Monitoring Objects Manipulations to Detect Abnormal Behaviors

Gabriele Civitarese   Claudio Bettini
Università degli Studi di Milano, Dept. of Computer Science, EveryWare Lab
Email: {gabriele.civitarese, claudio.bettini}@unimi.it

***Abstract*—With a growing population of elderly people the number of subjects at risk of cognitive disorders is rapidly increasing. Many research groups are studying pervasive solutions to continuously and unobtrusively monitor fragile subjects at their homes, reducing health-care costs and supporting medical diagnosis. Among the several behavioral aspects which clinicians are interested in monitoring, anomalous behaviors while performing activities of daily living are of great importance. In this work, we aim at improving the state of the art on this topic by enabling the recognition of fine-grained anomalies by detecting specific object manipulations. We attach tiny Bluetooth Low Energy accelerometers to several household objects in order to detect which manipulations are performed by the inhabitant on which object. Detected manipulations, combined with data from other environmental sensors deployed in the home, are used to infer ADLs and fine-grained abnormal behaviors. Preliminary results on a dataset with hundreds of complex activities captured in a smarthome environment show the effectiveness of the proposed method.**

I. Introduction

The increased availability of affordable and reliable sensing infrastructure is improving the ability to continuously monitor activities at home. In the context of eHealth, several research efforts have addressed the problem of identifying deviations in normal ways of performing activities of daily living (ADL) by the elderly (among recent ones [1], [2], [3]). A general approach is to build a model of the "regular" behavior in order to identify those activity patterns which diverge from the expected ones [1]. The main drawback of this type of approaches is that behavioral changes are detected without giving specific explanations of what happened. Other research groups tried to refine the identification by recognizing the general anomaly's category (e.g. omission, substitution, replacement, ...) using statistical methods [2]. However, the results show an high rate of false positives. In our previous work [3] we also started focusing on *fine-grained* anomaly detection, with the goal of providing more informative feedback to clinicians. Indeed, for each anomaly category we want to distinguish between several different ways of performing it, which may vary depending on the ADLs being performed, on the objects involved, on specific patterns of actions, etc..

This paper reports significant advances along this research direction by exploiting our recent results on using Bluetooth Low Energy (BLE) accelerometers attached to everyday objects in order to recognize performed manipulations [4]. Differently from other sensor-based approaches, which only monitor generic interactions with the objects [5], we apply machine learning to recognize the details of the manipulations performed (e.g., "*the bottle has been used to pour water*"). Moreover, the use of those devices addresses the obtrusiveness of wearable solutions [6], [7] and the privacy issues of audio/video-based approaches [8].

Our main contributions are threefold:

- We improved our experimental setup for ADL and anomaly recognition with a more accurate and reliable sensing infrastructure, which also includes the capability of monitoring the fine-grained manipulations of everyday objects.
- We propose a knowledge based reasoning framework to detect new fine-grained abnormal behaviors based on objects manipulations.
- We present preliminary results on a new dataset consisting of hundreds of complex/interleaved ADLs and anomalies.

The rest of the paper is structured as follows. Section II presents the general architecture of the proposed framework. Section III presents the technique to recognize fine-grained abnormal behaviors. Section IV summarizes preliminary experimental results. Finally, Section V concludes the paper.

II. General architecture

In this section we illustrate the framework we designed and implemented for detecting fine-grained abnormal behaviors.

A. System's architecture

The architecture of our recognition framework is shown in Figure 1.

Our framework consists of two main components. The Local data processing part is in charge of continuously collecting and pre-processing raw data from sensing devices. It runs within the smart-home environment. The Anomalies recognition platform component runs recognition algorithms on data provided by the Local data processing component. Currently, the recognition algorithms run periodically (e.g., on all the data collected in each day). This component could be deployed both in the smart-home environment or as a cloud service.

We consider a smart-home environment instrumented with two kinds of sensing devices: a) environmental sensors to monitor the inhabitant's interaction with the home environment, b) wireless accelerometers attached to a set of everyday objects in order to recognize the performed manipulations.

Raw data from environmental sensors are preprocessed by SEMANTIC INTEGRATION OF SENSOR DATA module, which applies simple inference rules to derive high-level events. The MANIPULATIONS DETECTION module collects the accelerometer data produced by the monitored objects and applies standard machine learning techniques to detect the performed manipulations. The MANIPULATIONS REFINEMENT module aggregates data from the underlying sensing subsystems and refines the manipulations classification. Aggregated and refined sensing data is used by the ADLS RECOGNITION module to detect activity instances with their timespans. Detected ADLs along with sensed data are transmitted to the FINE-GRAINED ABNORMAL BEHAVIORS RECOGNITION module, which applies knowledge-based reasoning to infer the occurred fine-grained abnormal-behaviors.

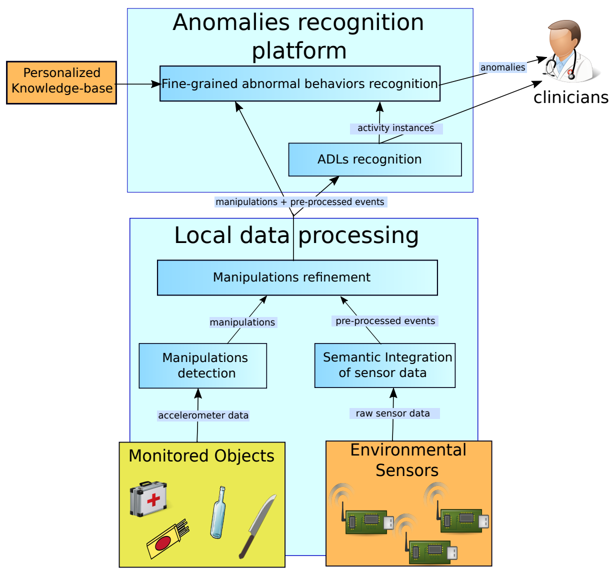

Fig. 1. The system's framework

B. Sensing subsystems

Two different modules are in charge of pre-processing sensor data.

1) Semantic integration of sensor data: The smart-home is equipped with a wide variety of environmental sensors, like pressure, presence, magnetic and power sensors. These sensors capture the interaction of the elderly with the home environment (e.g. stove usage, open/close of doors and repositories, pressure on chairs, presence at certain locations, ...) and transmit their raw data to the SEMANTIC INTEGRATION OF SENSOR DATA module. This module is in charge of reducing noisy measurements and to apply simple knowledge-based rules to derive timestamped pre-processed events. For example, if the energy meter attached to a stove cooker detects a substantial increase in power consumption (starting from no power consumption), the event "*the stove is turned on*" is inferred and timestamped appropriately.

We assume that a unique timestamp is assigned to each event, imposing a total order on event timestamps.

2) Manipulations detection: The MANIPULATIONS DETECTION module is in charge of recognizing the manipulations performed on the objects of interest. Every considered object is equipped with a tiny wireless accelerometer which transmits its data through a BLE interface. The data is first segmented to identify when exactly the object was manipulated. This is done with a simple thresholding method which detects when the object is actually in motion. From each segment, several statistical features are extracted. A machine learning algorithm is used to assign the fine-grained label of the performed manipulation (e.g., the bottle has been used to pour water). A starting and an ending timestamps are associated with the labeled manipulation. More details about this subsystem can be found in our previous work [4].

C. Manipulations refinement

The main task performed by the MANIPULATIONS REFINEMENT module is to combine the detected manipulations with information derived from environmental sensors to derive more precisely characterized manipulations. This process is done by a set of rules taking into account temporal and semantic relationships.

Example 1: Suppose that the system detected the manipulation of a medicine box and classified it as *significant displacement*. This can happen if the medicine box has been moved from a place to another[1]. If the system also detects that the medicine repository has been opened just before the start of the manipulation, the system can infer that the medicine box has been retrieved from the medicine repository. Hence, the manipulation class provided by the MANIPULATION DETECTION for that specific box in that timespan will be converted from *significant displacement* to *retrieved from repository*.

Refined manipulations along with pre-processed events are then temporally totally ordered.

D. Activity Recognition

The recognition of the specific activity being performed is sometimes a pre-requisite to detect an anomaly. This happens for abnormal actions within the activity (e.g., the subject forgets to use salt while preparing pasta), for anomalies related to sequence of activities (e.g., the subject prepared a meal but forgot to eat it), and for behavioral changes (e.g., significant variation of typical execution time of an activity).

The ADLS RECOGNITION module takes as input the pre-processed environmental sensor events and the refined manipulations, applies recognition algorithms, and produces as output a set of performed activity instances with unique identifiers, activity type, and timestamps denoting the temporal boundaries. Despite our experimental system uses specific machine learning algorithms [3], this framework does not impose any

[1]A displacement is characterized as significant when the movement is not just a minor involuntary change of position, like for example a medicine box moved within a drawer while grasping a different box

specific ADL recognition algorithm; the only requirement is the capability of capturing complex and interleaved situations.

E. Abnormal Behaviors Recognition

The FINE-GRAINED ABNORMAL BEHAVIORS RECOGNITION module takes as input the activity instances detected by the ADLS RECOGNITION module and the aggregated sensor data from the MANIPULATION REFINEMENT module. Since the anomaly recognition method is based on knowledge-based reasoning, this information is encoded into logic predicates. The knowledge base is also enriched with personal knowledge related to the monitored individual and logic rules are then used to infer the anomalies that have possibly occurred. The approach is explained in detail in the Section III. Detected anomalies along with recognized ADLs instances are sent to software modules that support data analysis and visualization for the clinicians.

III. RECOGNITION OF FINE-GRAINED ANOMALIES

In this section we describe how we formally model *fine-grained abnormal behaviors* and how our system can detect them by combining the information of ADLs recognition, environmental sensor observations, and object manipulations.

We consider *abnormal behaviors* (*anomalies* for short) as behaviors observed within a short time period (which can vary from few seconds to a day) that diverge from the expected ones according to models provided by domain experts. In particular, our work focuses on those anomalies which can be predictive of cognitive decline. We already investigated the recognition of several types of those anomalies in [3]. In this work we aim at refining the granularity of the anomalies that we can detect by considering the manipulations of objects in the home environment.

A. Categories of anomalies

Abnormal behaviors which are symptoms of cognitive decline have been investigated by neuropsychology researchers [9], which characterized several functional difficulties in achieving everyday tasks which can be predictive of serious cognitive disorders like MCI or Dementia. In Table I we summarize the categories of anomalies that we consider in this work. Note that they all involve manipulations of everyday objects. Each category corresponds to several different types of abnormal behaviors.

TABLE I
LIST OF CONSIDERED ABNORMAL BEHAVIORS CATEGORIES

Type of anomaly	**Description**	**Example**
Omission	An important step within an ADL is not performed	The medicine box has been retrieved but no medicine is taken
Substitution	A different object than appropriate is used or a different component action than expected is performed	Pouring sugar instead of salt to prepare pasta
Replacement	The subject replaces a correct action with a wrong one	Putting the medicine box in the fridge
Wrong activity	The subject performs an activity that should not be done	The subject takes a not prescribed medicine
Inefficient execution	The subject performs actions which slow down/compromise the execution of the ADL	The subject takes double of the usual time in watering plants
Repetition	The subject repeats an ADL that he/she already performed forgetting it already took place	A medicine which is prescribed once is taken twice
Searching	The subject actively searches through home's repository for an item	The subject forgets where he/she put the salt and he/she searches it in all the repositories

B. Activity-dependent anomalies

Abnormal behaviors may depend on how ADLs are executed (or non-executed), while others are more generic and not tight to a specific ADL. In our model we consider both types of anomalies, distinguishing the two categories:

- *Activity-dependent*: if the anomaly is contextual to the occurrence (or non-occurrence) of one or more activity instances (e.g., the elderly executes the activity of taking medicines but takes the wrong medicine)
- *Activity-independent*: if the anomaly is not contextual to the occurrence (or non-occurrence) of one or more activity instances (e.g., the elderly just keeps on wandering around, searching for something for an unusually long time)

In the case of *activity-dependent* anomalies, they could be related to one or multiple activity instances. For instance, the omission of the activity *Preparing the table* can be considered as anomalous only if the ADL is not performed before the activities *Eating lunch* or *Eating dinner*, while it is not anomalous to omit it before the *Preparing breakfast* ADL. Another example is the repetition of medicines intake, where the same medicine is taken twice in two different *taking medicines* instances within the same prescription time.

Activity-independent anomalies, on the other hand, only relies on the sensed information. For instance, a substitution like *the butter is inserted in a non-refrigerated repository* should be fired independently with respect of recognized ADLs.

C. Subject-dependent anomalies

Orthogonally with respect to activity dependency, we also consider subject-dependent abnormal behaviors. Indeed, a behavior which is abnormal for an individual could be normal for another one. First of all, this may depend on medical prescriptions (e.g. medicines to be taken, diet, ...). Considering this type of information it is possible to provide detailed anomalies like "*the patient forgot to take his/her morning medicine*".

In addition to medical prescriptions, another important aspect of personalized abnormal behaviors are personal habits. For instance, the usual time and duration of execution of an ADL may vary for each person. Whenever an individual changes significantly his/her habits (e.g. taking longer to perform ADLs), it may be a symptom of a cognitive decline. Hence, we define some rules which capture the deviation from a past "normal behavior", mining statistics on normal execution of ADLs.

In general, subject-dependent rules are dynamically generated by considering medical prescriptions and personal habits.

Example 2: Consider the case where the subject is searching for salt to prepare pasta. It may be a normal habit to open two or three repositories in order to effectively find and retrieve the salt shaker. Hence, until the subject keeps on behaving as usual, no anomaly is detected. However, if the subject wanders around the home opening several times different repositories (much more than the usual two or three times), it may be considered as an abnormal behavior. Hence, the anomaly related to *searching for an item* is dynamically defined based on the past normal behavior of the subject.

D. Anomaly representation and recognition

The abnormal behaviors are usually described in natural language by domain experts (i.e., clinicians). We use a first-order logic knowledge-base to model those descriptions in terms of temporal relations between detected ADLs, high-level events, manipulations and personalized knowledge of the monitored individual. Then anomaly recognition is performed using a logic programming engine.

In our model, an anomaly is represented with the predicate $anomaly(an, aid, obj, t_s, t_e)$ where an is the anomaly's type, aid is the identifier of the activity instance related to the anomaly (if any), obj is the object related to the anomaly (if any), and t_s and t_e are respectively the starting and ending time of the anomaly occurrence.

A manipulation occurrence is represented by the logic fact $manipulation(o, m, t_s, t_e)$, where o is the object being manipulated, m is the manipulation label provided by the machine learning algorithm, and t_s and t_e are the starting and ending time of the manipulation, respectively. A sensor event as, for example, the opening of a cabinet is represented with the logic fact $action$(*PreparingTable-Id, open, Silverware-Drawer*, 2016-11-12 12:01:34), where the first argument is an activity instance id (used only when the system has classified the event as part of an activity), the second argument is the event type, the third is the object/area involved, and the last is the timestamp. An activity, as for example preparing table, once recognized by the system is represented by the logic fact $activity$(*PreparingTable-Id*, *PreparingTable*, 2016-11-12 11:58:00, 2016-11 12-12:05:12), where the first argument is the activity instance identifier, the second the activity type, and the last two are starting and ending timestamps, respectively. Subject-dependent knowledge, like prescribed medication, is also added to the knowledge base in terms of logic facts. For example, if $medicineA$ has to be taken everyday between 8 and 9 am, the fact $prescribedMedicineTime(medicineA, 8am, 9am)$ is added to the knowledge base.

Table II illustrates some examples of first-order logic rules used to infer abnormal behaviors. The role of object manipulations in the process of recognizing fine-grained anomalies is highlighted by the first two rules. Both rules detect an *omission*: the subject didn't take a medicine which was prescribed in a specific time interval (e.g. in the morning). In the first case the subject completely forgets to take the medicine, while in the second case the subject actually retrieves the medicine from the repository but then forgets to take it. Even if the practical implication of the two anomalies is the same, they represent two different patterns which may be important to distinguish for devising appropriate intervention mechanisms and possibly also for the clinical evaluation.

The third example illustrates a rule for the recognition of an activity-independent anomaly. When the manipulation *return* of a refrigerated object is close to the interaction with a non-refrigerated repository, the anomaly is fired. Note that his particular behavior is anomalous regardless of the performed ADL.

The last rule is an example of subject-dependent and activity-independent abnormal behavior. The anomaly is fired when the subject consecutively opens and closes k repositories without retrieving or returning any item, which indicates confusion about where an item is placed. The number of repositories k which are consecutively accessed to fire the anomaly is subject-dependent, and it is mined by analyzing the past normal behavior of the subject. If the subject consecutively accesses more repositories than the usual, then the behavior is considered as abnormal. Hence, the rule is automatically generated based on the value of k. The value Δt indicates the maximum amount of time between the opening of two repositories.

It is important to note that the occurrences of considered anomalies are not intended to provide an automatic diagnosis of the patients cognitive status. Indeed, individual habits or personal traits can include execution of abnormal behaviors which are not related with cognitive declines. However, the frequency of detected anomalies and their temporal trend can be used to infer behavioral changes.

IV. Experimental evaluation

A. Experimental setup

We implemented the system's prototype within our smart lab, which is instrumented with several environmental sensors like magnetic, power, presence and plug sensors. Those sensing devices are used to capture the interaction of the inhabitant with the home environment (repositories, chairs, electrical stove, ...) and continuously communicate their readings to a smart-home gateway using Z-Wave protocol.

We also attached tiny BLE accelerometers to several objects which are interesting to monitor different ADLs. In particular, we considered medicine boxes, a liquid bottle, a knife, food/beverage packages and a watering can. Those devices transmits continuously their accelerometer data to an Android mobile application which runs the Manipulations detection module. Manipulations are classified in real-time and then transmitted to the gateway.

TABLE II
EXAMPLES OF RULES MODELING ABNORMAL BEHAVIORS

No.	Rule	Anomaly type
1	$anomaly(medicine\_not\_even\_retrieved, Aid, Medicine, T_s, T_e) \leftarrow isMedicine(Medicine) \wedge prescribedMedicineTime(Medicine, T_s, T_e) \wedge activity(Aid, takingmedicine, T_{sa}, T_{ea}) \wedge (T_{sa} > T_s) \wedge (T_{sa} < T_e) \wedge (T_{ea} > T_s) \wedge (T_{ea} < T_e) \wedge not(isManipulated(Medicine, T_s, T_e))$	Omission: the subject completely forgot to take a prescribed medicine
2	$anomaly(medicine\_retrieved\_but\_not\_taken, -, Medicine, T_s, T_e) \leftarrow isMedicine(Medicine) \wedge prescribed(Medicine, T_s, T_e) \wedge isRetrieved(Medicine, T_s, T_e) \wedge isReturned(Medicine, T_s, T_e) \wedge not(isOpened(Medicine, T_s, T_e))$	Omission: the subject retrieved a prescribed medicine from the repository but then forgot to take it
3	$anomaly(refrigeratedfood\_in\_wrong\_repository, Aid, Item, T_s, T_e) \leftarrow isRefrigeratedFood(Item) \wedge manipulation(Aid, return, Item, Repository, T_s, T_e) \wedge not(isRefrigeratedRepository(Repository))$	Replacement: the subject places an object which needs to be refrigerated in a non-refrigerated repository
4	$anomaly(searching, -, -, T_{o_1}, T_{c_k}) \leftarrow action(-, open, Repository_1, T_{o_1}) \wedge action(-, close, Repository_1, T_{c_1}) \wedge (T_{o_1} < T_{c_1}) \wedge not(returnedOrRetrievedObjectsBetween(T_{o_1}, T_{c_1})) \wedge \ldots \wedge action(-, open, Repository1, T_{o_1}) \wedge action(-, close, Repository_k, T_{c_k}) \wedge (T_{o_k} < T_{c_k}) \wedge not(returnedOrRetrievedObjectsBetween(T_{o_k}, T_{c_k})) \wedge ((T_{o_k} - T_{o_1}) < (\Delta t \times k))$	Searching: the subject opens and closes several repositories more than k times to find some item without returning and retrieving any object (k is subject-dependent) within a short time interval.

TABLE III
PRELIMINARY RESULTS

Anomaly	Precision	Recall	F_1 score
Medicine X not even retrieved	0.79	1.0	0.88
Retrieve Medicine X But Not Opened	0.67	1.0	0.80
Open Medicine X Twice	0.88	0.70	0.78
Wrong Medicine Opened	1.0	0.8	0.89
Object Retake Multiple Times	1.0	1.0	1.0
Wrong Repository	0.9	1.0	0.95
Repository Search	0.89	0.89	0.89
Overall	0.88	0.91	0.89

We configured a Raspberry Pi to act as the smart-home sensor gateway to collect environmental sensors observations and object manipulations. A NodeJS REST server is in charge of receiving sensor data and storing it on a SQLite database. Periodically (e.g. at the end of each day) the gateway executes the SEMANTIC INTEGRATION OF SENSOR DATA module on environmental sensors, and transmits the derived high-level events, along with objects manipulations, to the MANIPULATIONS REFINEMENT module. Both modules are written in Java language. The MANIPULATIONS REFINEMENT module produces several log files containing the refined and aggregated sensor data. The sensor logs are used by the recognition algorithms of the ANOMALIES RECOGNITION PLATFORM, which are executed off-line.

B. Dataset collection

In order to validate our system, we accurately designed the acquisition of a dataset of several ADLs and anomalies. Our target activities are related to the kitchen environment. In the specific we considered the following ADLs: *taking the prescribed medicines*, *preparing breakfast*, *preparing meal* (i.e. lunch or dinner), *laying the table*, *eating*, *cleaning up* (i.e. clear the table and washing dishes) and *watering plants*. We designed several realistic scenarios, where each scenario represents a whole day of ADLs and abnormal behaviors performed by a different subject in its kitchen. Activities execution is designed to be as realistic as possible, with complex and interleaved patterns. In order to obtain a dataset which is the more general and robust possible, we introduced in all the scenarios several levels of variability in performing the ADLs/anomalies:

- **Variability in how a task is performed**: the same ADL can be performed in several different ways. For instance, a medicine can be taken with or without drinking water. Another example is the preparation of the meal, which can significantly vary depending on the recipe.
- **Variability in the order of actions**: Even two different ADLs execution which consist in the same task can significantly vary. The order of execution of actions can be different. Suppose, for instance, the pasta preparation. The inhabitant can significantly vary the order at which he/she accesses to repositories to retrieve food items and cooking instruments.
- **Variability in how ADLs are interleaved**: ADLs are often performed in an interleaved fashion. Hence, we introduced in the scenarios different ways of interleaving the activities. For instance, while sitting at the table during lunch, the inhabitant stops eating for a while to take its medications.
- **Variability in how an anomaly occurs**: Abnormal behaviors can occur with different patterns just like ADLs. Suppose for example the anomaly "*forgetting to take a prescribed medicine*". This can be done by totally forgetting to take it (no interaction with the medicine box) or by retrieving it at its repository but then forgetting to take it. Thanks to the manipulation recognition, we can monitor and distinguish these two cases.

Moreover, we included additional realistic scenarios, asking to the actors to simulate ADLs without following pre-defined scripts. Of course, those scenarios do not contain abnormal behaviors.

In total we acquired 752 instances of ADLs and 150 different patterns of abnormal behaviors. Those ADLs and anomalies have been collected in 40 scripted and 20 unscripted scenarios executed by 19 different volunteers.

C. Results

In this work we focus on the accuracy in recognizing fine-grained abnormal behaviors mostly based on the execution of ADLs. This accuracy is affected by different factors: a) the propagation of errors from noisy sensing devices, b) mistakes in manipulations classification, c) mistakes in ADLs recognition and d) inaccuracy in the modeling of abnormal behaviors rules. Since the accuracy of activity recognition has been extensively evaluated in several previous works, we assume to have an ADL recognition system which is accurate at 100%. This allows us to have a better understanding on which errors are introduced by the anomaly recognition rules engine in combination with a possibly noisy sensing infrastructure. Moreover, since ADLs recognition systems are becoming very accurate, this assumption does not imply unrealistic results (e.g. running SmartFABER [3] ADLs recognition algorithm on the dataset described in this section resulted in an overall F1 score greater than 0.9).

Our preliminary results are summarized in Table III. Due to lack of space, we report the results related to a subset of the considered abnormal behaviors. The first two anomalies represent the first two rules in Table III, which have been already discussed. The anomaly "*Open Medicine X Twice*" represents the scenario where the subject takes the same medicine twice within the prescription time. "*Wrong Medicine Opened*" describes the situation where the subject takes a medicine which is not prescribed at the time of intake. The anomaly "*Object Retake Multiple Times*" indicates the scenario where the subject repeatedly interacts with an object without performing any useful action with it. "*Wrong Repository*" anomaly captures the situation where an object is placed in a not appropriate repository. Finally, "*Repository Search*" is the last anomaly in Table III and it has been already introduced.

From the results it is possible to understand how manipulations recognition impacts on anomaly detection. In the first four rules, the mistakes in manipulations classification have a greater impact. This is because those anomalies are based on the specific manipulation "*accessing the content of medicine box*", which is often confused with "*significant movement*". On the other hand, the remaining rules are just based on significant movements of the objects, which are more easy to detect. The overall results are very promising, showing an average F_1 score of almost 0.9. Moreover, the method produced a very low number of false positives. Consider that the dataset consists of over 700 activity instances which included in total over 6.000 sensor events, while the total number of actual instances of abnormal behaviors is just 150. The total number of false positives considering the whole dataset is very low, and this means that rarely an abnormal behavior is fired during the normal execution of ADLs (i.e. the true negative rate is very high).

However, from the results it emerges the well known inflexibility of symbolic approach with respect of a noisy sensing infrastructure.

V. Conclusions and discussion

In this paper we showed how, by detecting specific object manipulations, we could significantly refine and improve our previous work on the recognition of fine-grained abnormal behaviors.

Despite the promising results, the proposed method has several limitations. The BLE accelerometers used in this work are suitable for a controlled environment like ours, while they may be more challenging to deploy in a real home environment for both costs (using a device for each object of interest could be expensive), energy consumption (these devices are powered by batteries), and practicality (objects could get in water, high temperature or they may have to be re-attached frequently in case of one-time use objects). However, we are confident that technological innovations will soon mitigate these limitations. Regarding the anomaly detection methods future work includes a) the investigation of probabilistic methods to model and detect certain types of anomalies that are not suited to be modeled by logic, b) the real-time operation of our framework, and c) the acquisition of a dataset from seniors using monitored objects while performing the activities.

Acknowledgments

The authors would like to thank Alessandra Plati and Claudio Tappia for their excellent work in dataset acquisition and software implementation.

References

[1] E. Hoque, R. F. Dickerson, S. M. Preum, M. Hanson, A. Barth, and J. A. Stankovic, "Holmes: A comprehensive anomaly detection system for daily in-home activities," in *2015 International Conference on Distributed Computing in Sensor Systems*, June 2015, pp. 40–51.

[2] B. Das, D. J. Cook, N. C. Krishnan, and M. Schmitter-Edgecombe, "One-class classification-based real-time activity error detection in smart homes," *IEEE Journal of Selected Topics in Signal Processing*, vol. 10, no. 5, pp. 914–923, Aug 2016.

[3] D. Riboni, C. Bettini, G. Civitarese, Z. H. Janjua, and R. Helaoui, "Smart-faber: Recognizing fine-grained abnormal behaviors for early detection of mild cognitive impairment," *Artificial Intelligence in Medicine*, vol. 67, pp. 57 – 74, 2016.

[4] G. Civitarese, S. Belfiore, and C. Bettini, "Let the objects tell what you are doing," in *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct*, ser. UbiComp '16. ACM, 2016, pp. 773–782.

[5] D. J. Patterson, D. Fox, H. Kautz, and M. Philipose, "Fine-grained activity recognition by aggregating abstract object usage," in *ISWC '05: Proceedings of the Ninth IEEE International Symposium on Wearable Computers*. Washington, DC, USA: IEEE Computer Society, 2005.

[6] J. Clawson, J. A. Pater, A. D. Miller, E. D. Mynatt, and L. Mamykina, "No longer wearing: Investigating the abandonment of personal health-tracking technologies on craigslist," in *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, ser. UbiComp '15, 2015, pp. 647–658.

[7] J. Ranjan and K. Whitehouse, "Towards recognizing person-object interactions using a single wrist wearable device," in *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct*, ser. UbiComp '16. ACM, 2016, pp. 722–731.

[8] J. Wu, A. Osuntogun, T. Choudhury, M. Philipose, and J. M. Rehg, "A scalable approach to activity recognition based on object use," in *Proc. IEEE International Conference on Computer Vision (ICCV 07)*. IEEE Press, 2007, pp. 1–8.

[9] T. Giovannetti, B. M. Bettcher, L. Brennan, D. J. Libon, M. Burke, K. Duey, C. Nieves, and D. Wambach, "Characterization of everyday functioning in mild cognitive impairment: A direct assessment approach," *Dement Geriatr Cogn Disord*, vol. 25, pp. 359–365, 2008.