






# Ethical Surveillance: Applying Deep Learning and Contextual Awareness for the Benefit of Persons Living with Dementia

Steve Williams<sup>1</sup> , J. Mark Ware<sup>1</sup> , and Berndt Müller<sup>2</sup> 

<sup>1</sup> Faculty of Computing, Engineering and Science,  
University of South Wales, Pontypridd, UK

{steve.williams1, mark.ware}@southwales.ac.uk

<sup>2</sup> Department of Computer Science, Swansea University, Swansea, UK  
berndt.muller@swansea.ac.uk

**Abstract.** A significant proportion of the population has become used to sharing private information on the internet with their friends. This information can leak throughout their social network and the extent that personal information propagates can depend on the privacy policy of large corporations. In an era of artificial intelligence, data mining, and cloud computing, is it necessary to share personal information with unidentified people? Our research shows that deep learning is possible using relatively low capacity computing. When applied, this demonstrates promising results in spatio-temporal positioning of subjects, in prediction of movement, and assessment of contextual risk. A *private* surveillance system is particularly suitable in the care of those who may be considered vulnerable.

**Keywords:** Privacy · Deep learning · Assisted-living · Mobile computing · Ethics · mHealth · Wearable health · Dementia · Safer walking · GPS · LSTM · RNN

## 1 Background

Advancements in mobile devices that can be worn and carried, their interconnectivity, and the improvement of artificially intelligent tools provide a significant opportunity to assist in the care of the aged. In accordance with a human right to private life, we have examined methods to keep tracking information private unless there is a moral argument, such as risk to the person being monitored, that justifies a breach in privacy. In this scenario, safety is paramount and in the interests of beneficence and non-maleficence an ethical policy in terms of design is employed, which defines that personal information is precious and should therefore not be shared on the internet.

Dementia is a debilitating condition that is growing with the aging society. Continuance with life in the community is encouraged, since social interaction and physical activity stimulates a healthy mental state in the person with symptoms (PwS) along with the family carer. We seek bespoke artificially intelligent solutions for these persons living with dementia (PlwD) who wish to preserve independence of the PwS.

Initial system infrastructure and findings are published in [1], the suitability of a mobile computer technology in tracking PwS and ethical aspects are previously outlined in [2]. The work described here contributes to the ethical debate regarding the question at which point information gathered when monitoring a PwS should be shared. We investigate a technological solution that keeps data private until a threshold of risk is reached. AI is used to learn what is ‘normal’ for a person (based on individual habits), various metrics are then used in the decision making to change the default private state.

To this end, a monitoring system is designed that requires the PwS to carry a mobile phone and wear a fitness tracker. It is understood that some may not be comfortable with this and it is anticipated that the mobile technology component will ultimately be integrated in a single wearable device. This technology can be particularly useful for patients who have early-onset dementia, i.e., those of working age and therefore more likely to be used to carrying a phone or wearing a smart-watch.

## 2 The Problem

The onset of dementia has a profound effect on the PwS and the wider family unit. Diagnosis can bring with it a loss of role function, uncertainty about the future, fear of being a burden, and reduced mobility that can lead to social isolation [3]. The objective of this study is to create an ‘electronic safety net’ that can provide peace of mind to the carer, while preserving the rights and independence of the PwS. A key aim of the project is to delay residential care.

### 2.1 Dementia

Dementia is caused by several diseases of the brain. There is a wide spectrum of symptoms, some of which may manifest in a propensity to walk independently at inappropriate times [4]. Literature indicates that this can lead to premature mortality [4–6]. Actions to mitigate this risk can lead to increased dependence, to curtailment of social activities, and reduction in quality of life [7]. Elopement episodes are a major reason for nursing home admission [8]. A study in Finland reports that the latter may be delayed, using assistive technology, by an average of eight months [9].

### 2.2 Privacy

Online data privacy divides opinion. Many elect to share very varied information about their lives publicly on the internet, but this is not always a conscious decision – Terms and conditions regarding data sharing tend to be ignored by many users as they install applications and use online services. Nevertheless, consent given in this way is often referred to as informed when the potential for data propagation is mentioned in the supplied information, even though this information is rarely considered thoroughly.

Leaks of private information have recently been in the news headlines. Data stored on the internet, e.g. by cloud services, is often assumed to be safe, but human intervention and inadequate security measures allow breaches [10]. Advocates of privacy treat personal information very differently and avoid sharing their information with

people or organisations. This attitude is supported by cyber-security activists, e.g. in a report of vulnerabilities leading to 91 exploits of tracking service providers in January 2018 [11]. In the case of care for persons who may be considered vulnerable it seems ethically correct that a strict data protectionist policy should be the default.

### 2.3 A Human Rights-Based Approach

The World Health Organisation (WHO) advocates a human rights-based approach for PlwD [12]. In our study, almost two years of personal data was collected. This included location (derived using GPS and nearby Wi-Fi nodes), activity recognition, indoor movement, and logs of heart rate, steps, and sleep patterns. This kind of monitoring undoubtedly has the potential to invade a person's right to private life. The tracking was described by the subject as a big-brother bad dream. On reflection, the level of 'invasion' depends on who has access to the data.

## 3 Machine Learning (ML)

The aim is for an algorithm to learn human mobility patterns of an individual, and to assess the perceived risk against the learnt normality that is deemed to be 'safe'. A measure of risk is used to determine the level of protection required on the personal data collected. To protect privacy, propagation of this information is restricted to the secure home network. No interaction with the wearable or phone is required of the PwS.


To improve potential accessibility to many users in the long run, the equipment used in a working prototype is a standard smart-phone and a home-based 'hub', which is a credit-card sized computer with limited resources, such as a Quad-Core 1.2 GHz CPU and 1 GB RAM. Networking between the two in 'monitoring' mode is via on-board Bluetooth and Wi-Fi only while at home.

**Unconventional Deep Learning:** Deep learning (DL) discovers intricate patterns in large datasets by using multiple processing layers to learn representations of data [13]. Sequential and parallel information is processed in a cyclical (recurrent) fashion by modifying internal weightings of input signals to produce an expected output signal [14, 15]. The hardware platform described may seem restrictive for a DL task in an age where we are used to resources being server based and ubiquity being the norm. Convention says that DL requires large computing capacity, but this is not available for the present use case. Long Short-Term Memory (LSTM) networks [16] are a type of Recurrent Neural Network suitable for learning and predicting sequential patterns in timelines. Using accelerometers, as commonly found in modern mobile devices, LSTM are deployed in human activity recognition (HAR). X-Y-Z accelerometer readings are interpreted over a defined time-period and then compared to those taken in a laboratory to determine probability that a categorised activity is taking place [17]. We have assimilated this using GPS sensor data. A dataset suitable for learning using an LSTM neural network was developed, and the resultant tensor was deployed to an Android device to calculate the probability of being on a learnt trajectory or otherwise.

The novel concept that surveillance need not be *invasive* is introduced. There is a host of literature relating to HAR [18], there are indoor monitoring studies with AI, e.g. [19], and studies of wandering trajectories, e.g. [20]. None of these describe categorisation of the normal movements of a person together with discrete monitoring that keeps information private until anomalies are found.

### 3.1 ML Methodology

**Data:** GPS data is collected from one subject using a standard HTC-10 smartphone used solely for that purpose. Considerable data preparation is required using the minute-by-minute location coordinates. Data is first compartmentalised based on total movement to date ( $tm$ ). This is then divided by an increment ( $i$ ) giving sub-divisions as shown in Fig. 1 with  $i = 20$ .




**Fig. 1.** Boundaries of the extent of total movement for 3 months,  $i = 20$ . Map: © Google.


To optimise computation time, daily data is reduced to only the proportion that represents movement.

**Categorisation.** Points within each segment (or compartment) are assessed for each trajectory and each segment's points are compared using a kd-tree-based nearest-neighbour algorithm [21]. The degree of similarity is assessed giving a percentage and a threshold provides a similarity decision. There is difficulty in some trajectories where, for example, topographical, atmospheric, or networking issues used in test data collection lead to sparse and noisy data. Sparse data was dealt with using 1d-univariate interpolation [22]. This is particularly important in the early days of training where there are few trajectories to compare. Noisy data is essentially ignored at this time by

adding a tolerance to the similarity decision just described which is explained in more detail below. The result of the comparison algorithm is a segment chain (string) for trajectories with 1 or 0 signifying a match in each square (Fig. 2).





**Fig. 2.** A successful match of two segment chains. 29 segments,  $i = 10$ .



**Fig. 3.** Interpolation used to deal with sparse data causing accuracy issues. Map data: © Google.

Categorisation by comparison of trajectory segment chains by only comparing matched segments significantly reduces the computational capacity required in terms of processing and memory. If a match is found, interpolated point data is added to a master repository with which future comparisons are made. An encoded polyline [23] reduces database size requirements and allows trajectories to be stored as entities. In time, the necessity for interpolation is reduced as the repository trajectory density increases.

As seen in Fig. 3, interpolation may cause significant deviation from the route that is travelled, e.g. by cutting corners and using roundabouts, but this level of granulation is considered satisfactory as a ‘zone of safe movement’ is maintained. Matching segments rely on a nearest neighbour tolerance ( $nnT$ ) and merging with subsequent trajectories eventually creates a dense category master that is used to define this zone.




**Fig. 4.** (a) A comparison tolerance  $nnT$  leads improved matches while ignoring noise. (b) Interpolated points are merged to create a dense category master.


$mnT$  set at 0.005 in decimal degrees, equating to just over 500 m, is used in the experiments. This tolerance can be linked to  $tm$  in further work as the extent of movement defines the granularity required within the movement space. The resultant categories develop into a densely populated polyline seen in Fig. 4b. All movement within a data collection period are matched with destinations recognised in the initial cluster analysis.

In addition to our collected data, the comparison algorithm was tested using seven users’ data from the Geolife (GL) dataset [24]. This contains better quality GPS trajectories and includes higher variance in modes of travel. With  $mnT$  applied to nearest neighbour algorithm it is observed that small deviations from a route are not a significant problem. As can be seen in Fig. 5, four separate tracks converge on a destination and in the extent of this day’s movement all points are within one segment.

Noise, detours and differing distances included in two tracks taking Route 1 and Route 2 in Fig. 6, both arrive at the same place  $E^1$  and  $C^2$ .  $mnT$  allows for the eventuality of  $C^1$  and  $D^1$  not matching Route 2. Adding both to the master increases the possibility that subsequent trajectories match by widening the dataset.



**Fig. 5.** Detours and converging paths are handled using  $mnT$  and segment comparison. Map: © Google Maps.



**Fig. 6.** Widening the category master by allowing a nearest neighbour tolerance.

**Bearing.** Some GL users’ data highlighted the difficulty of recognising direction of travel in that only one-way trajectories are recorded. Experimentation with inclusion of direction of travel gave complex results, consequently movement is treated as omnidirectional; the category master is essentially an amalgamation of history on that route.

**Time Factor.** This is an important consideration in the study scenario, but the likelihood of a person travelling a recognised trajectory at the same time is low so prediction of this is not required. There are detours from a route, the method of travel may change, there may be traffic. These factors all have a significant impact on spatio-temporal data and following extensive experimentation, it is concluded that data-point true timestamps cause confusion. Instead, each category master is indexed sequentially.

**Predictability.** Major studies in human mobility patterns find that there is a high degree of temporal and spatial regularity [25]. In the datasets investigated, this study

concur; the number of places travelled to is surprisingly low. The three most regularly visited destinations are selected for demonstration; these are travel to University (south), to social visits (west) and to a supermarket (east) seen in Fig. 7.



Fig. 7. Three categories of travel overlaid with noise showing, interpolated;  $3 \times 10^4$  records.

**Pre-processing.** Category masters are exported and the number of records per category is equalised by interpolating (increased or reduced) to 10,000 records each. Noise is amplified where outliers are interpolated. These outliers will be removed in later versions of the system. The data is stacked and normalised. Train:test split is 80:20.

**Machine Learning.** Inspiration for this is credited to work using Convolutional Neural Network and LSTM RNN in mobile phone HAR applications. The solution selected for our application is Tensorflow 'BasicLSTMCell' stacked with 'Multi-RNNCell' with 64 hidden units. The neural network is expected to learn geo-spatial data to predict categorisation (of the trajectory) when it is given further blocks.

### 3.2 ML Results

Experimentation found that the number of time steps set at 10, in blocks of 10 gave an accuracy of 90–97% over 500 epochs in less than 1.5 h (Fig. 8).

**Deployment.** Using our dataset, the resultant tensor is imported to an Android application that sequentially passes arrays of 10 steps of a test trajectory in a timed fashion. A Tensorflow classifier returns the probability of the array being Category 1, 2 or 3 for the three trained classes. These predictions are logged on the phone (Fig. 9).

#### Mobile Results


Category 1: Correctly predicted with 98–99% certainty unless trajectories overlap.

Category 2: Correctly predicted with 55–86% certainty.

Category 3: Correctly predicted with 77–90% certainty.



**Fig. 8.** LSTM training session over 1.4 h. 90–97% accuracy.



**Fig. 9.** Android category prediction results: the vertical scale on these graphs range from 0 to 1 where 1 = 100% certainty. (a) Category 1 (dotted line). (b) Category 2 (dashed line). (c) Category 3 (solid line).

The Tensorflow classifier gives reliable prediction of a route being tested in all cases. These are very satisfactory results. Overlap between two categories returning a 50:50 result in the Category 1 test is perfectly acceptable since the routes do overlap.

**Public Dataset Results**


Training was carried out using matched trajectories in the GL dataset with similarly acceptable results (Fig. 10).

**Volunteer Test Results**

Six volunteers were recruited from a convenience sample that consented to be tracked by phone and fitness tracker for a period of three months. GPS data was collected by enabling Google Timeline and by configuring their phone accordingly. Places visited, and routes taken are stored on Google servers<sup>1</sup>. At the end of the period data was

<sup>1</sup> Note that Google Timeline is only used for data gathering in this initial feasibility study. The full solution uses GPS data stored only locally on the mobile device and processed on the home hub.





**Fig. 10.** LSTM training session for 7 users (in 9 tests) using Geolife dataset.

exported, and a bespoke script was developed to interpret the data. With very few exceptions, it was found that the main locations visited can be classified as attending a place of work or study, going shopping and visiting friends or family.

**Machine Learning Results:** As previously explained, the six datasets consist of latitude, longitude, and elevation. When subjected to the same neural network, training test results are shown in Table 1.

**Table 1.** Volunteer machine learning results

Volunteer	Age	Gender	Phone	% accuracy result
1	46	F	iPhone 5s	86
2	23	M	iPhone 6s	84
3	22	F	iPhone 5s	85
4	80	F	Android Galaxy S5	89
5	45	M	Android Galaxy J3	97
6	54	M	Android Galaxy S4	88

Volunteer 3 presented problems in that trajectories overlapped due to the topography of their home address. A revised method that change the way overlapping trajectories are categorised can be used overcome these issues. Volunteer 4 had noticeably more restricted movement and comparably reduced distance travelled causing specific issues of data sparsity. For these reasons both datasets required manual categorisation and matching.

**Data Augmentation:** Noisy data was dealt with using Google Snap-to-Road [26] and/or TrackMatching [27] and when sparsity occurred, gaps were filled using route finding techniques such as Google Directions API [28] or a variation of Open Street Maps routing [29]. Route finding methods of augmentation cause an element of subjectivity but provided data suitable to test the network.

**Deployment Results:** The Android simulator gives closely comparable results as that with our own data, for example volunteer 1:

*Category 1:* Correctly predicted with 42–54% certainty.\*

*Category 2:* Correctly predicted with 97–99% certainty.

*Category 3:* Correctly predicted with 98–99% certainty.

\* routes overlap

It is concluded from this series of tests that the developed machine learning model provides adequate accuracy in the categorisation of routes. The differing data collection methods give a valuable insight into how best to develop training data. The method of transferring learnt information to a mobile phone is particularly interesting as heavy processing can be carried out on a hub, thus preserving the restricted battery resource of the mobile device.

## 4 Situation Appraisal


Real-time appraisal of the situation of the person being monitored is key to ensuring their well-being. The system is designed with PlwD in mind, so apart from elopement, issues specific to PwS are considered:

**Sleep and Dementia.** Circadian rhythm disorders can present as an early component of the disease. They have significant impact on patients and caregivers and are a ‘major risk factor for early institutionalisation’ [30, 31]. Symptoms include sleep disturbances, sun-downing, and agitation. Instances of elopement regularly occur at night. Disturbance in sleep of the PwS has a significant effect on care-givers that can lead to their ill-health [32, 33].

Therefore, monitoring of sleep is highly relevant to this study. Not only should PwS’s safety outdoors be monitored, but a metric of well-being should be used to modify system sensitivity. The following section describes the approach and some technicalities of machine learning in this area. In addition, the initial metrics used in appraisal of the contextual situation the person being monitored is in are outlined. Factors such as sleep, and heart rate are here referred to as the ‘pre-disposition’ of the person. This may be understood as a metric for their well-being.

**Data Collection.** Although sharing data to a manufacturer’s server breaches the complete privacy rule, a FitBit fitness monitor is used in this study for convenience. A dedicated wearable with direct, local data access would allow to preserve privacy and will be used in the final prototype. A FitBit ‘Ionic’ is one of many devices that are popular with those who wish to monitor, for example, a keep-fit regime. While wrist actigraphy is customarily used in sleep research there is evidence that FitBit devices provide close estimation of total sleep time [34]. Over 2 years of data was collected from one subject using this device and more than three months of data from six volunteers using similar models. It was found that the data collected gives a good representation of actual sleep patterns. A secure authenticated oAuth2.0 API is used to access data from the FitBit servers, yielding daily data, when visualised is shown in Fig. 11.


**Machine Learning.** Machine-learning techniques have been developed that assess the data, which includes minute by minute heart-rate, steps and sleep records. The



**Fig. 11.** Graph of steps, heart rate and sleep records for one day.

requirement is that human activity is discretely monitored with automated realisation of trends. The daily situation is then evaluated and compared to what is deemed ‘normal’ for the individual.


**Sleep Period.** This is modelled using a Gaussian Distribution in order to give clarity on the expected duration of sleep. Long term changes to averages in a 3-month moving period for example using start, finish and duration of sleep may be used in appraisal (cf. Fig. 12).



**Fig. 12.** Gaussian distribution of sleep start, finish and duration for 3 months.

While the subject in this study does not suffer from disruption in diurnal rhythm per se, average sleep per week, and distribution, give an interesting illustration that show variance in the time-period (cf. Fig. 13).

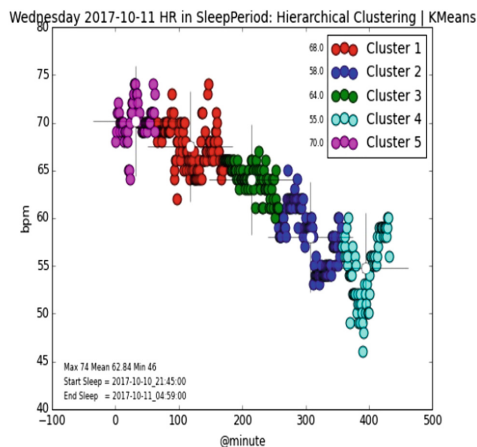
Results can be categorised using  $1\sigma$  or  $2\sigma$ , i.e. 68% or 98% of the norm ( $\mu$ ). Waking times can be defined as normal (up to  $1\sigma$ ), early or late (between  $1\sigma$  and  $2\sigma$ ), and very early or late ( $>2\sigma$ ). When visualised, trends are apparent, there are outliers that represent exceptional occurrences in this period.



**Fig. 13.** Trends in weekly sleep time for one year ignoring restlessness.

**Activity.** The same method was applied to step-count in daytime and night periods. This information is useful in recognising active periods during the night and possible association of these with daytime sedentary periods possibly in correlation with less sleep at night.

**Heart Rate While Asleep.** It is observed in our dataset that in the day, heart rate closely relates to physical activity such as steps, but while asleep, lack of movement can be used to provide a period in which it is possible to benchmark and provide reliable regression analysis. As illustrated in Fig. 11, minute by minute daily heart rate is collected. When heart-rate while asleep is extrapolated across days and polynomial regression compared, clear differences are evident. Centroids of five sleep periods are analysed. The sleep periods are *start*, *early*, *mid*, *late* and *finish*. In this way, varying sleep periods of different lengths are normalised. Agglomerative clustering with simple Euclidean affinity [35], and k-means cluster analysis [36] are used to give single centroids for each period. Having results for each period makes it possible to visualise clusters (cf. Fig. 14), and conclude  $\mu$  and  $\sigma$  in any defined period.



**Fig. 14.** Five clusters of heart rate readings in one night.  $\mu = 63$  bpm.

Centroid data for a date period gives an average of averages ( $\mu$ ). If a limit of, for example,  $(\mu + 0.5\sigma)$  is applied on data available then results can be categorised as being ordinary (0) or otherwise (1). Currently, only checks for high heart rate are introduced into the training data, but others could be included in future.

**Neural Network.** A neural network was developed using Keras and TensorFlow. The input layer and 2nd hidden layer uses a rectifier activation function with 6 nodes and 9 inputs. These are maximum, minimum,  $\mu$ ,  $\sigma$  and the five centroid results. A sigmoid activation function is used on the output layer. The classifier function is compiled using the ‘adam’ optimiser with ‘binary\_crossentropy’ loss, 500 epochs are used in training in batches of 10. The data is split so that 80% is used in training and 20% as test data. Experimentation using just over a year of raw-data were useful in that they were able to recognise the categorisation that had been applied to that data. 94.9% accuracy was measured for 1 year’s data in 105 s. Using 1 month of data, 85% accuracy was measured in 24 s, this was then improved to 87% by only supplying centroid measurements i.e. 5 inputs. It was concluded that pre-processing using hierarchical clustering then k-means and applying a fixed rule for categorisation is reliably recognised in a Neural Network; it is possible, for this subject, to predict with 87% accuracy given a month of data. The categorisation of data requires a rule to build training data, but once training has taken place the resultant tensor is deployed to the phone to process daily readings autonomously. The system successfully recognised exceptional heart rate events of the subject.

A fitness tracker that is worn 24 h a day provides an efficient way to collect information for this study. The product used is aesthetically pleasing and if a PwS is used to wearing a smart-watch it should not present a problem in use. Activity and indication of heart-rate levels may provide a useful indicator of well-being of a person at night.

## 5 Contextual Factors

When coupled with fundamental contextual factors such as time-of-day, distance from home, and weather conditions, contextual risk of being at a location outdoors can be used in decisions regarding preservation of privacy. The following sections summarise factors used in this study.

### 5.1 Time and Distance Metric

Time-of-day is easily determined on a computer and is an important factor when considering risk. Weightings  $w(t)$  of time  $t$  and distance from home are used for analysis. A time metric simply uses the hour of day, this is provisionally set as follows:

$$w(t) = \begin{cases} 1 & \text{if } 9:00 \leq t < 18:00 \\ 2 & \text{if } 6:00 \leq t < 9:00 \text{ or } 18:00 \leq t < 21:00 \\ 3 & \text{if } 21:00 \leq t \text{ or } t < 6:00 \end{cases} \quad (1)$$

with 1 representing low risk. A distance metric  $w'(d)$  for distance  $d$  is similarly set at:

$$w'(t) = \begin{cases} 1 & \text{if } 0 \text{ km} \leq d < 1 \text{ km} \\ 2 & \text{if } 1 \text{ km} \leq d < 5 \text{ km} \\ 3 & \text{if } 5 \text{ km} \leq d \end{cases} \quad (2)$$

These definitions are notional values for use in the experiments which should initially be decided by the user dyad after considering the questions of *when* and *how far* is normal for the individual PwS.

### 5.2 Weather Metric

In a similar way, scales of risk can be compiled using weather forecasting applications such as Dark Sky or OpenWeatherMap. The Dark Sky API [37] offers a full collection of meteorological conditions and is used in the study. The locality of the subject is known so forecast data is retrieved for temperature, precipitation and wind in that area. A rudimentary weather metric is defined using a matrix (see Table 2 where, again, 1 corresponds to low risk).

**Table 2.** Weather metric matrix

Weight	Temperature (°C)	Precipitation (mm/h)	Wind (Beaufort scale)
1	]15..25]	[0..1]	[0..5]
2	]10..15], ]25..27]	]1..4]	]5..11]
3	[∞..15], ]27..∞]	]4..∞]	]11.. ∞]

Figure 15 illustrates how this matrix can be used to conclude an accumulated weather measurement and how this can be weighted by time of day.

Time	TEMP °C		?	RAIN			WIND m/s			accumulated			
	Weight	Score		Weight	Result	Score	Weight	Result	Result	Weight	Result		
0	3	2	3	100.0	1	1	1	1	1	105	15	1575	
1	3	2	3	100.0	6	3	18	3	1	3	124	21	2604
2	3	7	3	14.3	3	2	6	5	1	5	28	18	509
3	3	8	3	12.5	4	2	8	7	2	14	38	21	788
4	3	3	3	100.0	5	3	15	10	2	20	138	24	3312
5	3	-1	3	100.0	6	3	18	12	3	36	157	27	4239
6	3	-1	3	100.0	7	3	21	15	3	45	169	27	4563
7	2	7	3	14.3	5	3	15	19	3	57	88	18	1589
8	2	8	3	12.5	4	2	8	23	3	69	92	16	1464
9	1	9	3	11.1	3	2	6	27	3	81	99	8	793
10	1	10	2	10.0	2	2	4	31	3	93	108	7	756
11	1	15	1	6.7	1	1	1	10	2	20	29	4	115
12	1	27	2	50.0	0	1	0	12	3	36	87	6	522
13	1	30	3	100.0	10	3	30	15	3	45	176	9	1584
14	1	27	2	50.0	0	1	0	19	3	57	108	6	648
15	1	16	1	6.3	0	1	0	23	3	69	76	5	381
16	1	10	2	10.0	0	1	0	15	3	45	56	6	336
17	1	12	2	8.3	0	1	0	12	3	36	45	6	272
18	2	9	3	11.1	0	1	0	10	2	20	33	12	397
19	2	8	3	12.5	0	1	0	9	2	18	33	12	390
20	2	5	3	100.0	0	1	0	5	1	5	107	10	1070
21	3	-3	3	100.0	0	1	0	4	1	4	107	15	1605
22	3	-5	3	100.0	0	1	0	3	1	3	106	15	1590
23	3	-10	3	100.0	0	1	0	2	1	2	105	15	1575
24	3	-3	3	100.0	0	1	0	1	1	1	104	15	1560

**Fig. 15.** Examples of a weather metric at different times of the day.


The weather metric, time, and categorised location result, when amalgamated with pre-disposition give a measure of risk. Put simply, if the monitored person is well rested and is outside on a sunny afternoon in a place which is defined as normal then perceived risk may be low, an accumulated score is used to derive the overall risk. When viewed together, these measures can be used to decide the point at which to override the privacy rule to ensure the monitored person is not harmed. The following section describes a working application designed to illustrate this.

## 6 Risk Analysis

Complex methods can be applied to calculate the perceived risk to the PwS; all the metrics described may be used to adjust the overall sensitivity of the system.

### 6.1 Inferring an Unknown Location

The Tensorflow Classifier, described in Sect. 3, is used in prediction of where the subject is in relation to normally visited places. If the subject moves to a new space, the contextual risk of that activity is assessed using time, distance from home, and forecasted weather conditions. In Fig. 16(a) movement along the test trajectory is outside known areas (shaded grey), distance and known temperature for the area is monitored (left graph above map). Risk is visualised in the right graph. As distance from home (start point) reduces, the system perceives this as returning and hence risk decreases. In Fig. 16(b) a detour outside a known path instigates appraisal and logs this as a new place, leading to an accumulation of risk, that is reset when the probability of normal movement increased, as shown in Fig. 16(c).



**Fig. 16.** (a) Graphed representation of accumulating risk reducing with distance. (b) Risk increasing when taking a detour from the trained path. (c) Correct categorisation of trajectory with 99% accuracy – risk is reset.

The described system is indicative of how machine learning can be used to assess the ‘normality’ of outdoor movement and changes in sleep and heart rate patterns. The attraction of using AI in this way is that data learnt can be transferred across platforms and re-training can take place using a private network overnight. The tensor allows the mobile device carried by the PwS (in our study a phone) to act as an autonomous agent that does not require the internet, which has a very significant positive impact on battery life.

This research introduces a method using an AI agent to continuously assess the situation and make an ethical decision on overriding the default level of privacy. The point at which a measure of risk translates into a decision to breach privacy for the good of the monitored person is the topic of significant and ongoing debate, to which this technical study contributes. The following section briefly touches on considerations in this field.

## 7 Ethics

The ethical debate regarding the point at which location data is shared, and with whom is an interesting area to which our findings contribute. If activity, time, place, or weather is appraised as high risk or ‘inappropriate’, a prior moral framework that rates *safety and risk* versus *privacy* can justify that recent movement and current location may be shared. A wellbeing metric can be used to determine system sensitivity. The sharing of location can take the form of an SMS alert, or an alert via the internet including a map showing the current position of the PwS to a trusted carer. Continuous updates can facilitate speedy recovery. In all other cases, the PwS may continue independently and all data collected is kept private.

Several questions arise:

- When applied to vulnerable persons, who may decide the threshold and who defines what is ‘inappropriate’?
- Is normality really ‘safe’?
- In production, would an AI-based algorithm implementing a definition of privacy be trusted?

Our work does not attempt to answer these questions, but provides investigation into the capabilities of technology. It is found in literature that technological solutions fail to offer a considered approach to resolve well-known privacy issues. Surveillance of those who may be deemed vulnerable is considered by many as ethically inappropriate, but ‘needs must’ and carers are taking DIY approaches [38], with systems that use technology not optimised for privacy. This exposes them to potential security vulnerabilities as described above. We have shown that a fitness tracker can be used to learn what is normal in terms of heart rate while asleep. This, and other contextual matrices can be used to modify system sensitivity. A *private* monitoring system that uses AI to determine out-of-the-ordinary movement is novel. Since it respects privacy, this surveillance is *not intrusive*. Development and implementation of such a system is likely to provide PlWD with an ethically robust ‘safety net’ that may be used to improve quality of life. It can increase independent living of the PwS, provide peace of mind to



the carer, while not requiring data sharing to call centres, or collection on central servers of tracking providers. Our system achieves secure data control and maintains data ownership.

## 8 Conclusion

The research presented shows promising results both in recognition of human geo-spatial activity and in prediction of movement along normally travelled routes. A wellness monitor discretely monitors both sleep and heart rate and this can be used to learn what is normal to enable it to flag exceptions. A cost-effective working prototype has been produced to demonstrate that deep-learning techniques can be applied to spatio-temporal data after programmatically categorising normally travelled trajectories. It has been found that when only part of a trajectory has been travelled, likely destinations can reliably be inferred. The application is designed to restrict personal information propagation to a home network and the limitations of computing capacity do not detract from the quality of results.

The World Health Organisation recognises that surveillance is intrusive, that the human rights of PwS are sometimes denied and that abuse of liberties is present. Locking doors to stop a person eloping violates their human right to liberty, but surveillance normally results in sharing of personal information, so is contrary to their human right to private life. Risk, when deviations from known places are sensed, is assessed automatically on a smart-phone in the context of time, extent and weather conditions.

Human rights (of private life and liberty) of the person with symptoms will be respected until the point at which it is judged that a prior moral argument of safety and risk supersedes the importance of privacy. If this happens, alerts containing location and recent movements are shared with an assigned carer, thus facilitating swift recovery.

The potential of the AI system described here is considerable. It is likely that many who value the importance of privacy highly will welcome a surveillance system that monitors but does not divulge detail. Predictions of likely trajectory of movement using real-time location data is novel, as is the concept of private surveillance as described. Availability of an internet connection or at least cellular coverage to deliver alerts is a requirement for implementation.

Ongoing work includes the processing of data from recruited volunteers, it is difficult to assess how the data-sets used differ from that which could be collected from PwS, subject to gaining the appropriate ethical approvals trials will embark with recruited PwS. The assessment of complex and intertwined trajectories and comparison of different scales of movement is currently under investigation. Findings will contribute to further refinement of the methodology after consultation with health professionals and PlwD. In an ideal scenario this would be used for prolonged independence of PwS, alleviation of a 24/7 burden of care, and could delay the necessity of moving the PwS to a care home.

**Acknowledgements.** Knowledge Economy Skills Scholarships (KESS) is a pan-Wales higher level skills initiative led by Bangor University on behalf of the HE sectors in Wales. It is part funded by the Welsh Government's European Social Fund (ESF) convergence programme for West Wales and the Valleys and is supported by the industrial partner SymIConnect Limited.

## References

- Williams, S., Müller, B.: Agents and dementia—smart risk assessment. In: Criado Pacheco, N., Carrascosa, C., Osman, N., Julián Inglada, V. (eds.) EUMAS/AT-2016. LNCS (LNAI), vol. 10207, pp. 277–284. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-59294-7\\_22](https://doi.org/10.1007/978-3-319-59294-7_22)
- Williams, S., Ware, J.M.: Is the use of 'mobile computer technology' appropriate for locating people with dementia? In: 2015 Proceedings of GIS Research UK (GISRUK), pp. 659–664. Figshare, Leeds (2015). <https://doi.org/10.6084/m9.figshare.1491375>
- Read, S., Toye, C., Wynaden, D.: Experiences and expectations of living with dementia: a qualitative study. *Collegian* **24**(5), 427–432 (2017). <https://doi.org/10.1016/j.colegn.2016.09.003>
- Silverstein, N.F.: *Dementia and Wandering Behavior: Concern for the Lost Elder*. Springer, New York (2006). ISBN 0-8261-0272-7
- McShane, R., et al.: Getting lost in dementia: a longitudinal study of a behavioral symptom. *Int. Psychogeriatr.* **10**(03), 253–260 (1998). <http://www.ncbi.nlm.nih.gov/pubmed/9785146>
- Ali, N., et al.: Risk assessment of wandering behavior in mild dementia. *Int. J. Geriatr. Psychiatry* **31**, 367–374 (2016). <https://doi.org/10.1002/gps.4336>
- Martyr, A., et al.: Living well with dementia: a systematic review. *Alzheimer's Dement. J. Alzheimer's Assoc.* **13**(7), 1567–1568 (2017). <https://doi.org/10.1016/j.jalz.2017.07.725>
- Cipriani, G., Lucetti, C., Nuti, A., Danti, S.: Wandering and dementia. *Psychogeriatrics* **14**, 135–142 (2014). <https://doi.org/10.1111/psyg.12044>
- Riikonen, M., Mäkelä, K., Perälä, S.: Safety and monitoring technologies for the homes of people with dementia. *Gerontechnology* **9**(1), 32–45 (2010). <https://doi.org/10.4017/gt.2010.09.01.003.00>
- Mathews, L.: Data From 540,000 GPS Vehicle Trackers Leaked Online *forbes.com* (2017). <https://www.forbes.com/sites/leemathews/2017/09/22/data-from-540000-vehicle-tracking-devices-leaked-online/#40b9c009274b>. Accessed 19 Sept 2018
- Trackmageddon website: Multiple vulnerabilities in the online services of (GPS) location tracking devices (2018). <https://0x0.li/trackmageddon/>. 18 Aug 2018
- World Health Organisation: Ensuring a human rights-based approach for people living with dementia (2015). The need for a human-rights based approach: [http://www.who.int/mental\\_health/neurology/dementia/dementia\\_thematicbrief\\_human\\_rights.pdf](http://www.who.int/mental_health/neurology/dementia/dementia_thematicbrief_human_rights.pdf). Accessed Mar 2018
- LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. *Nature* **521**, 436–444 (2015). <https://doi.org/10.1038/nature14539>
- Schmidhuber, J.: Deep learning in neural networks: an overview. *Neural Networks* **61**, 85–117 (2014). <https://doi.org/10.1016/j.neunet.2014.09.003>
- Brownlee, J.: How to Implement the Backpropagation Algorithm From Scratch in Python (2016). *Machine Learning Mastery: <https://machinelearningmastery.com/implement-backpropagation-algorithm-scratch-python/>*. Accessed Mar 2018
- Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Comput.* **9**(8), 1735–1780 (1997). <https://doi.org/10.1162/neco.1997.9.8.1735>
- Kwapisz, J.R., Weiss, G.M., Moore, S.A.: Activity recognition using cell phone accelerometers. *ACM SIGKDD Explor. Newsl.* **12**(2), 74–82 (2010). <https://doi.org/10.1145/1964897.1964918>

18. Ramamurthy, S.R., Roy, N.: Recent trends in machine learning for human activity recognition—a survey. *Data Min. Knowl. Discov.* **8**, 1–19 (2018). <https://doi.org/10.1002/widm.1254>
19. Arifoglu, D., Bouchachia, A.: Activity recognition and abnormal behaviour detection with recurrent neural networks. In: 14th International Conference on Mobile Systems and Pervasive Computing (MobiSPC 2017) (2017). *Procedia Comput. Sci.* **110**, 86–93. <https://doi.org/10.1016/j.procs.2017.06.121>
20. Batista, E., Borrás, F., Casino, F., Solanas, A.: A study on the detection of wandering patterns in human trajectories. In: 6th International Conference on Information, Intelligence, Systems and Applications (IISA), pp. 1–6. IEEE, Corfu (2015)
21. Maneewongvatana, S., Mount, D.M.: Analysis of approximate nearest neighbor searching with clustered point sets. CoRR [arXiv:cs/9901013v1](https://arxiv.org/abs/cs/9901013v1) (1999)
22. SciPy.Org: Interpolation (scipy.interpolate). <https://docs.scipy.org/doc/scipy/reference/interpolate.html>. Accessed 17 Mar 2018
23. Google: Encoded Polyline Algorithm Format. Google Maps Api: <https://developers.google.com/maps/documentation/utilities/polylinealgorithm?csw=1>. Accessed 17 Mar 2018
24. Zheng, Y., et al.: Geolife GPS trajectories 1.1. In: Geolife GPS Trajectory Dataset - User Guide. Microsoft Research (2011). <https://www.microsoft.com/en-us/research/publication/Geolife-gps-trajectory-dataset-user-guide>. Accessed 18 Sept 2018
25. González, M.C., Hidalgo, C.A., Barabási, A.: Understanding individual human mobility patterns. *Nature* **453**, 779–782 (2008). <https://doi.org/10.1038/nature06958>
26. Google Maps Platform: Snap to Roads. Roads API: <https://developers.google.com/maps/documentation/roads/snap>. Accessed 18 Sept 2018
27. Fabrice Marchal: TrackMatching API. TrackMatching Website: <https://mapmatching.3scale.net/mmswag>. Accessed 18 Sept 2018
28. Google Directions API: Google Directions API. Google Maps: <https://developers.google.com/maps/documentation/directions/start>. Accessed 18 Sept 2018
29. Lambertus: YOURS Routing\_API. Yet another OpenStreetMap Route Service: [https://wiki.openstreetmap.org/wiki/YOURS#Routing\\_API](https://wiki.openstreetmap.org/wiki/YOURS#Routing_API). Accessed 18 Sept 2018
30. Laure, P.-D., Yammine, P., Bastuji, H., Croisilef, B.: Sleep and Alzheimer’s disease. *Sleep Med. Rev.* **19**, 29–38 (2015)
31. Hope, T., et al.: Predictors of institutionalization for people with dementia living at home with a carer. *Int. J. Geriatr. Psychiatry* **13**(10), 682–690 (1998)
32. McCurry, S., Logsdon, R., Teri, L., Vitiello, M.: Sleep disturbances in caregivers of persons with dementia: contributing factors and treatment implications. *Sleep Med. Rev.* **11**(2), 143–153 (2007)
33. Brodaty, H.D.: Family caregivers of people with dementia. *Dialogues Clin. Neurosci.* **11**(2), 217–228 (2009). <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3181916/>. Accessed 17 Apr 2018
34. Keill, A.K., et al.: Validity of wearable fitness trackers on sleep measure. *Med. Sci. Sports Exerc.* **48**(5S), 10 (2016)
35. scikit-learn developers: sklearn.cluster.AgglomerativeClustering. <http://scikit-learn.org/stable/modules/generated/sklearn.cluster.AgglomerativeClustering.html>. Accessed 18 Sept 2018
36. sklearn.cluster developers: sklearn.cluster.KMeans. scikit-learn: <http://scikit-learn.org/stable/modules/generated/sklearn.cluster.KMeans.html>. Accessed 18 Sept 2018
37. Dark Sky Website: Dark Sky Weather API. <https://darksky.net/dev>. Accessed 18 Sept 2018
38. Gibson, G., et al.: The everyday use of assistive technology by people with dementia and their family carers: a qualitative study. *BMC Geriatr.* **15**, 89 (2015). <https://doi.org/10.1186/s12877-015-0091-3>