SeAct: Semantic Adaptive Segmentation of Sensor Data Streams for Human Activity Recognition

Amanda D. P. Venceslau[1,3], Vânia M. P. Vidal[2], Rossana M. C. Andrade[2,3], José Gilvan R. Maia[4], José Wellington F. Silva[1]

[1]Federal University of Ceará, Campus de Crateús, Crateús-CE, Brazil

[2]Department of Computing – Federal University of Ceará
Campus do Pici, Fortaleza-CE, Brazil

[3]Group of Computer Networks, Software Engineering and Systems (GREat)
Federal University of Ceara, Campus do Pici, Fortaleza-CE, Brazil

[4]Virtual University Institute – Federal University of Ceará
Campus do Pici, Fortaleza-CE, Brazil

{amanda.pires,rossana}@ufc.br, vvidal@lia.ufc.br, gilvanmaia@virtual.ufc.br

Abstract. *Pervasive computing delivers services based on user needs through smart environments that incorporate and integrate everyday objects discreet and non-intrusive. Personal applications provide the data collected by sensors for Human Activity Recognition. The main limitation is that these activities need to be continuously segmented for HAR. Furthermore, a growing problem is related to the disambiguation of activities since some actions generated by the same sensors belong to different activities. This paper proposes a hybrid method, SeAct, which dynamically adjusts segment size, combining machine learning and semantic inference. Experiments with CAD-120 datasets and a state-of-the-art hybrid method improve recognition accuracy and precision.*

1. Introduction

The number of elderly people is expected to more than double globally in the next three decades[1]. On the other hand, ubiquitous computing provides services according to the user's needs through smart environments that incorporate and integrate everyday objects in a non-intrusive, discrete manner [Postolache et al. 2009]. Ambient Assisted Living (AAL) equipped with sensors, wireless networks, and software applications, allows monitoring and collecting of user and environmental information to provide assertive technologies for healthcare, support, and assistance [Memon et al. 2014]. Based on the processing and analysis of actions, it is possible to interpret information and recognize users' Activities of Daily Living (ADL). Human Activity Recognition (HAR) is an emerging area of ubiquitous computing due to its importance in supporting monitoring and assistance applications [Patel and Shah 2021].

Many applications rely on the interpretation of continuously generated data streams. This scenario requires methods that segment the data stream into intervals, segments, or windows that can be analyzed in an insightful manner (e.g., perform an activity or anomaly detection), so applications can adjust their behavior accordingly or even warn

[1]https://unstats.un.org/sdgs/indicators/regional-groups/

about dangerous situations. According to [Sfar and Bouzeghoub 2019], sensor data segmentation is a key factor that can reduce or increase the accuracy of HAR methods. Data segmentation methods can be classified as static or dynamic [Fu 2011]. Static methods may choose a time window size that may be too long or too short for activity recognition, leading to errors. Dynamic methods use flexible size windows to adapt as they learn about the environment and user's behavior.

Previous research in this area [Okeyo et al. 2014, Wan et al. 2015, Sfar and Bouzeghoub 2019, Triboan et al. 2019] assumes that the resident's life routine changes, but these do not combine or process the events generated by the sensors to recognize activities. Most data segmentation techniques found in literature are *data-based* or *knowledge-based*. However, *hybrid approaches* propose to overcome the difficulties and incorporate the advantages of *data-based* and *knowledge-based* techniques. The hybrid methods use a semantic approach that determines the segment size according to simple events obtained (e.g., from objects used). However, such an approach does not consider the fact that, in real-world applications, simple event processing can result in the detection of numerous activities related to the same object of use. For example, when someone prepares a meal using a bowl, simplistic approaches can also report activities related to feeding and cleaning. Combining simple events into a Complex Event Processing (CEP) engine can support knowledge-based recognition systems capable of inferring activities according to events collected from user behavior patterns. CEP consists of associating or combining simple events to identify high-level situations based on defined rules [Chen et al. 2014]. As an example of events, we can mention the objects of use and posture during daily human actions.

In this article, we propose a new hybrid method called SeAct that combines a data-based and knowledge-based method. After learning the user's behavior during activities, our method allows dynamic selection of the best time window size. In an offline stage, a machine learning method learns the resident's behavior, and an ontology with domain-related concepts are built (the so-called ActiOn ontology) defining concepts such as sensors, events, and activities for later use by inference rules. This allows to adapt the time window according to semantic reasoning rules during the online stage, i.e., simple events are processed through inference rules to recognize the activities and to obtain their duration. These are the main contributions of this paper:

- SeAct, a new hybrid method for semantic adaptive segmentation of sensors data streams for HAR, so time window size is dynamic and based on complex event processing for disambiguation of activities.
- We evaluate our method through a reproducible experiment that uses the CAD-120 dataset and compares it to a state-of-the-art (SOTA) approach considering similar sensor data [Díaz-Rodríguez et al. 2014]. We obtained improved accuracy and precision results, concluding that our method is competitive.

The remainder of this paper is organized as follows: Section 2 reviews previous work on hybrid methods for segmentation of sensors data streams for human activity recognition. Section 3 describes the SeAct method and its offline and online steps. Section 4 discusses the evaluation results and Section 5 concludes the article.

2. Related Work

Sensor data segmentation techniques applied to HAR can be classified into *metrics-based*, *data-based*, *knowledge-based*, and *hybrid*. *Metric-based* techniques determine window size according to the result of calculated metrics [Wan et al. 2015]. However, treating incoming events according to their correlation with sensors' location can add irrelevant events to the segment and make it harder to disambiguate activities: correlation does not consider the user's temporal and contextual characteristics. *Data-based* techniques [Mallick et al. 2018, Dong et al. 2020, Pan et al. 2019, Rawashdeh et al. 2020, Krishnan and Cook 2014, Asghari et al. 2020] resort to machine learning or probability distribution to optimize segment size concerning sensor events. These are usually pre-segmented data with fixed window sizes. However, these models do not incorporate expressive user context information. *Knowledge-based* techniques [Triboan et al. 2017, Civitarese et al. 2018, Triboan et al. 2019] ensure interpretation and expressiveness to provide information enrichment. Many adopt a generic model built on common sense knowledge. Previous knowledge-based methods generally assume that the resident's routine does not change, which is not the case in smart environments. *Hybrid techniques* have been developed to overcome these difficulties by incorporating the main advantages of data-based and knowledge-based techniques. As we propose a hybrid approach, the following analysis of related work is limited to this category of methods.

[Salguero et al. 2018] use time and event-based, hybrid segmentation. First, time-based segmentation is applied to define the relevant events in a sliding time window for each activity, whose size is adjusted by statistical analysis. The approach is evaluated under real-world conditions in an experiment where the ontology and event-based segmentation automatically generates resources that serve as input to the supervised learning classification algorithms. The authors instantiate a plain ontology during the segmentation process, then combine concepts and properties to create new concepts. These concepts also provide knowledge to describe activities in a more comprehensive representation. Adopting the F-measure for the Singla and Ordoñez dataset, the authors obtained 0.86 compared to SOTA. In addition, the work correctly classified 85.25 instances for the Kasteren dataset. The authors compare the results by the average percentage of correctly classified windows.

DataSeg [Sfar and Bouzeghoub 2019] combines a clustering algorithm with an ontology-based method to dynamically choose the best time window size. A plain ontology is first created with a default activity classification in terms of their duration and sensors to alleviate the cold start problem. The clustering step is carried out offline once a training dataset is acquired from the user's activities. DataSeg uses the ontology and communicates with the activity recognizer module to determine the best size for the current time window for online recognition. DataSeg was evaluated in terms of accuracy result and F-score, surpassing a static window and a metrics-based technique that uses the measure of mutual information to calculate sensor correlation.

POLARIS [Civitarese et al. 2019] takes into account spatio-temporal dependencies, in addition to the correlations between activity and usage object of the data streams. Ontological inference derives these probabilistic dependencies between the types of sensor events and the classes of activities performed, and refers to these operations as semantic correlations. The output of the semantic correlation reasoning layer is directly

considered in the statistical analysis in a Markov Logic Network (MLN), assuming probabilistic and knowledge-based reasoning to recognize the activities. The authors evaluate the impact of the segmentation method by an F-score for activity recognition. Adopting the CASAS and SmartFaber datasets, the authors obtained 0.76 as an average recognition for all activities.

Finally, in [Díaz-Rodríguez et al. 2014] the authors performed an experiment using the well-known CAD-120 database [Koppula et al. 2013]. The objects of use and human posture extracted from this base provide information about the user's context when capturing the events or actions that define related activities. The proposed hybrid method uses a semi-supervised k-nearest neighbor (k-NN) algorithm to classify actions into activities. The segmentation strategy defined in [Díaz-Rodríguez et al. 2014] uses the information obtained by the k-NN algorithm and stores it in a hash table, so different activities are retrieved with similar actions. Finally, filters are used together with a fuzzy ontology to recognize an activity.

Adopting well-known datasets in the segmentation of sensor data streams is fundamental for an honest comparison between works. However, the number of public datasets in HAR is scarce [Junior et al. 2022] and the available datasets do not have many instances, which can be explained by the fact that building such datasets involves monitoring human beings.

Table 1. Hybrid sensor data streams segmentation methods for HAR. The columns adopt: *ST* – Semantic Technology. The lines adopt: *RSD* – Raw Sensor Data; *OOU* – Object of Use; *HP* – Human Posture; and *L* – Location

Work	Datasets	Feature	Window	ML Algorithm	ST
Salguero et al., 2018	Kasteren, Singla, and Ordoñez	RSD	Static	Probabilistic Learning	Ontology
Sfar; Bouzeghoub, 2019	Aruba	RSD	Dynamic	Clustering	Ontology
Civitarese et al., 2019	CASAS and SmartFABER	OOU+L	Dynamic	Markov Network	Semantic Inference
Diaz et al.,2014	CAD-120	OOU+HP	Dynamic	k-NN	Fuzzy Ontology
SeAct	**CAD-120**	**OOU+HP**	**Dynamic**	**k-NN**	**Semantic Inference**

In Table 1, we compare our method with the other methods considered relevant in hybrid sensors data stream segmentation to HAR in terms of: datasets; characteristics of the data collected by sensors; window types; processing types; description of the technique based on data and knowledge; and the use of an ontology to describe the segmentation result. Some SOTA approaches such as [Salguero et al. 2018] and [Sfar and Bouzeghoub 2019] manipulate raw sensor data, data from sensors implanted in objects (e.g., sensor M003 ON corresponds to an open refrigerator). Furthermore, these approaches use ontology for modeling and querying. Approaches like [Civitarese et al. 2019, Díaz-Rodríguez et al. 2014] use anno-

tated events (e.g., objects of use, location, and human posture) and process the events gradually [Civitarese et al. 2018] or querying the events from data structures [Díaz-Rodríguez et al. 2014] to define a dynamic window. Analyzing the knowledge-based techniques, in [Civitarese et al. 2018], the defined ontology and its constraints are used in the inference of activities, which the Markov Network filters. In [Díaz-Rodríguez et al. 2014], the ML algorithm learns user behavior patterns, and the segment is defined by querying the events represented by a data structure. The fuzzy ontology is used to filter the activities detected in the previous steps. In its turn, our approach proposes the correlation between simple events and processing complex events to obtain a dynamic window. The segment is determined by the duration of activities that are inferred by rules and semantic reasoning. To compose our rules, we use the learning algorithm.

3. SeAct: Semantic Adaptive Segmentation of Sensors Data Streams

The SeAct method has an architecture that integrates sensors data stream segmentation techniques into a new hybrid approach that processes complex events, providing sensor data stream segments to applications. The method is designed to be used in different applications, provided that a set of user monitoring data can be provided and other requirements defined in that section. Our method is divided into two steps (see Figure 1), namely, offline and online.

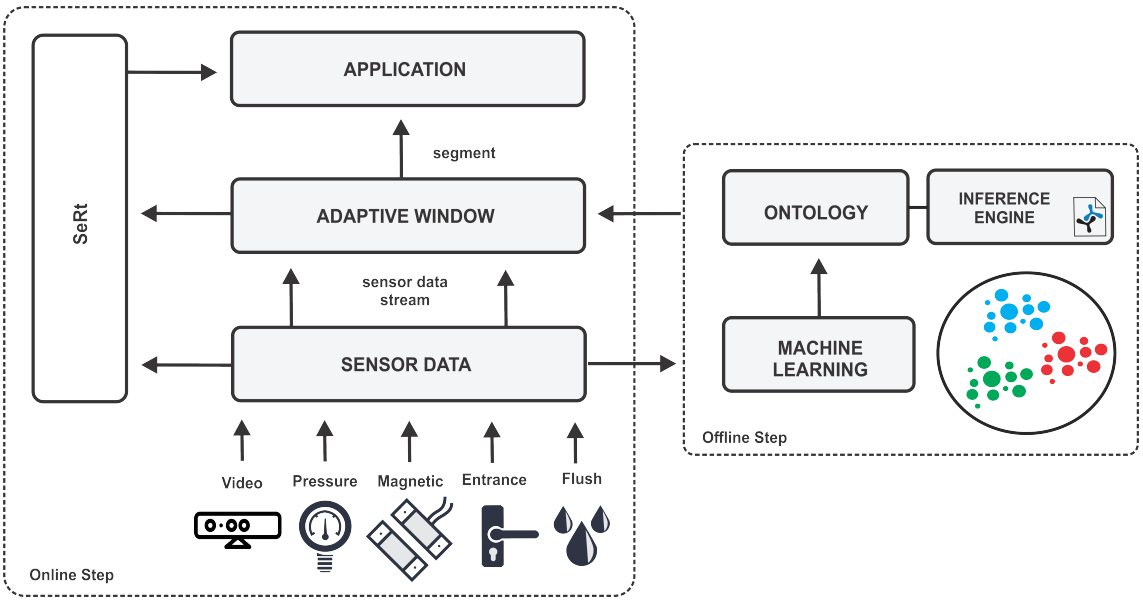

Figure 1. SeAct Architecture

A data-driven technique is needed to learn user behavior in the offline step. A set of monitoring data from a resident is acquired. This data is used as input for processing that uses machine learning. The learning result generates valid semantic inference rules for event detection and activity inference. The online step, a knowledge-based technique, uses ontology and reasoning about the definite rules to define the activity and its duration, used by the time window adaptation algorithm.

SeAct segmentation starts in the offline stage by acquiring a dataset of behavior patterns and activities of individuals (see Figure 1) through four sub-steps: (1) Sensor

Data Streams Acquisition; (2) Machine Learning Model; (3) ActiOn Ontology; and (4) Semantic Inference Rules. If the smart environment is a smart home, we can call this set of data as belonging to the residents. This helps to define the semantic inference rules used in the online step. An online step that performs the semantic adaptive segmentation of sensor data streams in time windows based on the knowledge provided by the ontology.

In the following subsections, we provide details of the proposed method, highlighting the need for a hybrid approach capable of segmenting sensor data streams. We use a smart home as a case study.

3.1. SeAct: Offline Step

The offline step is divided into four sub-steps: (1) Sensor Data Streams Acquisition; (2) Machine Learning Model; (3) ActiOn Ontology; and (4) Semantic Inference Rules.

3.1.1. Sensor Data Streams Acquisition

We adopted a case study in the domain of HAR to compare our method with SOTA, so we chose the public-domain RGB-D CAD-120 dataset [Koppula et al. 2013], in which body postures of users that appear in front of an RGB-D camera were extracted with the 3D location of objects in the scene using a Microsoft's Kinect, which provides depth and color information. Each detected posture corresponds to a sequence of acquired images, which also includes the objects participating in the activity. The user's body is represented by Kinect's 15 standard 3D points (head, neck, left and right torso and shoulders, elbows, hips, hands, knees, and feet). To understand relationships between variables, particularly, the relationship between posture and object events involved in the scene, we need to investigate how this relationship can impact the definition of actions and, consequently, the recognition of activities.

Initially, an individual must be monitored for a variable period of time since there is no insight into her behavior beforehand. The literature demonstrates that public datasets are available in different domains so that the community benefits, deepening knowledge in the area. We need to investigate how the relationship between posture and objects can impact the definition of actions and, consequently, the recognition of activities.

3.1.2. Machine Learning Model

ML-based methods resort to different formulations (e.g., probability, statistics, and linear algebra) to train classification models using historical behavior described using characteristics relevant to the problem [Akbar et al. 2019]. The proposed hybrid method uses machine learning to learn user behavior. Actions can be anomalies (e.g., falls, strokes, heart failures, etc), so assistants can make decisions and take actions as fast as possible. We use an ML model that learns the individual actions of participants using more than one action-related feature, i.e., multiple simple events are acquired for the same action. For example, in a fall detection application, postures can be combined with accelerometer data to define user actions. On the other hand, event sequences related to the posture when opening and closing the microwave can be insightful for a nutrition scenario. Attention

must be drawn to the fact that "atomic" actions themselves do not denote an activity like "eating": such recognition must analyze several actions.

SeAct improves segmentation results by adopting a sensor-dependent window extension, a strategy lacking in most SOTA works [Serpush et al. 2022]. Furthermore, we explore what we define as simple events, the information acquired by the sensors, processing it through complex events. However, some studies show promising results when using the dependency, such as [Díaz-Rodríguez et al. 2014]. A supervised learning algorithm, k-nearest neighbors (KNN), is used to classify activities belonging to a given dataset based on the monitored events. This is performed as soon as a new dataset is provided to learn the residents' different routines. For comparison purposes, our experiments use a data-based technique to classify activities based on the features, human posture (e.g., "movement") and object of use (e.g., "glass") proposed in [Díaz-Rodríguez et al. 2014]. So, RGB-D CAD-120 is used as input to the Piecewise Aggregate Approximation (PAA) time series algorithm [Keogh et al. 2001] to learn the sequence of postures adopted by users and the positions of objects concerning actions.

In the last step to learn different sub-activities or captured actions involving the user's use of objects, the authors adopt instance-based learning using k-nearest neighbors (KNN) to classify user's usage objects, postures, and activities. The classification outcome is used in our next sub-step and the main features of the CAD-120 dataset are adapted to ontological concepts and relationships.

3.1.3. ActiOn Ontology

Smart environments often require the description of their participants, sensors, and a representation suitable for decision-making. However, in this scenario, there is an exchange of information that requires a network infrastructure and communication between services, so addressing data volume is relevant for minimizing potential bottlenecks [Elsaleh et al. 2020]. Consequently, we adopt a core ontology to represent the concepts pertinent to the applications, adopting only domain-related concepts. Such ontology must have a T-Box describing the concepts related to the application domain. Moreover, applying semantic technologies to HAR systems can also automate, enrich information, and facilitate decision-making. To represent the data set with the behavior pattern of residents and thus represent human activities semantically, we propose an ontology that, due to the proposed generality, provides support for modeling and extending concepts and classes according to the application domain. Initial modeling considers the objects of use and the posture. The ontology contains information about the objects, postures, and high-level activities that must infer the duration of the activities.

In our case study, we designed ActiOn (Fig. 2) to semantically represent the information provided by the ML model. We added some instances with information related to activities, events, sensors, and a person as an example. Initially, there are four main classes (*Activity*, *Sensor*, *Event*, and *Person*), four main objects properties (*hasEvent*, *composedOf* and *partOf*), and a *DataProperty* (*hasValue*, *hasDurationRange*). The *Activity* and *Event* classes are respectively the domain and range of *compositeOf*, an inverse property of *partOf* – *Event* is the domain class and the range is the Activity class. So, *hasValue* depends on the *Event* class with a string value that represents the value of the

event (i.e., using the bowl object while the resident pours a liquid).

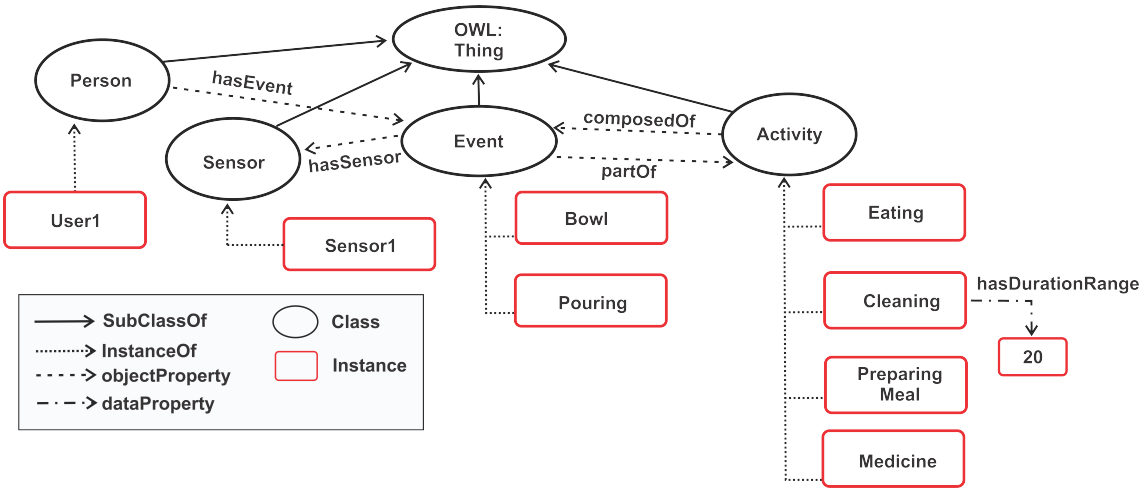

Figure 2. ActiOn ontology with instances.

3.1.4. Semantic Inference Rules

A core ontology enables defining semantic rules to represent and infer knowledge by using its T-Box and the behavior information obtained by the ML algorithm's output. In our case study, such resident behavior refers to events such as objects and human posture, in addition to the activities that were learned. Rules are generated according to the behavior information inferred using the ML model. So, only what was captured and classified as the behavior of that user in question will be used to generate the rules. The rule set is automatically updated after classification. To derive or infer new facts that do not exist in the knowledge base, it is necessary to use reasoning with rules based on first-order predicate logic or description logic to conclude a sequence of statements (premises) derived from predefined rules [Krötzsch et al. 2012] . Such software is known as a reasoning engine (or reasoner): it deals with RDFS and OWL vocabularies and infers facts from semantic data and ontologies. We use the Apache Jena reasoner, an open-source framework that provides a Java API to extract data obtained from files generating RDF knowledge graphs. We exemplify the use of rules used by our approach in an example of inference of the eating activity that receives two events, the glass object, and the user's posture or position, movement (see Figure 2).

In SeAct, object and posture (events) with their associated activities deriving from ML step are fed to the rules. This is significantly advantageous as it requires expert intervention only if no dataset is provided, automatic dynamic rule update can take place. Moreover, this method is flexible to consider the different routines of each resident.

3.2. SeAct: Online Step

The online adaptation sub-step results in dynamically adjustable windows. The adaptive window is based on the collected sensor data stream, which populates the ontology and then enables the reasoning of semantic inference rules that choose the window size. For our HAR case study, we modify the adaptive window Algorithm 1 to retrieve events from

Table 2. Example of Inference Rules in SeAct

Rule 1	[EatingMealRule1:: (?user SeAct:hasEvent ?event1), (?event1 rdf:type SeAct:Event), (?event1 SeAct:hasvalue “glass”), (?event1 SeAct:hasRelatedTime ?temp1), (?temp1 SeAct:hasValue ?tval1), (?user SeAct:hasEvent ?event2), (?event2 rdf:type SeAct:Event), (?event2 SeAct:hasvalue “movement”), (?event2 SeAct:hasRelatedTime ?temp2), (?temp2 SeAct:hasValue ?tval2), makeTemp(?activity) $\rightarrow$ (?activity rdf:type SeAct:activity), (?user SeAct:hasActivity ?activity), (?activity SeAct:hasValue “eating”), (?activity SeAct:compositeOf ?event1), (?event1 SeAct:partOf ?activity), (?activity SeAct:compositeOf ?event2), (?event2 SeAct:partOf ?activity)]

a data stream and recognize activities to extract their duration, which is used to define the segment. According to [Sfar and Bouzeghoub 2019], a sequence of events can be represented as $\{E_1, \ldots, E_n\}$, where $\{E_i\}$ refers to the i-th event captured by sensors, and each sensor event is encoded in the model of $(date, time, eventValue)$. This aims to divide event streams into s-sized time windows.

The time window size is initially set to $60s$. Then, the events that occurred under that size are extracted (*readOnline* at Line 4). Sensed values are validated to extract which object has that sensor and which posture is identified at that event (Line 5). These events are inserted into the ontology (Line 6) to apply rules that generate the related activities (*applyRulesActivities* at Line 7), so the activities resulting from some rule, i.e., that corresponds to some predefined rule (Section 3.1.4), will be inferred and can be extracted (Lines 9-11). After verification, the activity is added to the set of possible activities (Line 8). The segment (size of the time window) is adapted to a window corresponding to the activity duration extracted by our algorithm. Thus, as the window is adapted using the first events, it should collect the other possible events also related to the current activity.

4. Experiments and Discussion

SeAct was applied to the smart health scenario of a smart environment following the evaluation and experimentation methodology proposed by [Wohlin et al. 2012]. **Experiment Definition.** Monitoring environments provide the knowledge base needed for activity recognition applications used by caregivers, doctors, and hospitals that assess residents in their homes. Such HAR systems demand effective activity recognition. **Ex-**

```
Data: Ontology O, Rules R, Sensors data SD
Result: Set of possible resAct, Segment s
resAct ← NULL;
while SD ≠ NULL do
    interval ← 60;
    sensors ← readOnline(SD, interval);
    events ← validation(sensors);
    insertEvents(events, O);
    applyRulesActivities(O, R);
    resAct ← extractActivities();
    if resAct ≠ NULL then
        s ← getDuration(resAct, interval);
    end
    return s, resAct;
end
```

Algorithm 1: Adaptive window selection procedure.

periment Planning. We hypothesize that a semantic adaptive segmentation method that uses user context attributes can increase accuracy, precision, and recall compared to other approaches. **Experiment Operation.** We used CAD-120 dataset [Koppula et al. 2013], which includes recordings from 12 different ADLs (composed of several sub-activities) of 4 subjects, two males and two females (one of them is left-handed). User contexts need to be captured and used together to test our semantic solution to satisfy the hypothesis in our proposal. Our experiments were executed on iOS 11.5.2 with a Intel Core i5 Dual-Core (1.8 GHz, 1600 MHz DDR3) and 8 GB RAM. **Analysis and Interpretation.** Our results are compared to a SOTA [Díaz-Rodríguez et al. 2014] approach that uses the same HAR dataset and uses a hybrid method that stores inferences from ML model into a Hash Table. This approach does not use temporal semantics to relate events and infer activities. Our approach, in its turn, uses an ontology with semantic knowledge about events and activities through an ontological vocabulary. It also enriches and relates events with information about sensors, users, time of occurrence, and degree of uncertainty in detection. As shown in Tab. 3, SeAct outperforms the SOTA approach in terms of precision and accuracy while keeping a perfect recall, i.e., retrieves all genuine detections. **Replicability.** Our experiment uses a known dataset and a SOTA approach that has code available on the Github platform. We also provide instructions to replicate the experiment.

Table 3. Comparison of our approach to activity recognition using CAD-120.

Method	Precision	Recall	Accuracy
Algorithm proposed in [Díaz-Rodríguez et al. 2014]	0.60	1.00	0.38
SeAct	0.70	1.00	0.68

SeAct improves performance on CAD-120 by fair margins, achieving a precision of 0.70 and an accuracy of 0.68. This supports the hypothesis that our segmentation method is both effective and competitive. Individual results for each activity are shown in a confusion matrix (see Fig. 3): in SeAct, most activities and, consequently, segments are

chosen correctly. However, we can notice in [Díaz-Rodríguez et al. 2014] that activities that share the same objects (e.g., cereal, eating, and microwaving) are not chosen correctly. Furthermore, we must highlight that SeAct infers more than one possible activity, which is useful for critical HAR applications.

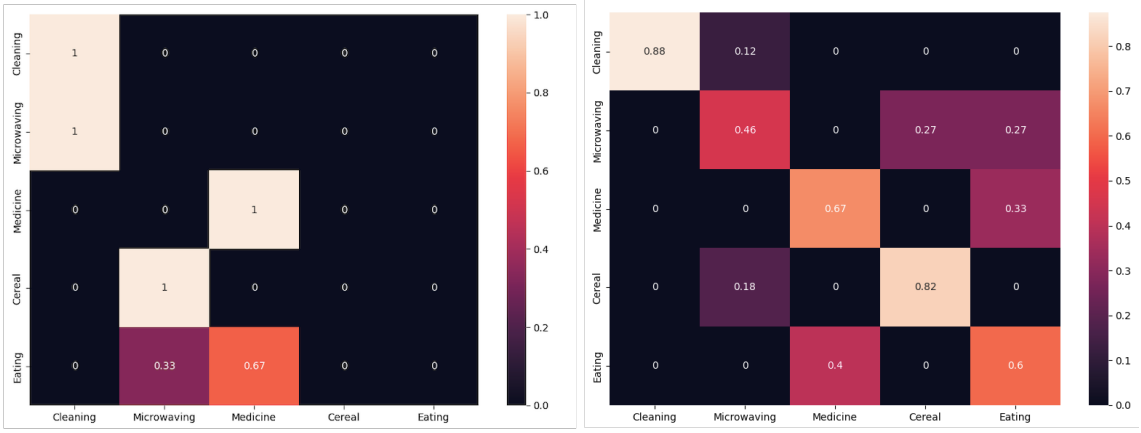

Díaz-Rodríguez et al. (2014) SeAct

Figure 3. Confusion matrices for activities segmented on CAD-120.

We can cite threats to validity: (i) we do not test different datasets; and (ii) other SOTA methods have not been tested as they are not reproducible. Finally, the SOTA method aims to recognize activities, with segmentation being part of the process. So we experimented with the segmentation step to get a fair comparison.

5. Conclusion and Future Work

While most hybrid sensor data segmentation methods adopt at least one user context attribute, in this work we propose the use of the semantic correlation of user context attributes to disambiguate activities in the reasoning step. So, our hybrid method combines machine learning and semantic inference. Experimental evaluation showed that the precision and accuracy of our approach is SOTA for hybrid methods. We also propose an ontology of the segmentation process that allows the semantic representation of the metadata involved, tracking, and researching this information by activity recognition systems. Finally, we successfully evaluated our ontology using competency questions. Our contributions can be applied for other domains, such as self-adaptive systems that take human movement and posture patterns into account. As future works, we plan to define a framework to integrate our method into a publish/subscribe system and carry out further experimentation with the fully implemented framework.

References

Akbar, A., Chaudhry, S. S., Khan, A., Ali, A., and Rafiq, W. (2019). On complex event processing for internet of things. In *2019 IEEE 6th International Conference on Engineering Technologies and Applied Sciences (ICETAS)*, pages 1–7. IEEE.

Asghari, P., Soleimani, E., and Nazerfard, E. (2020). Online human activity recognition employing hierarchical hidden markov models. *Journal of Ambient Intelligence and Humanized Computing*, 11(3):1141–1152.

Chen, C. Y., Fu, J. H., Sung, T., Wang, P.-F., Jou, E., and Feng, M.-W. (2014). Complex event processing for the internet of things and its applications. In *2014 IEEE International Conference on Automation Science and Engineering (CASE)*, pages 1144–1149. IEEE.

Civitarese, G., Bettini, C., Sztyler, T., Riboni, D., and Stuckenschmidt, H. (2018). Nectar: Knowledge-based collaborative active learning for activity recognition. In *2018 IEEE International Conference on Pervasive Computing and Communications (PerCom)*, pages 1–10. IEEE.

Civitarese, G., Sztyler, T., Riboni, D., Bettini, C., and Stuckenschmidt, H. (2019). Polaris: Probabilistic and ontological activity recognition in smart-homes. *IEEE Transactions on Knowledge and Data Engineering*, 33(1):209–223.

Díaz-Rodríguez, N., Cadahía, O. L., Cuéllar, M. P., Lilius, J., and Calvo-Flores, M. D. (2014). Handling real-world context awareness, uncertainty and vagueness in real-time human activity tracking and recognition with a fuzzy ontology-based hybrid method. *Sensors*, 14(10):18131–18171.

Dong, Y., Liu, J., Gao, Y., Sarkar, S., Hu, Z., Fagert, J., Pan, S., Zhang, P., Noh, H. Y., and Mirshekari, M. (2020). A window-based sequence-to-one approach with dynamic voting for nurse care activity recognition using acceleration-based wearable sensor. In *Adjunct Proceedings of the 2020 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2020 ACM International Symposium on Wearable Computers*, pages 390–395.

Elsaleh, T., Enshaeifar, S., Rezvani, R., Acton, S. T., Janeiko, V., and Bermudez-Edo, M. (2020). Iot-stream: A lightweight ontology for internet of things data streams and its use with data analytics and event detection services. *Sensors*, 20(4):953.

Fu, T.-c. (2011). A review on time series data mining. *Engineering Applications of Artificial Intelligence*, 24(1):164–181.

Junior, E. C., Andrade, R. M. C., Venceslau, A. D. P., Oliveira, P. A. M., Santos, I. S., and Oliveira, B. S. (2022). Where is the internet of health things data? In *Proceedings of the 24th International Conference on International Conference on Enterprise Information Systems (ICEIS) (to appear)*, pages 1–10.

Keogh, E., Chakrabarti, K., Pazzani, M., and Mehrotra, S. (2001). Dimensionality reduction for fast similarity search in large time series databases. *Knowledge and information Systems*, 3(3):263–286.

Koppula, H. S., Gupta, R., and Saxena, A. (2013). Learning human activities and object affordances from rgb-d videos. *The International Journal of Robotics Research*, 32(8):951–970.

Krishnan, N. C. and Cook, D. J. (2014). Activity recognition on streaming sensor data. *Pervasive and mobile computing*, 10:138–154.

Krötzsch, M., Simancik, F., and Horrocks, I. (2012). A description logic primer. *arXiv preprint arXiv:1201.4089*.

Mallick, M., Kodeswaran, P., Sen, S., Kokku, R., and Ganguly, N. (2018). Tsfs: An integrated approach for event segmentation and adl detection in iot enabled smarthomes. *IEEE Transactions on Mobile Computing*, 18(11):2686–2700.

Memon, M., Wagner, S. R., Pedersen, C. F., Beevi, F. H. A., and Hansen, F. O. (2014). Ambient assisted living healthcare frameworks, platforms, standards, and quality attributes. *Sensors*, 14(3):4312–4341.

Okeyo, G., Chen, L., Wang, H., and Sterritt, R. (2014). Dynamic sensor data segmentation for real-time knowledge-driven activity recognition. *Pervasive and Mobile Computing*, 10:155–172.

Pan, S., Berges, M., Rodakowski, J., Zhang, P., and Noh, H. Y. (2019). Fine-grained recognition of activities of daily living through structural vibration and electrical sensing. In *Proceedings of the 6th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, pages 149–158.

Patel, A. and Shah, J. (2021). Smart ecosystem to facilitate the elderly in ambient assisted living. In *Proceedings of International Conference on Recent Trends in Machine Learning, IoT, Smart Cities and Applications*, pages 501–510. Springer.

Postolache, O., Madeira, R. N., Correia, N., and Girão, P. S. (2009). Ubismartwheel: a ubiquitous system with unobtrusive services embedded on a wheelchair. In *Proceedings of the 2nd International Conference on PErvasive Technologies Related to Assistive Environments*, pages 1–4.

Rawashdeh, M., Al Zamil, M. G., Samarah, S., Hossain, M. S., and Muhammad, G. (2020). A knowledge-driven approach for activity recognition in smart homes based on activity profiling. *Future Generation Computer Systems*, 107:924–941.

Salguero, A. G., Espinilla, M., Delatorre, P., and Medina, J. (2018). Using ontologies for the online recognition of activities of daily living. *Sensors*, 18(4):1202.

Serpush, F., Menhaj, M. B., Masoumi, B., and Karasfi, B. (2022). Wearable sensor-based human activity recognition in the smart healthcare system. *Computational Intelligence and Neuroscience*, 2022.

Sfar, H. and Bouzeghoub, A. (2019). Dataseg: dynamic streaming sensor data segmentation for activity recognition. In *Proceedings of the 34th ACM/SIGAPP Symposium on Applied Computing*, pages 557–563.

Triboan, D., Chen, L., Chen, F., and Wang, Z. (2017). Semantic segmentation of real-time sensor data stream for complex activity recognition. *Personal and Ubiquitous Computing*, 21(3):411–425.

Triboan, D., Chen, L., Chen, F., and Wang, Z. (2019). A semantics-based approach to sensor data segmentation in real-time activity recognition. *Future Generation Computer Systems*, 93:224–236.

Wan, J., O'grady, M. J., and O'Hare, G. M. (2015). Dynamic sensor event segmentation for real-time activity recognition in a smart home context. *Personal and Ubiquitous Computing*, 19(2):287–301.

Wohlin, C., Runeson, P., Höst, M., Ohlsson, M. C., Regnell, B., and Wesslén, A. (2012). *Experimentation in software engineering*. Springer Science & Business Media.