Kozlowski, M., Byrne, D. B., Santos-Rodriguez, R., & Piechocki, R. J. (2018). Data Fusion for Robust Indoor Localisation in Digital Health. In *2018 IEEE Wireless Communications and Networking Conference Workshops (WCNCW 2018): Proceedings of a meeting held 15-18 April 2018, Barcelona, Spain* (pp. 302-307). Institute of Electrical and Electronics Engineers (IEEE).
<https://doi.org/10.1109/WCNCW.2018.8369009>

Peer reviewed version

Link to published version (if available):
[10.1109/WCNCW.2018.8369009](https://doi.org/10.1109/WCNCW.2018.8369009)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the author accepted manuscript (AAM). The final published version (version of record) is available online via IEEE at <https://ieeexplore.ieee.org/document/8369009> . Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Data Fusion for Robust Indoor Localisation in Digital Health

Michał Kozłowski¹, Dallan Byrne¹, Raúl Santos-Rodríguez², Robert Piechocki¹

Department of Electrical and Electronic Engineering¹

Department of Engineering Mathematics²

University of Bristol, UK

Email: {m.kozlowski, dallan.byrne, enrsr, r.j.piechocki}@bristol.ac.uk

Abstract—This paper offers an approach for the combining of signals from multiple sensors observing everyday activities in a digital health care monitoring context. The IoT environment presents a number of advantages for indoor localisation. The amalgamation of several passive sensors can be used to provide an accurate location. This location often bears unique signatures of activity especially when considering residential environments. However, it is only the basic human instincts, such as periodicity and routine, that make this possible. The fact that behaviours and actions recur naturally is an important assumption in this paper. The study proposes a method, whereby semantic information about the location is learned from an additional source. This method deals with the question of robust indoor localisation prediction by extracting additional activity information available from a wrist worn acceleration sensor. A number of different fusion models are considered, before choosing and validating the model which provides highest improvement in accuracy and robustness over the baseline example. The performance of the methods is examined on different unique datasets, which closely resemble residential living scenarios.

I. INTRODUCTION

In the field of Digital Health, indoor localisation is considered essential as a proxy to understand the behaviour of patients within their home environment. The localisation approaches can be divided into two categories – Range-Based and Range-Free. Range-Based schemes utilise, among others, Time-of-Arrival (TOA) [1], Relative Signal Strength (RSS) [2] and Ultra Wideband (UWB) [3] in order to infer position of the sensor node. Range-Free rely on proximities between sensor nodes [4].

RSS based techniques are arguably state-of-the-art in indoor localisation [5]. RSS is the measure of power received by the receiver and will vary depending on the medium and positions of obstacles in the environment [6]. RSS provides the best trade-off between system complexity and the need for training, as it is low cost to implement and deploy [7], [8]. Recent localisation competition [9] showed, that RSS based systems can achieve comparable performance to active methods, such as UWB or ultrasound without the necessity for expensive hardware.

This paper presents a Range-Based probabilistic method of localisation. It relies on the variations between distributions of RSS in different locations in the house [10]. The location is inferred through a temporal state-space model – Hidden Markov Model (HMM) [11]. Variations of the signal are encoded in the

locations themselves. Given any location, the distribution of the signals arriving there will be highly dependent on shadowing effects [12] and the user's current position indoors [6].

In localisation literature, the RSS signal is often complemented by other types of sensors [13]. In order to enrich the data in this study, a wrist-worn accelerometer is used as an additional source of information about the activity. Accelerometers provide a measure of acceleration in three dimensions - x , y and z - represented by gravitational forces acting in those three directions [14]. Processing the accelerometer data involves feature extraction and classification of predefined tasks. However, wrist-mounted accelerometer activity recognition is usually inferior to prediction from sensors mounted on different parts of the body [15]. Regardless, the wrist remains the least intrusive and most socially acceptable place to wear a sensor [16].

The probabilistic models presented in this paper are analysed on SPHERE Challenge Dataset [17]. The dataset was collected in a specially adapted two-floor residential house. It includes RSS and Accelerometer data, with richly labelled location and activity information. The performance of the optimal model is then examined using a separate, unique dataset. The High Resolution Localisation (HRL) dataset was collected specifically for this paper, using the equipment described in the paper by Pope et al. [18]. This data includes location information in higher resolution, but very sparse activity labelling.

In the context of residential health monitoring, accurate localisation of the patient in their environment is considered crucial [19]. This is the main motivation behind this work. The main contributions of this paper are the dataset, graphical models and the subsequent robustness analysis. We present and scrutinize the performance of novel models, linking activity information to passive RSS using unique data. We then demonstrate that this additional source of information is not only beneficial to location inference, but can also safeguard against noise and loss of data. We finally show the limits of these models, and provide reasons as to why they exist.

The paper is structured as follows: In Section II, the data collection, processing and feature extraction is discussed. Section III describes the fusion models. The models are then analysed on SPHERE Challenge data and the optimal model is further examined on the HRL dataset in Section IV. The points for future work are also discussed there.

II. DATA COLLECTION AND PROCESSING

A. SPHERE Challenge Dataset

For the collection of this dataset, a house was filled with 4 Access Points (APs) which provide the RSS information. The users were asked to wear a SPHERE wrist wearable [20], which served as a RSS anchor as well as accelerometer sensor. They then performed a number of scripted tasks in predefined areas of the house. The data was annotated using video, and included a variety of room-level location and activity labels. We refer the reader to [17] for further details.

The 4 APs available in the house were moved from their nominal positions in-between experiments. This meant, that learning a model on two affected experiments could produce different signatures of the RSS in the same locations. In addition to this, an AP would occasionally be out of commission for a period of time and would not be sending information about the signal strength. This, in localisation context, could be fallible, especially if the algorithm is based solely on RSS. The dataset however was chosen as an optimal platform to test models which are designed with robustness in mind. It was deemed that this dataset would closely resemble the data available in a real-world scenario, as well as reproduce the possible shortcomings which are likely to be encountered whilst rolling out this kind of IoT system.

Out of the 20 activities labelled in the dataset, there are only a handful which could help with localisation. A number of specific labels would be grouped together into a single class. The reason for their groupings stems from the sparsity of RF coverage in the test bed house. As the vast majority of the scripted experiments took place downstairs, the SPHERE dataset study included only one AP upstairs. The rooms with poor coverage included two bedrooms, a toilet and a corridor area. By distinguishing the activities performed in the bedrooms, such as 'sit-to-lie' and 'lie-to-sit' transitions, it was easier to predict the upstairs locations more accurately. This was because these particular movements are more often performed in these rooms and could be used to aid the RSS-only prediction of location.

The labels from SPHERE Challenge dataset were banded into 5 separate groups. These are tabulated in Table I. Group 1 helped with ambulation information. Group 2 was used to aid the localisation upstairs, as the tasks in that group were found to be most prevalent there. Group 3 aided with the staircase determination, in order to make the floor transitions more accurate. Group 4 only includes sitting, which was performed in a variety of rooms, much like squatting in Group 5.

B. High Resolution Localisation Dataset

The SPHERE Challenge Dataset lacked the granularity required to examine performance of localisation algorithms thoroughly. This was because only room-level labels were available. This necessitated the generation of a more diluted dataset, which could later be used for testing the robustness of the methods. Thus, the unique HRL dataset was collected.

In this dataset a different three-room dwelling was filled with an abundant amount of APs – 8 in total. The house floor was tessellated into $1\text{m} \times 1\text{m}$ tiles. A paper tag was placed on the floor tiles to label them. Each tile was assigned coordinates

in Euclidean space, relative to the 'tile zero'. The users then wore the updated version of the SPHERE accelerometer sensor [16], as well as a camera. The users performed three unscripted 'free-living' experiments of varying length. This dataset includes no comprehensive activity information. Again, this would mimic the real-life scenario. The collection and annotation of activity information is arduous for both the user and annotator.

TABLE I: Optimal label groups for activity recognition in SPHERE Challenge data

Group 1	Group 2	Group 3	Group 4	Group 5
Jump Walk-with-load Walk	Bending Kneeling Lying Standing Lie-to-sit Sit-to-lie Sit-to-stand Stand-to-kneel Stand-to-sit Straighten Turn	Ascend Descend Stand-to-bend Kneel-to-stand	Sitting	Squatting

C. Feature Extraction

There are a number of studies concerned with time series feature extraction, and accelerometer in particular. Common features include mean, mode and median, zero crossing rate and first five values of Short Time Fourier Transform [21], [22].

In order to extract the features from SPHERE Challenge data, a window of 6.4s as per Zhang et al. [23] was used. The windowing method was an overlapping rolling window, producing $K - N$ extraction samples, where K is the number of aggregated time bins. It segmented the data, sampled at 20Hz, into vectors of length $N = 128$, from which simple features were extracted based on direction-invariant magnitude. Each feature was then recorded and a number of different classifiers were used. Those classifiers were chosen on the basis of the state-of-the-art within the community [21], [22]. They include k-Nearest Neighbours (k -NN), Decision Trees, Linear Discriminant Analysis (LDA), Quadratic Discriminant Analysis (QDA) and HMM.

Not all of the features have the same relative impact over the classification accuracy. Minimum-Redundancy Maximum-Relevance (mRMR) [24] was used to choose the most effective subset of features based on the mutual information. The most dominant features were also the simplest – full list is shown in Table II.

TABLE II:
List of accelerometer features

Feature		
Variance	SD	Max/Min
Mean/Median/Mode	Range	Area
Sum	Kurtosis	25th Percentile
RMS	Skewness	

For HRL data however, temporal aggregation was required. Temporal aggregation is the accumulation and averaging of data points into respective temporal bins of specific duration, effectively down-sampling the data. HRL was sampled at 5Hz, outputting 5 separate unique values at each sample. Data was then aggregated into 0.2s time bins. It was found that a window of 1.2s performed best for feature extraction. This yielded $N = 6$ data points in each window. The better performance was likely due to the quality of data available.

III. MODELS

The notation in this section is as follows: L denotes location, RSS is the observation of the RSS, A is the inferred activity and Acc are the observations of the accelerometer features.

Fig. 1: Baseline

Fig. 2: Model 1

Fig. 3: Model 2

Fig. 4: Model 3

The model from Fig. 1 is used as a Baseline. It only uses RSS as its location observation. At a given time t , the trained model will compare the current observation of the RSS against all the location states. This method is widely accepted in literature [25] [26]. The stipulation in this model is that the distinctiveness of the signal in each room/tile is enough to localise a user in a residential environment. This model does not account for the user's activity information, nor does it make any contextual assumptions about the layout of the localisation environment.

Fig. 2 shows a first improvement on the Baseline. In addition to the previous RSS observations, it assumes that the location is also determined by the current activity of the user. This model came out of the belief that, for example, it would be more likely to assume the user is in the kitchen because they are cooking, instead of inferring the opposite. In order to infer the activity however, the feature observations are required.

The second model in Fig. 3 ignores the activity information. It instead relies on the fact that the user's raw accelerometer features are enough to distinguish specific location in the house. This model is the simplest of all three and only

considers observations to infer a single level network. This model is likely to be the most robust out of the three, mainly due to lesser complexity.

The final model in Fig. 4 does not directly link activities to locations, but the two nodes are nonetheless jointly dependent through observations. It is stipulated that the extra activity information might have some influence on how the location is inferred.

Inference was done through an HMM. Consider, that an observable state at a given time t is given by $v(t)$ and a hidden state given by $\omega(t)$. A vector representing a sequence of observable states is denoted as $\mathbf{V}^T$ and a sequence of hidden states ω^T , where T in both cases is the number of states. The model can then be evaluated by [27]:

$$P(\mathbf{V}^T) = \sum_{r=1}^{r_n} P(\mathbf{V}^T | \omega_r^T) P(\omega_r^T) \quad (1)$$

where r is the index for each particular sequence. This can further be described in terms of transition and emission probabilities given respectively as:

$$P(\omega_r^T) = \prod_{t=1}^T P(\omega(t) | \omega(t-1)) \quad (2)$$

$$P(\mathbf{V}^T | \omega_r^T) = \prod_{t=1}^T P(v(t) | \omega(t)) \quad (3)$$

The emission probabilities in this case were replaced with Gaussian Mixtures. It describes the symbols with a single Gaussian [28]:

$$b_j(RSS_{tk}) = \sum_{k=1}^M N(RSS_{tk} | \mu_{jk}, \Sigma_{jk}) \quad (4)$$

for RSS emissions, where $1 \leq j \leq T$ are the location states, and $1 \leq k \leq M$ is the number of APs. For the accelerometer feature case:

$$b_j(Acc_{tl}) = \sum_{l=1}^S N(Acc_{tl} | \mu_{jl}, \Sigma_{jl}) \quad (5)$$

where $1 \leq l \leq S$ is the number of features.

IV. RESULTS

A. SPHERE Challenge data analysis

The location was measured using 3 different metrics. The SPHERE Challenge data only labels rooms and as such the result is room level. The resolution was big enough, such that a direct comparison between a label and prediction at each t could be used. If $y(t)$ is our label and $h(t)$ the prediction, then our accuracy $E(t)$ is:

$$E(t) = \begin{cases} 1 & \text{if } y(t) = h(t) \\ 0 & \text{else} \end{cases}$$

Second metric is a basic Euclidean distance error between prediction and label rooms/tiles [9]. Finally, the path error

Fig. 5: Accuracy of each model. Experiments 6 and 7 were omitted due to under performance.

specifies the shortest Dijkstra path [29] between prediction and label.

The SPHERE Challenge Dataset was used to analyse how the models would perform on low-resolution but well labelled data. The dataset was separated into 10 identical experiments performed by different users. In order to test the model, 10-fold cross validation was used across all experiments. Two experiments - 6 and 7 - were under-performing. After removing those from the fold, the performance of the remaining bins increased. They have therefore been omitted from subsequent analysis.

Every enhancement model improved the nominal result, suggesting that the inclusion of accelerometer data is advantageous. The small deviations between the models were, in addition to their architectures, likely caused by their complexity. Model 3 is the most complex of the remaining two. Fig. 5 shows that it performed similarly to Model 2, never deviating for more than 5%. Those two models share similarities in the way they infer the location, but the prediction coming from a less complex Model 2 is more accurate. Model 1 did not follow any other method. It was more accurate when predicting the path error than Model 2. However, it required more elaborate pre-processing and inference methods, as it would be inferred on two levels. The increased number of inference steps are more likely to harbour inaccuracies and false positive activity predictions. This in turn translates into inaccurate location result. Table III shows the overall average result of the SPHERE Challenge data analysis.

B. High Resolution Localisation analysis

Model 2 was therefore chosen as the optimal network, by leveraging the result obtained to the complexity of the system. It was used to validate the hypothesis set out on the previous dataset. The data consisted of three separate 'free-living' experiments. Those experiments included 'everyday' behaviours and tasks which are likely to be found in any Digital Health data collection study. As with the SPHERE Challenge data, cross-validation was used train and test the

model. It is important to point out that the chosen method did not require any activity labels. The results can be seen in Table IV.

This paper will test the robustness of different models by presenting two experiments on the HRL data. Firstly, an artificial noise, disguised as packet drop rate, will be iteratively increased. This is to see how the baseline and the enhanced model will perform when faced with missing data. Secondly, the APs will be gradually removed. This will mean that there will be fewer sources of information. The experiment will check how the enhanced model will perform when faced with less data in a smaller indoor environment.

When using the HRL data, the resolution was reduced to $1m \times 1m$. This meant that the actual distance error could remain similar, whilst the tile-level accuracy metric would fail to provide a viable result. The finer resolution increased the overall temporal error. Consider Table IV, where the room-level accuracy is now perfect, but tile-level reduces to 15.88% in the best case. Due to that fact, only the path error and the distance error were considered during the robustness study.

C. Discussion

Our RSS-based system achieves comparable performance to the state-of-the-art RSS implementations in the Microsoft Localisation Competition [9]. The average distance error achieved by our method in Table IV (1.59m) is similar to the error achieved by Chen et al. [9], [30] (1.37m) using analogous infrastructure. However, our experimental scenario and the testing environment differs from the competition setup and as such the two cannot be directly compared.

Although the improvement over the Baseline is slight, one of the goals is to study the robustness of the Model against different types of perturbations.

Firstly, the packet loss rate between the APs and the wearable, which is naturally present with a value of 22.75%,

TABLE III: Results of testing the models on SPHERE Challenge data

	Room-level Accuracy (%)	Distance (m)	Path (m)
Baseline	70.4	1.28	2.11
Model 1	74.2	1.05	1.47
Model 2	75.9	1.01	1.51
Model 3	73.3	1.09	1.61

TABLE IV: Results of Model 2 with HRL

	Room-level Accuracy (%)	Tile-level Accuracy (%)	Distance (m)	Path (m)
Baseline	100	14.66	1.65	1.96
Model	100	15.88	1.54	1.95

Fig. 6: Path error per increasing packet drop rate.

Fig. 7: Distance error per increasing packet drop rate.

was increased. Fig. 6 shows the path error of the Baseline and Model to increasing packet loss rates. Similarly, Fig. 7 illustrates the distance error. Both of graphs show the average performance of $n = 57$ random injections of noise into the system together with the standard deviation. The Model's path finds a minimum at 50% of dropped packets. It is at that point that the result shows an improvement of 10cm over the Baseline. This happens as the accelerometer values, originating in the wearable, are invariant to range, whereas RSS are not. The former will appear the same or similar at each AP, whereas the latter will vary with each AP. This makes accelerometer information complementary and thus more immune to added noise. After 50% of noise however, the Baseline begins to outperform the Model. It appears that, again, the complexity plays an important role in the prediction. Since the Model requires the accurate estimation of more parameters than the Baseline, it is prone to overfit the data. This can also be confirmed by the standard deviation of the error at high packet drop percentage. Model's error fluctuates more broadly than the error for the Baseline. Additionally, the distance error confirms this, but to a lesser extent as it presents a much smoother increase. This is due to less rigorous distance measurement – any deviation from the label will be scaled linearly, as opposed to being a function of the layout of the environment.

Secondly, an experiment was devised to understand how the Model performs when APs are removed to simulate a scenario with reduced numbers of APs. This study involved taking the RSS distributions for each AP and ranking them according to their pairwise overlap, computed as the Weizman's measure (also known as the overlap coefficient). The APs were then removed one by one according to this criteria in order to reduce the total number of APs removing as little information as possible. Figures 8 and 9 show the performance for the Baseline and Model for both path and distance metrics. Similarly to the previous experiment, the behaviour is consistent with the hypothesis that the Model, given relatively noiseless data, will outperform the Baseline, even when faced with fewer sources of information.

Fig. 8: Path error given increasingly fewer Access Points.

Fig. 9: Distance error given increasingly fewer Access Points.

V. CONCLUSION

This paper proves that inferring the location of an individual in their own home could be improved by incorporating additional data sources. To that end, a specific accelerometer signature was associated with a specific location. By performing the same, or similar tasks, in the same places, the signatures are comparable enough between different free-living experiments, as to aid the RSS localisation technique. The results show that the localisation is robust even when noise is added to the system and if the sources of information are gradually being removed. However at higher noise levels, and at larger information absence, the Model performs more poorly, as compared to the Baseline.

Future work will involve incorporating additional sensors, including gyroscope and magnetometer, together with additional information pertaining to the layout of the house to explore how the layout complexity relates to packet drop and the location accuracy.

ACKNOWLEDGEMENTS

This study was supported by EPSRC and SPHERE IRC, grant number EP/K031910/1. The authors would like to thank all colleagues from the SPHERE group.

REFERENCES

- [1] I. Guvenc and Z. Sahinoglu, "Threshold-based TOA estimation for impulse radio UWB systems," in *2005 IEEE International Conference on Ultra-Wideband*, pp. 420–425, Sept. 2005.
- [2] X. Zhu and Y. Feng, "RSSI-based Algorithm for Indoor Localization," *Communications and Network*, vol. 05, p. 37, May 2013.
- [3] S. Gezici, Z. Tian, G. B. Giannakis, H. Kobayashi, A. F. Molisch, H. V. Poor, and Z. Sahinoglu, "Localization via ultra-wideband radios: a look at positioning aspects for future sensor networks," *IEEE Signal Processing Magazine*, vol. 22, pp. 70–84, July 2005.
- [4] T. He, C. Huang, B. M. Blum, J. A. Stankovic, and T. Abdelzaher, "Range-free Localization Schemes for Large Scale Sensor Networks," in *Proceedings of the 9th Annual International Conference on Mobile Computing and Networking, MobiCom '03*, (New York, NY, USA), pp. 81–95, ACM, 2003.
- [5] H. Liu, H. Darabi, P. Banerjee, and J. Liu, "Survey of Wireless Indoor Positioning Techniques and Systems," *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, vol. 37, pp. 1067–1080, Nov. 2007.
- [6] P. Bahl and V. N. Padmanabhan, "RADAR: an in-building RF-based user location and tracking system," *Proceedings IEEE INFOCOM 2000*, vol. 2, pp. 775–784 vol.2, 2000.
- [7] M. Mizmizi and L. Reggiani, "Design of RSSI based fingerprinting with reduced quantization measures," in *2016 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, pp. 1–6, Oct. 2016.
- [8] R. McConville, D. Byrne, I. Craddock, R. Piechocki, R. Santos-Rodríguez, and J. Pope, "Understanding the quality of calibrations for indoor localisation," *2018 IEEE 4th World Forum on Internet of Things*, 2018.
- [9] D. Lymberopoulos and J. Liu, "The Microsoft Indoor Localization Competition: Experiences and Lessons Learned," *IEEE Signal Processing Magazine*, vol. 34, pp. 125–140, Sept. 2017.
- [10] V. Ramadurai and M. L. Sichitiu, "Localization in Wireless Sensor Networks: A Probabilistic Approach," in *Proceedings of the 2003 International Conference on Wireless Networks*, 2003.
- [11] L. Rabiner and B. Juang, *Fundamentals of speech recognition*. PTR Prentice Hall, Apr. 1993.
- [12] A. A. Hussein, T. A. Rahman, and C. Y. Leow, "Performance Evaluation of Localization Accuracy for a Log-Normal Shadow Fading Wireless Sensor Network under Physical Barrier Attacks," *Sensors (Basel, Switzerland)*, vol. 15, pp. 30545–30570, Dec. 2015.
- [13] A. Conti, D. Dardari, M. Guerra, L. Mucchi, and M. Z. Win, "Experimental Characterization of Diversity Navigation," *IEEE Systems Journal*, vol. 8, pp. 115–124, Mar. 2014.
- [14] M. J. Mathie, A. C. F. Coster, N. H. Lovell, and B. G. Celler, "Accelerometry: providing an integrated, practical method for long-term, ambulatory monitoring of human movement," *Physiological Measurement*, vol. 25, pp. R1–20, Apr. 2004.
- [15] A. Mannini, S. S. Intille, M. Rosenberger, A. M. Sabatini, and W. Haskell, "Activity recognition using a single accelerometer placed at the wrist or ankle," *Medicine and science in sports and exercise*, vol. 45, pp. 2193–2203, Nov. 2013.
- [16] X. Fafoutis, A. Vafeas, B. Janko, R. S. Sherratt, J. Pope, A. Elsts, E. Mellios, G. Hilton, G. Oikonomou, R. Piechocki, and I. Craddock, "Designing Wearable Sensing Platforms for Healthcare in a Residential Environment," *EAI Endorsed Transactions on Pervasive Health and Technology*, vol. "3", Sept. 2017.
- [17] N. Twomey, T. Diethe, M. Kull, H. Song, M. Camplani, S. Hannuna, X. Fafoutis, N. Zhu, P. Woznowski, P. Flach, and I. Craddock, "The SPHERE Challenge: Activity Recognition with Multimodal Sensor Data," *arXiv:1603.00797 [cs]*, Mar. 2016. arXiv: 1603.00797.
- [18] J. Pope, R. McConville, M. Kozłowski, X. Fafoutis, R. Santos-Rodríguez, R. J. Piechocki, and I. Craddock, "Sphere in a box: Practical and scalable curvalve activity monitoring smart home kit," *2017 IEEE 42nd Conference on Local Computer Networks Workshops*, pp. 128–135, 2017.
- [19] T. Van Haute, E. De Poorter, P. Crombez, F. Lemic, V. Handziski, N. Wirstm, A. Wolisz, T. Voigt, and I. Moerman, "Performance analysis of multiple Indoor Positioning Systems in a healthcare environment," *International Journal of Health Geographics*, vol. 15, p. 7, Feb. 2016.
- [20] X. Fafoutis, B. Janko, E. Mellios, G. Hilton, R. S. Sherratt, R. Piechocki, and I. Craddock, "A Low-Maintenance Wearable Activity Tracker for Residential Monitoring and Healthcare Applications," in *Proceedings of the International Conference on Wearables in Healthcare (HealthWear)*, June 2016.
- [21] S. J. Preece, J. Y. Goulermas, L. P. J. Kenney, and D. Howard, "A Comparison of Feature Extraction Methods for the Classification of Dynamic Activities From Accelerometer Data," *IEEE Transactions on Biomedical Engineering*, vol. 56, pp. 871–879, Mar. 2009.
- [22] N. Ravi, N. Dandekar, P. Mysore, and M. L. Littman, "Activity Recognition from Accelerometer Data," in *ResearchGate*, vol. 3, pp. 1541–1546, Jan. 2005.
- [23] S. Zhang, A. V. Rowlands, P. Murray, and T. L. Hurst, "Physical Activity Classification Using the Geneva Wrist-worn Accelerometer," *Medicine & Science in Sports & Exercise*, vol. 44, pp. 742–748, Apr. 2012.
- [24] H. Peng, F. Long, and C. Ding, "Feature selection based on mutual information criteria of max-dependency, max-relevance, and min-redundancy," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 27, pp. 1226–1238, Aug. 2005.
- [25] Y. Ni, J. Liu, S. Liu, and Y. Bai, "An Indoor Pedestrian Positioning Method Using HMM with a Fuzzy Pattern Recognition Algorithm in a WLAN Fingerprint System," *Sensors (Basel, Switzerland)*, vol. 16, Sept. 2016.
- [26] M. K. Hoang, J. Schmalenstroer, C. Drueke, D. H. T. Vu, and R. Haeb-Umbach, "A hidden Markov model for indoor user tracking based on WiFi fingerprinting and step detection," in *21st European Signal Processing Conference (EUSIPCO 2013)*, pp. 1–5, Sept. 2013.
- [27] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification (2Nd Edition)*. Wiley-Interscience, 2000.
- [28] H. Tang, M. Hasegawa-Johnson, and T. S. Huang, "Toward robust learning of the Gaussian mixture state emission densities for hidden Markov models," in *2010 IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 5242–5245, Mar. 2010.
- [29] D. E. Knuth, "A generalization of Dijkstra's algorithm," *Information Processing Letters*, vol. 6, pp. 1–5, Feb. 1977.
- [30] Z. Chen, Q. Zhu, H. Jiang, H. Zou, Y. C. Soh, L. Xie, R. Jia, and S. Costas, "An iBeacon Assisted Indoor Localization and Tracking System," *Microsoft Indoor Localization Competition*, 2015.