arXiv:1411.7591v1 [cs.CV] 27 Nov 2014

Egocentric Video Biometrics

Yedid Hoshen Shmuel Peleg
The Hebrew University of Jerusalem
Jerusalem, Israel

Abstract

Egocentric cameras are being worn by an increasing number of users, among them many security forces worldwide. GoPro cameras already penetrated the mass market, and Google Glass may follow soon. As head-worn cameras do not capture the face and body of the wearer, it may seem that the anonymity of the wearer can be preserved even when the video is publicly distributed.

We show that motion features in egocentric video provide biometric information, and the identity of the user can be determined quite reliably from a few seconds of video. Biometrics are extracted by training Convolutional Neural Network (CNN) architectures on coarse optical flow.

Egocentric video biometrics can prevent theft of wearable cameras by locking the camera when worn by people other than the owner. In video sharing services, this Biometric measure can help to locate automatically all videos shot by the same user. An important message in this paper is that people should be aware that sharing egocentric video will compromise their anonymity.

1. Introduction

Head worn Egocentric cameras are becoming very popular. GoPro cameras are now used not only by extreme sports enthusiasts but also by law enforcement and military professionals. With the forthcoming release of Google Glass such cameras will further penetrate the mass market.

Egocentric video is different from hand-held video in several respects:

- The camera is always on the user, and video can be recorded at any time e.g. when the user is walking.
- The camera is attached to the user's head. It therefore follows user head and body motions.
- The camera does not record images of the user, which may imply that user identity is hidden. However, in this work we show that in many cases user identity can be recovered.

Determining the identity of the user can be important. Many portable devices have some biometric means to verify that the user is in fact allowed to use the device. A common visual verification method is face recognition [29], but it cannot be used in head mounted devices. It may be possible to use non visual means such as speaker verification [4] or fingerprint recognition [18], but they are not always effective and available. In this work we set to verify user identity from the recorded by head-worn cameras.

Many users feel secure that sharing their egocentric videos does not compromise their identities, as shown in Fig. 1. Police forces released footage of officer activity, and commando operations recorded by cameras on soldiers heads are widely published on YouTube. Some users have even recorded and published what appears to be their own crimes. A consequence of our work is that privacy of such videos is compromised in many cases.

It has been found that people can be distinctly identified by biometric characteristics such as height, stride length, walking speed etc. Much research has been performed to extract and compare gait information from a video observing a walking person [10, 15, 5]. Most successful video methods deal with a static camera and a tightly controlled environment. Such methods have been effective in cases such as airport security, but are less applicable in uncontrolled settings.

In this work we extract biometric information from egocentric video. We concentrate on video recorded when the user is walking. Each user has distinctive features: step size, height, walking style, leg length etc. User biometrics are expressed in specific camera motions, which are characteristic of the user.

We learn biometric features and classifiers using a Convolutional Neural Network (CNN) architecture which includes layers corresponding to biometric feature extraction and classification. The input to our network is a very coarse optical flow computed over multiple frames of a few seconds' duration. Different architectures are trained for person identification and for verification.

Experiments are performed on both a public dataset and on a new, more challenging dataset verifying the perfor-



a) b) c)

Figure 1. a) A GoPro video uploaded to YouTube allegedly capturing a crime from the POV of the robber. Can the robber be identified? b) A GoPro video uploaded by US soldiers in combat. Are their identities safe? c) Combating theft of egocentric devices.

mance of our method. We find that a very short video (4 seconds) can identify the user accurately. Using our trained short video classifiers as building blocks, we use simple schemes to significantly improve verification and identification performance.

The ability to determine the identity of the user quickly and accurately is important for prevention (e.g. camera theft) and for forensic analysis (who committed the crime). An additional application includes retrieving all the videos of a user on a video sharing website. Wearing a mask does not reduce recognition rate, of course.

2. Previous Work

Much work has been done on human biometrics. Some popular biometric measures are face recognition [29], speaker voice identification [4] and fingerprint recognition [18]. For human height and gender recognition from photos and videos see for example [20, 2]. A review of biometric measures is presented in [11].

An important and longstanding biometric measure is Gait, an analysis of a person's walking style. A pioneering work by Murray [22] showed that gait is a highly distinctive biometric measure. Initial work has been carried out using specialized equipment such as reflectors and moving lights [6]. Over the last few decades much work has been done on extracting gait information from videos obtained by static cameras ([10, 15, 5]). The spatial shape of the body is commonly used (e.g. [28]). This encodes measures such as height, width, leg length etc. Another popular class of features is temporal based features (e.g. [21]). Such information encodes information such as step velocity, acceleration and frequency.

Several works studied gait analysis from non visual sensors such as accelerometers [19] and pressure sensors [1]. In this work we concentrate on learning gait as expressed in egocentric video.

Little work has been done on gait obtained from moving cameras. In a pioneering work, Shiraga et al. [27] have studied identifying persons from backpack mounted stereo cameras. By estimating rotation and period of motion using 3D geometry they were able to identify users with great accuracy. This however has the disadvantage of requiring specialized equipment. We instead learn gait from widely used standard head-mounted cameras.

Using optical flow for activity recognition from head-mounted cameras has been done by [13, 24, 26, 16] and others. Poleg et al. [23] have used optical flow motion vectors for hiding away the identity of the camera user. We are not aware of published works that used optical flow pattern for identifying egocentric camera users.

Feature design for time series data has been extensively studied. It is particularly important for speech recognition systems ([25]). Speaker verification is a long standing problem which is related to this work. Features that have been found to be helpful for speaker recognition include Linear Predictive Coding (LPC) [9] and MFCC [7].

Instead of hand designing features, we learn features along with the classifier, end-to-end, using convolutional neural networks (CNN). For an overview of deep neural networks see [3]. Learned features are sometimes better than hand-designed features [14].

3. Learning Biometrics from optical flow

In this section we detail our method for learning biometrics from head-worn cameras. Egocentric video suffers from bouncy and unsteady motion caused by user head and body motion. Although usually a nuisance, we show that this information can be useful for biometric feature extraction and consequently for identifying the user.

Real world scenarios vary greatly with the amount of data available about each user. In collaborative settings we may have several hours of video of a user's walking motion. In non-collaborative settings we may have only very short videos (4-30 seconds) of the user's motion.

In this paper we design a CNN classification architecture. Similar networks can be used for user verification (user vs. world) and user identification (one of M known users). We train the network to differentiate between video sequences recorded by the user in the past and sequences recorded by other users. This classifier can then be used to determine if a new sequence was recorded by the user or not. The architecture combines jointly both feature learn-

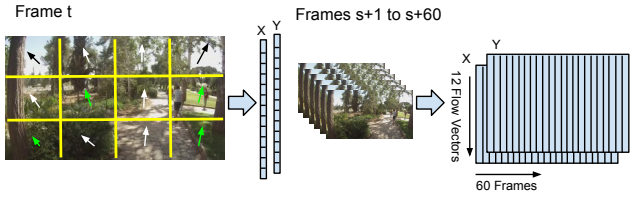

a) b)

Figure 2. a) 12 Optical flow vectors are calculated for each image. We obtain two columns (each of 12 values), for the x and y optical flow components. b) The optical flow columns are stacked over 60 frames into 12×60 grids separately for the x and y components. Each value is divided by the square root of its magnitude.

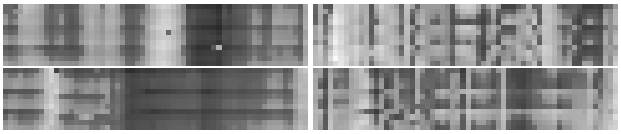

Figure 3. Two example of the front end features. The features contain 12 optical flow vectors per frame (vertical axis), computed for each of 60 frames (horizontal axis). The left and right images are the x and y optical flow features correspondingly.

ing and classification. The classifier is used as a building block in a simple high-level scheme that improves performance from longer sequences. This approach is detailed in Sec. 3.2.

3.1. Low Level Feature Extraction

We extract sparse optical flow features, and train the CNN on the extracted features. Although CNNs are sometimes used to extract features from raw pixel data, this is not applicable here. Our datasets do not contain enough examples to learn features from pixel-level data, and optical flow features are usually difficult to extract automatically even on large datasets.

In the following sections we assume that the video frames were preprocessed in the following way (see Fig. 2):

- Frames are partitioned into a small number of non-overlapping blocks of grid size $d_x \times d_y$. In the experiments we used $d_x = 3,\ d_y = 4$.
- The optical flow of each block between subsequent frames is calculated by the Lucas Kanade algorithm [17]. This can be done very fast.
- The magnitude of the optical flow vector of each block is passed through a compressive non-linearity (we use a square root non-linearity).

The video is divided into sequences of $d_x \times d_y$ optical flow vectors of length T seconds (we use $T = 4$ seconds unless otherwise specified). The sequences overlap by $T/2$. The data instances are therefore of dimension $d_x \times d_y \times T \times fps \times 2$ (we use 15 fps). The factor of 2 corresponds to the dimension of the optical flow vector. Visualization of the input data is shown in Fig. 3.

3.2. User Classification

In this section we assume that we have several minutes of walking egocentric video from each person. This is a reasonable assumption if the user gives us access to a long egocentric video or several short sequences recorded over a period of time. In this case we have enough data to train a classifier, that can determine if a certain video was recorded by the user.

3.2.1 Training a short sequence classifier

Our task in this section is to classify a feature vector x (of fixed length) obtained as described in Sec. 3.1 into a person ID in the identification scenario and target / non-target person in the verification scenario. We describe a CNN architecture to learn features and classifier end to end.

Due to the limited number of data points available in our datasets, we limit our CNN to only 2 hidden layers. Using more layers increases model capacity but also increases over-fitting. The architecture is illustrated in Fig. 4.

Due to the coarse sampling of the optical flow we do not assume much spatial invariance in the features. On the other hand the precise temporal offset of the user's actions is usually not important, e.g. the precise time of the beginning of a user's step is less important than the time between strides. Our architecture should therefore be temporally invariant. The first layer was thus designed to be convolutional in time but not in space. The kernel size spans all the blocks across the x and y components over K_T frames (we use $K_T = 10$). The convolutional layer consists of M kernels (we use $M = 16$). The outputs of the kernels $z_m^1 = W_m * x$ are passed through a ReLU non-linearity ($max(z_m^1, 0)$). We pool the outputs substantially in time, as the feature vector is of high dimension compared to the amount of training data available (we use kernel length of 20 and stride of 15). These values were chosen to roughly correspond to the typical time interval between steps.

The data is then passed through two fully connected layers each followed by a sigmoid non-linearity ($\sigma(z) = \frac{1}{1+e^{-z}}$). The first fully connected hidden layer has N_1 hidden nodes (we used $N_1 = 32$). The second hidden layer has the same number of nodes as output classes (we used 6 in the identification case and 2 in the verification case).

To motivate the use of a CNN, we note that Linear Predictive Coding (LPC), a common technique for speaker verification can be described as convolving the signal with a set of FIR filters. This can be realized by a set of convolutions. The convolutional layer also acts as non-linear PCA. The last two layers act as a non-linear classifier, much like RBF-

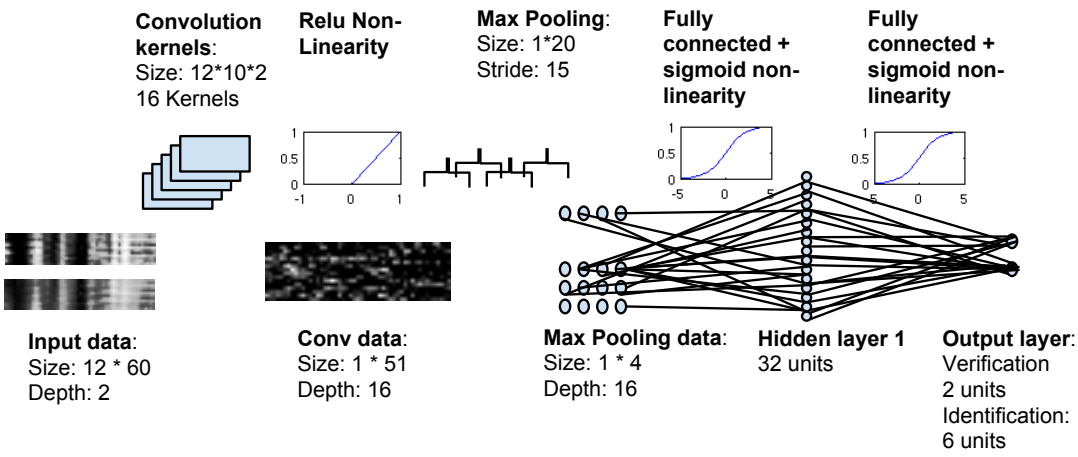

Figure 4. A diagram of the short sequence user recognition CNN architecture. The operations on the data are shown on top, the sizes of subsequent data layers are shown on the bottom. The Neural Network learns the features jointly with the classifier, therefore automatically creating features optimal to this task.

SVM often used in speaker verification. PCA, LPC and the final RBF-SVM classifier are typically applied separately. The CNN trains all operations jointly, possibly improving performance.

The network was trained by Stochastic Gradient Descent on a GPU using the Caffe package [12]. The specific parameter values used were: learning rate 0.001, weight decay 0.004 and momentum 0.9. The mini-batch size was 200.

3.2.2 User classification on longer sequences

Sec. 3.2.1 described a method to train an identity classifier on a short video sequence. The video that we are given to classify might be significantly longer than the duration of the input data instance. A simple scheme is introduced to classify the full video.

Low level features are first extracted as described in Sec. 3.1. The classifier is then run on each data instance x_t and yields probabilities $P(y_t = i)$ where y_t is the output label (person ID for identification, target / non-target person for verification). We have tested two methods of combining the predictions on multiple data instances:

- Mode classifier: Hard classify each data instance $\hat{y_t} = argmax_i P(y_t = i)$ and choose the label that fit the largest number of data instances $argmax_i \#\{\hat{y_t} = i\}$.

- MAP classifier: Classify the video into the globally most likely label, $argmax_i \prod_t P(y_t = i) = argmax_i \sum_t log(P(y_t = i))$. This assumes that data instances are IID, however we have found this requirement is not necessary for the success of the method. MAP classification has helped boost the identification performance on the FPSI dataset to around 95%.

The effectiveness of the various methods is compared in Sec. 5.

4. Experiments

Several experiments were performed to verify the effectiveness of our method . As there is no standard dataset for egocentric video biometrics, we use both a public dataset and a new dataset we collected for our experiments.

The First-Person Social Interactions (FPSI) dataset was collected by Fathi et al [8]. 6 individuals (5 males, 1 female) collected a day's worth of egocentric video each recorded by head-worn GoPro cameras. Due to battery and memory limitations of the camera, the users occasionally took the cameras off and put them on again, ensuring that camera extrinsic parameters were not kept constant.

In this work we deal with human biometrics while walking rather than sitting or standing. We extracted the walking portions of each video using manual labels. It is possible instead to use a classifier such as in [24] to find the walking intervals.

The FPSI dataset contains a large amount of video for a small number of users. Another common scenario is having a small amount of video (a few minutes) for a larger number of users. To this end, we have collected the Egocentric Video Biometrics dataset (EVB).

Our dataset consists of 47 head-mounted video sequences from 34 individuals. The dataset contains outdoor sequences, where individuals walked mostly within the same geographical area. The apparatus used was a GoPro camera (of models Hero 3 Black , Hero 3 White, Hero 3+ Silver and Hero 3+ Black) attached to baseball caps (see Fig. 5). Participants were told to walk normally for several minutes. 7 other participants recorded two long sequences



Figure 5. The apparatus used to record the EVB dataset.

(5-7 minutes) each, using a different cap for each sequence. The video sequences were manually processed to retain the walking sections only.

Features were extracted from the data using the steps described in Sec. 3.1:

1. Sub-sampling the frame rate to 15 FPS.
2. Dividing each frame into a 3×4 grid (each cell of size 360×480 pixels).
3. Computing optical flow vectors for each cell using Lucas-Kanade [17]. We used the code distributed by [24].
4. Reducing the dynamic range of the optical flow by taking the square root of optical flow magnitude.
5. We used the above optical flow vectors, 12 vectors per frames for each of 60 frames (4 seconds of video), as input features to the CNN.

Using the features calculated above, we have performed the following experiments:

4.1. User Identification

Given a number of individuals each with a large amount of training data, we classify a new video sequence into the user that recorded it. We tested our method on the FPSI dataset. We use for training the first 80% of sequences taken early in the day, and tested on the last 20% sequences recorded at the end of the day (for each individual). This is done to reduce overfitting to a particular time or camera setup. The data were randomly sub-sampled to ensure equal number of samples for each person in both training and testing. The results are described in Sec. 5.

4.2. User Verification

Given a target user with a few minutes of training data, and negative training data by other users (possibly with only a small amount for each user), we verify if a new short video sequence came from the target user or some other user. Verification on longer sequences is done by combining the predictions from subsequent short sequences. As the FPSI contains only 6 users it was not suitable for the verification task (this is elaborated upon in Sec. 5). Instead we used the new EVB dataset. For each experiment we selected as target user, a user that recorded sequences with two different caps. This ensured we did not fit to the extrinsic parameters of a particular camera. One target sequence was used for training and the other for testing. Each long sequence contains 100-200 input examples. The other sequences from the EVB dataset are divided equally between negative training and negative test data samples. Care was taken to ensure that all sequences of a non-target user would appear only in the training or test data but not in both. This was done to ensure we did not overfit to specific non-target users. We randomly sub-sampled the data to ensure equal numbers of negative and positive train and test data.

4.3. Baseline Comparison

To understand how the performance of the CNN compares to other high quality classifiers, we compared the CNN verification performance against an RBF-SVM trained on a set of predetermined features. The features used were:

1. Raw - The same features used by the CNN.
2. PCA - The dimensionality of each frame is reduced from 24 to 8 by PCA after optical flow vector calculation (stage 3). The rest of the stages are carried out normally. The dimensionality of the features is 8×60. This tests if reducing feature dimensionality to the most significant components improves generalization.
3. LPC - After stage 3, each of the 24 time series (optical flow values for a given block and x or y component over time), is encoded by 5 LPC values. No square root operation (stage 4) is carried out. The dimensionality of the features is 24×5. LPC is commonly used in speaker verification.
4. PCA-LPC - After stage 3, the dimensionality of each frame is reduced by PCA to 8 components, each time series is then encoded by 5 LPC components. No square root operation (stage 4) is carried out. The output dimensionality is 8×5.

We have found that subtracting the average optical flow vector for each frame, from all the optical flow in the frame (after stage 3) helped performance and present the results with such pre-processing applied. It is in fact akin to 2D stabilization. Additionally, as is standard practice, all feature vectors are mean and variance normalized across the training set before being used by the SVM. We have carried out the above experiments for linear SVM too, but the results for RBF-SVM are better in almost all cases.

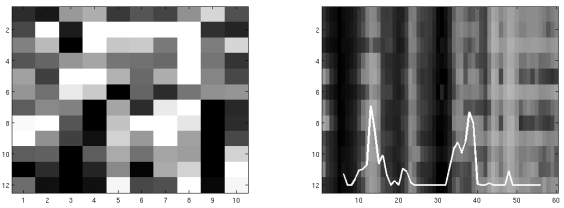

a) b)

Figure 6. a) The weights of a filter learned by the network for the y optical flow component (the convolution is in time only, the filter kernel has the same height and depth as the features). It looks for positive acceleration followed by negative acceleration 0.33 seconds later. b) The response of the filters superimposed on optical flow features (y component). It fires at suitable acceleration boundaries e.g. frame 15.

Actual // Pred	1	2	3	4	5	6
1	123	32	3	13	46	31
2	52	166	1	3	21	38
3	2	1	216	8	4	0
4	11	2	16	217	4	0
5	22	15	4	2	126	14
6	34	28	4	1	43	47

Table 1. Confusion matrix for identification performance on the FPSI dataset for 4 second long sequences. Higher accuracy is obtained on longer sequences.

5. Results

In this section we present the results of the experiments described in Sec. 4.

Feature learning and classification was done end to end using a Convolutional Neural Network. The network was kept small due to the limited amount of training data. Although neural networks are not usually easy to interpret, looking at the weights of the first layer can give some intuition about the features learned by the network. Fig. 6.a) shows a first layer filter learned by the network and the corresponding responses on an input example. For illustration purposes only the y component is shown. Looking at the weights, we see that the filter is tuned for positive acceleration in the y direction, followed by negative acceleration 5 frames (0.33 seconds) later. From Fig. 6.b) we can see the filter is triggered when such intervals between accelerations occur (e.g. frame 15) but not at shorter intervals (e.g. frame 24). This can correspond to classification by inter-step interval.

Table. 1 presents the person classification performance of our network. The table shows the confusion matrix between the 6 users for video sequences of 4 second duration. The system was able to discriminate between the users effectively although some users are more difficult to separate for example User 1 and User 6.

Most of the time, test videos are longer than 4 seconds.

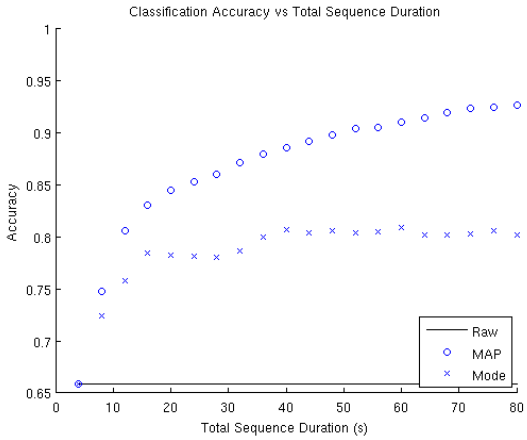

Figure 7. Classification accuracy vs Total sequence duration for $T = 4$ seconds. MAP classification consistently beats mode classification. They significantly improve on the short sequence identification performance.

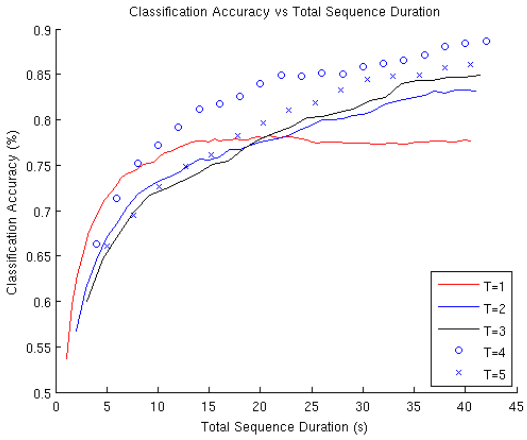

Figure 8. Classification accuracy vs. Short Sequence Duration used. More frames result in a more accurate classifier, but shorter sequences yield more feature points per video. $T = 4$ seconds results in the best overall performance.

There are several ways of using our previously trained classifier to identify the user. Two such strategies were described in Sec. 3.2.2. In Fig. 7 we compare the performance of the strategies. We see that longer videos do result in better identification for both strategies. We can also see that MAP classification significantly beats the mode classification rule (up to a score of around 95% compared to the mode score of below 80% and the random guess score of 16.7%).

To justify our choice of short sequence duration (4 seconds unless specified otherwise), we plot the average classification rate vs short sequence duration T in Fig. 8. It can be seen that $T = 4$ seconds presents a good tradeoff between model richness (due to having enough frames) and limiting overfitting (by having enough data points).

Fig. 9 shows the identification probabilities for each

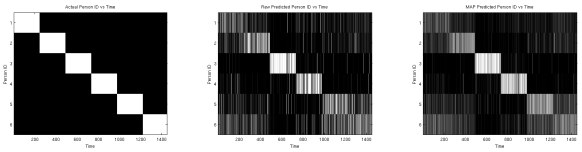

a) b) c)

Figure 9. a) Ground truth labels. b) Raw CNN probabilities. c) MAP rule probabilities (for $T = 12$ seconds.)

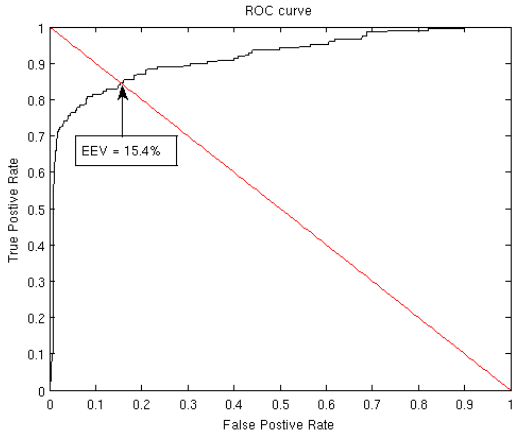

Figure 10. A graph illustrating the EER measure. When the line intersects the ROC curve, the False Acceptance Rate and False Rejection Rates are equal. The point of intersection is the EER.

sample in the test videos. All test input sequences for each user are presented in subsequent temporal order. The actual labels are presented in Fig. 9.a, the probabilities from the CNN are shown in Fig. 9b and the MAP results (for 12s) are shown in Fig. 9c. This effective temporal smoothing gives intuition for how the MAP rule improves classification performance.

Table. 2 presents the verification performance for 7 persons against the EVB dataset. Only individuals with sequences of length greater than 5 minutes were used as positive targets. In order to compare system performance by a single number it is common in biometrics to use the Equal Error Rate (EER), the error rate at which the False Acceptance Rate (FAR) and False Rejection Rate are equal. An illustration of the EER is presented in Fig. 10.

The CNN is compared against 4 baseline methods: SVM, PCA-SVM, LPC-SVM, PCA-LPC-SVM. The CNN methods perform better than the baseline methods on almost all cases. The closest competitor is LPC-SVM (which is commonly used in speaker verification). CNN has an average EER of 12.1% compared to 17.2% achieved by LPC-SVM. As LPC encodes the data by an FIR filter, it can be learned by convolutions. The convolutions are also able to use information across the frame and in both x and y flow components whereas LPC only looks at a single block along time. It is therefore less general than the CNN.

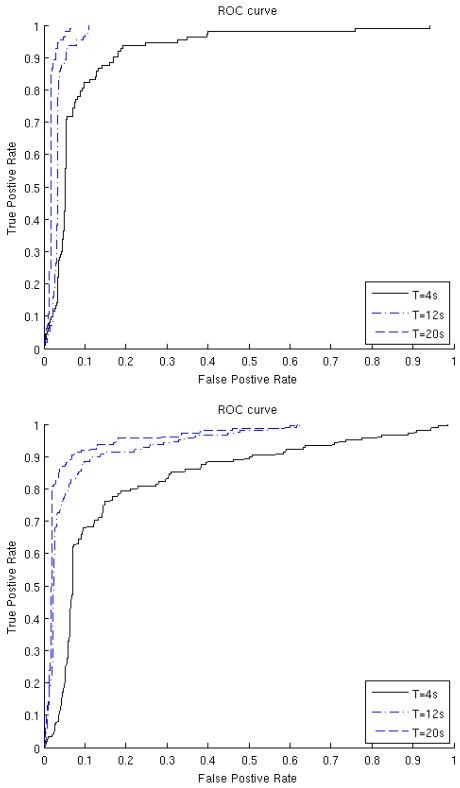

Figure 11. ROC curves for Person 1 (top) Person 2 (bottom) from the EVB dataset. The three lines in each plot show the verification performance for different total test sequence duration. Each 4 second sequence was classified using the CNN, the scores were aggregated using the MAP rule.

Plots of the ROC curve are presented for Person 1 and Person 2 from the EVB dataset in Fig. 11. It should be noted that the negative test examples are previously unseen users. It is crucial for verification as we cannot model each person in the world. By modeling the target person we can separate the person from the global population. Our results show the network is effective at learning the target person rather than attempting to model the global population, and is therefore able to generalize to unseen non-target users.

We tried learning a verification classifier by choosing as a target person from the FPSI dataset, and 4 other users as negative training data with no added users from the EVB dataset. The morning sequences of the target person were used for training and the afternoon for testing. We tested the verification performance between the afternoon sequences of the target user and the remaining 6th non-target user from the FPSI dataset. The network however, fit to the non-target users and has not been able to generalize to the new non-target user. We therefore conclude that a significant number of non-target users (such as present in the EVB dataset) is required for training a verification classifier.

To test the possibility that stabilization would take away

Person	1	2	3	4	5	6	7	Mean
SVM	14.5%	25.3%	10.6%	24.3%	26.3%	16.2%	41.0%	22.6%
PCA-SVM	22.2 %	40.1%	17.0%	30.9%	29.3%	15.7%	41.5%	28.1%
LPC-SVM	8.3%	21.5%	9.2%	23.9%	**10.5**%	12.0%	34.9%	17.2%
PCA-LPC-SVM	19.4 %	34.8%	21.5%	36.1%	22.8%	18.0%	42.9%	27.9%
CNN	**7.3** %	**11.6**%	**4.8**%	15.9%	14.2%	**8.7**%	22.5%	**12.1**%
CNN-Stable	9.6%	15.8%	6.8%	**15.8**%	19.7%	12.6%	**18.4**%	14.1%

Table 2. Verification Equal Error Rates for 7 different persons against the EVB dataset. The CNN result is compared with i) RBF SVM, ii) PCA dimensionality reduction to 8 components followed by an RBF-SVM, iii) coding each of the 24 series by 5 LPC components followed by an RBF-SVM. iv) Reducing the number of data-series to 8 by PCA, coding each of the 8 time-series by 5 LPC components followed by an RBF SVM. We show the CNN result with and without 2D stabilization. The CNN architecture significantly outperformed the baseline methods in most cases. Stabilization is shown to degrade results slightly. LBP-SVM is the closest competitor, its performance was improved by stabilization (only the stabilized case is shown). The total duration of each sequence used was 12 seconds.



a) b) c)

Figure 12. Common failure cases: a) When walking alongside another person, the user frequently turns resulting in sharp blurry motions. b) Turning sharply results in atypical fast motions. c) Large nearby occluding objects can cause atypical optical flow patterns.

some or all the biometric information in the frame motions, the verification experiments were also performed after subtracting the mean framewise optical flow vector from all optical flow in the frame. As motion between frames is small and some lens distortion correction was performed, this is similar to 2D stabilization. Table. 2 shows that stabilization degrades CNN performance only slightly (2%). It actually improved performance for the SVM based methods. We therefore conclude that simple stabilization does not remove the biometric information from the video. Various 3D stabilization schemes are more complex and can possibly remove more biometric information. This investigation is left for future work.

In Fig. 12 several failure cases are shown. It can be seen that failure is sometimes caused by occlusions, lack of features or drastic head movements (shown by the significant blur). It is possible that by filtering out such cases, higher recognition performance may be achieved.

The CNNs used in the above experiments have only 2 hidden layers and do not have many hidden units. In order to limit overfitting, we could not use larger networks due to the limited amount of data. If larger datasets will become available, larger models could be trained probably resulting in higher recognition performance.

6. Conclusion

A method to determine the identity of a user from head worn egocentric camera video has been presented. Our method relied on biometric information implicit in camera motion, which can identify users accurately. Biometric features were learned automatically using Convolutional Neural Networks. The CNN architectures were shown to generalize and improve on physically motivated hand designed features.

Experimental evidence has confirmed that our method can determine user identity with high accuracy. The experiments were carried out on a commonly used egocentric video dataset, and on a larger dataset collected by us.

We have tested the effects of 2D video stabilization on classification accuracy, and found only slight degradation in performance. It is possible that more elaborate 3D stabilization would have a more significant effect, we leave this investigation for future work.

The implication of our work is that users' head-worn egocentric videos give much information away. This information can be used benevolently (e.g. camera theft prevention, user analytics on video sharing websites) or maliciously. Care should therefore be taken when sharing such raw video.

Acknowledgement: This research was supported by Intel ICRI-CI, by Israel Ministry of Science, and by Israel Science Foundation.

References

[1] S. J. M. Bamberg, A. Y. Benbasat, D. M. Scarborough, D. E. Krebs, and J. A. Paradiso. Gait analysis using a shoe-integrated wireless sensor system. *Information Technology in Biomedicine, IEEE Transactions on*, 12(4):413–423, 2008. 2

[2] C. BenAbdelkader, R. Cutler, and L. Davis. Person identification using automatic height and stride estimation. In *Pattern Recognition*, 2002. 2

[3] Y. Bengio. Learning deep architectures for ai. *Foundations and trends in Machine Learning*, 2(1):1–127, 2009. 2

[4] F. Bimbot, J.-F. Bonastre, C. Fredouille, G. Gravier, I. Magrin-Chagnolleau, S. Meignier, T. Merlin, J. Ortega-García, D. Petrovska-Delacrétaz, and D. A. Reynolds. A tutorial on text-independent speaker verification. *EURASIP*, 2004:430–451, 2004. 1, 2

[5] A. F. Bobick and A. Y. Johnson. Gait recognition using static, activity-specific parameters. In *CVPR*, 2001. 1, 2

[6] J. E. Cutting and L. T. Kozlowski. Recognizing friends by their walk: Gait perception without familiarity cues. *Bulletin of the psychonomic society*, 9(5):353–356, 1977. 2

[7] S. Davis and P. Mermelstein. Comparison of parametric representations for monosyllabic word recognition in continuously spoken sentences. *ASSP*, 1980. 2

[8] A. Fathi, J. K. Hodgins, and J. M. Rehg. Social interactions: A first-person perspective. In *CVPR*, 2012. 4

[9] S. Furui. Cepstral analysis technique for automatic speaker verification. *ICASSP*, 1981. 2

[10] J. Han and B. Bhanu. Individual recognition using gait energy image. *PAMI*, 2006. 1, 2

[11] A. K. Jain, P. Flynn, and A. A. Ross. *Handbook of biometrics*. Springer, 2007. 2

[12] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014. 4

[13] K. M. Kitani, T. Okabe, Y. Sato, and A. Sugimoto. Fast unsupervised ego-action learning for first-person sports videos. In *CVPR*, 2011. 2

[14] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012. 2

[15] L. Lee and W. E. L. Grimson. Gait analysis for recognition and classification. In *Automatic Face and Gesture Recognition*. 1, 2

[16] Z. Lu and K. Grauman. Story-driven summarization for egocentric video. In *CVPR*, 2013. 2

[17] B. D. Lucas, T. Kanade, et al. An iterative image registration technique with an application to stereo vision. In *IJCAI*, volume 81, pages 674–679, 1981. 3, 5

[18] D. Maltoni, D. Maio, A. K. Jain, and S. Prabhakar. *Handbook of fingerprint recognition*. springer, 2009. 1, 2

[19] J. Mantyjarvi, M. Lindholm, E. Vildjiounaite, S.-M. Makela, and H. Ailisto. Identifying users of portable devices from gait pattern with accelerometers. In *ICASSP*, 2005. 2

[20] B. Moghaddam and M.-H. Yang. Gender classification with support vector machines. In *Automatic Face and Gesture Recognition*, 2000. 2

[21] H. Murase and R. Sakai. Moving object recognition in eigenspace representation: gait analysis and lip reading. *PRL*, 17(2):155–162, 1996. 2

[22] M. P. Murray. Gait as a total pattern of movement: Including a bibliography on gait. *American Journal of Physical Medicine & Rehabilitation*, 46(1):290–333, 1967. 2

[23] Y. Poleg, C. Arora, and S. Peleg. Head motion signatures from egocentric videos. In *ACCV*, 2014. 2

[24] Y. Poleg, C. Arora, and S. Peleg. Temporal segmentation of egocentric videos. In *CVPR*, 2014. 2, 4, 5

[25] D. A. Reynolds and R. C. Rose. Robust text-independent speaker identification using gaussian mixture speaker models. *Speech and Audio Processing, IEEE Transactions on*, 3(1):72–83, 1995. 2

[26] M. S. Ryoo and L. Matthies. First-person activity recognition: What are they doing to me? In *CVPR*, 2013. 2

[27] K. Shiraga, N. T. Trung, I. Mitsugami, Y. Mukaigawa, and Y. Yagi. Gait-based person authentication by wearable cameras. In *INSS*, 2012. 2

[28] L. Wang, T. Tan, H. Ning, and W. Hu. Silhouette analysis-based gait recognition for human identification. *PAMI*, 25(12), 2003. 2

[29] W. Zhao, R. Chellappa, P. J. Phillips, and A. Rosenfeld. Face recognition: A literature survey. *Acm Computing Surveys (CSUR)*, 35(4):399–458, 2003. 1, 2