Gait Representation Using Flow Fields

Khalid Bashir
khalid@dcs.qmul.ac.uk

Tao Xiang
txiang@dcs.qmul.ac.uk

Shaogang Gong
sgg@dcs.qmul.ac.uk

School of EECS,
Queen Mary University of London,
London E1 4NS, UK

Abstract

Gait is characterised by the relative motions between different body parts during walking. However, most recently proposed gait representation approaches such as Gait Energy Image (GEI) and Motion Silhouettes Image (MSI) capture only the motion intensity information whilst ignoring the equally important but less reliable information about the direction of relative motion. They thus essentially sacrifice discriminative power in exchange for robustness. In this paper, we propose a novel gait representation based on optical flow fields computed from normalized and centred person images over a complete gait cycle. In our representation, both the motion intensity and the motion direction information is captured in a set of motion descriptors. To achieve robustness against noise, instead of relying on the exact value of the flow vectors, the flow direction is discretised and a histogram based direction representation is formulated. Compared to the existing model-free gait representations, our representation is not only more discriminative, but also less sensitive to changes in various covariate conditions including clothing, carrying, shoe, and speed. Extensive experiments on both indoor and outdoor public datasets have been carried out to demonstrate that our representation outperforms the state-of-the-art.

1 Introduction

Gait is concerned with how people walk and can be characterised by the relative motions between different body parts during walking. However, most recently proposed gait representation approaches [1, 5, 10, 11, 14] focus solely on the motion intensity information. This is mainly because that although the relative motion direction is useful for representing gait, it is also more difficult to capture and less reliable. In particular, during a walking cycle, whilst the whole body is moving towards one direction, multiple independent motions also co-exist across different body parts. Furthermore, these independent motions constantly change in both direction and magnitude. One possible solution to this is to adopt a model based approach [8, 12, 15, 21] which models the human body configuration (e.g. 2D/3D skeletons) and the model parameters estimated over time encode the detailed relative motion information. Nevertheless, model-based approaches are sensitive to image noise, self-occlusion, shadows, and view changes, which leads to inferior performance on public gait datasets [5, 10, 11].

In this paper we propose a novel gait representation approach which captures both the intensity and direction information about the relative motions between body parts, and critically is robust to noise. More specifically, optical flow fields are computed for centred and normalised human figures extracted from a complete gait cycle. In order to extract robust gait features from the noisy optical flow field, we do not rely on the exact values of the magnitude and direction of the flow vectors. Instead, the flow direction is discretised and a histogram based direction representation is formulated, which gives rise to a set of spatio-temporal motion descriptors. In particular, our gait representation consists of a Motion Intensity Image (MII), which measures the intensity of relative motion at each pixel location, and four Motion Direction Images (MDIs), each of which represents the likelihood of the direction of motion being along one specific motion direction during a complete gait cycle. These motion images are then fused together during recognition.

Compared with existing model free approaches [1, 5, 10, 11, 14], our representation are more descriptive as more comprehensive motion information is captured. In addition, the fusion of both motion intensity and direction information makes our representation more robust to various covariate conditions that affect gait. Generally speaking, there are two types of covariate conditions: conditions that affect the shape appearance of human figure such as carrying and clothing conditions, and conditions that affect gait itself including shoe-wear type, surface and time. Similar to existing binary cumulative motion images such as Gait Energy Image (GEI) and Motion Silhouettes Image (MSI), our MII is less sensitive to covariates that affect gait itself. Our MDIs, on the other hand, capture little shape appearance information; they are thus less sensitive to shape appearance related covariates. A fusion of both MII and MDIs makes our representation insensitive to changes in both types of covariate conditions.

We perform extensive experiments using the CASIA [20] and Soton dataset [17] which contain both indoor and outdoor gait sequences to validate the effectiveness of our representation against the state-of-the-art. The results indicate that our flow field based gait representation outperforms existing alternatives, especially when there are changes in various covariate conditions.

1.1 Related Work

The idea of extracting gait features from optical flow fields has been exploited before in the late 90's [7, 13]. Little and Boyd [13] represent the distribution of each flow field using moments and the periodic structure of the flow distribution features is then exploited for gait representation. Treating the optical flow fields as spatial patterns, the extracted holistic features are weak in discriminative power for an intra-class object recognition task such as gait recognition, although they have been shown to be very useful for inter-class recognitions (e.g. action recognition [4, 18]). Instead of extracting holistic features, Huang *et al.* [7] proposed to use the magnitude of flow vectors directly as templates. These templates are then projected to a low dimensional canonical space for recognition. However, the useful motion information is neglected. In addition, relying on the exact value of the flow vector magnitude makes their representation sensitive to noise.

A number of approaches have been proposed recently to separate the motion and static shape appearance information in gait representation. In [10] Motion Silhouette Contour Templates (MSCT) and Static Silhouette Templates (SST) are formulated for representing the two types of information. On a similar note is the fusion of shape and kinematics features [3]. In both approaches, the shape features used are less discriminative for gait representa-

tion and prone to changes in appearance related covariate conditions such as clothing and carrying. As for the motion features, the MSCT does not contain motion direction information; although those kinematics features in [3] do, they have to be extracted based on a human body configuration model. In contrast, the proposed MII contains both motion intensity and shape appearance information whilst the MDIs are concerned solely with the motion direction information. Also, both MII and MDIs are robust to noise and easy to compute. More importantly a fusion of both makes our representation insensitive to various covariate condition changes.

2 Gait Motion Descriptors

Given a gait sequence, a figure-centric optical flow field is computed for each frame. This is followed by gait cycle segmentation, and computation of our gait motion descriptors for each cycle.

2.1 Computing Figure-centric Optical Flow Fields

A figure-centric spatio-temporal volume for each walking person is first extracted. To this end the walking person needs to be segmented from each image frame. This is achieved by background subtraction. Connected component algorithm is then applied to the segmented foreground regions and a silhouette image of the person is obtained. The silhouettes extracted from the sequence are used for two purposes: 1) They are used as the foreground masks to extract the original images of the person, which are then centred and normalised to a standard width and height to generate the figure-centric spatio-temporal volume. 2) The silhouettes are also used to extract gait cycles via the maximum entropy estimation [16] in the lower half of the silhouettes. An example of normalized silhouette and the corresponding figure-centric original image of the walking person from a carrying-bag sequence in the CASIA dataset are shown in Fig.1 (a) and (b) respectively.

(a) Silhouette   (b) $\mathbf{I}_{t-1}$ (c) $\mathbf{I}_t$ (d) $\mathbf{F}_t$

Figure 1: Computing figure-centric optical flow field. (a) shows the silhouette at frame $t-1$; (b)&(c) are the extracted figure-centric images of the walking person at $t-1$ and t respectively; (d) is the computed optical flow field.

Given a complete gait cycle consisting of T figure centric images $\{\mathbf{I}_1, .., \mathbf{I}_t, .., \mathbf{I}_T\}$, optical flow fields are estimated for each frame. The flow field for the t−th frame is denoted as $\mathbf{F}_t$ and computed using two consecutive frames $\mathbf{I}_{t-1}$ and $\mathbf{I}_t$. A recently proposed optical flow estimation algorithm is adopted [2]. In order to alleviate the effect of noise, the flow fields are smoothed by applying a 3X3 Gaussian filter. An example of the figure centric flow field is shown in Fig.1 (d).

2.2   Computing Motion Descriptors

Optical flow fields are inevitably noisy and error-prone due to image noise and self-occlusions. For intra-class object recognition, extracting robust motion descriptors from noisy flow field input is crucial. To this end, instead of using the exact values of the flow vectors as in previous methods [4, 7, 18], we extract very coarse yet informative motion intensity and direction information by taking the following steps:

1. The flow field $\mathbf{F}$ (the subscript t is omitted here for notational simplicity) is decomposed into two scalar fields $\mathbf{F}_x$ and $\mathbf{F}_y$ corresponding to the horizontal and vertical components of the flow. The two scalar fields are further decomposed into four non-negative fields denoted as $\mathbf{F}_x^+$, $\mathbf{F}_x^-$, $\mathbf{F}_y^+$, and $\mathbf{F}_y^-$. In particular, at pixel location (i, j), if $F_x(i, j) > 0$, we have $F_x^+(i, j) = F_x(i, j)$ and $F_x^-(i, j) = 0$; otherwise we have $F_x^+(i, j) = 0$ and $F_x^-(i, j) = -F_x(i, j)$. Similarly $\mathbf{F}_y^+$, and $\mathbf{F}_y^-$ are obtained from $\mathbf{F}_y$. Note that the direction of flow vector will become extremely unreliable when the flow magnitude is small. Therefore, if the magnitude $\|\mathbf{F}(i, j)\|$ is less than 1 pixel, all four non-negative fields at that pixel location are set to zero.

2. The four non-negative flow fields are discretised into binary images and denoted as $\hat{\mathbf{F}}_x^+$, $\hat{\mathbf{F}}_x^-$, $\hat{\mathbf{F}}_y^+$, and $\hat{\mathbf{F}}_y^-$. This is because the exact values of these non-negative fields are still too noisy to be trustworthy.

3. Given the 4 binary images obtained from the original flow field for each frame of the complete gait cycle of T frames, we now compute how the motion intensity and direction are distributed during the gait cycle. To achieve this, a 5-bin histogram is computed at each pixel location $\mathbf{B}(i, j) = \{B_I(i, j), B_x^+(i, j), B_x^-(i, j), B_y^+(i, j), B_y^-(i, j)\}$. More specifically, at each frame t, if all of the four flow fields ($\hat{\mathbf{F}}_x^+$, $\hat{\mathbf{F}}_x^-$, $\hat{\mathbf{F}}_y^+$, $\hat{\mathbf{F}}_y^-$) are zero, $B_I(i, j)$ is incremented by one. $B_I(i, j)$ therefore counts the total number of frames in the gait cycle where there is no relative motion at pixel (i, j). The count in the bin $B_x^+(i, j)$ will be incremented if the corresponding binary flow images has a non-zero value. We therefore have $B_x^+(i, j) = \sum_1^T \hat{\mathbf{F}}_x^+$. Similarly we obtain the values of the other three histogram bins.

4. The 5 histogram bins at each pixel location are normalised by the length of the gait cycle T, which give rise to 5 motion descriptors used in our approach for gait representation. Specifically, from $B_I(i, j)$ we obtain the Motion Intensity Image (MII), denoted as $\mathbf{M}$, which measures the motion intensity distribution over a complete gait cycle. The other four bins contribute to four Motion Direction Images (MDIs) which measure the distribution of the motion direction along four directions: right, left, up, and down, and are denoted as $\mathbf{M}_x^+$, $\mathbf{M}_x^-$, $\mathbf{M}_y^+$, and $\mathbf{M}_y^-$ respectively. The size of the 5 descriptors is identical to that of the figure-centric person image.

   An example of the 5 motion descriptors illustrated as grey-level images is shown in Fig.2. The Motion Intensity Image $\mathbf{M}$ is a time-normlised accumulative motion intensity image with lower value meaning that relative motion occurs more frequently during the gait cycle. As can be seen in Fig.2, low values are observed at the legs and arms area, whereas the static areas such as head and torso are represented as high values. Note that our MII is similar in spirit to the widely adopted Gait Energy Image (GEI) [5], which aims to capture the same information but termed differently as motion energy. However, since a GEI is computed using the binary silhouette as opposed to flow fields, the measurement of motion

(a) $\mathbf{M}$ (b)$\mathbf{M}_x^+$ (c) $\mathbf{M}_x^-$ (d) $\mathbf{M}_y^+$ (e) $\mathbf{M}_y^-$
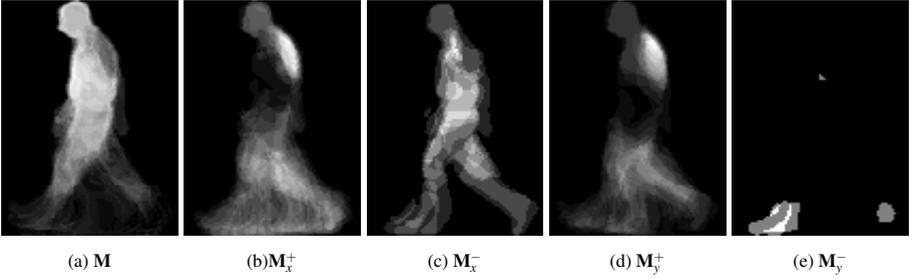
Figure 2: Example of our 5 motion descriptors.

intensity is indirect and less accurate. The four Motion Direction Images give more detailed information as for how likely relative motions of different directions can take place at each pixel location during a gait cycle. It can be seen from Fig.2(b)-(e) that, as expected, motions of different directions have different distributions at different locations. It is also noted that compared with the other three MDIs, all the high intensity areas in $\mathbf{M}_y^-$ are concentrated around the two feet. This is not surprising as downwards motion is rare for people walking on a flat surface. Nevertheless this suggests that the discriminative power of $\mathbf{M}_y^-$ is much weaker than that of the other three MDIs. Consequently, in this paper only three MDIs are used for gait representation together with the MII.

Fig.3 shows how different motion descriptors are affected by different covariate condition changes. It can be seen clearly that the appearance related covariate conditions such as carrying and clothing have a visible effect on the MII. This is because that the shape of the low motion-intensity areas in the MII are affected by the shape appearance changes caused by variations in these conditions. In contrast, Fig.3(b)-(d) indicate that the effect of those covariate conditions is minimal on the MDIs. Therefore the MDIs are more robust against shape appearance changes. On the other hand, the MDI would be more sensitive to covariate conditions such as shoe and surface that affect more on the gait itself. Overall, these two types of motion descriptors capture complementary information about gait and have different levels of sensitivity towards different types of covariate conditions. They are therefore fused together for recognition as described next.

3 Gait Recognition

The four motion descriptors ($\mathbf{M}$, $\mathbf{M}_x^+$, $\mathbf{M}_x^-$, $\mathbf{M}_y^+$) are used independently for computing the dis-similarity between the gait sequences of a probe subject and a gallery one. The dis-similarity scores are then fused for matching the two subjects. Instead of using the descriptors directly as templates, they are projected to a subspace for dimensionality reduction. This is achieved using Component and Discriminant Analysis (CDA) based on Principle Component Analysis and Multiple Discriminant Analysis to simultaneously achieve a good data representation and class separability [6]. The Euclidean distance of the descriptors in the CDA subspace is used to measure the dis-similarity between two subjects.

Once the dis-similarity scores have been computed for each of our four motion descriptors, they are fused together as follows. First, the dis-similarity scores need to be normalised as they fall into different value ranges. In particular, dis-similarity score for each descriptor is normalised by the averaged dis-similarity score between two different subjects in the training dataset. After normalisation, the final similarity score used for template matching is

(a) $\mathbf{M}$ (b) $\mathbf{M}_x^+$ (c) $\mathbf{M}_x^-$ (d) $\mathbf{M}_y^+$
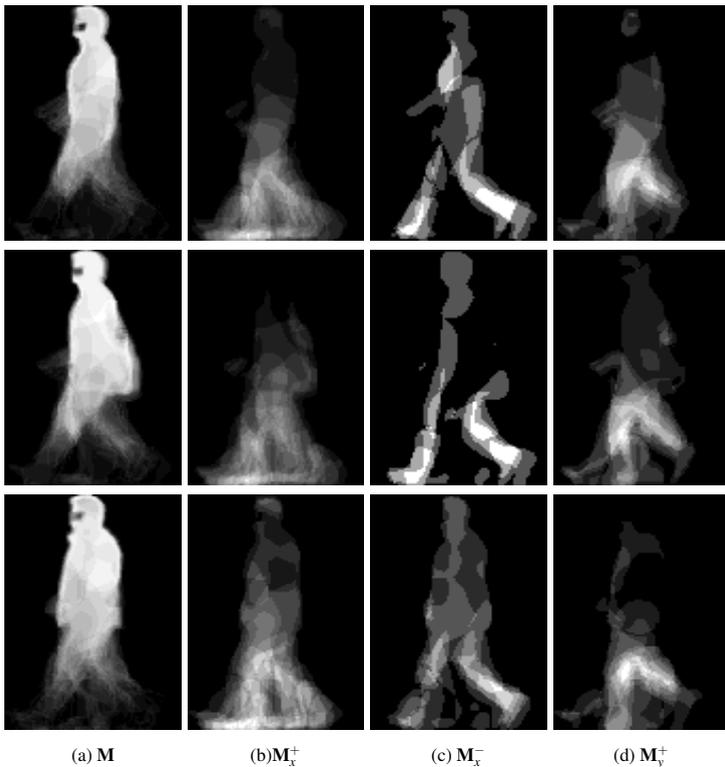
Figure 3: Effect of different covariate conditions on our motion descriptors. Top row: normal condition; middle row: the same person carrying a bag; bottow row: the same person wearing a bulky coat.

obtained as:

$$D = \lambda D_{\mathbf{M}} + (1 - \lambda)(D_{\mathbf{M}_x^+} + D_{\mathbf{M}_x^-} + D_{\mathbf{M}_y^+}) \tag{1}$$

where $D_{\mathbf{M}}, D_{\mathbf{M}_x^+}, D_{\mathbf{M}_x^-}, D_{\mathbf{M}_y^+}$ are the normalized dis-similarity score obtained using our gait descriptors, and λ is the fusion weight.

4 Experiments

4.1 Datasets and Settings

Two datasets, the CASIA and Soton datasets, were used for experiments which cover both indoor and outdoor scenarios and are amongst the most comprehensive public gait datasets.

The CASIA gait dataset [20] captures gait video sequences in an indoor environment. The dataset consists of 124 subjects. For each subject there are 10 walking sequences consisting of 6 normal walking sequences where the subject does not carry a bag or wear a bulky coat (CASIASetA), 2 carrying-bag sequences (CASIASetB) and 2 wearing-coat sequences (CASIASetC). Each sequence contains multiple gait cycles. The original image size of the database is 320x240. We used 4 of the 6 normal sequences as the gallery set and the rest of

the sequences in set CASIASetA (CASIASetA2) were used as the probe set together with CASIASetB and CASIASetC.

The Soton Large dataset is part of the Soton database [17]. It contains 116 subjects captured in both indoor and outdoor environment. The dataset has 6 subsets, SotonSetsetA to F. In our experiments the most widely used SotonSetsetA and SotonSetsetE were used. SotonSetsetA was captured in a controlled indoor environment whilst SotonSetsetE was captured outdoor with objects moving in the background thus closely resembling a real-world application scenario. Both SotonSetsetA and SotonSetsetE are featured with people walking under unchanged normal conditions. For both SotonSetsetA and SotonSetsetE, half of the sequences for each subject was used as the gallery set and the rest as probe set.

The Soton Small dataset consists of 11 subjects captured indoor under various covariate condition changes. Those covariate conditions include clothing condition (e.g. heavy/light clothes), carrying condition (e.g. carrying/without carrying bags), shoe (e.g. boots/trainers) and speed (quick/slow). In particular, for each subject there are two sequences under normal conditions (SotonSmallSetA), 3 carrying-bag sequences (SotonSmallSetB), 2 wearing-coat sequences (SotonSmallSetC), 5 sequences with different shoes (SotonSmallSetSh), and one sequence each for slower and quicker walking (SotonSmallSetS). We use one of the normal sequences from SotonSmallSetA as the gallery set. The other normal sequence (SotonSmallSetN) and all other subsets were used as the probe set. Both the Large and Small Soton datasets were captured at standard PAL resolution.

The size of normalized figure-centric images for all datasets was set to 128x88. For all experiments, the results obtained using $\lambda = 0.5$ (see Eqn. 1) are reported. The effect of different values of λ is also investigated.

4.2   Comparative Evaluation

Probe Set	TM	GEI	M	$\mathbf{M}_x^+$	$\mathbf{M}_x^-$	$\mathbf{M}_y^+$	Fusion
CASIASetA2	97.6%	99.4	**99.4%**	96.3%	84.9%	93.4%	97.5 %
CASIASetB	52.0%	60.2%	56.6%	57.3%	32.8%	46.7%	**83.6%**
CASIASetC	32.7%	30.0%	14.8	**50.9%**	29.1%	24.6%	48.8%
Overall	60.2%	62.8%	56.7%	68.2%	48.9%	54.9%	**76.6%**

Table 1: Results on the CASIA dataset. The proposed descriptors are compared with a template matching method in [20] and GEI [5].

Results on CASIA dataset - Table 1 shows the recognition rates obtained using the 4 different motion descriptors and the fusion result. The results shows that when there is no change in the covariate conditions (CASIASetA2) Motion Intensity Image (MII $\mathbf{M}$) gives slightly better results than the three Motion Direction Images (MDIs $\mathbf{M}_x^+,\mathbf{M}_x^-,\mathbf{M}_y^+$). When the shape appearance related covariate conditions change in the probe set (CASIASetB and CASIASetC), the opposite is observed. In particular, $\mathbf{M}_x^+$ achieves better result than the other three descriptors when used alone. It can also be seen in Table 1 that the fusion result is better than each descriptor alone or close to the best. Overall, considerable improvement can be obtained by fusing all 4 descriptors together. We also compare our results with those obtained using existing methods in Table 1. A single descriptor in our approach can achieve comparable results whist the fusion result is significantly higher than those of the alternative methods, particularly when there are changes in covariate conditions.

Probe Set	M	M_x^+	M_x^-	M_y^+	Fusion
SotonSetA	99.1%	96.5%	96.5%	93.1%	**99.1%**

Table 2: Results for our descriptors on SotonSetA.

Probe Set	MSI	Frieze	SVB Frieze	MSCT+SST	GEI	Fusion
SotonSetA	84.8%	96.0%	84.0%	84.0%	99.1%	**99.1%**

Table 3: Comparison with existing approaches including MSI [9], Frieze Patterns [14], SVB Frieze Patterns [11], MSCT+SST [10] and GEI on SotonSetA.

Probe Set	Wagg and Nixon [19]	M	M_x^+	M_x^-	M_y^+	Fusion
SotonSetE	67.0%	**100%**	93.1%	93.9%	93.9%	97.4%

Table 4: Results on the outdoor SotonSetE.

Results on Soton Large dataset -The results obtained on the indoor SotonSetA using the proposed motion descriptors are reported in Table 2. Our result based on fusing the 4 motion descriptors is also compared with the results reported in the literature in Table 3. Again, under the same covariate condition, both the MII and the fusion give excellent result. Compared with alternative approaches, both our approach and GEI yield near-perfect result and outperform the other model-free approaches. Table 4 shows the results on the outdoor dataset SotonSetE. The result suggests that our motion descriptors are robust to the lighting changes and moving background presented in the outdoor environment. In particular, our results are significantly higher than that of Wagg and Nixon [19].

Probe Set	GEI	M	M_x^+	M_x^-	M_y^+	Fusion
SotonSmallSetN	100%	100%	100%	100%	100%	**100%**
SotonSmallSetB	86.3%	76.2%	76.2%	66.7%	61.9%	**90.4%**
SotonSmallSetC	72.7%	54.6%	81.8%	63.6%	72.7%	**90.9%**
SotonSmallSetSh	100%	100%	92.3%	82.1%	87.2%	**100%**
SotonSmallSetS	100%	100%	85.7%	76.1%	76.2%	**100%**
Overall	94.2%	90.3%	87.3%	77.6%	79.6%	**97.1%**

Table 5: Results on the Soton Small dataset.

Results on Soton Small dataset -Table 5 shows the results on the Soton Small dataset which contain changes in shape-appearance related covariates (SotonSmallSetB and SotonSmall-SetC) and covariates that affect gait itself (SotonSmallSetSh and SotonSmallSetS). It can be seen from Table 5 that under shape appearance related covariate changes, MDIs give better result than MII. However, when the covariate conditions affect gait itself MII achieves better performance. This result is consistent with our analysis on the characteristics of different descriptors in Sec. 2.2. It is also obvious from the results that the descriptor fusion result outperforms that of using each descriptor alone. The performance of our descriptors is compared with GEI and Table 5 shows that our performance is better.

(a) CASIA

(b) SotonSetA

(c) SotonSetE

(d) Soton Small Set
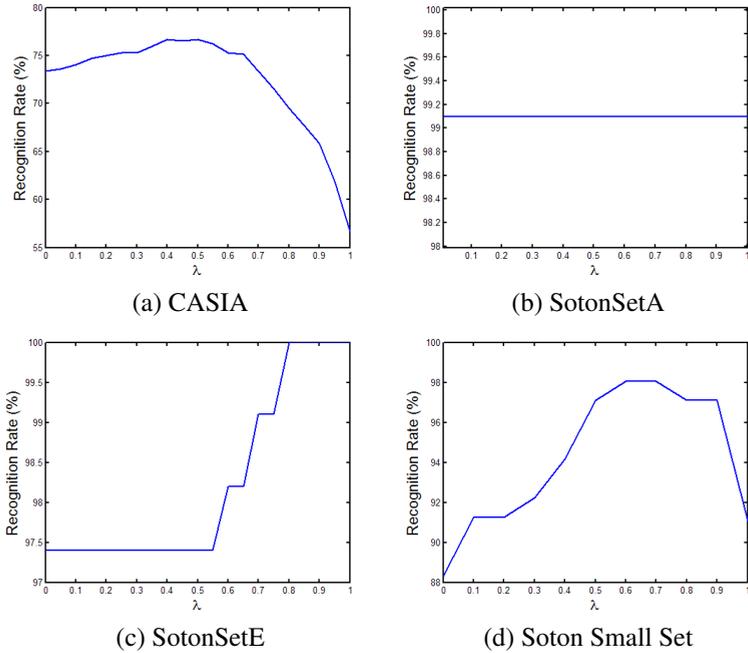
Figure 4: Effect of the fusion weight λ. Larger value of λ gives more weight to MII.

4.3 The Effect of the Fusion Weight λ

The results reported so far were obtained by setting the value of the fusion weight λ to 0.5, i.e. equal weight is given to the MII and the three MDIs. The effect of λ is investigated and the result is shown in Fig. 4. As can be seen in Fig. 4 when there are covariate changes, setting the value of λ either too low or too high will lead to worse performance. This is because that our MII and MDIs contain complementary information and have different levels of sensitivity towards different types of covariates. Overall, a value of 0.5 seems to be a safe choice.

5 Conclusion

We have proposed a novel gait representation based on optical flow fields computed from normalized and centred person images over a complete gait cycle. In our representation, both the motion intensity and the motion direction information is captured in a set of motion descriptors. The formulated Motion Intensity Image (MII) and Motion Direction Images (MDIs) have different levels of sensitivity toward different types of covariate conditions. The Fusion of them thus gives us a gait representation that is not only more discriminative, but also less sensitive to changes in various covariate conditions including clothing, carrying, shoe, and speed. Extensive experiments on both indoor and outdoor public datasets have been carried out to demonstrate that our representation outperforms the state-of-the-art.

References

[1] C. BenAbdelkader, R. Cutler, and L. Davis. Gait recognition using image self-similarity. *EURASIP J. Appl. Signal Process.*, (1):572–585, 2004.

[2] T. Brox, A. Bruhn, N. Papenberg, and J. Weickert. High accuracy optical flow estimation based on a theory for warping. In *ECCV*, pages Vol IV: 25–36, 2004.

[3] Y. Chai, Q. Wang, J. Jia, and R. Zhao. A novel gait recognition method via fusing shape and kinematics features. In *ISVC*, pages 80–89, 2006.

[4] A. Efros, C. Berg, G. Mori, and J. Malik. Recognizing action at a distance. In *ICCV*, page 726, 2003.

[5] J. Han and B. Bhanu. Individual recognition using gait energy image. *PAMI*, 28(2): 316–322, Feb 2006.

[6] P.S. Huang, C.J. Harris, and M.S. Nixon. Recognizing humans by gait via parametric canonical space. *Artificial Intelligence in Eng.*, 13:359–366, 1999.

[7] P.S. Huang, C.J. Harris, and M.S. Nixon. Human gait recognition in canonical space using temporal templates. *VISP*, 146(2):93, April 1999.

[8] D. Kim, Daehee Kim, and J. Paik. Model-based gait recognition using multiple feature detection. In *ACIVS*, pages 1018–1029, Berlin, Heidelberg, 2008.

[9] T.H.W. Lam and R.S.T. Lee. A new representation for human gait recognition: Motion silhouettes image (msi). In *ICB*, pages 612–618, 2006.

[10] T.H.W. Lam, R.S.T. Lee, and D. Zhang. Human gait recognition by the fusion of motion and static spatio-temporal templates. *PR*, 40(9):2563–2573, September 2007.

[11] S.K. Lee, Y.X. Liu, and R.T. Collins. Shape variation-based frieze pattern for robust gait recognition. In *CVPR*, pages 1–8, 2007.

[12] C.C. Lien, C.C. Tien, and J.M. Shih. Human gait recognition for arbitrary view angles. In *ICICIC*, page 303, 2007.

[13] J.J. Little and J.E. Boyd. Recognizing people by their gait: The shape of motion. *Videre*, 1(2), 1998.

[14] Y. Liu, R. Collins, and Y. Tsin. Gait sequence analysis using frieze patterns. In *ECCV*, May 2002.

[15] H. Lu and P. Venetsanopoulos. A layered deformable model for gait analysis. In *AFGR*, pages 249–254, April 2006.

[16] S. Sarkar, P. Phillips, Z. Liu, I. Vega, P. Grother, and K. Bowyer. The humanID gait challenge problem: Data sets, performance, and analysis. *PAMI*, 27(2):162–177, 2005.

[17] J. Shutler, M. Grant, M. Nixon, and J Carter. On a large sequence-based human gait database. In *RASC*, pages 66–71, 2002.

[18] H. Uemura, S. Ishikawa, and K. Mikolajczyk. Feature tracking and motion compensation for action recognition. In *BMVC*, 2008.

[19] D.K. Wagg and M.S. Nixon. On automated model-based extraction and analysis of gait. In *AFGR*, pages 11–16, 2004.

[20] S. Yu, D. Tan, and T. Tan. A framework for evaluating the effect of view angle, clothing and carrying condition on gait recognition. In *ICPR*, pages 441–444, 2006.

[21] G. Zhao, G. Liu, H. Li, and Pietikainen. 3D gait recognition using multiple cameras. In *AFGR*, pages 529–534, April 2006.