. on ICPR*, pages 2162–2165. IEEE Computer Society, 2010. 1, 2, 5
- [9] S. Sivapalan, D. Chen, S. Denman, S. Sridharan, and C. Fookes. Gait energy volumes and frontal gait recognition using depth images. In *Proc. Int. Conf. on IJCB*. IEEE, 2011. 1, 3, 6
- [10] S. Sivapalan, D. Chen, S. Denman, S. Sridharan, C. Fookes, et al. The backfilled gei - a cross-capture modality gait feature for frontal and side-view gait recognition. *Proc. Int. Conf. on DICTA*, 2012. 1, 2, 4, 5, 6
- [11] J. Wright, A. Y. Yang, A. Ganesh, S. S. Sastry, and Y. Ma. Robust face recognition via sparse representation. In *IEEE Trans. on PAMI*, 2008. 2, 4
- [12] S. Yu, D. Tan, and T. Tan. A framework for evaluating the effect of view angle, clothing and carrying condition on gait recognition. In *Proc. Int. Conf. on ICPR*, volume 4, pages 441–444, 0-0 2006. 1, 4