

Analysis of patient domestic activity in recovery from Hip or Knee replacement surgery: modelling wrist-worn wearable RSSI and Accelerometer data in the wild

Mike Holmes^{1*}, Hao Song¹, Emma Tonkin¹, Miquel Perello Nieto¹, Sabrina Grant¹ and Peter Flach¹

¹ University of Bristol

*mike.holmes@bristol.ac.uk

Abstract

The UK health service sees around 160,000 total hip or knee replacements every year and this number is expected to rise. Expectations of surgical outcome are changing alongside demographic trends, whilst aftercare may be fractured as a result of resource limitation or other factors. Conventional assessments of health outcomes must evolve to keep up with these changing trends. In practice, patients may visit a health care professional to discuss recovery and will provide survey feedback to clinicians using standardised instruments, such as the Oxford Hip & Knee score, in the months following surgery. To aid clinicians in providing accurate assessment of patient recovery a continuous home health care monitoring system would be beneficial.

In this paper the authors explore how the SPHERE sensor network can be used to automatically generate measures of recovery from arthroplasty to facilitate continuous monitoring of behaviour, including location, room transitions, movement and activity; in terms of frequency and duration; in a domestic environment. The authors present a case study of data collected from a home equipped with the SPHERE sensor network. Machine learning algorithms are applied to a week of continuous observational data to generate insights into the domestic routine of the occupant. Testing of models shows that location and activity are classified with 86% and 63% precision, respectively.

1 Introduction

The UK health service sees around 160,000 hip and knee replacements every year [National Joint Registry, 2018] within the National Health Service and this number is expected to increase. Hence, innovative approaches to evaluating surgical outcomes will be needed to respond to the increasing burden of joint replacement surgery. Health care interventions, such as surgeries, are only part of a patient's journey. Expectations of surgical outcome are changing alongside demographic trends [National Joint Registry Editorial Board, 2017]. Conventional assessments of health outcomes must

evolve to keep up with these changing trends. After joint replacement, up to 30% of patients report minimal improvement or their symptoms get worse and not all patients are satisfied with their outcome [Beswick *et al.*, 2012]. Poor outcomes include continuing pain, functional limitation and increased health care utilisation. Consequentially, improving outcomes after joint replacement is a key research priority.

Patients routinely receive a follow-up appointment approximately six weeks following surgery. However, this may not be with the surgeon, but with a registrar. This may complicate assessments. Various strategies have been proposed to increase efficiency whilst maintaining quality and patient acceptability, such as the use of 'virtual clinics' [Williams, 2014]. These rely on Patient Reported Outcome Measures (PROMs), such as the Oxford Hip or Oxford Knee Score and the EQ-5D, a measure of health status. These can assess various health outcomes including pain, function and aspects of quality of life, but have sometimes significant limitations. For example, PROMs may be subjective to a certain extent and may reflect the patient's level of pain [Senden *et al.*, 2011; Stevens-Lapsley *et al.*, 2011].

Previously, research has explored the relationship between PROMs and objective measures, notably performance-based tests such as timed walks or sit-to-stand tests [Bolink *et al.*, 2012]. Such objective measures are administered in controlled, laboratory style settings, and may not reflect levels of activity in daily life. Multimodal sensor systems present in domestic settings, such as those used in ambient assisted living scenarios [Rashidi and Mihailidis, 2013], allow assessment of behaviour and activity in a natural setting. Establishing a relationship between PROMS and multimodal sensor data permits us to develop effective methods of passive monitoring and recovery after surgery, providing a further data source that, if used alongside PROMS, may allow for relatively timely intervention in the event of complications, potentially improving patient outcomes.

1.1 Contribution

The contribution of this research is to make an initial evaluation of statistical method, from literature (section 2), which may provide measurement or classification of mobility information including location and room transitions, movement intensity and distance, and posture & ambulatory activities, using data gathered *in the wild*. Techniques detailed in litera-

ture (section 2) are applied to one week of continuous observational data recorded within a real residence (section 3.1).

Results of classifier training are presented in section 4, along with visualisation of measurements and classifications for location, movement and activity. An evaluation of methods using real-world data is presented in discussion in section 5, with conclusions in section 6.

A valuable outcome of this initial evaluative research has been to highlight the future work (section 7) necessary to develop algorithms for long-term measurement and classification which can be robust to the challenges presented when working with data gathered *in the wild*.

1.2 SPHERE: A sensor platform for health care in a residential environment

SPHERE is an interdisciplinary research project which aims to develop sensor technologies capable of supporting a variety of practical use cases, including healthcare and ambient assisted living outcomes. An additional goal of SPHERE is to build systems that are considered acceptable by the public and which are flexible and powerful enough to function well in a broad variety of domestic environments [Woznowski *et al.*, 2015; 2017].

‘Smart home’ systems development has primarily taken place in laboratory settings [Alam *et al.*, 2012], or, as in the SPHERE project, in a customised home [Tao *et al.*, 2015]. Research, development and testing of multimodal sensor technologies was completed in a home owned by the project, the SPHERE House. In 2017, the SPHERE project began to deploy a multimodal sensor network into dozens of homes in the South West of England.

The work reported here is part of a set of initial studies on data generated using the SPHERE sensor network in deployment. In particular, this study is intended to establish the behaviour of the sensor network and of the associated analytic infrastructure, including measurements of participant location, movement and activity, in a genuine deployment context ‘in the wild’.

2 Related work

Key indicators of relevance to PROMS include movement patterns (such as room to room transfers), patterns of improvement (quality of movement, distance walked, climbing stairs), activities undertaken (such as cooking or cleaning) and sleep (e.g. hours sleeping, quality of sleep).

This study focuses on three measurements of participant domestic behaviour including location, movement and activity. In this section, the authors provide a brief overview of research relating to each method employed.

2.1 Indoor localisation

Indoor localisation [Quan *et al.*, 2010; Wang *et al.*, 2014] is an important area of research for behavioural analysis in residential health care. The ability to predict the location of a patient not only gives insight into domestic routine and habitation, but allows other information to be physically contextualised. The SPHERE low-energy Bluetooth network provides a mesh of overlapping or interacting signal strength fields

which provide a pattern of distinct Received Signal Strength Indicator measures, based on proximity to each receiver in the network.

As in literature [Quan *et al.*, 2010; Wang *et al.*, 2014], RSSI has been used to *fingerprint* locations within a space by learning the discriminant RSSI vectors from a moving average [Quan *et al.*, 2010].

2.2 Movement

Measures of movement are often based on either measured acceleration [Preston *et al.*, 2012; Xiao *et al.*, 2016] or inferred positional change as represented by shift in Received Signal Strength Indicator (RSSI) [Krumm and Horvitz, 2004; Gansemer *et al.*, 2010b; 2010a]. Both approaches have advantages and disadvantages, considering different types of movement and the location of the accelerometer. A wrist-worn accelerometer as used in SPHERE may show spikes in magnitude based on ambulation (e.g. walking or running) but also for rapid hand / arm movements (e.g. chopping vegetables), and acceleration magnitude may not show spikes for low acceleration positional change (e.g. transcending stairs with aid of a stair lift). RSSI will highlight positional change regardless of acceleration but will not show movement that is non-positional (e.g. walking or running on a tread mill), and may overestimate small movements which block line of sight to RSSI receivers (e.g. rolling over in bed).

In this work both magnitude and RSSI based movement calculations are shown for comparison.

2.3 Activity

Activity recognition using wearable and mobile devices has been a major focus for the recent years [Bao and Intille, 2004; Kwapisz *et al.*, 2011; Siirtola and Rönning, 2012; Ravi *et al.*, 2005; Janidarmian *et al.*, 2017].

From a device prospective, mobile phones, smart watches and wrist bands have the dominant source of data, which normally captures the acceleration signal around the body of the users. In this paper we also focus on the 3-axes acceleration data obtained from a wrist band, which is one of the standard choice in the field.

3 Case study

In this study the authors present initial results and data visualisations for a single case study participant home of the SPHERE cohort over the first week of system installation.

The case study presents the authors *work in progress* in developing methods of analysis to monitor, visualise and validate key indicators of recovery from surgery, such as hip or knee arthroplasty. The experiments in this paper focus on experimental evaluation of methods for measurement of movement within the home.

3.1 Methods

In this section the authors present an overview of methods used to generate the three classification metrics: in-door localisation, movement and activity classification.

Data Collection

The case study home has been selected as it represents a simple use-case for SPHERE technology *in the wild*. The residence has a single occupant with few rooms and all rooms located on the same floor. Figure 1 shows a graphical representation of the layout of the residence.




Figure 1: Connection between different rooms in the house

The SPHERE system has been installed in the residence for five months. In this paper the authors focus on analysis of the first week of installation, so as to give an overview of the methods used for analysis and visualisation of data.

Figure 2 details the physical architecture of one subsystem of the SPHERE sensor network, the wrist-worn Bluetooth Low-Energy (BLE) wearable device. The wrist wearable harbours a tri-axial accelerometer and broadcasts over Bluetooth at 25Hz.




Figure 2: Diagram of wearable Bluetooth LE subsystem of SPHERE

Data collected from a second SPHERE subsystem; the environmental sensor network; will provide passive infra-red activation data, of use in validating the predictions made using the wearable data.

To develop a localisation training set for the home, during installation of the SPHERE sensor network, a technician performs an annotation procedure called a 'technician walk-around'. The technician carries the wearable device to each room in the home, annotating the start and end times in each labelled location. The technician walk-around was repeated prior to the sensor network being removed from the home. Figure 3 visualises the technician walk around.

The participant is asked to perform their daily routine in a fast-forward manner. That is, the participant starts from the location of the bed, visit corresponding locations according to their usual routine, while performing simple activities like "making a cup of tea". During the experiment, the participant

wears a head mounted camera to record the view, which can be later annotated for the activities performed.




Figure 3: Visualisation of technician walk around

Ethics: data collection and publication

The data used in this study has been collected as part of the SPHERE project [Woznowski *et al.*, 2015; 2017]. The participant in this case study has provided consent for data to be recorded within their home. Participation in the SPHERE project is voluntary and participants are at liberty to exit the experiment at any time.

Due to the sensitive nature of data collected within a real-world residential environment, data used in this study is not being made public alongside this paper. A data set of activity and location annotated SPHERE sensor data, recorded during short scripted experiments in the SPHERE House (*The SPHERE Challenge*) [Twomey *et al.*, 2016], is available online.

Classifying location using RSSI fingerprints

RSSI levels between the wrist-worn wearable and each installed receiver within the home have been recorded. Using a 3-second sliding window, a vector of RSSI values is constructed to represent the position of the participant. For each second and each gateway, the sum, mean, minimum, maximum and variance in RSSI are calculated across the second. Each second in the sliding window was then concatenated to produce a vector of length $n = 75$.

A multi-layer perceptron artificial neural network (MLP) with three hidden layers, 100 nodes per hidden layer, was trained to classify the location of the wearable based on RSSI vectors. To train the classifier, the annotations taken during the technician walk-around activity (figure 3) were used to label the training and test set of vectors. The set of labelled vectors were shuffled and split 50/50 between training and testing sets. The MLP was trained using the 'adam' [Kingma and Lei, 2015] algorithm. Results of training and testing are presented in the section 4.

Location predictions are used to visualise room occupancy over time and to localise movement metrics, providing a view on where movement happens within the home and at what times of day.

In addition, location predictions are used to visualise the frequency of predicted transition from room to room. Predicted room transitions are expected to abide by the adjacency of rooms, as given in the residence layout in figure 1.

Classifying movement intensity

Movement intensity is calculated by the magnitude of acceleration (equation 1), as given by the mean tri-axial accelerometer readings from the wrist-worn wearable, over a 1-minute window. This approach has been successfully demonstrated in [Xiao *et al.*, 2016]. The wearable device transmits acceleration in x, y and z dimensions at 25Hz.

$$A = \sqrt{a_x^2 + a_y^2 + a_z^2} \quad (1)$$

Acceleration magnitude was calculated for each 1-minute of accelerometer data. For each 1-second window, the standard deviation of magnitude was calculated. Movement intensity is here defined as the sum of magnitude standard deviations (per second) over a given time window.

Movement is also calculated by the Euclidean distance between consecutive RSSI vectors. Similarly to the method described in [Muthukrishnan *et al.*, 2009], an RSSI vector (sliding window) was calculated for each observation window, in this case 1-minute. For each 1-minute window, the Euclidean distance between the current and previous window was calculated.

Movement classification is compared to activations of passive infra-red sensors (PIR) installed in each room of the residence. PIR sensors activate when movement is detected in a room. As the case study home is a single occupancy residence, PIR activations are anticipated to occur inline with increased acceleration magnitude.

Classifying activities using RSSI and tri-axial acceleration

From the fast-forward experiment, three models are obtained for predicting standing, lay down, and walk with a one-vs-rest strategy.

As both the RSSI and acceleration are collected at a relatively high frequency and hence are not synchronised to each other, the sliding window (with 2 seconds length) approach has been applied on the raw signal to obtain the standard feature vectors.

Since each activity is only collected for a couple of seconds during the fast-forward experiment, here we strategically avoid features that require a higher amount of training data, leaving features only involving calculating the mean values, median values, and standard deviations for each individual acceleration and RSSI, as well as the overall acceleration and RSSI readings on the target wearable device.

Regarding the model, we apply the Logistic regression for obtaining probabilistic outputs, which can then be corrected via Beta probability calibration.

With calibrated probabilistic outputs, one can easily visualise the uncertainty with each prediction, as well as calculating the overall time spent on these activities.

3.2 Data

In this section the authors present a brief overview of the data generated using methods described in section 3.1.

Location: 379,234 location classifications were made at 1 second intervals. **RSSI based movement:** 6,375 estimates of

distance travelled were generated at 1-minute intervals. **Accelerometer based movement:** 6,363 accelerometer magnitude observations were calculated, at 1-minute intervals. **Activity recognition:** 302,400 estimates of activity were generated from RSSI and accelerometer data, at 2-second intervals.

4 Results

In this section the authors present initial results applying localisation, movement and activity classification algorithms to the first week of data from the case study home.

4.1 Indoor localisation

Table 1 shows the test set performance of the trained MLP indoor localisation classifier. Table 4.1 shows the training and testing split for each class.

Location	Train	Test
Bedroom 1	37	42
Bathroom 1	30	22
Kitchen 1	30	29
Living Room 1	57	49
Hall 1	37	50

Table 1: Location classifier test-set results

Class	Precision	Recall	f1-score	Support
bathroom 1	0.86	0.83	0.84	23
bedroom 1	0.79	0.94	0.86	35
hall 1	0.80	0.83	0.82	48
kitchen 1	1.00	0.91	0.95	32
living room 1	0.88	0.80	0.83	54
avg / total	0.86	0.85	0.85	192

Figure 4 visualises room occupancy in 2-hour windows, as predicted by the localisation classifier. Figure 5 shows location transitions from location to location, within the home.

4.2 Movement intensity and distance

Figure 6 shows movement intensity in 2-hour windows, over the observed week. Figure 7 shows movement as calculated by mean euclidean distance between RSSI vectors in each 2-hour window. Figure 8 shows passive infra-red sensor activations over the same period.

Table 2 shows movement intensity for each day in each residential location.

Figure 9 shows the intensity of movement calculated by accelerometer magnitude in each domestic locations, across the first week of observation.

4.3 Activity Recognition

As we train the corresponding model for each activity with the one-vs-rest strategy, the results can be simply evaluated as in table 3. Table 3 shows test set performance of the classifier. Figures 10, 11 and 12 show activity classifications aligned to location classifications (figure 4) across the observed week.

Table 2: Movement (magnitude) per day per location. Sub-index indicates the rankings from high to low values.

Day	Bathroom	Bedroom	Hall	Kitchen	Living room	Total
Mon	0.43	0.76	10.35	0.06	8.74	20.35 ₍₇₎
Tue	0.14	2.47	9.20	1.74	22.43	35.99 ₍₃₎
Wed	1.78	11.25	1.07	3.68	7.08	24.87 ₍₆₎
Thu	4.40	10.04	1.13	8.28	9.84	33.69 ₍₄₎
Fri	1.55	11.26	2.66	5.28	6.66	27.41 ₍₅₎
Sat	5.12	11.95	4.89	10.33	6.71	39.00 ₍₂₎
Sun	5.37	21.67	5.98	29.45	27.03	89.50 ₍₁₎
Total	18.79 ₍₅₎	69.4 ₍₂₎	35.28 ₍₄₎	58.82 ₍₃₎	88.49 ₍₁₎	




Figure 4: Location by 2-hour window across week 1




Figure 6: Movement intensity (magnitude) by 2-hour window




Figure 5: Location transitions across week 1 comparing non-overlapping intervals of 1 minute. Outer squares indicate: one-step adjacency in red and two-steps adjacency in dotted orange.

5 Discussion

Localisation predictions appear accurate on the small sample set used for training and testing the classifier. Results of classifier testing (Table 1) show average precision of 86% and recall 85% over the five classes. However, given the limited size of the data set, the assumption of independence between samples in the training and test set does not hold. For




Figure 7: Movement distance (RSSI) by 2-hour window

Table 3: Activity classifier 3-folds cross validation results

Class	Precision	Recall	f1-score	Support
lay down	0.75	0.92	0.83	13
stand	0.61	0.71	0.66	35
walk	0.53	0.47	0.50	36
avg / total	0.63	0.70	0.66	84




Figure 8: Passive infra-red sensor activation by 2-hour window




Figure 11: Predictions of the activity 'stand' in the predicted locations across one week




Figure 9: Movement (magnitude) by location across week 1




Figure 12: Predictions of the activity 'lay down' in the predicted locations across one week




Figure 10: Predictions of the activity 'walk' in the predicted locations across one week

that reason, in our current analysis we may expect lower levels of accuracy than estimated. From the location prediction data across the week (figure 4), a regular routine emerges. The participant occupies the bedroom from around 22:00 on most nights, with occupancy transitioning through the hall-

way, bathroom, kitchen and living room between 04:00 and 06:00 on most mornings. The participant leaves the residence between 06:00 and 08:00 between Tuesday and Friday. The living room is occupied in the evenings between 16:00 and 22:00, with the longest periods of occupancy on the Monday, Tuesday and Sunday.

The room transitions (figure 5) show that living room and kitchen are most frequently moved between. However, the transition matrix calculated by majority class in each pairwise minute causes the hallways to be under represented. Allowing for second order adjacency; essentially allowing for hops over the hallway; reduces the error.

Movement intensity (figure 6) by accelerometer magnitude supports the location predictions, showing that movement intensity decreases at 22:00 on most days and remains low until 04:00, a time when the participant is located in the bedroom. The most intense movement within the home was recorded on the weekend.

Movement intensity measurements by accelerometer magnitude are supported by the PIR activation (figure 8) data. With the exception of the installation period, between 08:00 and 12:00 on Monday, when there were SPHERE technicians

in the home, the individual participants' movement from the wearable maps well to the movement detected by the PIR sensors.

Movement intensity (accelerometer magnitude) by location highlights where activity occurs within the home. Table 2 shows that over the entire week the living room was where most movement occurred, followed by the bedroom, kitchen, hallway and finally the bathroom.

Figure 9 shows how movement intensity varied in locations over time. The visualisation shows low intensity movement during night time hours, when the participant is located in the bedroom. More intense movement is detected in both the kitchen and living room each day in the late afternoon and evening. The most intense and sustained movement occurred in the kitchen on Sunday morning.

A comparison of magnitude measurements in figure 6 and RSSI measurements in figure 7 during sleeping hours highlights a potential problem with using RSSI measurement. RSSI signal change can be by obscuring the wearable. During sleep the RSSI signals can be modified by a participant changing sleeping position and obscuring the wearable, resulting in a perceived movement in position.

Walking activity, shown in figure 10, appears mostly in the evening, overlapping with time spent moving in the kitchen and living room (Figure 6).

Standing activity, shown in figure 11, occurs in the hall, kitchen and living room and routinely at the beginning and end of each day, overlapping with time spent in the kitchen and living room (figure 4).

Laying down, shown in figure 12, indicates sleep during hours of low activity in the bedroom, as shown in figure 6.

6 Conclusions

The case study has demonstrated that RSSI and tri-axial accelerometer data can be used to measure key indicators of recovery from Hip and Knee replacement surgery, such as daily routine, sleep patterns, location transitions and movement intensity over a short period of time.

RSSI *fingerprints* collected during the technician walk-around activity were sufficient (Table 1) to model distinct locations within the home. However, transitions such as kitchen to bathroom, which appear in figure 5, are not physically possible according to the layout of the home in figure 1. The likely cause is the adjacency of kitchen and bathroom meaning the *fingerprints* may converge dependent on factors such as radio-frequency interference.

To improve the classification algorithm and reduce erroneous location transitions, a representation of possible transitions given a prior must be included in the model. This improvement will be a focus of future work.

Routine of activity and passivity is highlighted in movement estimates, with accelerometer magnitude providing the clearest view of true movement levels. RSSI tended to over estimate movement.

Activity classifications, shown in figures 10, 11 and 12, show a pattern of activity in locations. Laying down occurs mostly in the bedroom, with standing occurring mostly in the living room and kitchen, and walking predominantly occur-

ring in the bedroom and living room. The results suggest that the method of classification has produced meaningful activity classifications and should provide a basis for an expanded activity set in future work.

7 Future work

In future work, a longer period of observation will be analysed using methods identified in this paper. It is anticipated that with a longer period of observation issues such as concept shift or hardware failure may reduce the effectiveness of classifiers for periods of time.

Particularly, the method of RSSI fingerprinting used in this initial case study would not be robust to hardware failure or removal. In future work it will be necessary to develop a continuous retraining strategy, such that should a gateway be unplugged or suffer failure then location classification can recover and learn from the new gateway topology.

SPHERE and HemiSPHERE participants have completed additional surveys to help with annotation of routine and behaviour. In analysing a longer time period, future work will incorporate feedback from a sleep quality survey, daily diaries and Social Rhythm Metric (SRM).

We are currently collecting data from additional participants and pending to apply the same analysis on the collected data. We need to consider how our current analysis pipeline generalises to other house layouts, participants and number of house occupants.

One possible research direction is about how to validate the predicted locations when multiple participants are in the household. In our current analysis it was possible to validate with the PIR activation. However, we will need to incorporate different sensors and heuristics to differentiate between the house occupants.

In our current work activities are limited to walking, stand and lay down. In future work, these activities are to be extended to include sitting down, and climbing and descending stairs. Further activities are facilitated by additional annotation of head mounted camera data recorded during the *day in fast-forward* activity.

To further develop a view of domestic routine, data from smart-meter attached devices such as microwave, toaster, radio and television will be integrated with participant location predictions to highlight patterns of user interaction with domestic appliances.

References

- [Alam *et al.*, 2012] Muhammad Raisul Alam, Mamun Bin Ibne Reaz, and Mohd Alauddin Mohd Ali. A review of smart homes - past, present, and future. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 42(6):1190–1203, 2012.
- [Bao and Intille, 2004] Ling Bao and Stephen S Intille. Activity recognition from user-annotated acceleration data. In *International Conference on Pervasive Computing*, pages 1–17. Springer, 2004.
- [Beswick *et al.*, 2012] Andrew David Beswick, Vikki Wyld, Rachael Gooberman-Hill, Ashley Blom, and Paul

- Dieppe. What proportion of patients report long-term pain after total hip or knee replacement for osteoarthritis? a systematic review of prospective studies in unselected patients. *BMJ open*, 2(1):e000435, 2012.
- [Bolink *et al.*, 2012] SAAN Bolink, SN van Laarhoven, M Lipperts, IC Heyligers, and B Grimm. Inertial sensor motion analysis of gait, sit-stand transfers and step-up transfers: differentiating knee patients from healthy controls. *Physiological measurement*, 33(11):1947, 2012.
- [Gansemer *et al.*, 2010a] S Gansemer, S Puschel, R Frackowiak, S Hakobyan, and U Grossmann. Improved RSSI-based Euclidean distance positioning algorithm for large and dynamic WLAN environments. 9(1):9, 2010.
- [Gansemer *et al.*, 2010b] Sebastian Gansemer, Uwe Grossmann, and Syuzanna Hakobyan. RSSI-based Euclidean Distance algorithm for indoor positioning adapted for the use in dynamically changing WLAN environments and multi-level buildings. pages 1–6. IEEE, September 2010.
- [Janidarmian *et al.*, 2017] Majid Janidarmian, Atena Roshan Fekr, Katarzyna Radecka, and Zeljko Zilic. A comprehensive analysis on wearable acceleration sensors in human activity recognition. *Sensors*, 17(3):529, 2017.
- [Kingma and Lei, 2015] Diederik P Kingma and Jimmy Lei. Adam: A Method for Stochastic Optimization. page 15, 2015.
- [Krumm and Horvitz, 2004] J. Krumm and E. Horvitz. Loco-dio: inferring motion and location from wi-fi signal strengths. In *The First Annual International Conference on Mobile and Ubiquitous Systems: Networking and Services, 2004. MOBIQUITOUS 2004.*, pages 4–13, Aug 2004.
- [Kwapisz *et al.*, 2011] Jennifer R Kwapisz, Gary M Weiss, and Samuel A Moore. Activity recognition using cell phone accelerometers. *ACM SigKDD Explorations Newsletter*, 12(2):74–82, 2011.
- [Muthukrishnan *et al.*, 2009] Kavitha Muthukrishnan, Berend Jan van der Zwaag, and Paul Havinga. Inferring Motion and Location Using WLAN RSSI. In Richard Fuller and Xenofon D. Koutsoukos, editors, *Mobile Entity Localization and Tracking in GPS-less Environments*, volume 5801, pages 163–182. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009.
- [National Joint Registry Editorial Board, 2017] National Joint Registry Editorial Board. National joint registry 14th annual report 2017, 2017.
- [National Joint Registry, 2018] National Joint Registry. Joint replacement statistics, 2018.
- [Preston *et al.*, 2012] Timothy Preston, Wendy Baltzer, and Stewart Trost. Accelerometer validity and placement for detection of changes in physical activity in dogs under controlled conditions on a treadmill. *Research in Veterinary Science*, 93(1):412–416, August 2012.
- [Quan *et al.*, 2010] Michael Quan, Eduardo Navarro, and Benjamin Peuker. Wi-Fi Localization Using RSSI Fingerprinting. page 6, 2010.
- [Rashidi and Mihailidis, 2013] Parisa Rashidi and Alex Mihailidis. A survey on ambient-assisted living tools for older adults. *IEEE journal of biomedical and health informatics*, 17(3):579–590, 2013.
- [Ravi *et al.*, 2005] Nishkam Ravi, Nikhil Dandekar, Preetham Mysore, and Michael L Littman. Activity recognition from accelerometer data. In *Aaai*, volume 5, pages 1541–1546, 2005.
- [Senden *et al.*, 2011] Rachel Senden, Bernd Grimm, Kenneth Meijer, Hans Savelberg, and Ide C Heyligers. The importance to including objective functional outcomes in the clinical follow up of total knee arthroplasty patients. *The Knee*, 18(5):306–311, 2011.
- [Siirtola and Röning, 2012] Pekka Siirtola and Juha Röning. Recognizing human activities user-independently on smartphones based on accelerometer data. *IJIMAI*, 1(5):38–45, 2012.
- [Stevens-Lapsley *et al.*, 2011] Jennifer E Stevens-Lapsley, Margaret L Schenkman, and Michael R Dayton. Comparison of self-reported knee injury and osteoarthritis outcome score to performance measures in patients after total knee arthroplasty. *PM&R*, 3(6):541–549, 2011.
- [Tao *et al.*, 2015] Lili Tao, Tilo Burghardt, S Hannuna, Massimo Camplani, Adeline Paiement, Dima Damen, Majid Mirmehdi, and Ian Craddock. A comparative home activity monitoring study using visual and inertial sensors. In *E-health Networking, Application & Services (Health-Com), 2015 17th International Conference on*, pages 644–647. IEEE, 2015.
- [Twomey *et al.*, 2016] Niall Twomey, Tom Diethel, Meelis Kull, Hao Song, Massimo Camplani, Sion Hannuna, Xenofon Fafoutis, Ni Zhu, Pete Woznowski, Peter Flach, and Ian Craddock. The SPHERE challenge: Activity recognition with multimodal sensor data. *arXiv preprint arXiv:1603.00797*, 2016.
- [Wang *et al.*, 2014] Yujian Wang, Bin Zhao, and Zhaohui Jiang. RSSI-Based Smooth Localization for Indoor Environment. *The Scientific World Journal*, 2014:1–8, 2014.
- [Williams, 2014] Dan Williams. Virtual follow up of hip and knee replacement patients, 2014.
- [Woznowski *et al.*, 2015] Przemyslaw Woznowski, Xenofon Fafoutis, et al. A multi-modal sensor infrastructure for healthcare in a residential environment. In *IEEE Int. Conf. on Communication Workshop (ICCW)*, 2015.
- [Woznowski *et al.*, 2017] Przemyslaw Woznowski, Alison Burrows, Tom Diethel, et al. SPHERE: A sensor platform for healthcare in a residential environment. In *Designing, Developing, and Facilitating Smart Cities*, pages 315–333. Springer International Publishing, Cham, 2017.
- [Xiao *et al.*, 2016] Luo Xiao, Bing He, Annemarie Koster, Paolo Caserotti, Brittney Lange-Maia, Nancy W. Glynn, Tamara B. Harris, and Ciprian M. Crainiceanu. Movement prediction using accelerometers in a human population: Movement Prediction Using Accelerometers in a Human Population. *Biometrics*, 72(2):513–524, June 2016.