

Unsupervised adaptation for acceleration-based activity recognition: Robustness to sensor displacement and rotation

Ricardo Chavarriaga · Hamidreza Bayati · José del R. Millán

Received: date / Accepted: date

Abstract A common assumption in activity recognition is that the system remains unchanged between its design and its posterior operation. However, many factors affect the data distribution between two different experimental sessions. One of these factors is the potential change in the sensor location (e.g. due to replacement or slippage) affecting the classification performance. Assuming that changes in the sensor placement mainly result in shifts in the feature distributions, we propose an unsupervised adaptive classifier that calibrates itself using an online version of expectation-maximisation. Tests using three activity recognition scenarios show that the proposed adaptive algorithm is robust against shift in the feature space due to sensor displacement and rotation. Moreover, since the method estimates the change in the feature distribution it can also be used to roughly evaluate the reliability of the system during online operation.

Keywords Activity recognition · Sensor displacement · Unsupervised adaptation · Linear discriminant analysis · Expectation-maximisation ·

Based on “Unsupervised adaptation to on-body sensor displacement in acceleration-based activity recognition”, by H. Bayati, J. d. R. Millán, and R. Chavarriaga which appeared in the Proceedings of the International Symposium on Wearable Computers, San Francisco, California, June 2011. © 2011 IEEE.

R. Chavarriaga and H. Bayati contributed equally to this work.

Chair on Non-Invasive Brain-Computer Interface (CNBI)
Ecole Polytechnique Fédérale de Lausanne
CH-1015 Lausanne, Switzerland.
E-mail: ricardo.chavarriaga@epfl.ch

1 Introduction

Activity recognition from on-body sensors is largely being studied in applications like gaming [10], industrial maintenance [26] and health monitoring [29]. In particular, acceleration sensors have been applied for recognising different activities ranging from modes of locomotion [30] to complex daily living activities [19]. Typically, the design of these systems (e.g., feature selection, classification) assumes that the characteristics of the sensor network will not change. However, during system operation body-worn sensors may slip or rotate. Similarly, it is unrealistic to expect users to precisely re-attach the sensors at the same location from day to day. These changes may degrade the recognition performance. In order to address the issue of sensor location variability, we propose a self-calibrating approach based on probabilistic classifiers. The method tracks changes in the feature distribution in an unsupervised manner using an online implementation of the Expectation-Maximisation algorithm.

Several approaches have been proposed to cope with those changes in activity and gesture recognition using body-worn and ambient sensors (e.g., using vision-based recognition [31,32,9]). Some of them try to exploit the specific characteristics of the change they want to address. For example, Kunze et. al. used the combined information of gyroscope and accelerometers to distinguish between rotation and translation [12]. Their work suggests that, in contrast to sensor rotation, translation does not affect significantly the acceleration signals. Based on this, they proposed a heuristic method that yielded higher recognition rates for displaced sensors on body segments. Other approaches focus on the selection of displacement-invariant features [25]. Förster et. al. use genetic programming to find invariant fea-

tures on acceleration-based gesture recognition [7]. In this case, they left one sensor out from training and used evolving features of other sensors to train a classifier. In another work, the same group proposed an online unsupervised self-calibration algorithm [8]. Using online adaptation they adjusted a nearest class centre classifier (NCC). They applied the method on synthetic data in addition to two real life datasets corresponding to gesture-recognition scenario mentioned above and a fitness scenario dataset. Alternative methods to find invariant features is to train classifiers using data recorded at different body locations. Lester and colleagues used sensors placed on the shoulder, wrist and the waist of subjects performing daily life activities [14]. Then, they compare the performance of the classifiers trained on data from individual sensors against classifiers trained on data from the three sensors altogether. This approach relies on the recording of enough data from all available positions. However, this may imply a costly setup for the collection of the training data, while at the same time it might be difficult to fully cover all potential displacements that a sensor may suffer.

A third approach assumes that changes in the sensor placement affect the signal feature distributions in a particular manner. A particular case, termed *covariate shift*, assumes that the training and testing feature distributions change but the conditional distribution of the classifier output given an input is the same. Based on this assumption, Sugiyama et. al. proposed a modification of cross validation technique called importance weighted cross validation (IWCV) that can be used for model and parameter selection in classification tasks [27]. They used IWCV to select the parameters of an *importance weighted LDA* (IWLDA) where the weights are the ratio of the test and train pattern distributions in the calibration session. In experimental studies this ratio is replaced by its empirical estimates, using either Kullback-Leibler importance Estimation Procedure (KLIEP) or unconstrained least square importance fitting (uLSIF) [28, 15]. However, it should be noticed that this method requires a calibration session for estimating the ratio of the distributions between training and test session.

We propose a method that assumes that the main change in the feature distribution corresponds to a shift. This allows us to propose an expectation-maximisation algorithm that can be used online to estimate the feature distributions on an unsupervised manner. This yields a mechanism to estimate the distribution shift and adapt the original classifier. In this paper, we extend our previous studies of the method regarding sensor displacement in order to assess its performance upon rotational noise.

The rest of this paper is structured as follows: In section 2 we describe the proposed method. Then, we revisit previously reported tests in the case of sensor displacement [2]. For this, we use the same scenarios introduced by Förster and colleagues: a Human-computer-Interaction (HCI) and a fitness scenario (Sections 3.1.1 and 3.1.2, respectively) [8]. In section 3.2, we further test our method in the case of sensor rotation. To this end, we use the HCI and a daily living scenario. We then conclude and discuss the presented method and results, as well as future directions of research.

2 Unsupervised adaptive classifier

The underlying assumption of a typical recognition system is that the feature distributions estimated from the training data will remain unchanged during system operation. However, this assumption may not hold in real-life applications, resulting in a decrease in performance. In order to deal with that, we propose an unsupervised approach aimed at estimating changes in the feature distribution and thus allowing the adaptation of probabilistic classifiers. We particularly study this adaptive classifier in the case of changes in the sensor placement.

In the proposed approach, knowing that sensor displacement may result in changes in the overall feature distribution, we assume that these changes can be fully characterised by a shift of an unknown magnitude and direction. Given this, we estimate the distribution shift using an online version of the expectation-maximisation algorithm. Once the shift vector has been estimated, incoming samples can be *shifted back* and classified using the original classifier (i.e., the one trained in the original feature distribution).

Specifically, let $\mathcal{C}(\mathbf{x})$ be a classifier trained on data with feature distribution $p(\mathbf{x})$. Given the assumptions described above, the distribution of new incoming samples $p(\mathbf{y})$ will be equal to the original distribution shifted by a vector θ ,

$$p(\mathbf{y}) = p(\mathbf{x} + \theta) \quad (1)$$

Therefore, the classification performance will not be affected if samples are shifted back before classification: $\mathcal{C}(\mathbf{y} - \theta)$. In consequence, self-adaptation can be achieved by estimating the shift vector θ in an online, unsupervised manner, as described below.

Given the training feature distribution $p(\mathbf{x})$,

$$p(\mathbf{x}) = \sum_{i=1}^I P(\mathbf{z} = \omega_i) P(\mathbf{x} | \mathbf{z} = \omega_i) \quad (2)$$

where \mathbf{x} represents the features, $P(\mathbf{z} = \omega_i)$ is the prior probability of class i , I is the number of classes, and the class-conditional distribution is a normal distribution with mean μ_i and covariance matrix Σ_i ,

$$P(\mathbf{x}|\mathbf{z} = \omega_i) \equiv N(\mathbf{x}|\mu_i, \Sigma_i) \quad (3)$$

Let \mathbf{y} be the samples recorded during system operation. As described above, the shifted distribution $p(\mathbf{y} - \theta)$ should correspond to the same distribution as the training samples (Eq 1). Given a matrix \mathbf{Y} where the j -th column represents the j -th observation, \mathbf{y}_j and \mathbf{Z} be a matrix of labels, with corresponding \mathbf{z}_j that are latent variables. We can define the log-likelihood for a specific value of θ ,

$$\ln p(\mathbf{Y}|\theta) = \ln \sum_{\mathbf{Z}} p(\mathbf{Y}, \mathbf{Z}|\theta) \quad (4)$$

We use an expectation-maximisation (EM) algorithm to maximise the likelihood over θ [3]. Given the previous estimation θ^{old} the *E-step* corresponds to computing the posterior probabilities given the shift vector $p(\mathbf{Z}|\mathbf{Y}, \theta^{old})$. For the j -th observation, it is computed as:

$$P(\mathbf{z}_j = \omega_s | \mathbf{y}_j, \theta^{old}) = \frac{P(\mathbf{z}_j = \omega_s) P(\mathbf{y}_j - \theta^{old} | \mathbf{z}_j = \omega_s)}{\sum_{i=1}^I P(\mathbf{z}_j = \omega_i) P(\mathbf{y}_j - \theta^{old} | \mathbf{z}_j = \omega_i)} \quad (5)$$

The *M-step* corresponds then to evaluating θ^{new} ,

$$\theta^{new} = \arg \max_{\theta} Q(\theta, \theta^{old}) \quad (6)$$

where

$$Q(\theta, \theta^{old}) = \sum_{\mathbf{Z}} p(\mathbf{Z}|\mathbf{Y}, \theta^{old}) \ln p(\mathbf{Y}, \mathbf{Z}|\theta) \quad (7)$$

$$Q(\theta, \theta^{old}) = \sum_{j=1}^J Q_j(\theta, \theta^{old}) \quad (8)$$

where J is the number of patterns and $Q_j(\theta, \theta^{old})$ is defined as follows:

$$\sum_{i=1}^I P(\mathbf{z}_j = \omega_i | \mathbf{y}_j, \theta^{old}) (\ln P(\mathbf{z}_j = \omega_i) + \ln N(\mathbf{y}_j - \theta | \mu_i, \Sigma_i)) \quad (9)$$

Algorithm 1 Online shift estimation

```

Initialise  $\theta = \theta^0$ 
for every new sample  $\mathbf{y}_j$  do
  Compute posterior probability of the shifted sample
  using Eq. 5.
  Classify the pattern based on maximum posterior rule.
  Compute shift update,  $\Delta\theta$ 
  if ( $|\Delta\theta| > \Theta$ )
    Update the shift  $\theta$  (Eq. 10).
  endif
end for
  
```

In order to have a run time estimation of the distribution shift we use an online version of Levenberg-Marquardt algorithm [18]. This yields an on-line update rule that maximises Eq. 9 using its gradient (\mathbf{g}) and Hessian (H),

$$\theta^{new} = \theta^{old} + \Delta\theta \quad (10)$$

where,

$$\Delta\theta = (H + \lambda\mathcal{I})^{-1} \mathbf{g} \quad (11)$$

$$\mathbf{g} = \sum_{i=1}^I P(\mathbf{z}_j = \omega_i | \mathbf{y}_j, \theta^{old}) \Sigma_i^{-1} (\mathbf{y}_j - \theta^{old} - \mu_i) \quad (12)$$

$$H = \sum_{i=1}^I P(\mathbf{z}_j = \omega_i | \mathbf{y}_j, \theta^{old}) \Sigma_i^{-1} \quad (13)$$

The λ term in Eq. 11 is a small positive number and \mathcal{I} is identity matrix. This regularisation term prevents from inverting a singular matrix.

To sum up, given a trained probabilistic classifier—Linear or Quadratic Discriminant Analysis (LDA or QDA respectively)—shifts in the feature distribution can be estimated online using Algorithm 1. In order to avoid small oscillations in the estimation when there are small changes in the feature distribution, the shift θ is only updated when the magnitude of the estimated change exceeds a threshold (Θ). Note that at the beginning of the operation, an initial value Θ^0 has to be set. Having no knowledge about how the distribution may have changed since training, we set this value to be zero, thus assuming no change.

3 Results

We test the performance of the adaptation approach on real data recorded in activity recognition scenarios. We emulate changes in the sensor location (Section 3.1), as well as sensor rotation (Section 3.2). In the first case, we use data recorded simultaneously by sensors located at different places, using data from one sensor for training and testing on data from another one. In the latter case, signals from each sensor are artificially rotated by a given angle. The use of artificial rotation allows us to evaluate different conditions and better characterise the performance of the method. In the rest of the paper the adaptation mechanism detailed in the previous section is used to adapt a LDA classifier (henceforth termed aLDA).

3.1 Robustness to sensor displacement

We use two activity recognition scenarios to test the performance of the adaptive classifiers upon changes in the sensor location. The first one corresponds to a gesture based human-computer interaction (HCI) scenario, while the second one corresponds to fitness activities where aerobic movements are performed by the subject. In both scenarios, several acceleration sensors were placed on the subject limbs in order to simultaneously record activity at different body locations, see Figures 1(a), 1(b). This allows to emulate sensor displacement by testing classifiers trained at one location using data from a sensor placed at a different location. These scenarios and testing procedure were previously introduced by Förster and colleagues [8].

We report the classification performance of a static LDA classifier, as well as the proposed adaptive version of LDA (aLDA). Moreover, we also evaluate the performance of *Importance Weighted LDA*, (IWLDA). This method relies on the covariate shift assumption and requires a calibration dataset to estimate the distribution shift (c.f. Section 1, [28]). In the reported simulations for IWLDA we used all test samples as calibration dataset, therefore corresponding to the performance of an off-line recognition system. KLIEP was applied for the importance estimation (for IWLDA we set $\lambda = 1$ and for KLIEP we set $\delta = 0.01$ and three Newton iterations). It should be noticed that in the reported results for IWLDA the feature distribution change is first estimated and then kept fixed for estimating the accuracy on the testing set. On the contrary, for aLDA we report the accuracy of the classification while the adaptation process takes place, therefore emulating the online performance.

A more detailed analysis of the method performance upon sensor displacement, including a comparison with the adaptive NCC method proposed by Förster and colleagues [8], has been previously reported in [2].

3.1.1 HCI Gesture Scenario

The HCI scenario deals with the recognition of five different hand gestures: a triangle, an upside-down triangle, a circle, a square, and an infinity symbol [7, 8]. Six USB accelerometers are placed at different positions on the right lower arm of the subject, and aligned to minimise rotational variation (c.f. Figure 1(a) and [8]). We use data while one subject performs 50 repetitions of each gesture. Data are manually segmented to contain only a single action with duration between five to eight seconds. We created training and testing sets containing two thirds and one third of the data respectively, with classes equally distributed on both sets.

We assess the performance of the adaptive approach using the mean value, the standard deviation, min, max, energy in addition to magnitude of acceleration signals and correlation between each pair of three axes of each sensor. The dimensionality of the feature space is reduced using *canonical variate analysis* (CVA), also known as multiple discriminant analysis, leading to a four-dimensional feature space (i.e., corresponding to the number of classes minus one) [6, 11]. The update threshold Θ was set to 1.5, as this corresponds to the maximum estimated shift when the method is applied to the *training dataset*, and λ was set to the absolute value of the smallest non-positive eigenvalue of $H+0.01$.

The classification performance of the adaptive LDA and IWLDA classifiers is shown in Figure 2(a). In these plots, the performance of each approach (vertical axis) is compared to the performance of the fixed LDA classifier (horizontal axis). Each point corresponds to one of the tested sensor combinations. Red circles show the performance when there is no change in the sensor location (i.e., the classifier is tested on data from the same sensor it was trained). Points above the diagonal line correspond to an improvement due to the adaptation process with respect to the static classifier. It shows that the adaptive LDA outperforms the static classifier in most cases, while the accuracy remains similar when there is no change in the sensor location. Moreover, IWLDA results in very small improvement over the LDA classifier.

We also show the average performance for the three classifiers with respect to the sensor change (Figure 2(a), rightmost plot). The performance of the LDA classifier decreases significantly when tested with data recorded at a different location. In contrast, aLDA consistently

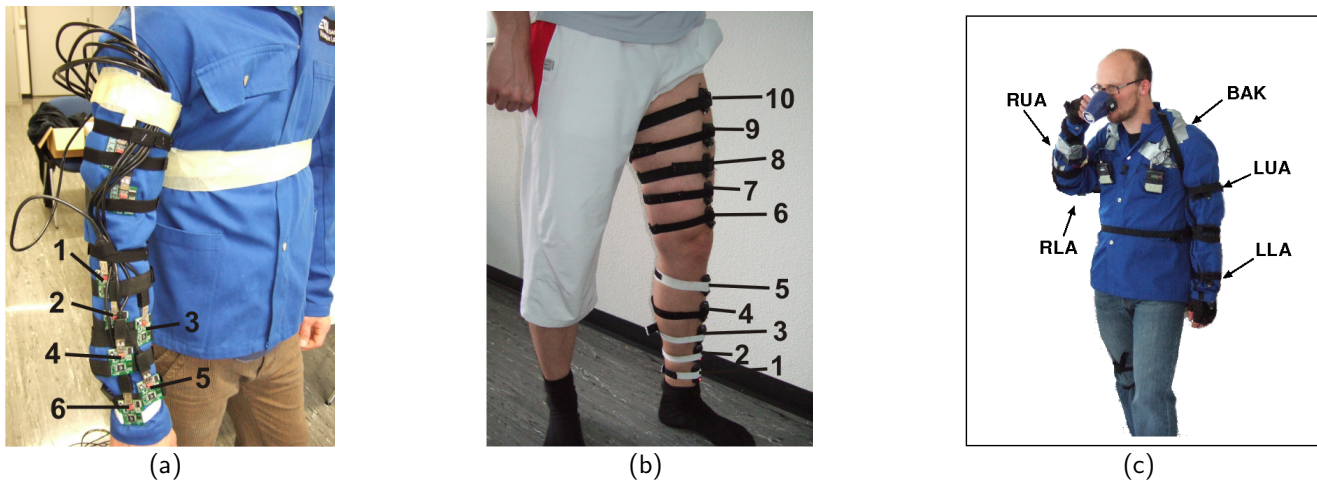


Fig. 1 Sensor placement for the different experimental setups. (a) HCI gesture recognition scenario. (b) Fitness Scenario. (c) Daily living scenario.

outperforms both the LDA and IWLDA classifiers. Surprisingly, IWLDA does not allow any improvement with respect to the LDA classifier when tested on another sensor location. This suggests that the recalibration process does not provide enough information to properly estimate the new feature distributions. Further discussion on this issue is presented in section 4.

Since the adaptation process relies on the estimation of changes in the feature distribution, one may expect that it performs better when there are small changes in the sensor location. In the case of no sensor location change ($t = s$), the aLDA adaptive mechanism yields a small decrease in performance with respect to the static classifier. In contrast, aLDA average performance is about 20% higher than LDA when tested in sensors located next to the training sensors ($|t - s| = 1$). Similarly, aLDA also improves performance in the other sensor combinations ($|t - s| > 1$). In particular, we observe that the aLDA is quite robust for the location sensors 3 to 6 (i.e., sensors located closer to the wrist). Indeed the average performance after displacement among of these positions is equal to 75.2% and 86.9% for the two simulated sets of features (c.f. Figure 3).

3.1.2 Fitness Activity Dataset

The second scenario corresponds to a fitness scenario where five different aerobic movements of the leg were recorded using 10 bluetooth acceleration sensors located on the subject's leg [8]. Five of the sensors were placed on the lower leg and the other five on the thigh (c.f. Figure 1(b)). Sensors were located equidistantly and roughly with the same orientation so as to model only translation. During the experiment the subject performs the movements shown in a video by an instructor. The

video contains all movement classes equally represented, and is presented five times.

For each sensor, the mean and variance of the acceleration magnitude based on a sliding window with two thirds of overlap are used as features. As in the previous application, the data was divided into a training and a testing set containing two thirds and one third of the data respectively, and simulation parameters for aLDA were the same as before. Similar to previous studies, we tested separately the sensors located on different leg segments (i.e., thigh or lower leg), as preliminary results show that little adaptation can be achieved for location changes between different limb segments.

In contrast to the previous scenario, in this case the performance of the aLDA and IWLDA classifiers does not significantly differ from the static LDA (c.f.

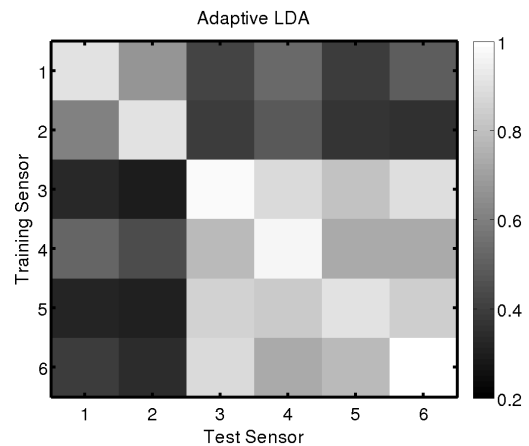


Fig. 3 HCI gesture scenario. Classification accuracy (encoded by grey levels) for each training-testing combination. Each row denotes the sensor used for training and each column represents the sensor used for testing the method.

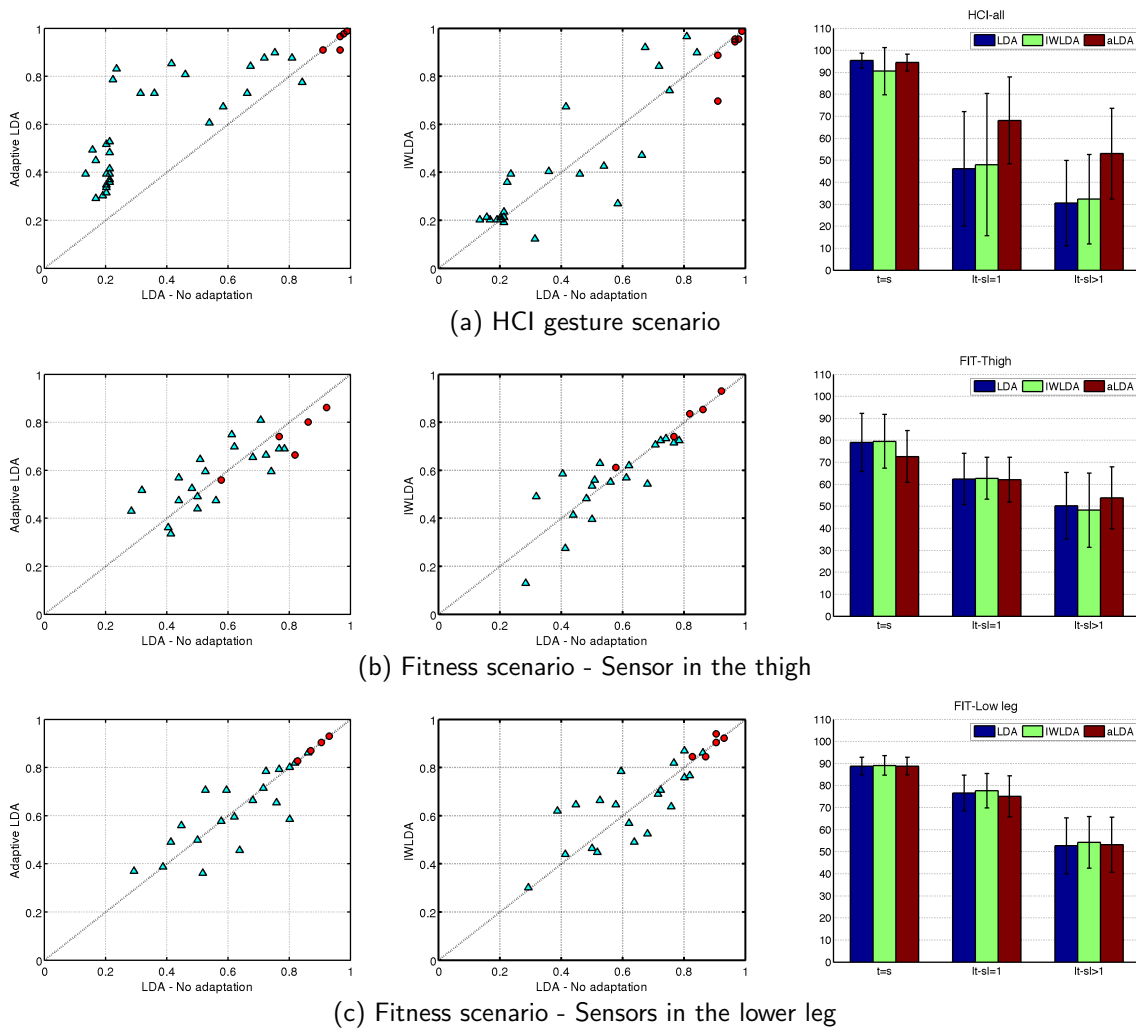


Fig. 2 Method performance upon sensor displacement - Classification accuracy. *Left*, aLDA. *Middle*, IWLDA. Each plot shows the accuracy of the adaptive classifier vs. the LDA classifier. Red circles show the cases when the classifier is tested at the same location it was previously trained. *Right*, Average performance for the three classifiers.

Figures 2(b) and 2(c). A performance increase is only observed when there is a large change in the sensor location ($|t - s| > 1$), especially for sensors located on the thigh. Indeed, the performance decrease of the static LDA classifier when tested in other locations is not as steep as in the HCI scenario. The average performance of the static LDA when testing in the closest sensor to the training one ($|t - s| = 1$) is about 62% and 76% for sensors on the thigh and lower leg respectively. Actually, the static LDA performs better than the previously proposed adaptive NCC for the sensors in the lower leg, suggesting that there is little room for performance improvement given the classifier characteristics [2].

3.2 Robustness to sensor rotation

3.2.1 HCI gesture scenario

We further tested the proposed method by emulating sensor rotation. Using the HCI gesture scenario, the sensor signals were artificially rotated in the range -90° to 90° . We report results for rotations around the x- and z- axis since (given the characteristics of the task, rotations around the y-axis had a small effect on the recognition performance). The data preprocessing and classifier training is performed as described in the previous section. During online adaptation the regularisation parameter λ was set to 0.005.

The performance of the fixed and adaptive classifier is shown in Figures 4 and 5 for rotations over the x- and z-axis, respectively. As expected, the performance of the LDA classifier drops drastically even after small sen-

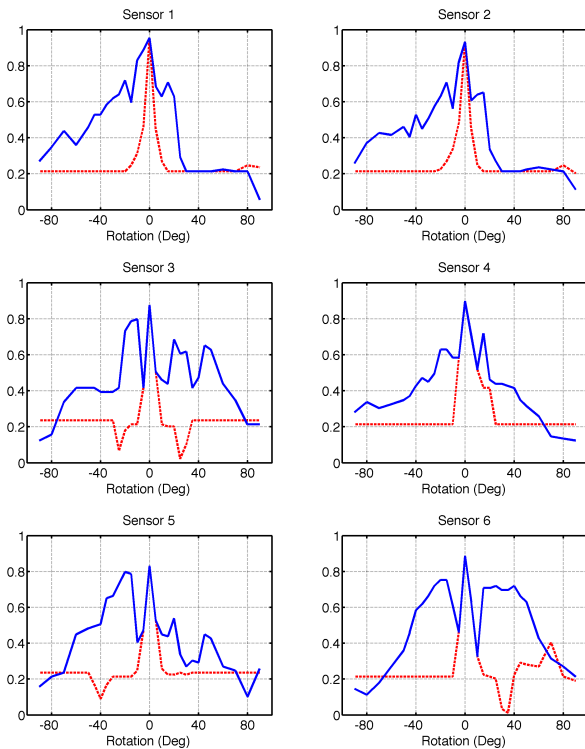


Fig. 4 HCI scenario. Rotations over the x-axis. Classification accuracy of the static LDA (dashed red line) and adaptive LDA (continuous blue line).

sensor rotations, especially in the case of rotations around the x-axis. The decrease in performance of the adaptive classifier is significantly smaller for all the tested sensors, thus yielding graceful degradation upon such type of sensor change.

3.2.2 Daily living scenario

A third dataset was used to test robustness against sensor rotation. The data is a subset of a larger recording performed in a rich-sensor environment [16, 20]. It corresponds to a daily living scenario where the subject performs a morning activities. During the recordings, each subject performed 5 times a run with activities of daily living (ADL) and one *drill* run. During the ADL run subjects freely performs the activities following a loose description of the overall actions to perform (i.e. wake up, make breakfast and take a walk), without precise instruction about more specific actions. During the drill runs, they performed 20 repetitions of a predefined sequence of activities including open doors and drawers, turn on/off the lights or drink.

We evaluate activity recognition using five acceleration sensors located on the back, right upper arm, left upper arm, right lower arm, and left lower arm as shown in Fig 1(c). Acceleration values for the three axis of

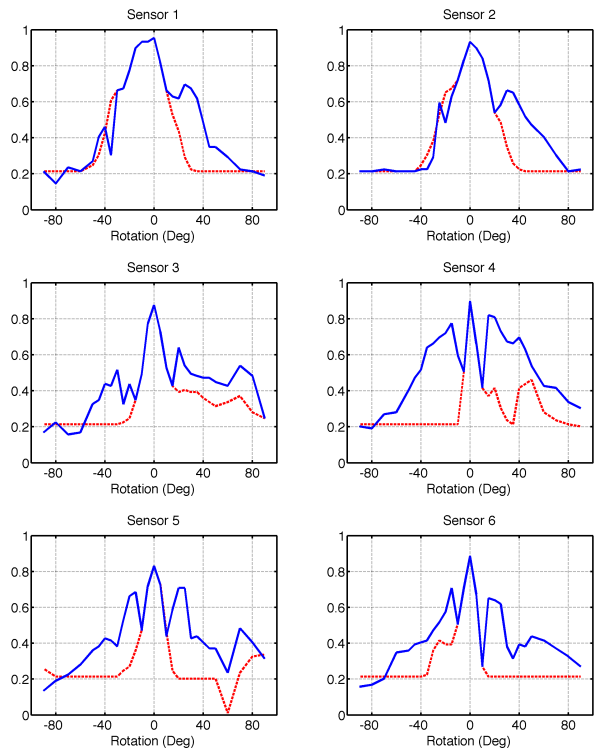


Fig. 5 HCI scenario. Rotations over the z-axis. Classification accuracy of the static LDA (dashed red line) and adaptive LDA (continuous blue line).

the five sensors were taken into account and features were again extracted using CVA. The projected features were feed either into the the LDA classifier or its adaptive version. As before, we emulate rotation of one sensor in the range $[-90^\circ, 90^\circ]$ in the y- and z- axis. We report results from 4 different subjects, using 3-fold cross-validation for each one of them.

We performed two sets of simulations using different sets of classes to be recognised. In the first one, termed *experiment A*, the set is composed of four activities: toggle switch, drink, clean table and close drawer. These classes exhibit a large discriminability that leads to high classification accuracy when there is no noise (0.81, 0.81, 0.79, and 0.84 for subjects 1 to 4, respectively). The second set, *experiment B*, was chosen to be more challenging and is composed of six activities (open door, open dishwasher, open drawer, clean table, drink and toggle switch). The classification accuracy for all subjects in this experiment, when no noise is added, is 0.63, 0.66, 0.66, and 0.64 for subjects 1 to 4, respectively.

Analysis of the features extracted using CVA shows that the most informative sensors are those located on the back and on the right arm. Indeed, in both experiments rotation of sensors in the left arm has no influence on the classification accuracy (results not shown).

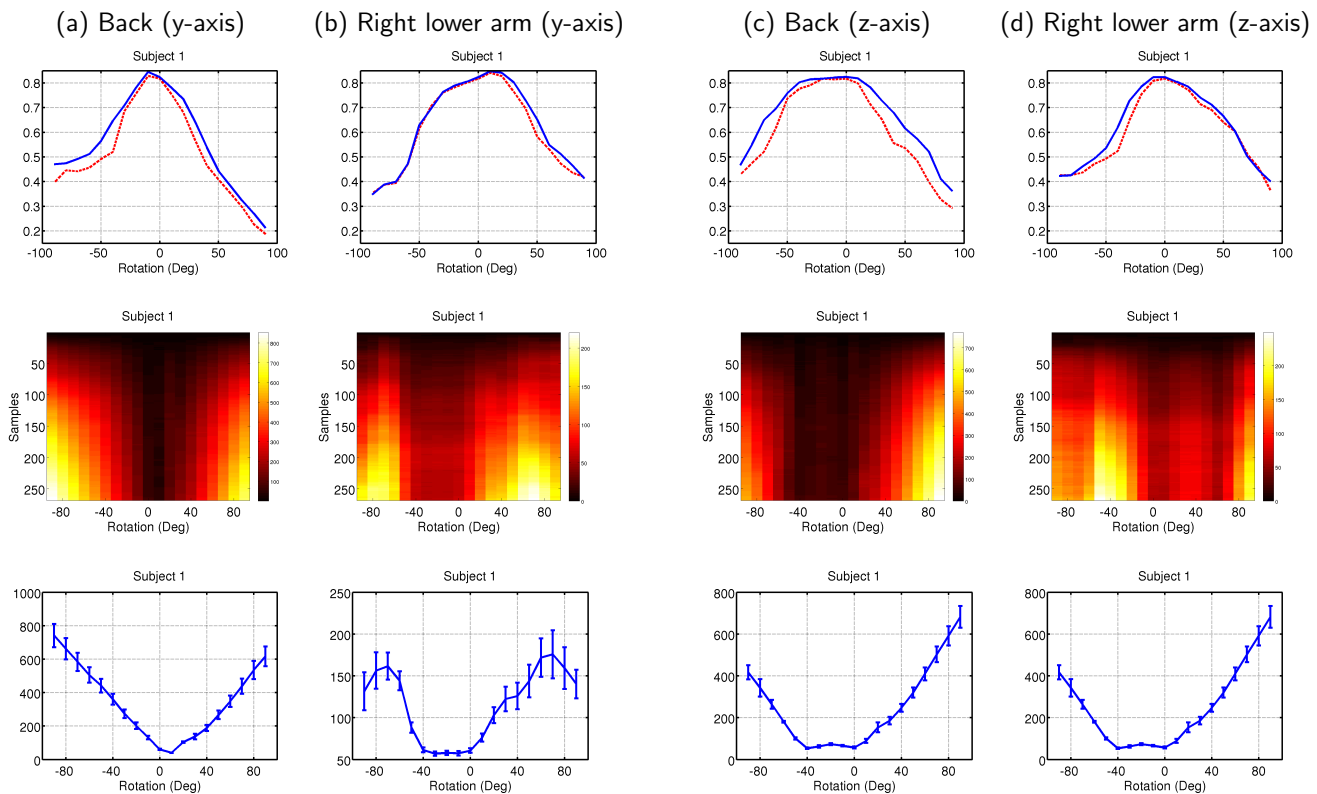


Fig. 6 Daily living scenario - Experiment A (4 classes). Performance on subject 1 when sensors on the back and right lower arm are rotated. *Top row*, Classification performance of the static LDA and adaptive LDA (dotted and continuous traces, respectively). *Middle row*, Evolution over time of the estimated shift. *Bottom row*, Final value of the estimated offset. Error bars show the standard deviation computed over the last 100 samples.

Figure 6 shows the performance for one representative subject after rotation of sensors on the back and the right lower arm for the experiment A. It can be seen that the performance of the LDA (dashed red trace) decreases after rotations of the sensor on the back, particularly around the y-axis. This decrease is less marked than in the previous database since all sensors are used for classification and only one of them is affected with noise. Once more, the adaptive process (continuous blue trace) outperforms the static approach for all subjects resulting in a more robust system against sensor rotation. Such effect is larger in the sensor located in the back than the one in the arm.

The figure also shows the evolution of the estimated shift as new samples are acquired (the values are color coded, where dark colors correspond to smaller values). It can be seen that for small sensor rotation the method quickly converges towards small values. In contrast, for large rotations ($> 40^\circ$) after 50 samples the value of the estimated shift starts to increase. We also show the final estimated shift (i.e., after 250 samples), as well as its standard deviation on the last 100 samples as a measure of the convergence of the adaptive mechanism.

A similar pattern was found in experiment B, as seen in Figure 7. In this case, sensors in the back and right upper arm are the most discriminative. Again, for small sensor rotations, the method is able to estimate the distribution shift and outperforms the static LDA classifier. Figures 8 and 9 show the performance increment, in terms of percentage of the LDA performance, for all subjects. It shows that the adaptive mechanism generally increases the performance of the original classifier in both experiments for all subjects.

4 Discussion

Deployment of activity recognition systems requires them to be able to cope with different factors that appear in real-life. One of those is the case of sensor displacement, especially for long-term running applications. Several approaches have been proposed to tackle this issue either by taking into account the type of change we want to be robust to, or by redundant training on the system using several body locations. However, the first approach can only address specific types of change, while

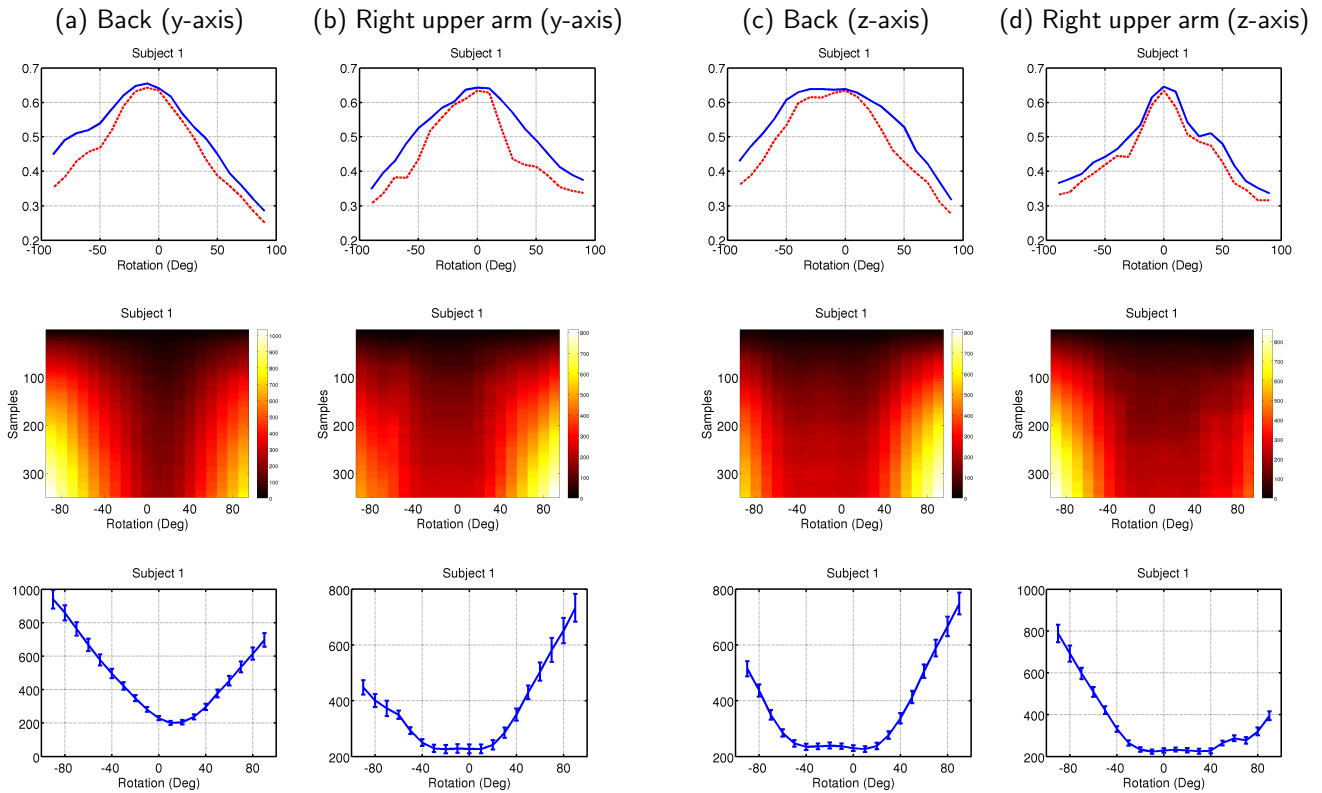


Fig. 7 Daily living scenario - Experiment B (6 classes). Performance on subject 1 when sensors on the back and right upper arm are rotated. *Top row*, Classification performance of the static LDA and adaptive LDA (dotted and continuous traces, respectively). *Middle row*, Evolution over time of the estimated shift. *Bottom row*, Final value of the estimated offset. Error bars show the standard deviation computed over the last 100 samples.

the second one imposes a significant overhead on the system design and calibration.

Alternatively, we propose an unsupervised adaptive mechanism that tracks changes in the feature distribution in an online manner. The proposed method extends probabilistic Gaussian classifiers assuming that changes in the sensor placement mainly result in a shift of the overall feature distribution. Given this assumption, unsupervised adaptation is achieved by estimating the feature distribution shift by means of an online version of expectation maximisation using the Levenberg-Marquardt algorithm.

Reckoning that such an assumption is unlikely to fully hold in real applications, we perform several simulations using activity recognition of realistic scenarios emulating both sensor translation and rotation. Experiments using body-worn accelerometers support the idea that this method is able to compensate for strong performance decrease without compromising the performance when the original classifier performs well (e.g., fitness scenario). We emulate sensor displacement using an experimental setup using sensors located at different positions of the upper and lower limbs, and test-

ing the classifier in a sensor located at a different position than the one used for training. Moreover, we assessed the method in the case of rotational changes in the sensor position by artificially rotating the recorded signals. This allows us to characterise its performance upon changes of different magnitude.

Regarding sensor displacement, we further compare with another adaptation technique, i.e., IWLDA. Our results show that aLDA performs as well as the IWLDA without requiring the availability of calibration data. Indeed, in the specific case of the HCI scenario, it significantly outperforms IWLDA. Furthermore, it should be taken into account that for the aLDA we report the testing performance while the adaptation process is taking place, thus providing an estimation of the online performance of the system. In previous work we have also shown that this method also outperforms another adaptive approach based on NCC classifiers [2].

In the HCI scenario the performance of the LDA classifier is strongly affected by sensor displacement. This effect is reduced by the adaptive mechanism (cf. Figure 2(a)). In particular, aLDA performance remains remarkably high for sensors located close to the wrist.

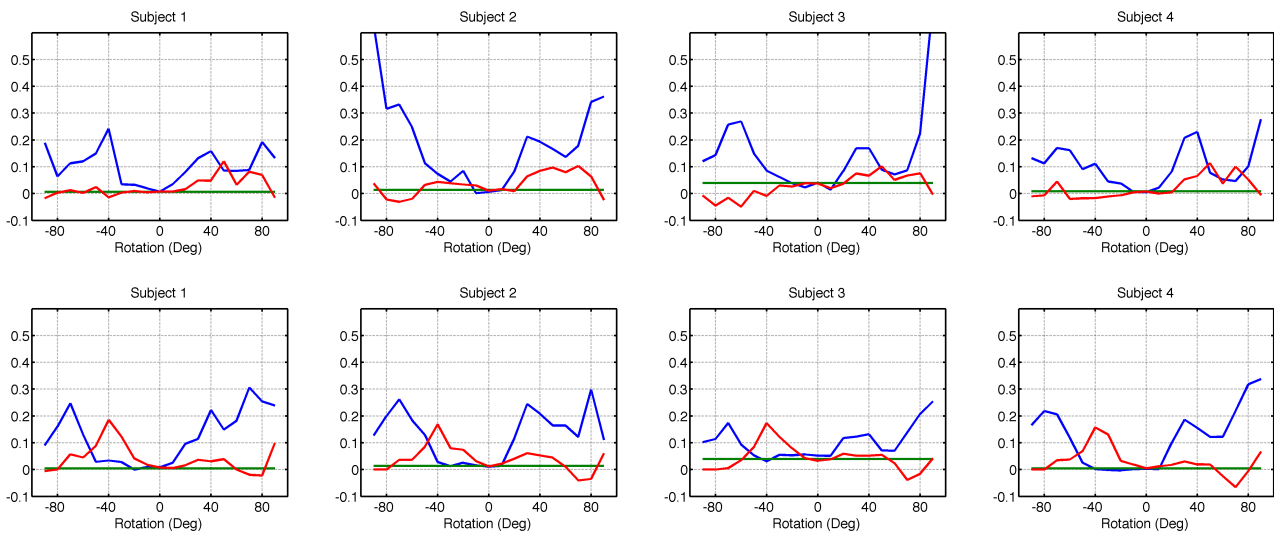


Fig. 8 Daily living scenario - Experiment A (4 classes). Performance increase (with respect to the LDA performance) when using aLDA upon rotation of signals of one sensor. Each line corresponds to the rotation of one of the three discriminative sensors (i.e., sensors on the back and right upper and lower arm). (Top) Rotation over the y-axis. (Bottom) Rotation over the z-axis.

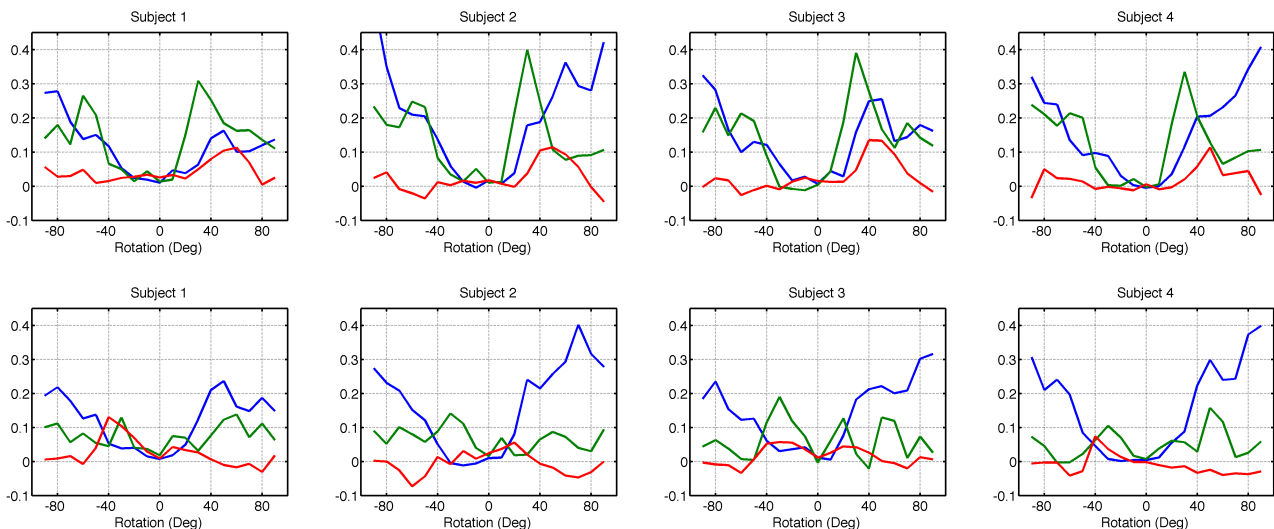


Fig. 9 Daily living scenario - Experiment B (6 classes). Performance increase (with respect to the LDA performance) when using aLDA upon rotation of signals of one sensors. Each line corresponds to the rotation of one of the discriminative sensors (i.e., sensors on the back and right upper and lower arm). (Top) Rotation over the y-axis. (Bottom) Rotation over the z-axis.

In contrast, the IWLDA is not able to capture the changes in the feature distribution despite the availability of the calibration process. In the fitness scenario, the adaptive mechanism does not have a significant impact, as the performance of the aLDA does not differ from the static classifier. This may be due to the fact that the LDA classifier already seems robust to small sensor displacements in this application thus leaving less opportunity for adaptation. A similar performance pattern was observed for the IWLDA, showing that our approach converges to the same estimation than the calibration process of this method.

Results suggest that the adaptive approach is more robust to sensor rotation than the fixed classifier (c.f. Section 3.2). In particular in the HCI scenario, the performance of the LDA classifier drops to chance level after rotation of about 15° , whereas the decrease of the adaptive approach is considerably smaller. In the daily living scenario, where several sensors are taken into account for classification, sensor rotation has a smaller impact on the performance of the LDA. However, even in these cases the aLDA consistently performs better than the static approach for the two sets of target classes that we presented.

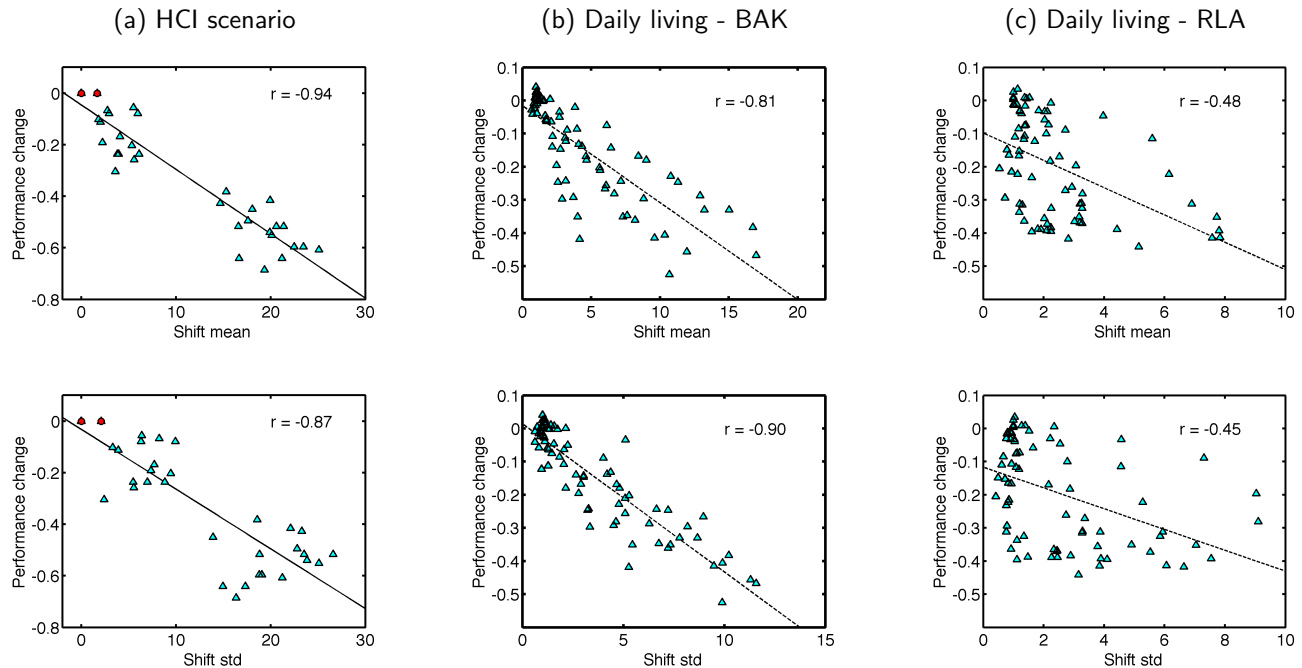


Fig. 10 Sensor displacement - Normalised change in performance (testing accuracy minus ac training accuracy) with respect to the mean (*Top*) and standard deviation (*Bottom*) of the estimated shift. Values are normalised with respect to the estimated values on the training set. *Left*, Gesture recognition scenario. *Middle*, Daily living scenario - Experiment A, rotations over z-axis. Sensor on the back. *Right*, Daily living scenario - Experiment A. Sensor on the right lower arm.

The rationale of the adaptation method is the estimation of changes in the feature distribution. Other approaches try to detect these changes in order to identify anomalous behaviour (e.g., sensor failure). They are mainly based on the characterisation of the feature distributions [4, 1, 17, 24] or the the classification output in classifier ensembles [23, 5]. In our case, the estimated shift provides a direct estimation of the changes in the feature distribution [2]. Such measure can be used to infer an online estimation of the system reliability, a critical point for systems that have to deal with dynamic changing environments [13, 21]. For example, if a sensor is considered non reliable (e.g., when the estimated shift Θ exceeds a given threshold), compensatory actions can be taken, such as its removal from a sensor network [13, 23, 22].

Figure 10 shows how the mean and standard deviation of the estimated shift correlates with the change in performance with respect to the original location, for both sensor displacement and rotation. In general larger estimated shifts correspond to a decrease in accuracy, meaning that it can provide information about the sensor reliability. However, in some cases the performance decreases though the estimated change is small, suggesting that in these cases the assumptions of the method are not satisfied.

Reported results suggest that despite the strong assumptions of the method, it is able to effectively capture changes in the feature distribution of the upcoming samples. Indeed, the presented approach results in graceful performance degradation upon sensor displacement. Furthermore, this is achieved in an unsupervised manner without requiring a calibration phase and using only two free parameters (λ and Θ). At this stage, it is limited by the types of changes it can effectively estimate, i.e., shifts on the feature distributions. Further work is being undertaken to extend it to also cope with other types of transformations (e.g., allowing for scaling and rotations). However, this may require iterative processes relying on a larger amount of data, and more free parameters that may compromise its application on run time applications. A trade-off then should be found between the performance increase that can be achieved and complexity of the method that should be used.

Acknowledgements We would like to thank K. Förster and D. Roggen from ETH Zurich, and H. Sagha from EPFL, Lausanne for providing the experimental data and insightful discussions. This work was supported by the EU-FET project ICT-225938 (*Opportunity: Activity and Context Recognition with Opportunistic Sensor Configuration*). This paper only reflects the authors' views and funding agencies are not liable

for any use that may be made of the information contained herein.

References

1. Alippi, C., Roveri, M.: An adaptive cusum-based test for signal change detection. In: ISCAS. IEEE (2006)
2. Bayati, H., Millán, J.d.R., Chavarriaga, R.: Unsupervised adaptation to on-body sensor displacement in acceleration-based activity recognition. In: IEEE International Symposium on Wearable Computers, ISWC (2011)
3. Bishop, C.M.: Pattern Recognition and Machine Learning. Springer (2007)
4. Chandola, V., Banerjee, A., Kumar, V.: Anomaly detection: A survey. ACM Comput. Surv. **41**(3), 1–58 (2009)
5. Chavarriaga, R., Sagha, H., Millán, J.d.R.: Ensemble creation and reconfiguration for activity recognition: An information theoretic approach. In: IEEE Int Conf Systems, Man, and Cybernetics (IEEE SMC) (2011)
6. Duda, R.O., Hart, P.E., Stork, D.G.: Pattern Classification, 2 edn. Wiley-Interscience (2000)
7. Förster, K., Brem, P., Roggen, D., Tröster, G.: Evolving discriminative features robust to sensor displacement for activity recognition in body area sensor networks. Int Conf Intelligent Sensors, Sensor Networks and Information Processing, ISSNIP pp. 43–48 (2009)
8. Förster, K., Roggen, D., Tröster, G.: Unsupervised classifier self-calibration through repeated context occurrences: Is there robustness against sensor displacement to gain? In: IEEE Int Symposium Wearable Computers (2009)
9. Hasanuzzaman, M., Zhang, T., Ampornaramveth, V., Gotoda, H., Shirai, Y., Ueno, H.: Adaptive visual gesture recognition for human-robot interaction using a knowledge-based software platform. Robotics and Autonomous Systems **55**(8), 643–657 (2007)
10. Kang, H., Lee, C.W., Jung, K.: Recognition-based gesture spotting in video games. Pattern Recognition Letters **25**(15), 1701–1714 (2004)
11. Krzanowski, W.J.: Principles of multivariate analysis. Oxford University Press, Oxford (1998)
12. Kunze, K., Lukowicz, P.: Dealing with sensor displacement in motion-based onbody activity recognition systems. In: Proc int conf on Ubiquitous computing, pp. 20–29. ACM, New York, NY, USA (2008)
13. Kurz, M., Hölzl, G., Ferscha, A., Calatroni, A., Roggen, D., Tröster, G., Sagha, H., Chavarriaga, R., del R. Millán, J., Bannach, D., Kunze, K., Lukowicz, P.: The opportunity framework and data processing ecosystem for opportunistic activity and context recognition. International Journal of Sensors, Wireless Communications and Control (2011)
14. Lester, J., Choudhury, T., Borriello, G.: A practical approach to recognizing physical activities. In: Lecture Notes in Computer Science : Pervasive Computing (2006)
15. Li, Y., Kambara, H., Koike, Y., Sugiyama, M.: Application of covariate shift adaptation techniques in Brain Computer Interface. IEEE Trans Biomedical Engineering **57**(6), 1318–1324 (2010)
16. Lukowicz, P., Pirkel, G., Bannach, D., Wagner, F., Calatroni, A., Förster, K., Holleczeck, T., Rossi, M., Roggen, D., Tröster, G., Doppler, J., Holzmann, C., Riener, A., Ferscha, A., Chavarriaga, R.: Recording a complex, multi modal activity data set for context recognition. In: 1st Workshop on Context-Systems Design, Evaluation and Optimisation at ARCS (2010)
17. Manly, B.F.J., Mackenzie, D.: A cumulative sum type of method for environmental monitoring. Environmetrics **11**(2), 151–166 (2000)
18. Marquardt, D.W.: An algorithm for least-squares estimation of nonlinear parameters. Journal of the Society for Industrial and Applied Mathematics **11**(2), pp. 431–441 (1963)
19. Ravi, N., D, N., Mysore, P., Littman, M.L.: Activity recognition from accelerometer data. Proc 17th conf Innovative applications of artificial intelligence **3** (2005)
20. Roggen, D., Calatroni, A., Rossi, M., Holleczeck, T., Förster, K., Tröster, G., Lukowicz, P., Bannach, D., Pirkel, G., Ferscha, A., Doppler, J., Holzmann, C., Kurz, M., Holl, G., Chavarriaga, R., Sagha, H., Bayati, H., Creatura, M., Millán, J.D.R.: Collecting complex activity data sets in highly rich networked sensor environments. In: Seventh International Conference on Networked Sensing Systems (2010)
21. Roggen, D., Förster, K., Calatroni, A., Holleczeck, T., Fang, Y., Tröster, G., Lukowicz, P., Pirkel, G., Bannach, D., Kunze, K., Ferscha, A., Holzmann, C., Riener, A., Chavarriaga, R., Millán, J.: OPPORTUNITY: Towards opportunistic activity and context recognition systems. In: Third IEEE WoWMoM Workshop on Autonomic and Opportunistic Communications (2009)
22. Sagha, H., Millán, J.d.R., Chavarriaga, R.: Detecting and rectifying anomalies in opportunistic sensor networks. In: International Conference on Body Sensor Networks, BSN. Dallas, Texas, USA (2011)
23. Sagha, H., Millán, J.d.R., Chavarriaga, R.: Detecting anomalies to improve classification performance in an opportunistic sensor network. In: IEEE Workshop on Sensor Networks and Systems for Pervasive Computing, PerSens. Seattle (2011)
24. Smyth, P.: Hidden Markov models for fault detection in dynamic systems. Pattern Recognition **27**(1), 149–164 (1994)
25. Steinhoff, U., Schiele, B.: Dead reckoning from the pocket—an experimental study. In: IEEE Int Conf Pervasive Computing and Communications (PerCom 2010). Mannheim, Germany (2010)
26. Stiefmeier, T., Roggen, D., Tröster, G., Ogris, G., Lukowicz, P.: Wearable activity tracking in car manufacturing. IEEE Pervasive Computing **7**(2), 42–50 (2008)
27. Sugiyama, M., Krauledat, M., Müller, K.R.: Covariate shift adaptation by importance weighted cross validation. J. Mach. Learn. Res. (2000)
28. Sugiyama, M., Suzuki, T., Nakajima, S., Kashima, H., von Bnau, P., Kawanabe, M.: Direct importance estimation for covariate shift adaptation. Annals of the Institute of Statistical Mathematics **60**(4), 699–746 (2008-12-01)
29. Tentori, M., Favela, J.: Activity-aware computing for healthcare. IEEE Pervasive Computing **7**(2), 51–57 (2008)
30. Van Laerhoven, K., Cakmakci, O.: What shall we teach our pants? IEEE Int Symposium on Wearable Computers, IWSC pp. 77–83 (2000)
31. Westeyn, T., Vadas, K., Bian, X., Starner, T., Abowd, G.: Recognizing mimicked autistic self-stimulatory behaviors using HMMs. Wearable Computers Proceedings. Ninth IEEE International Symposium on pp. 164–167 (2005)
32. Wilson, A.D., Bobick, A.F.: Realtime online adaptive gesture recognition. Pattern Recognition, International Conference on **1**, 1270 (2000)