

“THEME ARTICLE”, “FEATURE ARTICLE”, or “COLUMN” goes here: The theme topic or column/department name goes after the colon.

Lifelong Learning in Sensor-based Human Activity Recognition

Juan Ye and Simon Dobson
School of Computer Science,
University of St Andrews,
UK

Franco Zambonelli
Dipartimento di Scienze e
Metodi dell'Ingegneria,
Universita' di Modena e
Reggio Emilia, Italy

Human activity recognition systems will be increasingly deployed in real-world environments and for long periods of time. This significantly challenges current approaches to human activity recognition, which have to account for changes in activity routines, evolution of situations, and of sensing technologies. Driven by these challenges, in this article we argue the need to move beyond learning to *lifelong machine learning* – with the ability to incrementally and continuously adapt to changes in the environment being learned. We introduce a conceptual framework for lifelong machine learning to structure various relevant proposals in the area, and identify some key research challenges that remain.

A key enabler for many critical ubiquitous computing applications, in areas such as healthcare (e.g., mental well-being assessment, and clinical assessment on cognition and mobility) and home automation (e.g., automatic heating configurations) is sensor-based human activity recognition (HAR): the ability automatically to recognise and predict users' current and future activities from data collected from a wide range of wearable and ambient sensors. Based on inferred user activities, applications can deliver customised, in situ services in an automatic and unobtrusive manner.

Activity recognition has been extensively studied, and significant progress has been made through the use of modern data-driven techniques, including Hidden Markov Models, Conditional Random Fields, Support Vector Machines, and – more recently – deep neural networks. These techniques have achieved promising results in learning complex correlations between sensor data and activities of interest: their success in recognising human activities in lab and testbed settings is now enabling the move towards large-scale, in-the-wild, and long-term deployment of activity recognition systems [17]. This move however comes with its own challenges, notably

that neither the environment being observed, nor the population or pattern of activities being recognised, remain constant. This creates a need for HAR systems capable of embedding the means of automatic adaptation to changes, *i.e.*, lifelong learning.

On the one hand, lifelong learning in HAR covers the lifespan of a user, by continuously learning and adapting to changes in the user's health condition, activity routine, or living environments; on the other, it also covers the lifespan of a HAR system itself, by responding to the changes in sensor technologies, in sensor deployments, and in the operational environment.

Lifelong learning in HAR will be necessary to enable large-scale, long-term, and sustainable deployment of activity-aware applications. It directly tackles personalised, adaptable service delivery. For example in mobile health applications, it would be desirable to design care plan, intervention or treatment that is customised to individuals' health conditions and more importantly can dynamically evolve responding to the change of their conditions; e.g., the onset of a disease's symptoms or the recovery of a disease. Without such adaptation, an activity recognition system may fail to detect or misclassify activities, leading to ineffective treatment or to undesirable consequences. For example, failing to adjust rehab exercise in time for patients suffering knee osteoarthritis can cause more damage on the injured parts of the knee [5].

Lifelong learning will also of course bring with it a whole new set of challenges and opportunities that have not been properly studied before. These range from sensing technologies, system infrastructure, and privacy and security. In this article, we concentrate on how such challenges impact on the learning algorithms used for HAR, from which we argue the need for lifelong machine learning as a field of study, present opportunities and discuss open research questions in this new field of learning.

EMERGING ISSUES FOR REAL-WORLD AND LONG-TERM HAR

In its simplest form, human activity recognition can be phrased as a classification problem: mapping collections of raw sensor data or sensor features (X) into activity labels (y) through a learned or specified correlation model ($h: X \rightarrow y$), typically based on datasets containing ground truth information.

The specific challenges in real-world and long-term HAR are driven by the *constant, unavoidable, and unpredictable* changes in sensors, environments, and human activities. These changes result in the evolution in class space, in the correlation model, and in the feature space model, which refer to class evolution ($y \rightarrow y'$), concept drift ($h \rightarrow h'$), and feature evolution ($X \rightarrow X'$), respectively, also implying the need to continuously collect new ground truth information (Figure 1).

Class Evolution

Most of the current HAR approaches follow a well-established methodology [14]:

- **Deployment:** pre-define a closed set of activities of interest, and select a range of ambient and/or wearable sensors that can potentially detect the activities.
- **Model training:** collect sensor data for a short period of time, annotate them with activity labels, and build a computational model to correlate sensor data with activities by defining expert knowledge or training a machine learning technique.
- **Activity recognition:** recognise current activities from real-time sensor data.

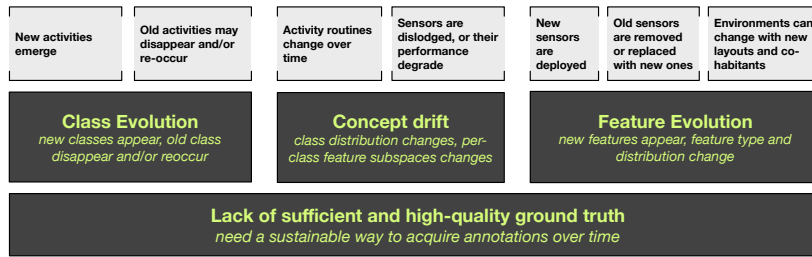


Figure 1: Emerging issues for HAR driven by constant and unpredictable changes in real-world and long-term deployment

A key assumption in such process is that *the set of activities is fixed*, which does not necessarily hold when we are starting to deploy the system in open and versatile environments for a longer period of time. The non-deterministic and dynamic nature of human behaviour can lead to the occurrence of new activity patterns or even new types of activities [4]; for example, a user may start playing a Kinect fitness game to be more active, or he/she can start having bath rather than simply getting a shower. Indeed, the whole point of the HAR system may be to monitor (in the case of deteriorating capability), or provoke (in the case of training applications), changes in the ways users conduct themselves. These changes in activities can give rise to class evolution – the emergence of new activity classes, and the disappearance and recurring of old activities, which requires systems with the ability to recognize and incorporate new classes, and possibly to dismiss old ones; e.g., users might not be able to run any more due to their knee joint problem.

Concept Drift

Another general assumption in the traditional activity recognition methodology is that *the computational model for activity recognition is static once defined or trained*. However, individuals can perform the same type of activities in very different ways over the long term. For example, degrading health conditions may affect the way a person walks, or different dietary requirements may affect the way they cook dinner.

These changes in the way activities are performed induce concept drift, *i.e.*, changes in activity distribution. The lack of adaptation to it will lead to a degradation of performance in the long run, and in the inability to correctly classify existing classes, or misclassify them as new, unknown, classes, resulting in dissatisfaction of the system as a whole.

Feature Evolution

Finally, short-term, HAR can generally assume that the characteristics of the environment, and their detection, are both stable and reliable. In real-world environments, sensors are exposed to all sorts of unexpected “hazards”, including being interacted with by non-primary users or pets, moved, misused, and broken, all of which can lead to unpredictable readings [19]. Also, we might deploy new sensors and remove or replace old sensors over time. Environments can change with the re-arrangement of furniture; *e.g.*, moving a sofa to make more space, or the change with co-inhabitants or visitors, *e.g.*, more frequent visit from carers, or newly introduced night “sleep-in” service. All such changes in sensors and environments will lead to unexpected sensor stream, which requires the activity recognition model to respond to. These changes often lead to feature evolution – the inclusion of new sensor features, and the change in feature distribution to each activity class.

Lack of Sufficient and High-quality Ground Truth

To build a robust activity recognition model, most existing data-driven techniques require a large amount of training data; that is, annotated sensor data with activity labels. For all the above changes, if we could continuously collect sufficient and high-quality annotations, we could re-

train the model all the time and thus the changes are less a challenge. However, acquiring annotations has been well acknowledged as a time- and effort-consuming task, which either relies on users’ constant self-reporting on what they are doing or recording users’ activities via videos that are later annotated by the others. Neither of these approaches are a sustainable solution for a high volume of users in long-term deployment.

TOWARDS LIFELONG LEARNING IN HAR: FRAMEWORK AND APPROACHES

Driven by the need to learn a variety of tasks over time and improve learning performance, lifelong machine learning has been identified as a future challenge in the field of machine learning [13] and is attracting a rapidly increasing amount of interest. It is defined as “systems that can learn many tasks over a lifetime from one or more domains. They efficiently and effectively retain the knowledge they have learned and use that knowledge to more efficiently and effectively learn new tasks” [18]. Putting into the context of HAR, we define

lifelong learning in human activity recognition as being characterised by incremental and continuous learning of new types or variations of activities while also adapting to the change in circumstances of users, sensors and environments.

Lifelong learning in HAR supports continuous learning and is the critical enabler for allowing widespread deployment of sensor systems and their use to support human activities and other critical applications. It exhibits a substantial challenge due to the combination of challenges in change management and the lack of sufficient ground truth presented in Figure 1.

Lifelong learning in HAR is relevant to a range of machine learning fields, especially online learning and transfer learning. Online or incremental learning refers to the ability to learn models incrementally from data in a sequential manner. It can be considered part of lifelong learning and directly targets concept drift in Figure 1. But the lifelong learning in HAR proposed here has a wider scope than online learning, which also needs to tackle ground truth challenge; that is, how to acquire annotations on changes in data in order to enable model update.

Transfer learning – how to transfer and adapt an activity model from a source domain to a target domain [8] – has recently been widely applied to HAR. The main difference is that lifelong learning in HAR puts emphasis on dynamic adaptation of a model over time, which requires *continuous* model adaptation and automatic identification of when and what adaptation to be needed.

General Reference Framework

Taking a holistic view on tackling the above challenges in Figure 1, we identify three key components: change detection, change annotation, and change adaptation, and propose a conceptual architecture that extends a traditional activity recognition framework and enables to frame these key components to enable lifelong learning in Figure 2.

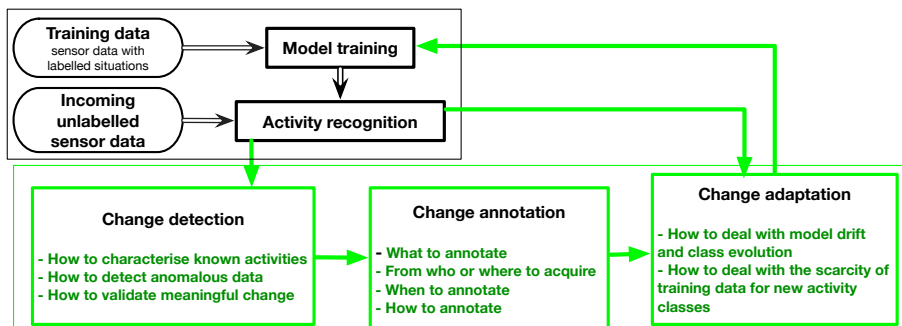


Figure 2 A general reference framework for lifelong learning in human activity recognition. It presents three key components (in green boxes) and shows how they are fit with the classical data-driven activity recognition framework (in black boxes).

- **Change detection** – to detect changes in feature space, classes, and correlation. This often involves
 - characterising sensor data for existing activities;
 - detecting anomalous data; *i.e.*, those that do not conform to the existing characteristic patterns; and
 - identifying data for a potentially new activity from all the anomalous data; *e.g.*, distinguishing from sensor noise.
- **Change annotation** – to learn semantics or labels of detected changes. This concerns a wide range of questions, including
 - what data to annotate; *i.e.*, how to select representative data for annotation, assuming that we cannot query each identified data point;
 - from who or where to get annotations; *i.e.*, whether we query users for their own activities, domain experts, or other sources (such as other users’ data or web resources);
 - when to annotate; *i.e.*, if we assume to query users, what is the right moment to query them so that they are most likely to attend the query;
 - how to annotate; *i.e.*, what form to query, such as free text or multimedia input; whether to enable multi-labels for one activity instance, or restrict to one label; how to guarantee semantic consistence and correctness of user annotations.
- **Change adaptation** – to incrementally adapt the activity recognition model to recognise and respond to the discovered changes. This often involves
 - dealing with concept drift and class evolution;
 - working with imbalanced, insufficient data for new activity classes; *i.e.*, building a reliable activity recognition model for new emerging activities, when at the early stage there often exists a limited number of data points for a new activity class, which can be much fewer than those for existing classes.

These three components form the core of lifelong learning in HAR. More importantly, they enable an iterative and incremental update to the original activity model with newly discovered and annotated activity data, rather than one iteration from training to recognition. The framework in Figure 2 can serve as a reference to design future sensor-based human activity recognition systems. In the extreme situation, the system could start with no pre-defined activities at all and incrementally learn a personalised set of activities for each user. To understand the existing research activities towards the emerging field of lifelong machine learning, we discuss some representative works devoted to directly targeting sensor-based *lifelong, unbounded, unknown, unseen* activity discovery or recognition.

Change Detection

Change detection has been mainly regarded in terms of an anomaly detection problem. There exist many anomaly detection techniques [6], among which clustering, one-class classifiers, and sequential pattern mining are the most popular approaches.

Clustering. Clustering is one of the commonly used techniques for unsupervised anomaly detection; *i.e.*, without the need of training data on anomaly classes. Its basic principle is to partition normal data points into multiple clusters, and any data point that falls outside and is far away from the clusters is considered as “abnormal”. These abnormal data points can potentially map to new activities.

Abdallah *et al.* propose a system called AnyNovel that applies incremental clustering model to discover and learn new physical activities, including sitting, walking, standing and lying from accelerometer data [1]. An Expectation-Maximisation (EM) clustering technique is used to cluster real-time data streams and identify a newborn cluster that lies outside the existing cluster decision boundary. Then a cohesion and separation criterion is utilised to validate whether this newborn cluster forms into a cohesive and dense group while keeping a considerable separation from the existing classes. The criterion is assessed through the cluster's density, weight, and temporal characteristics.

Similarly, Gjoreski *et al.* [10] and Ye *et al.* [14] have used a clustering technique to enable real-time clustering of streaming data. To validate clusters for potential new activities, Gjoreski *et al.* have proposed two temporal assumptions on human activities; that is, a human activity usually lasts for a certain period of time and there should not be frequent transitions between activities. With these assumptions, they have filtered short outliers and been able to more accurately discover meaningful clusters.

One-Class Classification (OCC). OCC is another common technique to detect new activities. The main difference from the above clustering technique is that with the support of sophisticated kernel functions, OCC has a capacity of modelling complex, nonlinear class boundary.

Yin *et al.* apply One-class Support Vector Machines (OSVMs) with a Radian Basis Function (RBF) kernel to discover abnormal activities such as falling or slipping on a wet floor from body-worn sensors that collect light, temperature, sound, and acceleration data [20]. An OSVM is trained to find a hypersphere that contains most data points corresponding to pre-defined normal activities. Then a Hidden Markov Model (HMM) is built to estimate the maximum likelihood of the current data trace by the existing model. If the likelihood is below a threshold, the new data trace is considered abnormal.

Sequential Pattern Mining. Sequential pattern mining is to discover sequences of frequent occurred itemsets by searching through the whole dataset. Rashidi *et al.* have combined a compression-based sequence mining with clustering algorithms to discover new activity patterns on binary sensor events collected from a smart home environment [16]. Using the minimum description length principle, they have found patterns that best compress the dataset and with sufficient frequency and length. A new activity is detected, if the current sensor sequence is not similar to the mined patterns, where the similarity is measured in edit distance.

Change Annotation

Change annotation concerns acquiring annotations to discovered changes so that we can update the activity model. Most of the current approaches focus on selecting the right data points for annotation.

Active Learning. Active learning is to alleviate the labelling effort by selecting the most informative instances to annotate, which can be the most uncertain instances; *e.g.*, lying between the boundary of classes, or the most representative instances; *e.g.*, lying right in the centre of a class.

Alemdar *et al.* have applied active learning strategies to select the most uncertain instances for annotation and iteratively update a HMM to infer daily activities [2]. Extending the basic activity learning strategies, Cheng *et al.* have proposed a density-weighted method to ensure the selected instances are representative of the underlying data distribution [7]. Hossain *et al.* have also extended the basic strategies with Silhouette Coefficient so as to select the instances that have formed into a cohesive cluster [12]. Both extensions have demonstrated an improved recognition accuracy with reduced labelling effort.

Mixture Model. Nguyen *et al.* propose a supervised activity discovery system, called *SuperAD*, which can identify both uncertain and unseen instances from the unlabeled data pool [15]. A supervised classifier is trained on labeled data and then predicts temporary labels on unlabeled data. Then a Gaussian Mixture Model (GMM) is trained on the pool of temporally labeled instances and estimates their likelihood by their corresponding submodel. A low likelihood on an

instance suggests the uncertainty of its temporary label and this instance will be selected to be re-annotated as either pre-defined activities or new activities. *SuperAD* has been demonstrated outperforming active learning strategies and random selection on selecting instances for annotation.

Change Adaptation

Once the changes are annotated, an activity recognition model needs to be extended and updated to incorporate these changes. For clustering-based techniques, it is less a problem in that a new activity is just another cluster(s) and it is easier to extend the activity recognition model. Most classifiers do not have the capacity of accommodating new classes. There exist classifiers that support incremental or online learning, including decision tree, naive Bayes, and neural networks. They allow updating and tuning parameters of their models from new pieces of data; that is, changing the distribution of classes or weights on features to predict classes. Also constrained programming has been applied to enable self-adaptive sensor configuration responding to the changes in environmental context [11]. Rarely do they support the evolution of their structure with new classes.

Evolving Classifiers. Andreu *et al.* have applied evolving fuzzy-rule-based classifiers, also known as eClass family, to support real-time, online learning of human activities from accelerometer data collected from wearable sensors [3]. Each activity is composed of a set of fuzzy rules, each of which captures a correlation with specific features generalised from accelerometer data. The rules are not fixed and can evolve over time with the change in correlations based on dynamic streaming data. Also it is flexible to introduce new rules for new activities or remove one if the activities are no longer relevant.

Zero-Shot Learning. Learning a model for new activities often suffers from the scarcity of training data. Zero-shot learners can be applied to deal with these problems by learning previously unseen activities with a few labelled instances or even without any training data. They are built by combining data-driven and knowledge-driven approaches. The data-driven model can be first trained on data of pre-defined classes and transferred to recognise a new class based on the knowledge-driven model, where the knowledge about new classes may come from text corpus [7], or domain experts.

Cheng *et al.* have adapted a zero-shot learner to recognise a new activity with limited training data [7]. A knowledge-driven model encodes the semantic relationship between high-level activities and low-level sensor attributes generated from accelerometer data. To include a new activity, domain experts and developers need to manually add new attributes and update the activity-attribute matrix with manually specified relationship between attributes and this new activity.

Summary

From the above overview, we can see change detection is the most well-studied with a similar principle; that is, characterising existing data patterns either in the form of a cluster, a hyper-space, or frequent sequential itemsets, and identifying any data that does not conform to the patterns. Few works have directly targeted at change annotation, only with a focus on selecting the most informative data points for annotation to reduce the annotation burden. However, most of them assume that users are always available to provide accurate annotations. None of them has conducted real-world experimentation on annotation acquisition. In change adaptation, the problem with imbalanced, insufficient training data for newly discovered activities has been acknowledged, but few targets concept drift and class evolution. This gives the motivation of this paper to present an overview of where the research is towards a lifelong learning system.

CHALLENGES AND OPPORTUNITIES

Beside existing proposals, many open research questions exist towards lifelong learning in HAR, pointing to potential research directions.

Distinguish Meaningful Change from Noise

Due to the pervasive problem with sensor noise – which is something that it’s impossible fully to engineer away – the key challenge in change detection is how to differentiate noise from “meaningful” changes in patterns. Beyond the attempted validation-based approaches, another possible direction to address this challenge is extension of decision boundary of existing classes. A boundary can be as simple as the union of existing clusters’ boundaries; that is, any data lying within this delineated global boundary is considered as normal. The global boundary might work well if the already identified clusters are close enough [6]. Another approach is to draw a slack space around the boundary of each cluster. The challenge is how to decide the threshold on the width of a slack space. A general approach is to adapt the threshold based on the false alarm rates; that is, the percentage of outliers are incorrectly inferred within an existing cluster [18]. However, this can be difficult for emerging activity recognition when there might not be enough ground truth to measure false alarm rates.

Informed Change Discovery

Finding out the cause of any anomaly is another significant difference between change detection in HAR and a general anomaly-detection problem. “Cause” here can refer to sporadic sensor noise or systematic sensor fault due to sensor recalibration or dislodgement [19], changes in an environment, new patterns for an existing activity class, and new types of activities. Each of these specific causes to anomaly might trigger different system-level actions; e.g., repairing or replacing sensors, or updating an activity recognition model. This brings us to another direction, i.e., *informed change discovery*. To achieve this, we might look into how to make use of the knowledge mined on both sensor data and activity patterns to supervise and explain the newly discovered anomaly. Existing work has demonstrated the value of domain knowledge such as temporal characteristics on human activities and sensor events in detecting new activities [1, 10] and sensor fault [19]. However, their main limitations are that (1) the knowledge is either manually crafted or based on simple statistical techniques, and (2) the knowledge is assumed to be fixed and will not change over time. To overcome these two limitations, we could investigate evolving knowledge mining techniques. For example, with the support of dynamic topic models [6] to track evolution of activity routine and sensor performance over time, we might be able to locate the root cause of a change in sensor features either as sensor failure or the change of routine.

Opportune Moment for User Annotation

When to annotate can be a rather difficult task, given that users might not be able to attend annotation requests all the time. To address this problem, it might be useful to look at the recent intelligent notification work; that is, finding the opportune moment when users are most likely to attend the query. The prediction of the opportune moment can be based on contexts including the time of day, location, the current physical activities, and the social company. Activity annotation is different from the message notification application, as it does not only require the user to attend to the query, but also to be able to provide the right activity label, which might incur extra effort from them. One way to integrate a notification system with activity annotation is to select the instance that is not only most informative but also closest to the query time so that the user might still be able to remember their activities. In addition, we can further leverage the context management in a notification system by presenting context as a clue to remind the users of their activities.

Quality of Annotations

It is challenging to allow users to annotate their own activities in free-text style. On the one hand, simply asking users “what are you doing at the moment” might not guarantee a relevant answer, which is different from annotation at the training phase where we have a set of predefined labels and simply ask the user to select whatever applies. The difficulty is that we might

not be able to foresee what new activities users are doing, so we cannot provide them with the predefined label set. Users might come up with all sorts of inconsistent expressions for their current activities; e.g., “boiling water”, “cooking a pasta”, or “preparing dinner”. On the other hand, users often multi-task where one activity can be interleaved with another activity concurrently; e.g., if a user is reading a book on a train, then possible answers to their current activities can be “reading a book”, “on a train”, or “reading a book on a train”. Therefore, the questions are how to guide or constrain annotations from users, how to extract relevant labels from their input, or both. The approach to check semantic correctness of user annotations might help but more thoughts should be put into the direction; e.g., should we allow multiple labels for each activity instance? This practice has been used in computer vision, which encourages capturing all the visual concepts in an image [4]. We might consider it in activity recognition to capture relevant contextual or situational information and concurrent activities. In addition, in real-world deployment, we randomly sample annotations from users [17], which are often not bounded to a temporal interval. Then the challenge is how to correlate these annotations with sensor data and to what temporal granularity.

Multiple Sources for Annotations

Rather than users themselves, we could consider gaining annotations from other users or other datasets with similar sensing technologies and activity recognition missions. Various transfer learning techniques have also been applied to HAR [8], which allows transferring an activity recognition model from one user to another user, or from one type of sensors to another type of sensors. This direction has a potential to reduce annotation burden on each user but the performance of these transfer learning approaches still needs a significant improvement, especially when faced with heterogeneous sensor feature sets and heterogeneous activity sets. To address the challenges, one core idea is to learn feature transformations so as to project feature spaces from different domains onto a common space to enable model adaptation. However, the majority of these techniques need some sort of ground truth to bridge the feature spaces; e.g., the same output labels or the co-occurrence of labels, or are built on strong assumptions; e.g., the domains share the same distribution. None of these always hold in the context of activity recognition. To address this problem, we could explore Structural Correspondence Learning (SCL) techniques, which link class labels through their semantic meaning [9], and instance transferring techniques, which borrow examples from similar but not the same activities [9].

Ensemble Learning for Evolving and Emerging Activity Classes

Managing emerging classes can be a challenging task. Ensemble learning techniques on streaming data classification have demonstrated its potential in dynamically including new classes without major changes to the already-learned model. In terms of HAR, the main challenges will involve accounting for the variety and dynamism of patterns in each type of activity. The key to address these challenges is to maximise the flexibility of an ensemble, which can be achieved through diverse ways of constructing base learners and decomposable structures of constructing an ensemble. For example, each base learner can be built on different sets of features; i.e., the ones with the most discriminating and predictive power, which will reduce the impact of the change in sensor deployment and thus significantly enhance the flexibility in feature space. A hierarchical ensemble can be designed to support dynamic composition of base learners with different semantics; e.g., temporal or causal dependencies. This will allow for flexibly adding a new emerging pattern or activity class, and/or remove a disappearing pattern or activity class.

CONCLUSION

Activity recognition has long been a popular research topic in pervasive computing, and lifelong learning in HAR moves towards dynamically and intelligently adapting to feature evolution, concept drift, and class evolution in a constantly changing environment. We have identified research challenges centred around learning, which are synthesized from the limitation of the re-

cent attempts. However, taking lifelong learning to reality and performing the real-world evaluation will set out more challenges to the interdisciplinary nature of pervasive computing, including user interaction on how to annotate and de-annotate activities based on their own preference and privacy concerns, system design on how to adapt applications responding to the new activities without much re-engineering effort, to name a few.

ACKNOWLEDGEMENT

We thank for the constructive feedback and comments from Prof Cecilia Mascolo at the University of Cambridge.

REFERENCES

1. Z. S. Abdallah, M. M. Gaber, B. Srinivasan, and S. Krishnaswamy. Anynovel: detection of novel concepts in evolving data streams. *Evolving Systems*, pages 1–21, 2016.
2. H. Alemdar, T. L. van Kasteren, and C. Ersoy. Using active learning to allow activity recognition on a large scale. In *Proceedings of AMI '11*.
3. J. Andreu, R. D. Baruah, and P. Angelov. Real time recognition of human activities from wearable sensors by evolving classifiers. In *Proceedings of FUZZ '11*.
4. N. Bicocchi, M. Mamei, and F. Zambonelli. Detecting activities from body-worn accelerometers via instance-based algorithms. *Pervasive and Mobile Computing*, 6(4):482–495, 2010.
5. Chen, K.-H.; Chen, P.-C.; Liu, K.-C.; Chan, C.-T. Wearable Sensor-Based Rehabilitation Exercise Assessment for Knee Osteoarthritis. *Sensors* **2015**, *15*, 4193–4211.
6. V. Chandola, A. Banerjee, and V. Kumar. Anomaly detection: A survey. *ACM computing surveys (CSUR)*, 41(3):15, 2009.
7. H.-T. Cheng, F.-T. Sun, M. Griss, P. Davis, J. Li, and D. You. Nuactiv: Recognizing unseen new activities using semantic attribute-based learning. In *Proceedings of MobiSys '13*, pages 361–374. ACM, 2013.
8. D. Cook, K. D. Feuz, and N. C. Krishnan. Transfer learning for activity recognition: a survey. *Knowledge and Information Systems*, 36(3):537–556, 2013.
9. O. Day and T. M. Khoshgoftaar. A survey on heterogeneous transfer learning. *Journal of Big Data*, 4(1):29, Sep 2017.
10. H. Gjoreski and D. Roggen. Unsupervised online activity discovery using temporal behaviour assumption. In *Proceedings of ISWC '17*.
11. P. Sawyer, R. Mazo, D. Diaz, C. Salinesi and D. Hughes. Using Constraint Programming to Manage Configurations in Self-Adaptive Systems. *Computer*, vol. 45, no. 10, pp. 56-63, Oct. 2012.
12. H. S. Hossain, N. Roy, and M. A. A. H. Khan. Active learning enabled activity recognition. In *Proceedings of PerCom '16*, pages 1–9. IEEE, 2016.
13. M. I. Jordan and T. M. Mitchell. Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245):255–260, 2015.
14. Juan Ye, Lei Fang, and Simon Dobson. Discovery and recognition of unknown activities. In *Proceedings of Adjunct UbiComp '16*.
15. L. T. Nguyen, P. Tague, M. Zeng, and J. Zhang. SuperAD: Supervised activity discovery. In *Proceedings of UbiComp '15 Adjunct*, pages 1463–1472. ACM, 2015.
16. P. Rashidi, D. J. Cook, L. B. Holder, and M. Schmitter-Edgecombe. Discovering activities to recognize and track in a smart environment. *IEEE transactions on knowledge and data engineering*, 23(4): 527–539, 2011.
17. S. Servia-Rodríguez, K. K. Rachuri, C. Mascolo, P. J. Rentfrow, N. Lathia, and G. M. Sandstrom. Mobile sensing at the service of mental well-being: A large-scale longitudinal study. In *Proceedings of WWW '17*, pages 103–112, 2017.
18. D. Silver, Q. Yang, and L. Li. Lifelong machine learning systems: Beyond learning algorithms. In *Proceedings of AAAI Spring Symposium on Lifelong Machine Learning*, 2013.

19. J. Ye, G. Stevenson, and S. Dobson. Detecting abnormal events on binary sensors in smart home environments. *Pervasive and Mobile Computing*, 33:32 – 49, 2016.
20. J. Yin, Q. Yang, and J. J. Pan. Sensor-based abnormal human-activity detection. *IEEE Transactions on Knowledge and Data Engineering*, 20(8): 1082–1090, 2008.

ABOUT THE AUTHORS

Juan Ye is a lecturer in the School of Computer Science at the University of St Andrews. Her research interests centre around adaptive pervasive systems, specialising in sensor-based human activity recognition, sensor fusion, context awareness, ontologies, and uncertainty reasoning. She has a PhD in computer science from University College Dublin. Contact her at juan.ye@st-andrews.ac.uk.

Franco Zambonelli is a professor of Computer Science at the University of Modena and Reggio Emilia. He got his PhD in Computer Science and Engineering from the University of Bologna in 1997. His research interests include: pervasive computing, multi-agent systems, self-adaptive and self-organizing systems. He is in the editorial board of the *ACM Transactions on Autonomous and Adaptive Systems*, *Elsevier Journal of Pervasive and Mobile Computing*, *IEEE Society & Technology Magazine*, the *BCS Computer Journal*. He is ACM Distinguished Scientist, member of the *Academia Europaea*, and IEEE Fellow. Contact him at franco.zambonelli@unimore.it.

Simon Dobson is a professor of Computer Science in the School of Computer Science at the University of St Andrews. He works on complex and sensor systems, especially on sensor data analytics and the modelling of complex processes on networks. He holds a BSc from the University of Newcastle upon Tyne and a DPhil from the University of York, both in computer science, is a Chartered Fellow of the British Computer Society, a Chartered Engineer, and Senior Member of the IEEE and ACM. Contact him at simon.dobson@st-andrews.ac.uk.
