Chapter 14
SPHERE: A Sensor Platform for Healthcare in a Residential Environment

Przemyslaw Woznowski, Alison Burrows, Tom Diethe, Xenofon Fafoutis, Jake Hall, Sion Hannuna, Massimo Camplani, Niall Twomey, Michal Kozlowski, Bo Tan, Ni Zhu, Atis Elsts, Antonis Vafeas, Adeline Paiement, Lili Tao, Majid Mirmehdi, Tilo Burghardt, Dima Damen, Peter Flach, Robert Piechocki, Ian Craddock and George Oikonomou

14.1 Introduction

It can be tempting to think about smart homes like one thinks about smart cities. On the surface, smart homes and smart cities comprise coherent systems enabled by similar sensing and interactive technologies. It can also be argued that both are broadly underpinned by shared goals of sustainable development, inclusive user engagement and improved service delivery. However, the home possesses unique characteristics that must be considered in order to develop effective smart home systems that are adopted in the real world [37].

The home is the quintessential personal space and, therefore, there is a greater expectation of privacy at home than in public spaces. People are likely to behave differently when they are at home, in the knowledge that information about their behaviour belongs only to themselves and perhaps others who share this environment with them. But these lived experiences in real-life settings are of great interest to many research fields, in particular those aiming to improve well-being and healthcare provision. Studies conducted in living labs and prototype smart homes can produce information concerning system functionality and usability, but they represent a compromise in terms of the complexity of everyday life. The shift of smart home technologies into real-life contexts has begun to expose a number of

P. Woznowski (✉) · A. Burrows · T. Diethe · X. Fafoutis · J. Hall · S. Hannuna · M. Camplani · N. Twomey · M. Kozlowski · B. Tan · N. Zhu · A. Elsts · A. Vafeas · A. Paiement · L. Tao · M. Mirmehdi · T. Burghardt · D. Damen · P. Flach · R. Piechocki · I. Craddock · G. Oikonomou
Faculty of Engineering, University of Bristol, Merchant Venturers Building, Woodland Road, Bristol BS8 1UB, UK
e-mail: p.r.woznowski@bristol.ac.uk

V. Angelakis et al. (eds.), *Designing, Developing, and Facilitating Smart Cities*,
DOI 10.1007/978-3-319-44924-1_14

barriers to a wider adoption of these systems, including high cost of ownership, inflexibility, poor manageability and difficulty achieving security [9].

In reality, homes are very dynamic environments. They vary in typology as well as layout, and can change sporadically as residents refurbish interior spaces. This can create a variety of challenges ranging from the initial system installation to signal propagation and data analytics. Homes contain material possessions, but they also have personal meaning for the people who live in them. They can be home to a single person or to multiple people, with different individual characteristics. If we think about several generations living under a single roof, it is likely that each individual has different abilities and motivations in terms of interacting with a smart home system. Previous research about how people create a smart home identified three roles of household members in relation to their smart homes: home technology drivers, who were primarily involved in planning and maintaining the technology; home technology responsibles, who did not deal with the technology directly but wanted it installed and would outsource necessary repairs and adjustments; and passive users who were removed from any phase of the home automation, but had learned the basics about the system by using it [36]. Overall, people have a greater expectation of control over technology at home than in public spaces. The home can also be a place of conflict and negotiation between the people who live there, which can affect how domestic technologies are adopted and used.

The need to restructure healthcare services is widely acknowledged and has led to the home being viewed as a key setting for health and care. Perspectives on the role of the patient in preventing and managing chronic illness have shifted from the self-management approach of conventional medicine towards an approach that sees patients, healthcare professionals and the wider community working together to develop holistic and personalised care plans. Within this context, smart home technologies have emerged as one way to empower patients to actively engage in the management of their well-being. Achieving a system that is technically feasible and clinically effective requires a multidisciplinary approach that combines the expertise of various stakeholders, including end users who will steer the design towards an acceptable outcome. The domestic end users of such a smart home system may comprise healthy individuals, individuals who are living with one or more chronic health conditions, and individuals who experience an acute disease or injury. Furthermore, people's health and care needs change suddenly or progressively throughout their lives.

The challenges of developing smart home technologies for health and care become evident as we begin to break down the various facets of the home and the diversity of its residents. This type of rich understanding of potential users has been used to develop inclusive design resources to inform the development of appropriate smart home technologies for health and care [10]. The remainder of this chapter details the development of a smart home system, which was designed to be retrofitted into real homes. We begin with a description of the SPHERE project, under which this research was conducted.

14.1.1 Overview of SPHERE

SPHERE (Sensor Platform for Healthcare in a Residential Environment) is an EPSRC-funded interdisciplinary research project, led by the University of Bristol in collaboration with the University of Reading and the University of Southampton. The overall aim is to develop a smart home platform of non-medical networked sensors, capable of gathering and integrating multiple types of data about the home environment and the behaviours of its residents to better understand a range of healthcare needs. Rather than targeting subsets of the population based on demographics or health conditions, the project takes an inclusive approach with a view to generating rich data sets. We anticipate that analysis of these data sets will generate evidence-based insights into factors that affect health and well-being, thus informing more appropriate and effective interventions.

The system comprises various sensors that can be broadly grouped into three categories: environmental sensors, which monitor temperature, humidity, luminosity, noise level, air quality, room occupancy, electricity metering and cold and hot water consumption; vision sensors, which are able to track people and provide information about quality of movement; and wearable sensors, primarily a low-power wrist-worn device that uses accelerometers to measure patterns of movement. A prototype of this system is installed in a two-bedroom Victorian residential property in Bristol, which serves as a test house for short- to long-term user studies. We felt it was important to test the system in a realistic setting, in a familiar and otherwise unremarkable domestic environment. Among other things, this has allowed researchers to experience some of the technical challenges of retrofitting this system into real homes. Once this system has been thoroughly tested and iterated, we plan to deploy it in up to 100 homes in Bristol for long-term and 'in the wild' studies.

14.2 Enabling Technologies

This section briefly describes the different sensing modalities used by the SPHERE system, namely, the wearable sensor, the environmental sensors and the video monitoring system.

14.2.1 The Wearable and Environmental Sensors

The SPHERE architecture monitors the residential environment using a custom environmental sensor board, named SPES-2. SPES-2 (shown in Fig. 14.1) is a battery-powered sensor board that is based on the Texas Instruments CC2650 System-on-Chip (SoC) for processing and wireless communication. The CC2650 is a multi-standard 2.4 GHz wireless system; it supports Bluetooth Low Energy (BLE) and IEEE 802.15.4. A meandered 2.4 GHz monopole antenna is printed on

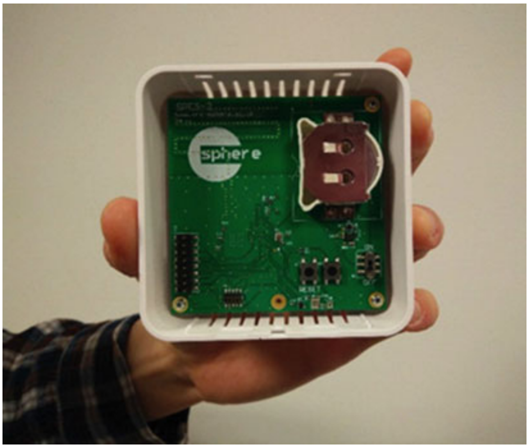

Fig. 14.1 The SPHERE environmental sensor

the board. For processing, the CC2650 incorporates an ARM Cortex-M3 microcontroller unit (MCU).

The SPES-2 incorporates a series of sensors for monitoring the residential environment. These include: a temperature and humidity sensor (HDC1000); a light sensor (OPT3001); a barometer (BMP280); a passive infrared motion sensor (EKMB1101); and a microphone (SPH0641LU4H-1) used for noise level sensing. Moreover, SPES-2 exposes an interface for connecting external sensors. Any low-power analogue or digital 3.3 V sensor is compatible. The board is powered by a 3 V CR2477 coin cell battery (typical capacity of 1000 mAh). The physical dimensions of the board are 75 × 75 × 1.6 mm, enclosed in an off-the-shelf casing (dimensions 85 × 85 × 25 mm). In addition to the in-house developed environmental sensors, the SPHERE system also incorporates a commercial electricity monitoring system by CurrentCost.

The SPHERE infrastructure also incorporates a custom activity tracker. The SPW-2 (Fig. 14.2) is a wearable sensor board, mounted on the user's wrist. Similar

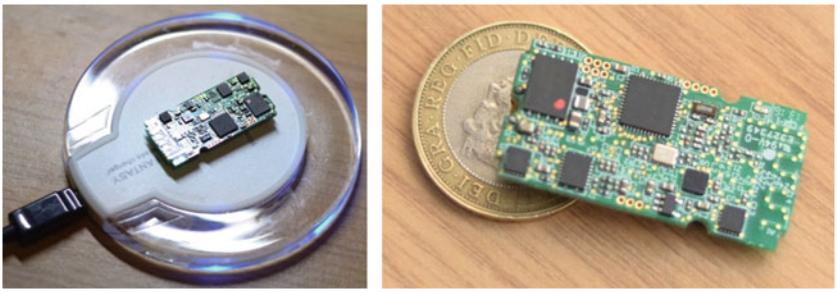

Fig. 14.2 The SPHERE wearable. *Left* While charging. *Right* Physical dimensions

to the environmental sensor board (SPES-2), SPW-2 is based on the CC2650 multi-standard 2.4 GHz SoC. Similarly to its predecessor [15], it employs two ADXL362 accelerometers. The accelerometers are separated by 30 mm, and are aligned with the user's limp in such a way, so that differential measurements provide rotational information on the limp's movements [57]. Additionally, the board incorporates a low-profile 2.4 GHz inverted-F antenna, and a flash memory of 512 MB for temporary data storage, whilst the user is out of the house.

The board is powered by a 100 mAh Lithium–Polymer (Li–Po) rechargeable battery (3.7 V). A Qi compatible inductive contact charging circuit [27] is incorporated to wirelessly charge the battery. Commercial off-the-shelf charging pads are compatible. The physical dimensions of the board are 39 × 20 × 1.6 mm, enclosed in a custom wristband (dimensions 41 × 22 × 8 mm).

The final component of the SPHERE infrastructure for environmental and body sensing is a mains-powered gateway, named SPG-2. SPG-2 employs two CC2650 sub-systems, which can be used either for simultaneous BLE and IEEE802.15.4 support or for implementing antenna diversity on the same standard to improve wireless coverage, as proposed in [17].

The gateway board incorporates a subset of the sensors of the environmental sensor. These include: a temperature and humidity sensor (HDC1000); a light sensor (OPT3001); a barometer (BMP280); and a microphone (SPH0641LU4H-1) used for noise-level sensing. The board is mains powered via USB and the same interface can also be used for programming and software debugging. Its physical dimensions are similar to those of the SPES-2.

14.2.2 Video Monitoring

The video monitoring component of the SPHERE system is a real-time multi-camera system, which is tasked not only with tracking individuals navigating their home environment, but also with providing continuous quality of movement information. Tracking information is provided in the form of 3D bounding boxes in world coordinates and quality of movement information is currently in the form of the log-likelihoods of a particular movement being 'normal' [43].

In the SPHERE platform, integration with other sensing modalities, user acceptance and deployment budget are primary considerations. To make deployment into the local community financially viable, it has been necessary to limit hardware selection to the low cost consumer RGB-D camera, Asus Xtion.[1] This camera needs to be coupled with a machine with suitable processing capacity, minimal intrusion to the user and minimal cost. The Intel Next Unit of Computing (NUC) with an i5 processor and 8 GB of RAM fulfils these requirements (Fig. 14.3).

[1] https://www.asus.com/3D-Sensor/Xtion/.

Fig. 14.3 NUC and Asus Xtion sensor in the SPHERE house

The small size and relatively low cost of the NUC, when compared to most workstations, allows it to be strategically placed in close proximity to other sensors. Another key feature of the NUC is its four USB 3.0 ports, which provide enough bandwidth to capture from four depth cameras simultaneously. This gives the system far more flexibility in deployment. Specifically, a range of configurations is possible: one NUC operating all of the cameras, one NUC per camera, or some configuration in between, depending on user preference as well as the specific circumstances of each individual deployment.

Another critical consideration for all of the SPHERE sensors and particularly the video subsystem is installation overhead and long-term reliability. Consider that SPHERE currently plans to deploy the system into up to 100 homes in the local community. With up to three NUCs and ASUS Xtions per home, the setting up and management of such a large enterprise rapidly becomes intractable without streamlined installation protocols and extremely reliable subsystems.

In light of these requirements, each of the NUCs runs GNU/Linux with the video system configured as a service, which is automatically launched when the machine boots. This makes the video system resistant to temporary power loss. A standardised system image is used to configure the system so that a single NUC can be unboxed and setup ready for deployment in around 10 min.

As mentioned previously, software reliability is critical to facilitate extended operational periods without being physically accessed. To aid this, we have used object-oriented design principles and common software design patterns where appropriate.

To support the collection of video data we have designed a pair of classes which capture all of the functionality needed in the video system: the Camera class and the CameraObserver class. The first of these encapsulates camera functionality providing depth, colour, bounding box and skeleton information. Instances of the CameraObserver class register themselves with a Camera to be notified when a new frame is available, with its associated data. The observer is then free to choose what to do with the data. This provides the system with a unified method to implement specific video experiments without understanding the underlying hardware configuration.

To integrate the video system into the SPHERE platform, we implemented a subclass of the CameraObserver which takes the video data, excluding the frames themselves, and serialises them into a JSON string. These strings are then transmitted over MQTT (Message Queue Telemetry Transport) protocol on the appropriate topic. This allows the video system to act as an IoT device to the rest of the SPHERE sensing platform.

Even considered in isolation, the output of the video system is a rich source of information. The velocity, aspect ratio and location of the bounding boxes provide important clues for measuring activities and behaviours. For example, when combined with environmental contextual information, such as the location of kitchen appliances, activities such as cooking, washing up and watching television can potentially be identified. Moreover, it may also be possible to determine general activity levels including the amount of time spent sedentary. Individuals are tracked using state-of-the-art algorithms such as [38]. The bounding boxes obtained may be further exploited to yield coarse estimates of human pose. For example, an approximately square bounding box is indicative of sitting, whilst a vertical elongated rectangle implies standing.

Quality of movement assessment is based on the skeletons provided by the PrimeSense[2] middleware. These are normalised for the global positioning and orientation of the camera and height variation. The relatively high dimensionality of the normalised skeletons is reduced using a modified version of Diffusion Maps [11], where Gerber's [22] method for addressing outliers in Laplacian Eigenmaps is exploited. The resulting high level feature vector, obtained from the normalised skeleton at one frame, represents individual poses and is used to build a statistical model of normal movement. Abnormal movement patterns are detected by their deviation from this model.

14.3 Overall System Architecture

The SPHERE system will be deployed in up to 100 homes in and around Bristol for long-term and 'in the wild' studies. Each deployment will consist a number of environmental, wearable and video sensors (Sect. 14.2), accompanied by a number of devices required for (i) data storage, (ii) network connectivity among sensors, and (iii) system management and monitoring.

The SPHERE system also consists a back end (called the "SPHERE Data Hub"), which is made up of a number of virtual machines, servers and storage devices physically situated at the University of Bristol. The Data Hub is used for data analytics and it also provides a system administration dashboard. It will lastly be used for long-term storage of all data collected from the 100 properties.

Each deployment property will be connected with the Data Hub through a secure Virtual Private Network (VPN) over a 3G, 4G, or fixed broadband link. On the deployment property, this VPN tunnel will terminate on a device called 'SPHERE Home Gateway', an Intel NUC PC identical to the one used by the video monitoring subsystem (Sect. 14.2.2). The overall system architecture is illustrated in Fig. 14.4.

[2]http://www.i3du.gr/pdf/primesense.pdf.

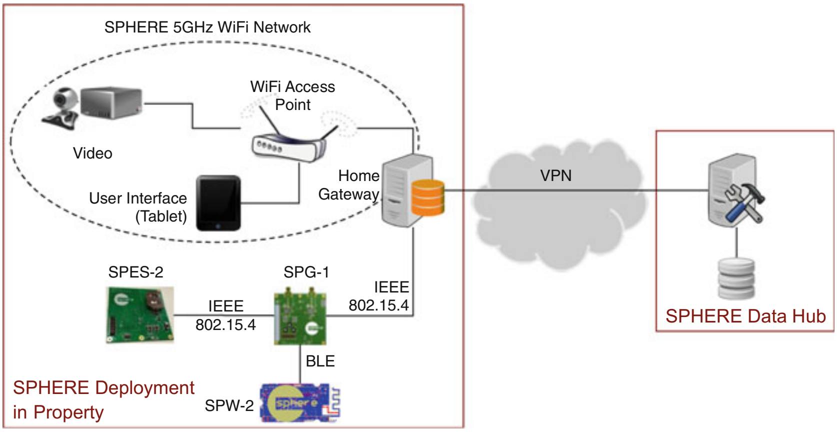

Fig. 14.4 Overall architecture of the SPHERE system

In addition to terminating the VPN tunnel between the University and the property, the Home Gateway serves a number of purposes:

- It provides a reliable, redundant and secure data storage medium for sensor data collected by the wearable, environmental and video sensors,
- It provides a time synchronisation source for all other SPHERE systems in the property,
- It hosts a dashboard that can be used by participants to visualise data, monitor and manage (e.g. start or stop) the system. The dashboard is presented to participants over a web-based interface through a pre-configured tablet computer.

The SPHERE system generates data that falls within one of two categories:

- **Sensor Data**: Sensor measurements collected by the wearable, environmental and video sensors. This category, for example includes wearable acceleration data, ambient light levels, environmental temperature, water or electricity usage, presence detection data and video bounding boxes.
- **Control and Monitoring Data**: Data used to monitor the system's overall state and to manage its individual components. This category includes, for example, network statistics, version of the software running on each device, device uptime, battery levels for battery-powered devices and more. In terms of managing the system, data within this category can be used to, for example, stop and restart the system or to install software updates.

Noteably, *Sensor Data* are stored on a Mongo database hosted at the Home Gateway, but—due to the sheer volume and also for security and data privacy purposes—they are *not* transmitted to the Data Hub over the network. Conversely, Control and Monitoring Data are transmitted to the Data Hub over the VPN link, in order to allow remote system management, administration, monitoring and generation of alerts to notify about system malfunctions.

In terms of network connectivity, a SPHERE deployment is conceptually broken down into two logical segments within each property. The wearable and environmental sensors communicate among themselves and with the Home Gateway using a combination of BLE and IEEE 802.15.4 wireless links. Both of those technologies operate in the 2.4 GHz radio frequency band. A 5 GHz WiFi network is used to ensure communications among the Home Gateway, Video NUCs and the tablet that provides users with the SPHERE User Interface. The choice of using the 5 GHz band was made for two reasons: (i) increased network performance, which is required to accommodate the large volume of video data; and (ii) prevention of interference with the 2.4 GHz low-power networking used by the wearable and environmental sensors. The entire SPHERE system is thus fully isolated from the user's own home network.

At the application layer, the SPHERE system makes extensive use of the MQTT protocol for data collection as well as for system monitoring, with an MQTT broker installed on the Home Gateway at every property. An MQTT client at the Data Hub is used to collect monitoring data and to issue control commands to all systems. Lastly, an MQTT client is used on every Video NUC in order to: (i) publish sensor data to the Home Gateway, (ii) publish monitoring data to the Home Gateway and the Data Hub, and (iii) receive management commands from the Home Gateway or the Data Hub.

For data collection from and management of the environmental sensor (SPES-2), we use the Constrained Application Protocol (CoAP) instead. Translation between MQTT and CoAP is achieved with a SPHERE-developed application layer proxy running on the Home Gateway. Lastly, the SPHERE wearable uses BLE to communicate with the rest of the infrastructure via an SPG-2 device. This SPG-2 encapsulates wearable sensor and monitoring data into CoAP packets before submitting them to the Home Gateway for processing and storage. The inverse process is followed in order to send control requests to the wearable.

14.4 Data Analytics and Interpretation

Ambient Intelligence (AmI) spaces process large quantities of *sensor data* and require robust and accurate Activity Recognition (AR) strategies. This is needed to infer activities of interest to monitor health, well-being or other personal benefits such as fitness level. A typical approach to the problem is to define a hypothesis to test which informs experimental design. Quantification of the hypothesis can be done either by simulation or by processing real sensor data. Experimental design also involves the selection of appropriate data collection methods, ground truth acquisition methods and annotation strategies.

14.4.1 Ground Truth

The validation of AR strategies involves comparing their output against ground truth/benchmark data. But before this can take place, this data has to be acquired. Ground truth data acquisition involves three stages: collecting data upon which to base the ground truth, deciding what labels are appropriate to describe the data and applying these labels annotating the data to obtain the ground truth. These stages are often developed as a result of an iterative process to determine the best data collection methods, most appropriate labels to use, and a suitable way to annotate the data to produce consistent and informative ground truth.

Obtaining annotations from self-reported diaries is imperfect as they rely purely on participant's compliance and their subjective perception and memory, which in general, becomes inaccurate with time. This is particularly the case with annotation which requires accurate temporal precision in order to be maximally effective, so it is unrealistic to expect detailed activity diaries with the exact timings of the activities. Researchers approach this problem in many different ways. Allen et al. [1] collected unsupervised activity data in the home using a computer set up to take participants through a routine. Input from the user was in the form of a button press from which the data was annotated. Van Kasteren et al. [52] asked participants to wear a Bluetooth headset that used speech recognition to label ground truth data. This method is inexpensive, but is of limited utility because it does not capture enough detail and contextual information.

Another strategy is for the researcher to record the activity and context during data collection [35, 42]. Pärkkä et al. [42] adopted a semi-overseen approach to collecting data for AR classification based on realistic activities. A single researcher followed the participant during the experiment and used an annotation app to record the activities. Even with this approach it was noted that there were annotation inaccuracies that were most likely correlated to predictive inaccuracy.

Methods using video recordings provide an objective reflection of participant's activities enabling a far more accurate and detailed activity ground truth data, however, these require additional attention in the form of ontology. For example, Atallah et al. [2] used video to annotate activities during laboratory experiments to train AR classifiers. Data can also be collected in an unobserved environment, encouraging natural behaviour; however, it can also be perceived as intrusive and will only capture actions with no room for participant interpretation.

Video annotation can be costly and time-consuming. Active learning is a technique that can reduce the amount of annotated data needed for training a classifier. In this approach, classifiers are to be trained with a minimal set of annotated data. Active learning algorithms attempt to quantify the utility of obtaining labels for new instances tensioned against the financial cost of querying an oracle for the true label. Only the instances that are deemed to yield maximal utility are selected. In particular, when coupled with transfer learning techniques (i.e. transferring knowledge from different contexts to a new context), active learning can dramatically reduce the quantity of labelled data required [12]. Hoque and Stankovic [26] employed a clustering technique to group smart home environmental data into activities and the user labelled the clusters. Another application for active learning, is to update classifiers or personalise them [33]. While attractive, these methods rely on a collaborative effort on the part of the user.

There are a number of available software tools which are suitable for video annotation, such as the ANVIL video annotation tool [29] or ELAN [8] developed by the Max Planck Institute for Psycholinguistics, The Language Archive, Nijmegen, The Netherlands.

The labels used to annotate data is another annotation consideration. Labels are often application specific, e.g. [42] used a hierarchical list of labels, aimed at capturing the context of the activities, whereas [49] focused purely on a specific disease and the associated symptoms. Logan et al. [32] presented a detailed activity ontology for the home using a custom tool that enabled annotators to label foreground and background activities for when the participant's attention is focused on another activity, addressing the fact that humans naturally multitask. Roggen et al. [44] used a four 'track' annotation scheme for annotating human activities based on video data including tracks for locomotion, left- and right-hand activities (with an additional attribute that indicates the object they are using), and the high-level activity.

In the SPHERE project, machine learning algorithms are initially trained and validated against recordings from a head-mounted video camera worn by participants. The data originates from the three different sensing modalities (depth cameras, wearable accelerometer and environmental sensors) deployed in a real house, which constitutes the testbed [56, 58]. The same SPHERE ontology of ADLs, underpins system-generated activity data and the controlled vocabulary used in video annotation.

14.4.2 Activity Recognition

In order to make instantaneous or longitudinal inferences about the health status of individual residents, a necessary first step is to be able to recognise normal Activities of Daily Living (ADL). To this end, we need a framework that allows us to take the inputs from multiple heterogeneous sensor sources, such as those described in Sect. 14.3, and make informed decisions that are tailored both to the individual and the context. Naturally, the SPHERE setting presents many sources of uncertainty. First, we are dealing with multiple sensor modalities (environmental, body-worn, video), each of which will have different noise profiles and failure modes. Second, as described in Sect. 14.4.1, we are dealing with a situation where annotated or labelled data is expensive and intrusive to acquire, and the resulting labels are potentially noisy and inaccurate (indeed in some cases there may be no *ground truth* in the classical sense, and we need to resort to modelling annotator disagreement explicitly). Lastly, patterns of human behaviour are subject to many factors (internal and external) that may or may not be attributed to the particular health context of a given individual.

Faced with such a situation, the most sensible approach would be to use *white box* modelling methods where possible. Model-based machine learning [6, 54] attempts to follow this ideal by encoding assumptions about the problem domain explicitly in the form of a model. Indeed, the model can be viewed simply as this set of assumptions, expressed in a precise mathematical form. These assumptions include the number and types of variables in the problem domain, which variables affect each other, and what the effect of changing one variable is on another variable. The result is that any decisions made by the system can be inspected, so that if the model is performing poorly, the solution is to re-examine the assumptions being made.

In the Bayesian paradigm, degrees of belief in states of nature are specified through the use of probabilities, which through the construction of probabilistic graphical models [30] allow us to apply a principled mathematical framework of the quantification of uncertainty to perform model-based machine learning. On the basis of the models we build, Bayesian decision theory tries to quantify the trade-off between various decisions, making use of probabilities and costs [3, 4].

A typical problem that we face is that the differences between individuals are too large to be captured by a single model. Hierarchical Bayesian models [21] allow us to simultaneously generalise over communities of residents whilst also learning personalised models. In addition, they allow us to be more flexible with our priors, by specifying *hyper-priors*, and then performing inference over the priors instead.

Adapting to multiple operating contexts

However, transparently dealing with degrees of belief does not solve all modelling challenges posed by the SPHERE project. Our models and inferences have to be applied in multiple contexts, and indeed any given context is liable to both gradual and abrupt shifts. Let us consider the example of modelling daily patterns of behaviour. A common approach for coping with the temporal aspect of daily patters is to introduce an *hour of day* feature to the classification model [28, 51]. However, when summarising the temporal nature of an activity into a coarse feature such as this, not only is information lost after discretisation, but also the strength of the periodicity of the action is ignored. Bayesian approaches have been ascribed to such periodic data [13] and these models can not only capture the complex multi-modal aspects of periodic patterns, but the resulting posteriors are interpretable and may be studied to increase practitioners understanding of the nature of their data.

In such situations, it will be crucial that we are able to trust the probabilities coming from the system. A machine learning system is well *calibrated* if the predicted probabilities it gives correspond to observed frequencies. This is natural in forecasting where we would expect it to rain in 60 % of days where a weather forecaster predicts a 60 % chance of rain [39] but carries over to machine learning as well. If a system is poorly calibrated then it suggests a problem either in the model (such as an overly restrictive assumption) or in the inference.

Different operating contexts call for different performance metrics, which perhaps incorporate a different notion of expected loss [25]. If the goal is to minimise loss, for example for the case of classification, a systematic approach would be that given a model, threshold choice methods that correspond with the available information about the operating condition should be applied, followed by comparison of their expected losses.

The operating contexts of smart environments can be affected by many factors. For example, it is well known that when multiple occupants reside in a smart home, models that do not adapt to these contexts will generally yield poor predictions. This happens because 'confusing' sensor data can arise when co-occurring activities are performed in different locations of the home. However, if the topology of the residence can be learnt, not only will predictive performance be boosted, but predictive confidence and calibration will likewise be improved [16, 50].

Different classification performance metrics such as F-score also imply a different notion of calibration [18]. More generally, the choice of the performance metrics in use should be seen as another modelling assumption rather than being independent from the model. Given that we expect the end users of our systems to include medical professionals as well as the residents themselves, we can easily see how the types of decision we would want to surface should be adaptable.

Explicitly modelling context change also favours domain adaptation and model reuse. We are building on the results of the REFRAME project,[3] which developed a general methodology for model reuse in machine learning called *reframing* (Hernández-Orallo, Prudêncio et al. (in press) [24]). The setting is exemplified by the recent ECML-PKDD'15 Discovery Challenge *MoReBikeS: Model Reuse with Bike rental Station data,*[4] which encouraged participants to build predictive models for new bicycle rental stations making use of previously trained models on other stations (for which the training data was however no longer available).

14.4.3 Localisation

Research [31] has shown that human activities in residential areas are highly correlated with their corresponding locations. Activities pertaining to SPHERE's research interests mostly occur indoors, thus rendering the traditional global navigation satellite systems like GPS or Galileo is redundant. Instead, a localisation solution which can indicate the relative indoor position is essential for future research opportunities.

Academia and industry have both tried various approaches to tackle indoor localisation problem. GE[5] and Philips[6] offer the LED-based indoor localisation

[3]http://reframe-d2k.org/.

[4]http://reframe-d2k.org/Challenge.

[5]http://www.gelighting.com.

[6]http://www.lighting.philips.co.uk/systems/themes/led-based-indoor-positioning.html.

system for use in retail outlets or hospitals. Video systems, such as MS Kinect or Intel RealSense, can also provide tracking information when the target is recognised within the effective region, but only if proper light condition is preserved. In scenarios where accurate location information is not necessary, the outputs from passive infrared [40] or sound sensors [5] are used as an indication of the subject's presence in certain areas. In addition, accelerometer and gyroscope [23, 55] data are also used for keeping track of the subjects, if the initial location is known—this method is known as one embodiment dead reckoning. It cannot, however, be considered an ideal solution to provide indoor location information yet, much like GPS system is now, for outdoor applications. The above approaches are used for different purposes in different contexts, and as such have drawbacks and limitations. These include arduous installation and deployment, low accuracy, high cost, limited coverage, and depending on the mentioned differing contexts, intruding upon subject's privacy. Thus, in relation to SPHERE, other indoor localisation methods are considered.

Nowadays, with an ever-increasing use of wireless systems, radio frequency (RF) signals are present everywhere penetrating all living spaces, including residential homes. Modern RF receivers allow for the parameters such as time delay, power strength, profile distortion, and even experienced reflections, to be accurately measured. These parameters can be used to estimate the distance between the transmitter and receiver. The distance estimations from spatially differentiated locations then are used for localising the target by triangulation or multilateration. In Bose and Foh [7] and Wang et al. [53] received signal strength (RSS) based ranging methods are described in detail. Generally an RSS-based method requires a relatively simple propagation environment, in order to avoid signal superimposition caused by multipath propagation. Thus, a high-density residential area creates a very challenging scenario for the application of this approach. Time is another parameter that can be used to indicate the distance. In both, academia and industry, researchers have tried to extract timing by specially designed wide band signals, such as those shown in Sahinoglu et al. [1, 19, 45]. Günther and Hoene [20] and Ciurana et al. [34] introduced a round trip measuring method based on the ACK mechanism of the 802.11 protocol with standard commercial off-the-shelf (COTS) devices. This method avoids multipath propagation problems, but is still limited by low clock resolution and stability of COTS devices. Exel [14] presents an improvement to the clock resolution problem by extracting the wireless communications signals from a dedicated receiver—this, however, is much more expensive than using COTS devices. By now it is apparent, that the above mentioned methods are either too expensive or inaccurate for widespread deployment.

Regarding the applications and requirements of SPHERE, the technology which can leverage existing wireless signals and cheap COTS equipment are very much preferred. Also, differing from the industrial applications such as the storehouse or assembly line robotics, which have hard requirements on localisation accuracy and resolution, there are no exact firm requirements, when referring to human activity research. Therefore, the development of the indoor localisation system for SPHERE includes two stages: the *premier* stage and the *advanced* stage. The Premier stage takes advantage of resources already deployed in the house—for example Bluetooth (BLE) access points (APs) and PIR sensors in each room. This provides us with room-level location information given, for example the optimum amount of BLE AP's spread around the house. According to the SPHERE's signal propagation research, three access points in the house at any one time are sufficient for this purpose [48]. Room-level information can provide only limited data to differentiate between activities. For instance, the detection of presence in the living room is associated with activities such as watching TV, reading, or chatting, and excludes actions not normally associated with this location, such as making tea, washing or taking shower. However, room-level information is still not fine-grained enough to aid recognition of some ADLs. In the context of multiple sensor platforms used in the SPHERE system, the localisation can be further strengthened by fusion of the other sensors' data—for example the on-body accelerometer, video or even electrical and water metres. Another approach is to deploy extra RF sensors (BLE in the SPHERE's context) in the house in order to better calculate the location information (*advanced* stage). The aim for this stage is to provide between $1 \times 1\ \text{m}^2$ and $1.5 \times 1.5\ \text{m}^2$ resolution location information based on the high density of RF receivers.

The implementation of the premier and advanced stages of the indoor localisation in the SPHERE prototype testbed are as follows:

Premier Stage: The SPHERE localisation system in this stage includes one custom wearable [15] and three distributed receivers [17]. The wearable broadcasts BLE advertisements in channel 37, 38 and 39 with 4 dBm emission power. Each receiver is equipped with one horizontal polarised folded dipole antenna and one vertical polarised folded dipole antenna for error correction purpose. On each receiver, wearable BLE advertisements are sniffed, timestamped, and subsequently saved into the database. In the database, the RSSs of the same BLE advertisements captured by different receivers are extrapolated together using the timestamps. Not all BLE advertisements can be received successfully by all receivers due to the signal attenuation caused by extended distance and obstacles in the propagation path. Thus, the received RSS samples are resampled using Pandas library. The data used for localisation experiments was acquired through scripted data collection. This script consists of a representative sample of ADLs occuring in every location in the testbed house. Classification methods, such as k-nearest neighbours (KNN) and support vector machine (SVM), were applied to this dataset. These classifiers provide around $80 \sim 85\ \%$ correct room-level localisation recognition with only three BLE receivers in the house [16]. Subsequently, the same data set was analysed using a Hidden Markov Model (HMM) approach which builds the time sequential relation between locations. HMM provides similar levels of correction rates. By comparing the characteristics of the errors in classification methods and the HMM, we surmise that the erroneous classifications occurs in a random spike manner while the errors in HMM method manifested themselves mostly as bursts. If we define the change of the location as one event, HMM shows much less spurious events than classification methods. Hence, mutual calibration between the classification and time sequential method is necessary to improve the localisation performance.

Advance Stage: In this stage, addtional BLE APs are deployed. Approximately two additional APs are installed in each room. Each AP is constituted by one Raspberry Pi and one COTS Bluetooth dongle. The APs are installed around the ceiling to mitigate the body shadow effect [41] and loitering personnel interference. The house is divided and labelled into 82 grids which are roughly 1×1 m^2. As the grid size is very small, it is very difficult for classification methods to cope with slight RSS difference between neighbouring grids. Thus, the localisation in this situation is more reliant on the time sequential relation.

Passive Sensing

There are multiple ways of approaching the localisation challenge in a sensor-rich, multi-modal setup. Methods based on simple binary sensors (e.g. PIR) are limited to single occupancy scenarios and can only provide room-level accuracy. On the other hand, RF-based approaches relying on change in RSSI, not only depend on participant wearing an RF transmitter (e.g. SPHERE wearable) but also require high-density infrastructure of RF receivers. Thus, passive sensing [47] is considered as one of the most likely candidates to provide fine-grained location information in a residential context. It originates from radar technology, and can then be extended to civilian applications by taking advantage of the RF signal already present in residential areas. As demonstrated in [46], passive sensing technology can quantify the precise distortion of human reflected RF signal which is related to the object moving along its trajectory. By synthesising the quantified signal distortion into the predefined classifiers or machine learning modules, pose and location information can be extracted. The successful implementation of passive sensing will lead to a device free solution for collecting location and even activity information in spaces where the wireless signals are presented. Passive sensing is thus an important, albeit budding, ambition of academia and industry, and requires large amounts of testing and research to be fully utilised. Incomplete as it may be though, it is nonetheless a rising research topic due its potential performance and applications in healthcare, security and entertainment.

Acknowledgement This work was performed under the SPHERE IRC, funded by the UK Engineering and Physical Sciences Research Council (EPSRC), Grant EP/K031910/1.

References

1. Allen FR, Ambikairajah E, Lovell NH, Celler BG (2006) Classification of a known sequence of motions and postures from accelerometry data using adapted Gaussian mixture models. Physiol Meas 2006:935
2. Atallah,L, Lo B, Ali R, King R, Yang G-Z (2009) Real-time activity classification using ambient and wearable sensors. IEEE Trans Inf Technol Biomed Publ IEEE Eng Med Biol Soc 13(6), 1031–1039

3. Berger JO (1993) Statistical decision theory and Bayesian analysis, 2nd edn. Springer-Verlag, New York, p 1993
4. Bernardo JM, Smith AFM (2008) Bayesian Theory. John Wiley & Sons, Hoboken, NJ, p 2008
5. Bian X, Abowd GD, Rehg JM (2005) Using sound source localization in a home environment
6. Bishop CM (2013) Model-based machine learning. Phil Trans R Soc A
7. Bose A, Foh CH (2007) A practical path loss model for indoor WiFi positioning enhancement. Inf Commun Signal Process
8. Brugman H, Russel A (2004) Annotating multi-media/multi-modal resources with ELAN. In: Proceedings of the 4th International Conference on Language Resources and Language Evaluation (LREC 2004). Lisbon, 2004, pp 2065–2068
9. Brush AJ, Lee B, Mahajan R, Agarwal S, Saroiu S, Dixon C (2011) Home automation in the wild: challenges and opportunities. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, pp. 2115–2124. ACM
10. Burrows A, Gooberman-Hill R, Coyle D (2015) Empirically derived user attributes for the design of home healthcare technologies. Pers Ubiquit Comput 19(8):1233–1245
11. Coifman RR, Lafon S (2006) Diffusion maps. Appl Comput Harmon Anal 5–30, (Elsevier)
12. Diethe T, Twomey N, Flach P (2016) Active transfer learning for activity recognition. In: 24th European Symposium on Artificial Neural Networks. Bruges: ESANN
13. Diethe T, Twomey N, Flach P (2015) Bayesian modelling of the temporal aspects of smart home activity with circular statistics. Mach Learn Knowl Discov Databases, 279–294. Springer International Publishing, Porto
14. Exel R (2012) Receiver design for time-based ranging with IEEE 802.11b signals. Int J Navig Obs
15. Fafoutis X, Janko B, Mellios E, Hilton G, Sherratt S, Piechocki R, Craddock I (2016) SPW-1: a low-maintenance wearable activity tracker for residential monitoring and healthcare applications. Int Conf Wearables Healthc (HealthWear). EAI
16. Fafoutis X, Mellios E, Twomey N, Diethe T, Hilton G, Piechocki R (2015) An RSSI-based wall prediction model for residential floor map construction. In: Proceedings of the 2nd IEEE World Forum on Internet of Things (WF-IoT). IEEE
17. Fafoutis X, Tsimbalo E, Mellios E, Hilton G, Piechocki R, Craddock I (2016) A residential maintenance-free long-term activity monitoring system for healthcare applications. EURASIP J Wirel Commun Netw 2016, 23
18. Flach PA, Kull M (2015) Precision-recall-gain curves: PR analysis done right. In: Proceedings of the Twenty-Ninth Annual Conference on Neural Information Processing Systems. NIPS
19. Fontana RJ, Gunderson SJ (2002) Ultra-wideband precision asset location system. Ultra Wideband Systems and Technologies, Baltimore
20. Günther A, Hoene C (2005) Measuring round trip times to determine the distance between WLAN nodes
21. Gelman A, Carlin JB, Stern HS, Rubin DB (2013) Bayesian data analysis, 3rd edn. Chapman and Hall, London
22. Gerber S, Tasdizen T, Whitaker R (2007). Robust non-linear dimensionality reduction using successive 1-dimensional Laplacian eigenmaps. In: Proceedings 24th International Conference on Machine learning, pp 281–288. ACM
23. Harle R (2013) A survey of indoor inertial positioning systems for pedestrians. IEEE Commun Surv Tutor 15(3):1281–1293
24. José H-O et al, Reframing in context: a methodology for model reuse in machine learning. AICOM, (in press)
25. Hernández-Orallo José, Flach Peter, Ferri Cèsar (2012) A unified view of performance metrics: translating threshold choice into expected classification loss. J Mach Learn Res 13 (1):2813–2869
26. Hoque E, Stankovic J (2012) AALO: Activity recognition in smart homes using Active Learning in the presence of Overlapped activities. In: Proceedings of the 6th International Conference on Pervasive Computing Technologies for Healthcare. IEEE, pp 139–146

27. Hui SY (2013) Planar wireless charging technology for portable electronic products and Qi. Proc IEEE 101(6):1290–1301
28. Kim E, Helal S, Cook D (2010) Human activity recognition and pattern discovery. Pervasive Comput. 48–53
29. Kipp M (2012) Annotation facilities for the reliable analysis of human motion. In: Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC). Istanbul, pp 4103–4107
30. Koller D, Friedman N (2009) Probabilistic graphical models: principles and techniques. MIT Press, Cambridge, Massachusetts
31. Lao L (2006) Location-based activity recognition. University of Washington
32. Logan B, Healey J, Philipose M, Tapia EM, Intille S (2007) A long-term evaluation of sensing modalities for activity recognition. In: Proceedings of the 9th International Conference on Ubiquitous Computing (UbiComp'07). Berlin: Springer-Verlag, pp 483–500
33. Longstaff B, Reddy S, Estrin D (2010) Improving activity classification for health applications on mobile devices using active and semi-supervised learning. In: 2010 4th International Conference on Pervasive Computing Technologies for Healthcare (PervasiveHealth), NO PERMISSIONS. IEEE, pp 1–7
34. Ciurana M, Barcelo-Arroyo F, Izquierdo F (2007) A ranging system with IEEE 802.11 data frames. In: IEEE Radio and Wireless Symposium. Long Beach
35. Maurer U, Smailagic A, Siewiorek DP, Deisher M (2006) Activity recognition and monitoring using multiple sensors on different body positions. Wearable Implant Body Sensor Netw. IEEE, Massachuset, pp 113–116
36. Mennicken S, Huang EM (2012) Hacking the natural habitat: an in-the-wild study of smart homes, their development, and the people who live in them. Pervasive Computing. Springer Berlin Heidelberg, pp 143–160
37. Mennicken S, Vermeulen J, Huang EM (2014) From today's augmented houses to tomorrow's smart homes: new directions for home automation research. In: Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing. ACM, pp 105–115
38. Munaro M, Menegatti E (2014) Fast RGB-D people tracking for service robots. Auton Robots, pp 1–16
39. Murphy Allan H, Winkler Robert L (1984) Probability forecasting in meteorology. J Am Stat Assoc 79:489–500
40. Narayana S, Prasad RV, Rao VS, Prabhakar TV, Kowshik SS, Iyer MS (2015) PIR Sensors: Characterization and Novel Localization Technique
41. Obayashi S, Zander J (1998) A body-shadowing model for indoor radio communication environments. IEEE Trans Antennas Propag 46(6):920–927
42. Pärkkä Juha, Ermes Miikka, Korpipää Panu, Mäntyjärvi Jani, Peltola Johannes, Korhonen Ilkka (2006) Activity classification using realistic data from wearable sensors. IEEE Trans Inf Technol Biomed 10(1):119–128
43. Paiement A, Tao L, Camplani M, Hannuna S, Damen D, Mirmehdi M (2014) Online quality assessment of human motion from skeleton data. In: Proceedings British Machine Vision Conference 2014
44. Roggen D et al (2010) Collecting complex activity datasets in highly rich networked sensor environments. In: 2010 Seventh International Conference on Networked Sensing Systems (INSS). IEEE, pp 233–240
45. Sahinoglu Z, Gezici S, Guvenc I (2008) Ultra-wideband positioning systems
46. Tan B (2015) Wi-Fi based passive human motion sensing for in-home healthcare applications. In: IEEE 2nd World Forum on Internet of Things. Milan
47. Tan B, Woodbridge K, Chetty K (2014) A real-time high resolution passive WiFi Doppler-radar and its applications. In: International Radar Conference. Lille
48. Tsimbalo E, Fafoutis X, Mellios E, Haghighi M, Tan B, Hilton G, Piechocki G, Craddock I (2015) Mitigating Packet Loss in Connectionless Bluetooth Low Energy. In: 2nd IEEE World Forum on Internet of Things (WF-IoT). Milan: IEEE. pp 291–296

49. Tsipouras MG, Tzallas AT, Rigas G, Tsouli S, Fotiadis DI, Konitsiotis S () An automated methodology for levodopa-induced dyskinesia: assessment based on gyroscope and accelerometer signals. Artif Intell Med 127–135. Elsevier
50. Twomey N, Diethe T, Flach P (2016) Unsupervised learning of sensor topologies for improving activity recognition in smart environments. Neurocomputing
51. Twomey N, Flach P (2014) Context modulation of sensor data applied to activity recognition in smart homes. In: Workshop on Learning over Multiple Contexts, European Conference on Machine Learning (ECML'14). Nancy, France
52. van Kasteren T, Noulas A, Englebienne G, Kröse B (2008) Accurate activity recognition in a home setting. In: 10th International Conference on Ubiquitous Computing—UbiComp'08. New York: ACM Press, pp 1–9
53. Wang, Y, Yang X, Zhao Y, Liu Y, Cuthbert L (2013) Bluetooth positioning using RSSI and triangulation methods. Las Vegas
54. Winn John, Bishop Christopher M, Diethe Tom R (2015) Model-based machine learning. Microsoft Research, Cambridge, p 2015
55. Woodman, O, Harle R (2008) Pedestrian localisation for indoor environments. In: Proceedings of the 10th International Conference on Ubiquitous Computing (UbiComp'08)
56. Woznowski P, et al (2015) A multi-modal sensor infrastructure for healthcare in a residential environment. In: IEEE ICC Workshop on ICT-enabled services and technologies for eHealth and AAL. London: IEEE, pp 271–277
57. Tsai Y-L, Tu T-T, Bae H, Chou PH (2010) EcoIMU: a dual triaxial-accelerometer inertial measurement unit for wearable applications. 2010 International Conference on Body Sensor Networks (BSN), Singapore
58. Zhu, Ni, et al. "Bridging e-Health and the Internet of Things: The SPHERE Project." Intelligent Systems, IEEE (IEEE), 2015: 39–46