NFC based dataset annotation within a behavioral alerting platform

Joseph Rafferty, Jonathan Synnott, Chris Nugent

School of computing and Mathematics
University of Ulster
Jordanstown, Northern Ireland
Email: {j.rafferty, j.synnott, cd.nugent}@ulster.ac.uk

Gareth Morrison

The Lava Group
17 Falcon Road
Belfast, Northern Ireland
Email: g.morrison@thelavagroup.co.uk

Elena Tamburini

I+ S.r.l
Piazza Puccini 26
Florence, Italy
Email: e.tamburini@i-piu.it

Abstract— **Pervasive and ubiquitous computing increasingly relies on data-driven models learnt from large datasets. This learning process requires annotations in conjunction with datasets to prepare training data. Ambient Assistive Living (AAL) is one application of pervasive and ubiquitous computing that focuses on providing support for individuals. A subset of AAL solutions exist which model and recognize activities/behaviors to provide assistive services. This paper introduces an annotation mechanism for an AAL platform that can recognize, and provide alerts for, generic activities/behaviors. Previous annotation approaches have several limitations that make them unsuited for use in this platform. To address these deficiencies, an annotation solution relying on environmental NFC tags and smartphones has been devised. This paper details this annotation mechanism, its incorporation into the AAL platform and presents an evaluation focused on the efficacy of annotations produced. In this evaluation, the annotation mechanism was shown to offer reliable, low effort, secure and accurate annotations that are appropriate for learning user behaviors from datasets produced by this platform. Some weaknesses of this annotation approach were identified with solutions proposed within future work.**

Keywords— annotation; NFC; smart environment; pervasive computing; machine learning; ubiquitous computing; behavior detection;

I. INTRODUCTION AND BACKGROUND

Pervasive and ubiquitous computing is a paradigm which, in part, encapsulates the process of embedding sensing, communication and computational services within environments [1], [2]. One application of this paradigm is Ambient Assistive Living (AAL). AAL is an approach to provide computational systems that are designed to assist persons within their environment [3], [4]. Many of these solutions operate by modeling and recognizing activities and offering assistance when an issue has been encountered.

Increasingly such assistive AAL solutions incorporate data-driven approaches to modelling and recognizing activities. In data-driven approaches, activity models are generated from datasets though statistical and machine learning. These datasets contain sensor records of user behavior and annotations/labels indicating when activities/behaviors of interest have been performed - such as wandering, agitation or failing. In order to produce useful activity models, annotations must be as accurate as possible [3], [5]–[12].

Traditionally, collection of annotations may be an arduous process [3], [5]–[10]. This process may involve retroactive analysis of collected information with subsequent definition of annotations or live annotation [3], [5]–[12].

Retroactive processes include using video/audio recordings, use of indirect observations of sensor data and self-recall.

The use of video/audio recordings involves reviewing a video/audio record of the environment that has generated a dataset [11]–[13]. This approach has shown to provide accurate results; however, it is time consuming and offers a potential invasion of privacy, reducing its utility.

Generating annotations [11], [12] though indirect observations of sensor data involves reviewing sensor records and identifying where likely events of interest have occurred. This approach has shown to be time consuming and error-prone due to the nature of manual sensor data analysis.

Use of a subject's self-recall [11], [12] has shown to be a less time consuming annotation process than reviewing video/audio recordings and use of indirect observations. However, this approach has shown to be the most error-prone of all the retroactive approaches. Such reduction in accuracy will limit the potential of generating usable activity models.

Generally, these retroactive processes are time consuming and so do not provide a scalable approach for generating annotations across a large number of smart environments.

Live annotation processes may include use of direct observations, a time diary, experience sampling or prompted labelling.

Direct observations [11], [12] involve assignment of a dedicated human observer to generate annotations. This approach produces accurate annotations but is extremely time consuming. This approach is not scalable as it requires a dedicated observer to be assigned to a subject or environment.

A time diary [11], [12] is a log where a subject records annotations manually. This approach is less time consuming than use of a human observer but is prone to less accurate annotation. This reduced accuracy may be attributed to a subject not immediately recording an annotation, issues related to incorrect time synchronization or a subject not being able to record an annotation when necessary.

Experience sampling [14] is a process where subjects are polled frequently to determine if an annotatable event has occurred in a previous window of time. This approach is more error prone than use of a time diary but is also less time consuming. However, such frequent polling has the potential to introduce fatigue and so reduce user engagement.

Prompted labelling [12] is a process where coarse changes in activity are detected and a subject is prompted to provide an annotation. This approach enables time accurate annotations to be generated in a manner that does not consume a large amount of time. However, application of this technique is currently limited and its reliance on a coarse change detection algorithm does not always provide reliable prompting.

Currently, direct observations and review of video/audio recordings produce the most accurate annotations but are time consuming and not scalable [11], [12]. Indirect observations and self-recall have been shown to provide the least accurate annotations [11], [12]. Experience sampling, use of a time diary and prompted labelling have moderate accuracy with moderate to low time consumption.

In order to increase quality and accuracy of annotation, while simultaneously reducing the effort and time required to generate them, a novel live annotation approach was employed in this study. Specifically, annotations are supplied, in-situ, by persons/caregivers that operate within an environment.

These annotations are provided through interaction with Near Field Communication (NFC) tags placed within an environment. If an activity/behavior of interest has occurred, a person in the environment may touch their smartphone, containing an annotation app, to a relevant NFC tag. This app will read the identification of the tag and record an annotation. This provides an intuitive annotation approach that enables accurate annotations to be generated by caregivers, while not without requiring any intensive training and enabling high levels of user engagement.

The approach was implemented within an AAL based evaluation platform to determine its utility. Specifically, detection of two behaviors that have health related implications were used as evaluation scenarios. These behaviors were falling [15] and sedentary activity [16].

The organization of the remainder of this paper is as follows; Section II presents the implemented platform, Section III discusses an evaluation of the annotation approach within an implemented platform. Section IV provides concluding statements and discussion on future studies.

II. NFC BASED ANNOTATION WITHIN A GENERIC BEHAVIORAL ALERTING PLATFORM

The implemented platform perceives an environment through ceiling-mounted thermal vision and sound level sensors. These sensors monitor factors such as environmental conditions, sound level, occupant count, occupant size and occupant movement. In addition to sensors, NFC tags are deployed to the environment to offer a mechanism to provide annotations when interacted with by a smartphone which has an annotation app installed.

The alert generation platform operates though a hybrid machine learning and rule based process. This hybrid process enables generic behavioral monitoring and alert generation, within the constraints of the metrics generated. The following Subsection presents an overview of this platform. This overview provides insight into the sensor platform, metrics produced, the annotation requirements and the integration of the annotation mechanism.

A. An overview of the implemented platform

The implemented platform has components located within an environment and on a cloud-based service. The environmental components are sensors, sensor-listeners and NFC tags.

Sensors are placed in each location which is to be monitored. The thermal vison sensor perceives the environment through a low-resolution grid of thermal data. Sensor-listeners transmit sensor data to the cloud service though use of a REST endpoint. On receipt of sensor data, the cloud service stores the raw data in a Time Series Database (TSDB) and a limited in-memory cache. All records stored within the TSDB are given a unique time index to a nanosecond resolution.

The data stored in the TSDB and in-memory cache is subsequently consumed by a feature extraction and fusion process. This process uses computer vision techniques to identify several features from the current frame, such as occupant count and occupant location. These thermal features are fused with appropriate sound level measurements and stored within the TSDB.

Generated features are subsequently consumed by a process that produces windows of metrics, on a per location basis. These windows of metrics are consumed by both a model training process and the alert generation mechanism.

The model training process uses a supervised machine learning process which relies on annotations, indicating when an event of interest is encountered. These annotations are provided by caregivers/persons who interact with NFC tags within an environment though use of a smartphone. Once these NFC tags are interacted with a relevant annotation is stored within the TSDB via the REST endpoint.

A single TSDB instance was used to store all metrics, sensor data and annotation data as this would provide a single,

coherent, source of time. This ensures that all metrics, sensor data and annotations are correctly synchronized. This synchronization allows annotations to accurately correspond with relevant sensor data and so are usable in training models. This TSDB may be clustered to offer a solution that scales with large volumes of data. In such a scenario, it's important to guarantee that all nodes have a synchronized clock to ensure time is modelled coherently. The chosen TSDB supports such clustered time synchronization.

The alert generation process operates on a window of live metrics which are classified by trained models and compared to caregiver-defined rules. Caregiver defined rules are specified through a web interface which interacts with the REST endpoint. The use of machine learning allows dynamic behavior detection models to be produced without the need to specify rules. A graphical representation of the of this platform is presented in Fig. 1.
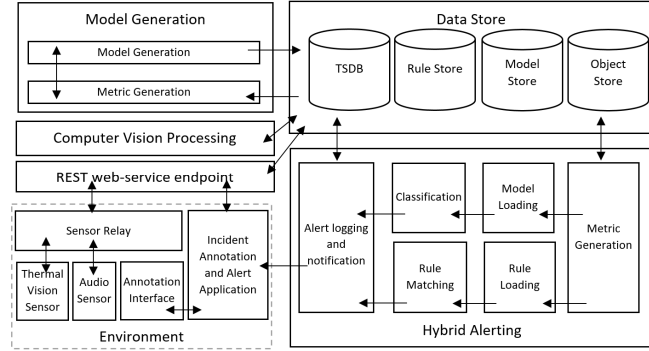

Fig. 1. The architecture of the implemented generic behavioral monitoring and alerting platform.

The sensors employed by this platform are detailed in the following Subsection.

B. Sensing infrastructure

In this platform, sensors observe an occupant's environment. These sensors are thermal vison and sound level monitors.
In this evaluation two different thermal vison sensors were used, sensor A and sensor B.

Sensor A can perceive the environment in an array of 16*16 pixels at a refresh rate of 4Hz. This sensor, and its perception of its environment is presented in Fig. 2.
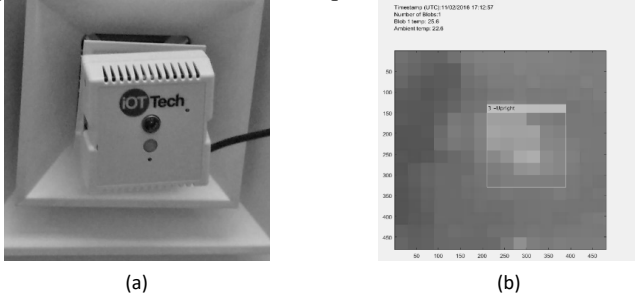

Fig. 2. Sensor A used in by this platform (a). The environment as perceived by sensor A, in a 16*16 array of thermal pixels (b).

Sensor B is able to perceive the environment through an array of 32*31 pixels at 20Hz. This sensor, and its perception of its environment is presented in Fig. 3.
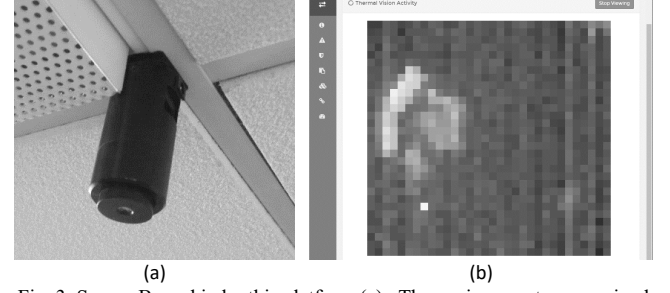

Fig. 3. Sensor B used in by this platform (a). The environment as perceived by sensor B, in a 32*21 array of thermal pixels (b).

These sensors have approximately a 90-degree field of view enabling observation of a large area, as determined by deployment height.

Both sensor A and B were used within the evaluation of this annotation mechanism. As such, the feature extraction and metric generation process can flexibly handle variation in size of frames of thermal data. Sensor A is a legacy component of this platform and is currently being replaced with the more capable sensor B. Future studies will only use sensor B. Both sensors send raw sensor data to the REST endpoint.

A final sensor is deployed to the environment, sensor C; this is a sound level monitor. This sensor samples sounds levels at a rate of 6 Hz. The sound level, in Root Mean Square (RMS), of these samples is calculated and forwarded to the REST endpoint. The RMS value accurately represents the volume of audio for a given sample [17].

All three sensors have a unique ID which is associated with a specific location in an environment, allowing fusion of generated features on a per location basis.

To accurately train data-driven activity recognition models from this sensor data, annotations must be gathered. These annotations have several requirements which must be catered for in this platform, these are presented in the following Subsection.

C. Annotation requirements

The hybrid alert generation component uses a supervised machine learning based process to partially provide its function. This supervised process requires annotations to be specified that indicate occurrence of behavioral events that alerts are desired for. These annotations are subsequently used to identify windows of metrics that are used to train models for classification of behaviors. These models are used to classify live windows of metrics, indicating whether a likely event of interest has happened.

This platform is intended to be deployed into a facility where caregiving/observing employees would be able to supply annotations when a behavior of interest occurred. However, manually supplied annotations may be inaccurate or neglected. To address this a low effort, secure and accurate mechanism of supplying these annotations was required.

In this platform, annotations have two classes of requirements to be catered for, functional and logistical. Functional requirements relate to the suitability of the annotation while logistical requirements concern the process of annotation.

Functionally, annotations to be used in this platform must offer an accurate timestamp and incorporate sensor location. Incorporating sensor location into the annotation ensures that windows of sensor metrics will only include relevant sensor data. Accurate timestamps allow identification of relevant windows of metrics.

Logistically, annotations must be provided with low effort, in a secure manner and through an assured process. A low effort process increases the likelihood that annotations will be provided. A secure mechanism ensures that only authorized individuals may provide an annotation, helping to ensure accuracy. Finally, an assured process safeguards that annotations are not provided in error.

These requirements have been considered and addressed by the annotation process described in the following Subsection.

D. NFC based dataset annotation

The NFC based annotation mechanism devised uses NFC tags deployed to an environment to serve as annotation interfaces. To function, these NFC tags first need to be placed in a location and enrolled through use of an administrative app deployed to a smartphone. During enrollment, a list of sensor locations is retrieved from the cloud service and presented to the app user. The app user can then select the sensor location that the NFC tag will be associated with. Once the location is specified a unique tag ID will be generated by the cloud service and recorded against the specified sensor location, the app's user will be prompted to touch the NFC tag with their smartphone. When the smartphone is touched to the NFC tag, the unique ID will be written to it and tag deployment is complete.

Caregivers are provided a smartphone equipped with a caregiver app. If a behavior of interest had occurred, caregivers may touch their smartphone to a relevant NFC tag. The app will then send the unique ID of that tag to the cloud service which will store an annotation within the TSDB. This annotation contains the sensor location associated with the NFC tag, the ID of the NFC tag and a server side timestamp. A single timestamp is required in this application, reducing annotation effort. The use of a smartphone app ensures that only authorized users provided can provide annotations. The apps that interact with the NFC tags are presented in Fig. 4.
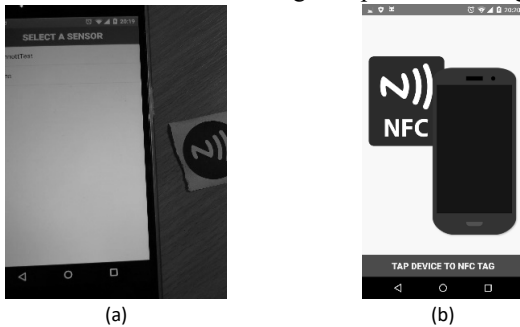

Fig. 4. (a) depicts an NFC tag being associated with a sensor though an admin app, (b) depicts an NFC tag explicitly being read by the app to record an annotation though a caregiver app.

To reduce the complexity of the annotation process, the caregiver app is designed so that it does not need to have the app active to operate. Instead the device only needs to be in a woken state when it is touched to the tag. Requiring the device to be woken has the benefit of reducing the chance that an annotation would be provided in error, offering an assured process. If required, this app can be modified to incorporate an onscreen confirmation of annotation at the cost of some additional effort. Additionally, annotations may be manually specified and modified though use of a web interface.

These annotations are used in conjunction with windows of sensor metrics to train classification models. The feature extraction and metric generation process is described in the following Subsection.

E. Feature extraction and metric generation

The thermal vison sensors used within this platform produce raw data which requires features to be extracted. Specifically, a computer vison process is employed to identify potential occupants that have been observed.

A computer vision process retrieves raw thermal vison data from the cloud service and performs comparative blob and entity detection against it.

This process performs initial filtering which eliminates thermal pixels that have gradually increased temperature or are outside of the likely human emissive temperature range. The filtered pixel readings are compared to historic thermal vison data to determine if any pixels are substantially warmer than pixels in that location, within the reference frame.

A bobbling process is then performed against these identified thermal pixels. Blobs that have a thermal pixel count less than 6 are discarded as these would be too small to represent a human at the intended sensor deployment heights. The remaining blobs represent occupants and provide several features including thermal pixel count, blob emissive temperature, suspected fall incident duration, blob width, blob height, blob aspect ratio and blob coordinates. In addition, Further information on this computer vison process may be obtained from [18]–[20] .

The features extracted from the computer vison process are combined with sound level data to produce windows of metrics. These windows of these metrics are consumed by the alert generation process. Further information on this alerting process is given in the following Subsection.

F. Hybrid supervised machine learning and rule based alert generation

The alert generation process uses both supervised machine learning and rule based alert generation.

Rules are specified by specially trained users on a per-window and per sensor location basis. These rules are a collection of comparative operators that are applied to metrics. These rules are applied to windows of live metrics and if they are met an alert is generated. These rules are defined through use of a web interface. Additionally, annotations can be used to indicate metrics that may be used within rules.

The supervised machine learning based process requires annotations to correlate metrics and events of interest, to train

a model used later in classification. Specifically, annotations are used to produce a set of *positively correlated* windows of metrics that extend from the point of annotation into the past, as dictated by static window sizes configured at a system-wide level. Windows of metrics that don't intersect with any annotations are segmented and placed in a set of *negatively correlated* metrics. These positive and negative sets of metrics are used to train detection models for each sensor location. Additionally, positive and negative sets of metrics from all sensors are combined to produce a global model. This global model functions as a fallback model for sensor locations that have insufficient trained data. The alert generation process classifies live windows of metrics against the trained model, where they exist for that window size.

The employed supervised machine learning process uses a Multi-Layer Perceptron and was integrated into the evaluation platform through used of the Java-based Neuroph library [21].

The hybrid model employed by this alert generation process allows alerts to be generated though rules while the machine learning process is being trained though annotation. Additionally, alert levels may be specified for the machine learning process and on a per-rule basis. These levels may be alert, warn and log.

An evaluation of this annotation mechanism within this platform is presented in the following Section.

III. EVALUATION

To evaluate the utility of this annotation mechanism several factors must be considered; suitability of annotations generated, time required to perform annotation and reliability of annotation.

In order to gauge the suitability of annotations generated, a real-world evaluation of the generic behavioral monitoring and alerting platform was performed with a focus on its machine learning based process. Two behaviors of interest were targeted; falling and sedentary activity.

To enable this evaluation, sensors and NFC tags were deployed to two environments. Once deployed real world data was collected over the course of three weeks. These environments were a kitchen and living room. Sedentary behavior detection was evaluated in the kitchen environment and fall detection was evaluated in the living room environment. Sensor A served as the thermal sensor in the kitchen environment. Sensor B was deployed to the living room environment where no sound level monitor was deployed.

Two machine learning based models were produced to generate alerts for each target behavior detection. The sedentary behavior model was trained with 150 windows/annotations and the fall detection model was trained with 30 windows/annotations. These trained models were applied to real-world data generated by observing an occupant. During this period of observation, several incidents were simulated in addition to a large period where no incidents occurred.

The results of this evaluation are presented in Table 1.

TABLE I.   RESULTS OF EVALUATION OF THE BEHAVIORAL DETECTION AND ALERTING PLATFORM

	Number of Incidents	*Detected Incidents*	*False positives*	*Detection rate*	*Sensitivity/ Specificity*
Machine learning					
Fall Detection	15	12	2	80%	80%/83.33%
Sedentary Behaviour	10	10	0	100%	100%/100%

This evaluation shows that the platform can detect and generate alerts for a range of behaviors with reasonable accuracy, proving the suitability of annotations generated by this mechanism. However, issues related to false positives and missed detections were present.

False positives were present in the fall detection evaluation. These false positives are unrelated to the annotation or training process but instead were due to inadequacies in the computer vison based process during live detection. Specifically, these errors were generated by uncommon postural changes, such as rapidly laying on a sofa from a standing position. Additionally, fall detections were related to a subject not being within the thermal sensor's field of view or having body part being partially occluded by environmental elements.

In order to evaluate the time required to perform annotation and reliability of recording annotations, a set of 60 additional annotations were performed. When performing these annotations, the time taken to perform the annotation was recorded. Following this the number of recorded annotations was verified and the average annotation time was calculated. When producing this set of annotations, a smartphone was placed within the pocket of the observing individual, emulating a realistic usage scenario. This individual was in the vicinity of a NFC tag.

In this evaluation, the average annotation process took 4.3 seconds and all 60 intended annotations were successfully recorded showing this mechanism be reliable.

This evaluation shows the annotation process can produce annotations that are suitable for use in training models. Additionally, this annotation process performs reliably and not as time consuming as existing annotation approaches [3], [5]–[12] .

This platform requires further evaluation and improvement, this is discussed within the following Section.

IV. CONCLUSION AND FUTURE STUDIES

This paper presents an NFC based annotation mechanism that was incorporated into a hybrid generic behavioral alerting platform. The alerting component incorporates machine learning and rule based processes to offer flexible specification of behaviors that require generation of alerts.

Behavioral rules are specified though a web interface and annotations, used to train the supervised machine learning process, were provided in-situ though the combination of NFC tags and a smartphone app.

An evaluation showed good accuracy in behavioral detection when a machine learning based processes

incorporated annotations produced by this NFC based annotation process. Additionally, this evaluation showed the annotation process produced usable, reliable and accurate annotations.

However, deficiencies are present related to the NFC based annotation platform. The most notable of which is the current requirement for the smartphone generating annotations to have an active internet connection. This limitation is due to the use of time synchronization though the TSDB, meaning the local time of the mobile device can't be relied upon for cached/offline annotations. This issue could be addressed by incorporation of logic to determine the time offset between cached/offline annotations and the TSDB. This logic would operate when the smartphone app regains an active internet connection and uploads the cached annotations.

Currently the annotation platform only supports a single NFC tag per location and so models can only be trained to detect one behavior. This is a limitation of the current implementation and support for enrollment of multiple classes of NFC tags, representing different behaviors, will be added to address this.

Future evaluation is planned within a real end-user environment. Specifically, a facility with over 100 occupants has been selected and plans for a trial are being agreed upon. This facility employs several caregivers, which will supply annotations to the machine learning component of the alert generation process. Additionally, a small number of rules will be implemented to provide fallback support for alert generation. The occupants of this facility exhibit multiple behaviors that will require specific alert generation, as such support for multiple annotations will be incorporated.

Finally, there are plans to adopt the NFC based annotation mechanism to future studies. In these studies, its utility will be evaluated compared to other mechanisms of dataset annotation.

REFERENCES

[1] M. Weiser, "The computer for the 21st century," *Sci. Am.*, vol. 265, no. 3, pp. 94–104, 1991.

[2] M. Satyanarayanan, "Pervasive computing: Vision and challenges," *IEEE Pers. Commun.*, vol. 8, no. 4, pp. 10–17, 2001.

[3] N. Roy, A. Misra, and D. Cook, "Ambient and smartphone sensor assisted ADL recognition in multi-inhabitant smart environments," *J. Ambient Intell. Humaniz. Comput.*, pp. 1–19, 2015.

[4] European Commission, "The Ambient Assisted Living (AAL) Joint Programme." [Online]. Available: http://ec.europa.eu/information_society/activities/einclusion/docs/ageing/aal_overview.pdf.

[5] L. Chen, J. Hoey, C. D. Nugent, D. J. Cook, and Z. Yu, "Sensor-Based Activity Recognition," *IEEE Trans. Syst. Man, Cybern. Part C (Applications Rev.*, pp. 1–19, 2012.

[6] P. Rashidi and A. Mihailidis, "A survey on ambient-assisted living tools for older adults.," *IEEE J. Biomed. Heal. informatics*, vol. 17, no. 3, pp. 579–90, May 2013.

[7] G. Acampora, D. J. Cook, P. Rashidi, and A. V Vasilakos, "A Survey on Ambient Intelligence in Health Care.," *Proc. IEEE. Inst. Electr. Electron. Eng.*, vol. 101, no. 12, pp. 2470–2494, Dec. 2013.

[8] M. P. Poland, C. D. Nugent, H. Wang, and L. Chen, "Smart Home Research:Projects and Issues," *Int. J. Ambient Comput. Intell.*, vol. 1, no. 4, pp. 32–45, Jan. 2009.

[9] H. Alemdar, H. Ertan, O. D. Incel, and C. Ersoy, "ARAS human activity datasets in multiple homes with multiple residents," in *2013 7th International Conference on Pervasive Computing Technologies for Healthcare and Workshops*, 2013, pp. 232–235.

[10] F. J. Aparicio-Navarro, K. G. Kyriakopoulos, and D. J. Parish, "Empirical study of automatic dataset labelling," in *Internet Technology and Secured Transactions (ICITST), 2014 9th International Conference for*, 2014, pp. 372–378.

[11] M. Stikic, D. Larlus, S. Ebert, and B. Schiele, "Weakly supervised recognition of daily life activities with wearable sensors," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 33, no. 12, pp. 2521–2537, 2011.

[12] I. Cleland, M. Han, C. Nugent, H. Lee, S. McClean, S. Zhang, and S. Lee, "Evaluation of prompted annotation of activity data recorded from a smart phone," *Sensors (Switzerland)*, vol. 14, no. 9, pp. 15861–15879, 2014.

[13] T. Plötz, C. Chen, N. Y. Hammerla, and G. D. Abowd, "Automatic synchronization of wearable sensors and video-cameras for ground truth annotation - A practical approach," *Proc. - Int. Symp. Wearable Comput. ISWC*, pp. 100–103, 2012.

[14] T. Goetz, A. C. Frenzel, N. C. Hall, U. E. Nett, R. Pekrun, and A. A. Lipnevich, "Types of boredom: An experience sampling approach," *Motiv. Emot.*, vol. 38, no. 3, pp. 401–419, 2014.

[15] S. S. Khan and J. Hoey, "Review of Fall Detection Techniques: A Data Availability Perspective," no. May, 2016.

[16] J. P. Buckley, A. Hedge, T. Yates, R. J. Copeland, M. Loosemore, M. Hamer, G. Bradley, and D. W. Dunstan, "The sedentary office: an expert statement on the growing case for change towards better health and productivity," *Br. J. Sports Med.*, p. bjsports–2015, 2015.

[17] E. Wold, T. Blum, D. Keislar, and J. Wheaten, "Content-based classification, search, and retrieval of audio," *MultiMedia, IEEE*, vol. 3, no. 3, pp. 27–36, 1996.

[18] J. Synnott, C. Nugent, and P. Jeffers, "A Thermal Data Simulation Tool for the Testing of Novel Approaches to Activity Recognition," in *Ambient Assisted Living and Daily Activities SE - 2*, vol. 8868, L. Pecchia, L. Chen, C. Nugent, and J. Bravo, Eds. Springer International Publishing, 2014, pp. 10–13.

[19] J. Rafferty, J. Synnott, and C. Nugent, "A Hybrid Rule and Machine Learning Based Generic Alerting Platform for Smart Environments. Engineering in Medicine and Biology Society," in *Engineering in Medicine and Biology Society (EMBC), 2016 38th Annual International Conference of the IEEE*, 2016.

[20] J. Rafferty, J. Synnott, C. Nugent, G. Morrison, and E. Tamburini, "Fall Detection Through Thermal Vision Sensing," in *Ubiquitous Computing and Ambient Intelligence: 10th International Conference, UCAmI 2016*, C. R. García, P. Caballero-Gil, M. Burmester, and A. Quesada-Arencibia, Eds. Springer International Publishing, 2016, pp. 84–90.

[21] "Neuroph." [Online]. Available: http://neuroph.sourceforge.net/.