Recognizing Social Gestures with a Wrist-Worn SmartBand

Jonathan Knighten, Stephen McMillan, Tori Chambers, and Jamie Payton
Department of Computer Science
University of North Carolina at Charlotte
Email: {jknigh28, smcmil11, tchamb15, payton}@uncc.edu

***Abstract*—The ability to recognize social gestures opens the door for the development of enhanced pervasive computing applications that are responsive to users' social interactions. In this paper, we explore the feasibility of using a smartband for social gesture recognition. We apply logistic regression, a supervised machine learning technique, to accelerometer data collected in a study of 32 users performing 12 social gestures. Our experimental results show promise for recognizing social gestures with a smartband; our simple approach achieves an average accuracy of 86% for classification of social gestures.**

I. Introduction

Smartbands offer new opportunities for pervasive computing applications that are responsive to the perceived state of the user and the surrounding environment. These commodity mobile devices support multiple forms of network connectivity (e.g, Bluetooth, WiFi) that can be exploited to share information and feature an array of on-board sensors that can be used to detect the current context. For example, smartbands currently on the market feature a barometer, heart rate monitor, light sensor, UV sensor, gyroscope, magnometer, and an accelerometer[1][2].

The use of accelerometer data collected by a smartband is particularly useful for identifying actions as they are performed by a user. Activity recognition (AR) is an important research topic in pervasive computing, and has been widely studied [1], [2], [3], [4], [5], [6]. Applications that rely on activity recognition have been developed for a wide range of purposes from judging rock-climbing competitions [7], to understanding canine behaviors [8], to addresssing freezing of gait episodes for people with Parkinson's disease [9]. However, many of the previous approaches have required multiple accelerometers per user [10] or were applied to broad postures and motions (e.g., sitting, standing, walking, going up/down stairs, running) [4], [6]. In this paper, we explore the feasibility of using a single smart band for recognition of fine-grained social gestures. The ability to recognize social gestures opens opportunities for new pervasive computing applications that can help to combat social isolation for vulnerable populations, act as assistive technology to provide the blind with information about non-verbal conversation cues, and provide support for the study of social interactions. We apply a supervised machine learning approach to smartband accelerometer data collected in a study of 32 participants performing 12 social gestures. Using logistic regression, we achieve an average classification accuracy of 86% across our gesture set.

[1]Samsung Gear sensor specs: http://www.samsung.com/global/microsite/gear/

[2]LG G Watch R sensor specs: http://www.lg.com/us/smart-watches/lg-W110-g-watch-r

II. Social Gesture Recognition Use Cases

We believe that creating solutions for fine-grained social gesture recognition on smartbands can enhance existing social networking applications (e.g., Facebook, LinkedIn) and can open the door for the creation of new applications that can have significant benefits for society. Below, we discuss two potential applications for smartband-based social gesture recognition.

A. Promoting Social Connectedness

Social isolation is a significant risk factor for stroke, myocardial infarction, and chronic diseases, including cancer and diabetes [11]. Those who perceive themselves as being socially isolated also report higher levels of anxiety, negative mood, and hostility and lower levels of happiness and life satisfaction [12]. A wearable computing application to promote social connectedness can serve as a mechanism to combat the effects of social isolation.

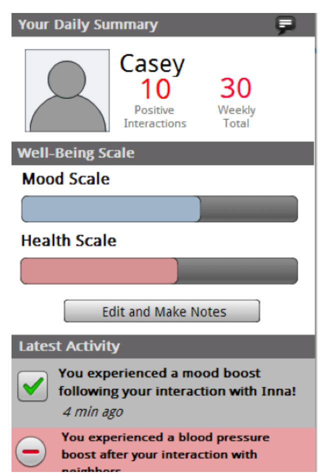

Fig. 1: Smartwatch app to promote social interaction

A prototype of a smartwatch application that we have begun to develop to encourage social interactions is shown in Figure 1. Such applications can be designed to detect the performance of social gestures, comparing data collected from the accelerometer on the smartwatch to a learned model of social gestures. Gamification can provide an incentive for use, with the application awarding points, badges, and rewards for social interactions. Users can also provide self-reported measures of mood and perceived health, review the quantity and kinds of interactions that they have performed, and explore trends in their mood and health reports. Such applications could also be designed to allow users to download archives of their activities that they could then share with health professionals.

978-1-4799-8425-1/15/$31.00 ©2015 IEEE

B. Assistive Technology for the Blind and Visually Impaired

When participating in social interactions, people who are blind or visually impaired are at a disadvantage in perceiving the physical movements of others as part of non-verbal communication. Assistive technologies have been proposed to support the blind in detecting and interpreting non-verbal communication, many of which use a camera and microphone to localize speakers in a conversation and require special-purpose haptic equipment to give vibrotactile feedback to users (e.g., the Social Interaction Assistant [13]). As more people begin to adopt smartbands, it may be possible to also use a social gesture recognition application deployed on smartbands as a complementary assistive technology. We can imagine a solution in which an application deployed on a smartband uses a microphone to detect the start of a conversation and uses the microphone, local network connections, and accelerometer data in a group membership algorithm, similar to the SocialWeaver application for smartphones [14]. On each group member's smartband, the local accelerometer data would be used to detect non-verbal conversation cues; these cues are communicated via a local communication link (e.g., Bluetooth) to a blind participant in the conversation. The blind person can be alerted either through separate haptic devices or through the smartband for various conversation cues.

III. Classifying Social Gestures

Our ultimate goal is to deploy a system for social gesture recognition on a smartwatch that can support context-aware applications. As such, in selecting a machine learning algorithm for recognizing social gestures, we focus on linear classification approaches. Labeled training data is given to a linear classifer as a feature vector (i.e., a collection of characteristics of the data that capture some property of the phenomenon under observation), and linear classification schemes learn a model represented as a weighted linear combination of those features. Abstractly, linear classifiers define a hyperplane that "splits" an input space into two spaces; simple extensions can support classification for more than one class. Since linear classification is performed by applying a simple function f that maps the dot product of the weight and feature vectors to a value (i.e., $f(\vec{w} \cdot \vec{v}) \rightarrow \mathbb{R}$), classification decisions can be made quickly.

Initially, we considered three candidate linear classifiers: linear SVM, Linear Discriminant Analysis (LDA), and logistic regression. Linear SVM is a non-probabilistic classification model, which seeks a hyperplane that maximizes the margin of separation between two classes. A set of points is collected to represent the margin between classes; these "support vectors" are then used to make classifications decisions. In LDA, each class is assumed to be normally distributed with equal covariance matrices. To classify a new point, a weighted distance between each class mean is calculated to determine which class is closest to the point. Logistic regression is a probabilistic model for classification. The key idea is to transform the problem of binary classification to a continuous criterion and then apply linear regression; this transformation is applied by taking the natual log of the odds that a particular input belongs to a given class.

Due to our desire to implement our system on a wearable device, we desire a classifier that has a small memory footprint. For this reason, linear SVM was ruled out because its memory requirements depend on the number of support vectors required. In practice, logistic regression and LDA perform similarly [15]. However, LDA makes assumptions about the normality and inner-class variance [15]; the data collected in our study of 32 users performing 12 gestures was non-normal and has non-equal class variance. Therefore, we select logistic regression as our primary approach.

A. Data Collection

The platform chosen for recording social gestures was a single, AX3 Watch manufactured by Axivity. The AX3 Watch contains a tri-axial MEMS accelerometer and 512MB of on board flash for data storage, and associated tools allows developers to control watch parameters such as sampling rate and sensor gravity threshold. For every experiment the accelerometer was sampled at 50Hz with a $\pm$4G gravity threshold [16].

A total of 32 volunteers participated in our data collection procedure. The population demographics of participants is reflective of the population demographics of the college in which our research lab is housed. Out of the 32 participants, 28 were male and 4 were female. Ages of the participants ranged from 18 to 63 years old; the average age was 25. Each participant was asked to perform a series of social gestures while wearing the AX3 device accelerometer on the right-hand wrist, oriented as illustrated in Figure 1.

Specifically, each user was asked to perform the gestures in Table I.

ID:	Gesture:	Description:
1	Celebratory Fist Pump	Users perform a fist pump gesture equated with excitement or celebration
2	High Wave	Users wave their right hand above their head
3	Hand Shake	Users shake hands with the research coordinator
4	Fist Bump	Users bump a closed fist with the research coordinator
5	Low Wave	Users perform a wave located close to the waist area
6	Point Straight	Users point straight ahead
7	Point Left	Users point to their left
8	Point Right	Users point to their right
9	Point Up	Users point up towards the ceiling
10	Motion Over	Users perform a gesture indicating come closer
11	High-Five	Users perform a high five with the research coordinator
12	Applause-Clap	Users clap as one would after a performance

TABLE I: User Study Gesture Descriptions

Each volunteer performed each gesture three times; each gesture performance was separated by a brief pause with hands in resting position.

B. Preprocessing: Windowing Data

A major challenge of activity recognition is extracting useful features from sensor data to be used in machine learning algorithms. Different types of sensor modalities have been used in activity recognition, but body-worn accelerometers are used often [7], [8], [18], [4], [5], [19], [20], [21], [22]. Sensor data from most modalities can be considered as a signal or a time series; for the case of the accelerometer, three signals are generated for each axis. To handle such data it must be

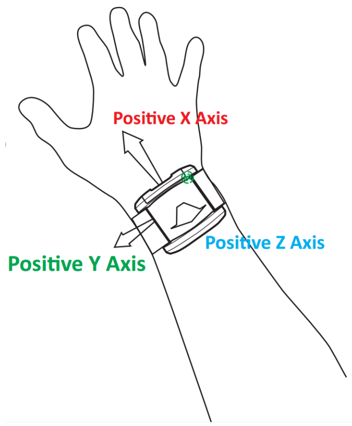

Fig. 2: Axivity AX3 Right Handed Axis Alignment [17]

discretized into sliding windows/frames, from which features are extracted; most modern techniques use overlapping sliding rectangular windows [7], [8], [18], [4], [5], while others use other techniques such as hamming windows [19], [7].

Table II gives the mean duration for each gesture type.

Gesture	Mean Duration 95%CI	Instances
1. Celebratory Fist Pump	116.064±6.973	297
2. High Wave	150.174±11.499	304
3. Hand Shake	118.961±7.627	306
4. Fist Bump	96.403±5.818	295
5. Low Wave	123.670±9.347	303
6. Point Straight	122.840±7.845	300
7. Point Left	132.468±10.188	295
8. Point Right	119.172±8.346	290
9. Point Up	124.107±9.392	300
10. Motion Over	122.911±9.704	303
11. High Five	107.617±6.179	290
12. Applause/Clap	130.503± 11.347	306

TABLE II: Gesture Durations in Number of Data Points

Given the average duration of the gestures in Table II, we applied a window size of 128 points to our data set, applying padding and trimming where necessary. Since we apply a supervised machine learning algorithm to recognize gestures, the windowed accelerometer data collected from the user study was annotated with labels that indicate the gesture being performed. Each accelerometer data file was manually annotated by a researcher with the starting point, end points, and a label that describes the type of gesture performed.

C. Preprocessing: Feature Extraction

A major decision that must be made for feature extraction is the domain from which features will be pulled from. The raw data generated from a sensor is in the time domain, and present the relationship of sensor level and time. A basic approach is to keep the data in the time domain, and extract statistical features from each window [3], [18], [23]. One advantage of this method is that it avoids the overhead of applying a transform such as the Discrete Fourier Transform (DFT), which would be desired on devices with low computational resources like the smartbands used in [3]. The key advantage of the frequency domain is that it captures periodic patterns in data that the time domain may not show. Activities such as waving, walking, and running, have a periodic trend or rhythm that is more apparent in the frequency domain than in the time domain. More recent works [18], [23] have evaluated the strength of using different preprocessing techniques and feature representations, such as convolving data with their respective empirical cumulative distribution functions (ECDF) [23], a deep learning approach using Restricted Boltzmann Machines (RBM) [18], and Principle Component Analysis (PCA) [18]. However, these more sophisticated approaches come with additional data processing costs and in a preliminary evaluation over our particular data set, offered little benefit.

We extract two sets of features, one related to the time domain and the other related to the frequency domain. The features were separated into these two groups due to the frequency domain features requiring a DFT, which is additional computational overhead. For each 128 point frame of data, each of the features in Table II were calculated for each of the x, y and z axis.

Feature:	Feature Set:
Min	Time Domain
Max	Time Domain
Mean	Time Domain
Standard Deviation	Time Domain
Pairwise Correlations	Time Domain
Zero Crossing Rate	Time Domain
Mean Crossing Rate	Time Domain
Skewness	Time Domain
Kurtoisis	Time Domain
Area Under Curve	Time Domain
Signal to Noise Ratio	Time Domain
Signal Energy	Frequency Domain
First 8 DFT Coef	Frequency Domain

TABLE III: Features Extracted

In total, we apply our classification approach to 60 features, 33 features of which are in the time domain and 27 in the frequency domain.

D. Classification with Logistic Regression

Given N training observations $\{(x^{(i)}, y^{(i)}), i = 1, 2, \ldots, N\}$, where $x^{(i)} \in \mathbb{R}^p$ and $y^{(i)} \in \{0, 1\}$, logistic regression operates by estimating the weights of a linear equation that expresses the likelihood of belonging to one of two classes [15]:

$$log\frac{p(y^{(i)}=1|x^{(i)})}{p(y^{(i)}=0|x^{(i)})} = \beta_0 + \beta_1 x_1^{(i)} + \cdots + \beta_p x_p^{(i)} = \vec{\beta}^T \begin{bmatrix} 1 \\ x_1^{(i)} \\ x_2^{(i)} \\ \vdots \\ x_p^{(i)} \end{bmatrix}$$

In our implementation, we use scikit-learn 0.15 [24] logistic regression in a one-versus-all scheme, which builds a binary classification model for each class [25]. This means for each class, we must estimate $\vec{\beta}$, the estimated weights for each class's logistic regression linear model. To find $\vec{\beta}$, maximum likelihood estimates (MLE) are used. Finding the estimates is an optimization problem [15], [26]:

$$\min_{\vec{\beta}} \sum_{i=1}^{N} -log(p(y^{(i)}|\vec{x}^{(i)};\vec{\beta}))$$

To prevent over-fitting we use LASSO, which introduces the L_1-norm regularization term to the MLE optimization problem [27], [26], [15]:

$$\min_{\vec{\beta}} \sum_{i=1}^{N} -log(p(y^{(i)}|\vec{x}^{(i)};\vec{\beta})) + \lambda \left\|\vec{\beta}\right\|_1,$$

$$where \left\|\vec{\beta}\right\|_1 = \sum_{i=1}^{p} |\beta_i|$$

A data item is classified by identifying the logistic regression model that maximizes the log odds.

IV. Results

To analyze the results from our created models we employed leave one subject out (LOSO) cross validation [28]. The advantage of this method is it shows how well a model generalizes to a new unseen user.

Feature Set:	Precision:	Recall:	F1-Score:	Accuracy:
Time Domain	0.84	0.84	0.84	0.836876
Frequency Domain	0.74	0.74	0.74	0.742479
All Features	0.87	0.86	0.87	0.864753

TABLE IV: Classification of Social Gestures with Logistic Regression

Table IV highlights the performance of the time domain features in isolation, the frequency domain features in isolation, and their combination.These results illustrate that reasonably accurate gesture classification is possible without applying a computationally intensive transformation to each window of data; reducing overhead associated with pre-processing of data is desirable for resource-constrained devices like smartbands.

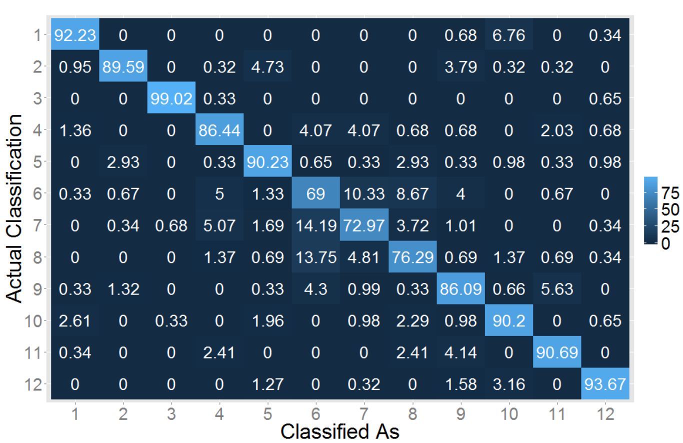

Fig. 3: Confusion Matrix for Logistic Regression Applied to 12 Social Gestures. Entries are percentage of examples that are correctly classified.

Precision, recall, and F1 results for each gesture are reported in Table V. The confusion matrix in Figure 3 highlights where misclassications occur. The most significant amount of misclassifications happen between the four directional pointing gestures (gestures 6-10). For many social gesture recognition applications, knowing the specific direction a user is pointing may not be important. Therefore, we also show results for a single pointing gesture in Figure 4.

Fig. 4: Confusion Matrix for Logistic Regression Applied to 9 Social Gestures. Directional pointing gestures are considered as a single general pointing gesture in this data set. Entries are percentage of examples that are correctly classified.

With a general point gesture, the model's accuracy increased to 92% and its overal precision, recall, and F1-score also increased to 92%.

	Separate Point Classes			*Single Point Class*		
Gesture:	**Prec.:**	**Recall:**	**F1**	**Precision:**	**Recall:**	**F1:**
1	0.94	0.92	0.93	0.94	0.94	0.94
2	0.95	0.90	0.92	0.94	0.90	0.92
3	0.99	0.99	0.99	0.99	0.99	0.99
4	0.85	0.86	0.86	0.93	0.90	0.91
5	0.88	0.90	0.89	0.88	0.91	0.90
6	0.66	0.69	0.67	0.84	0.88	0.86
7	0.77	0.73	0.75			
8	0.78	0.76	0.77			
9	0.83	0.86	0.84			
10	0.87	0.90	0.89	0.90	0.92	0.91
11	0.90	0.91	0.90	0.94	0.93	0.93
12	0.96	0.94	0.95	0.96	0.95	0.95

TABLE V: Precision, Recall, and F1 for Social Gesture Recognition with Logistic Regression

V. Related Work

Previous work has made strides in recognizing social interactions using sensors embedded in mobile phones. Two examples are particularly noteworthy. Cenceme [29] detects social interactions (e.g., having a conversation or dancing with friends) using a split-level classification approach, in which classification "primitives" are computed on the mobile phone and shipped to a server, where they are used in higher-level classification tasks. EmotionSense [30] lowers the barrier for the creation of tools that support collection of sensor data from mobile phones for the purpose of studying human interaction, allowing social scientists and psychologists to provide declarative specifications of sensing tasks for detecting emotion and properties of conversations. Ultimately, we aim to provide similar support for a framework that simplifies the development of pervasive computing applications that incorporate social gesture recognition; the work in this paper represents a small step towards that goal.

Perhaps most closely related to our effort is the use of a single accelerometer to classify social actions, such as drinking, stepping, laughing, gesturing, and speaking [31]. Although similar to our goals, their approach uses a stand-alone accelerometer worn around the neck and, more importantly, does not discriminate between different gestures. In contrast, our approach can potentially be deployed on a commodity smartwatch and discriminates between more fine-grained social gestures with relatively high precision and recall.

VI. Conclusions and Future Work

In this paper, we demonstrated the feasibility of applying machine learning algorithms to accelerometer data from a smartband for social gesture recognition. Applying logistic regression to a collection of time domain and frequency domain features, we were able to achieve an average accuracy of 92% across all gestures. These results indicate potential for future development of a smartband-based applications that are responsive to the performace of social gestures. However, we acknowledge that these results may be limited in their application in real-world settings due to the nature of our data collection method. Given the challenges in capturing the ground truth (i.e., gesture labels) associated with accelerometer data for gestures that are performed by users "in the wild", this initial study instead relied on a labratory study in which users performed a scripted set of gestures. As a result, our data set fails to capture some of the movement dynamics that would be present in a real-world setting. In the future, we plan to evaluate the application of supervised machine learning techniques for social gesture recognition over data sets collected from users in more natural settings. We also plan to deploy a classification system in its entirety on the smartband; we will evaluate its feasiblity and utility for supporting pervasive computing applications through an empirical evaluation of smart health and assistive applications that encourage and support social interactions.

Acknowledgments

We gratefully acknowledge Ashton Maltie for his initial efforts in collecting data and conducting preliminary analysis for this project. This work was partially supported by NSF under grant CNS-1156822. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the sponsor.

References

[1] N. Ravi, N. Dandekar, P. Mysore, and M. L. Littman, "Activity recognition from accelerometer data," in *AAAI*, vol. 5, 2005, pp. 1541–1546.

[2] L. Bao and S. S. Intille, "Activity recognition from user-annotated acceleration data," in *Pervasive Computing*, ser. Lecture Notes in Computer Science, A. Ferscha and F. Mattern, Eds. Springer Berlin Heidelberg, 2004, vol. 3001, pp. 1–17.

[3] U. Maurer, A. Smailagic, D. Siewiorek, and M. Deisher, "Activity recognition and monitoring using multiple sensors on different body positions," in *International Workshop on Wearable and Implantable Body Sensor Networks*, 2006.

[4] J.-g. Park, A. Patel, D. Curtis, S. Teller, and J. Ledlie, "Online pose classification and walking speed estimation using handheld devices," in *Proceedings of the ACM Conference on Ubiquitous Computing*, ser. UbiComp, 2012, pp. 113–122.

[5] H.-T. Cheng, M. Griss, P. Davis, J. Li, and D. You, "Towards zero-shot learning for human activity recognition using semantic attribute sequence model," in *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp)*, 2013, pp. 355–358.

[6] A. M. Khan, Y.-K. Lee, S. Y. Lee, and T.-S. Kim, "A triaxial accelerometer-based physical-activity recognition via augmented-signal features and a hierarchical recognizer," *IEEE Transactions on Information Technology in Biomedicine*, vol. 14, no. 5, pp. 1166–1172, 2010.

[7] C. Ladha, N. Y. Hammerla, P. Olivier, and T. Plötz, "ClimbAX: Skill assessment for climbing enthusiasts," in *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing*, ser. UbiComp, 2013, pp. 235–244.

[8] C. Ladha, N. Hammerla, E. Hughes, P. Olivier, and T. Ploetz, "Dog's life: Wearable activity recognition for dogs," in *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing*, ser. UbiComp '13, 2013, pp. 415–418.

[9] S. Mazilu, U. Blanke, M. Hardegger, G. Tröster, E. Gazit, and J. M. Hausdorff, "GaitAssist: a daily-life support and training system for parkinson's disease patients with freezing of gait," in *Proceedings of the ACM Conference on Human Factors in Computing Systems*, 2014, pp. 2531–2540.

[10] U. Maurer, A. Smailagic, D. P. Siewiorek, and M. Deisher, "Activity recognition and monitoring using multiple sensors on different body positions," in *International Workshop on Wearable and Implantable Body Sensor Networks*, 2006.

[11] J. S. House, K. R. Landis, D. Umberson *et al.*, "Social relationships and health," *Science*, vol. 241, no. 4865, pp. 540–545, 1988.

[12] J. T. Cacioppo, J. M. Ernst, M. H. Burleson, M. K. McClintock, W. B. Malarkey, L. C. Hawkley, R. B. Kowalewski, A. Paulsen, J. A. Hobson, K. Hugdahl *et al.*, "Lonely traits and concomitant physiological processes: the macarthur social neuroscience studies," *International Journal of Psychophysiology*, vol. 35, no. 2, pp. 143–154, 2000.

[13] T. McDaniel, S. Krishna, V. Balasubramanian, D. Colbry, and S. Panchanathan, "Using a haptic belt to convey non-verbal communication cues during social interactions to individuals who are blind," in *IEEE International Workshop on Haptic Audio visual Environments and Games*. IEEE, 2008, pp. 13–18.

[14] C. Luo and M. C. Chan, "Socialweaver: Collaborative inference of human conversation networks using smartphones," in *Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems*. ACM, 2013, p. 20.

[15] T. J. Hastie, R. J. Tibshirani, and J. H. Friedman, *The elements of statistical learning : data mining, inference, and prediction*, ser. Springer series in statistics, 2009, autres impressions : 2011 (corr.), 2013 (7e corr.).

[16] "Axivity," http://axivity.com/, accessed: 2014-05-22.

[17] *AX3 User Guide*, http://axivity.com/v2/products/AX3/AX3-User-Guide-v1-3.pdf, Axivity.

[18] T. Plötz, N. Y. Hammerla, and P. Olivier, "Feature learning for activity recognition in ubiquitous computing," in *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence*, ser. IJCAI'11, 2011, pp. 1729–1734.

[19] H. Hung, G. Englebienne, and J. Kools, "Classifying social actions with a single accelerometer," in *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing*, ser. UbiComp, 2013, pp. 207–210.

[20] A. Parate, M.-C. Chiu, C. Chadowitz, D. Ganesan, and E. Kalogerakis, "Risq: Recognizing smoking gestures with inertial sensors on a wristband," in *Proceedings of the 12th Annual International Conference on Mobile Systems, Applications, and Services*, ser. MobiSys '14, New York, NY, USA, 2014, pp. 149–161.

[21] A. Reiss, D. Stricker, and G. Hendeby, "Towards robust activity recognition for everyday life: Methods and evaluation," in *Proceedings of the International Conference on Pervasive Computing Technologies for Healthcare (PervasiveHealth)*, 2013, pp. 25–32.

[22] H.-T. Cheng, F.-T. Sun, M. Griss, P. Davis, J. Li, and D. You, "Nu-activ: Recognizing unseen new activities using semantic attribute-based learning," in *Proceeding of the 11th Annual International Conference on Mobile Systems, Applications, and Services*, ser. MobiSys '13, 2013, pp. 361–374.

[23] N. Y. Hammerla, R. Kirkham, P. Andras, and T. Ploetz, "On preserving statistical characteristics of accelerometry data using their empirical cumulative distribution," in *Proceedings of the 2013 International Symposium on Wearable Computers*, ser. ISWC, 2013, pp. 65–68.

[24] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.

[25] R. Rifkin and A. Klautau, "In defense of one-vs-all classification," *J. Mach. Learn. Res.*, vol. 5, pp. 101–141, Dec. 2004.

[26] S.-I. Lee, H. Lee, P. Abbeel, and A. Y. Ng, "Efficient l˜ 1 regularized logistic regression," in *Proceedings of the National Conference on Artificial Intelligence*, vol. 21, no. 1, 2006, p. 401.

[27] R. Tibshirani, "Regression shrinkage and selection via the lasso," *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 267–288, 1996.

[28] G. Xu and J. Z. Huang, "Asymptotic optimality and efficient computation of the leave-subject-out cross-validation," *The Annals of Statistics*, vol. 40, no. 6, pp. 3003–3030, 12 2012.

[29] E. Miluzzo, N. D. Lane, K. Fodor, R. Peterson, H. Lu, M. Musolesi, S. B. Eisenman, X. Zheng, and A. T. Campbell, "Sensing meets mobile social networks: the design, implementation and evaluation of the CenceMe application," in *Proceedings of the 6th ACM Conference on Embedded Network Sensor Systems*, 2008, pp. 337–350.

[30] K. K. Rachuri, M. Musolesi, C. Mascolo, P. J. Rentfrow, C. Longworth, and A. Aucinas, "Emotionsense: A mobile phones based adaptive platform for experimental social psychology research," in *Proceedings of the 12th ACM International Conference on Ubiquitous Computing*, ser. UbiComp, 2010, pp. 281–290.

[31] H. Hung, G. Englebienne, and J. Kools, "Classifying social actions with a single accelerometer," in *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp)*, 2013, pp. 207–210.