

Attributing Recognised Activities in Multi-Person Households Using Ontology-Based Finite State Machines

Stijn Verstichel¹, Tom Tourwé², Jelle Nelis¹, Nicolás González-Deleito²,
Elena Tsiporkova², Femke Ongenaë¹, and Bruno Volckaert¹

¹ Ghent University - imec, IDLab, Department of Information Technology,
iGent Tower, Technologiepark-Zwijnaarde 15, B-9052 Ghent, Belgium,
stijn.verstichel@UGent.be,

WWW home page: <http://IDLab.UGent.be>, <http://IDLab.technology>

² Sirris, Bd. A. Reyerslaan 80, B-1030 Brussels, Belgium,
tom.tourwe@sirris.be,

WWW home page: <http://www.sirris.be>

Abstract Continuous person monitoring systems usually assume single-person households. Presence of visitors, such as family members and friends, and multi-person households are therefore not taken into consideration, although they can still influence results. However, in order to reason about a specific monitored person's evolution and trends, it is needed to accurately recognise the different activities that these persons perform throughout a day. This paper considers multi-person households and focuses on the problem of assigning activities detected by a continuous monitoring system to the person that has performed these activities. The proposed solution consists of modelling the domain terminology into an ontology, as well as that person's typical habits by means of a Finite State Machine – represented in its own ontology – using the concepts from the domain ontology. Sequences of recognized activities are then compared to the Finite State Machines associated to the different persons in a household, and assigned to the person with the most similar modelled sequence.

1 Introduction

Our aging society is facing huge challenges on a variety of domains, not least in the home-care domain. A lot of research is performed into how people can be continuously monitored remotely, so as to keep track of how they are doing, whether they stick to their treatment protocol, etc. In order to be able to reason about the evolution and trends, one first needs to understand a person's habits, and, to this end, activity recognition is needed. One of the driving factors for the research presented in this paper, is the fact that most of the current State-Of-The-Art (SOTA) is focused on individuals. Oftentimes, the fact that human beings live and grow old together is neglected. One might argue that technology is not or less necessary when people are still living together, as they can monitor

one another, but this is certainly not always the case. Some people already are mentally or physically less able to care for themselves, let alone for a partner or someone else living in.

A wearable cannot uniquely identify a person, as, per se, it is unaware of who is wearing it. The system should therefore allow to link a wearable with a person (e.g. by its unique name or an ID). Environmental sensors – sensors placed in the living environment of a resident, e.g. Passive Infrared Sensors (PIR) sensors, ambient temperature sensors, etc. – that complement the wearable can pick up signals from other (non-)monitored persons. Such sensors can also identify a person, in case the person wears a tag, when its (unique) signal is captured by a tag reader, e.g. in a specific room of the household. Current monitoring solutions do not deal with this multi-person situation. Even in single-person households, visitors can be present. In order to correctly analyse a person’s daily living habits, these potential problems need to be properly catered for.

This paper investigates multi-person scenarios and more specifically focuses on the problem of assigning a recognised activity to one of the persons present in that environment. Conceptually, the following approach has been taken to tackle the aforementioned research challenge. Residents in a household have been asked to keep a diary of their daily activities. This diary is then used to generate a model that captures the typical activities that follow one another, with appropriate guards (e.g. the resident only takes a shower when he/she has gotten out of bed). In this work, a Finite State Machine (FSM) has been chosen as supporting technology. In principle, this FSM can be generated based on other data, where such patterns of daily living have been identified from historical data. The FSM-based approach does not assign single activities to a single resident directly, but considers whether an activity that is observed fits a (set of) FSMs that are linked to (a) user(s). Since each user has its own FSM, one strives to end up with a single FSM linked to that activity, which basically then assigns that activity to the resident corresponding to that FSM.

Using the fixed scenario below, i.e. the morning scenario at an elderly couple’s place, the developed algorithms and services can be validated and demonstrated and is used throughout this paper as a running example. The scenario requires a number of sensors to be installed at the household to detect the correct activity (later on referred to as **ACTIVITY PROFILE**) performed by the persons in that household. The goal of the activity attribution is to automatically distinguish which person is doing which activity in a multi-person household setting.

1. One person wakes up at 7:00,
2. and starts preparing breakfast at 7:02.
3. The second person wakes up five minutes later,
4. and takes a shower at 7:07.
5. They have breakfast together at 7:20.
6. One of them starts reading the newspaper at 8:00, and
7. the other one takes a shower at 8:00.

The research presented in this paper was conducted in close collaboration with partners in the SMARTpro project.³

In this section the context of the research challenge has been sketched. The adopted technologies, the actual ontology decomposition and the software architecture of the presented solution are detailed in Section 3 to 5. The next section, Section 2 introduces relevant related work. This paper is concluded in Section 6.

2 Related Work

Research presented at the Pervasive Health Conference in Oldenburg, 2014 [6], lists a number of algorithms and platforms for household activity monitoring. The authors rightly claim that those systems, however, are either limited to the tracking of a single-person household or otherwise are complex to install. Two studies have been presented, namely *(i)* a feasibility study on the usage of binary sensor networks and *(ii)* a test on the effectiveness of the multi-hypothesis tracking algorithm, based on data from two people residing in a living lab from the Washington State University. In future work, the focus lies on reliable generation of correct topology graphs as well as on the improvement of the algorithm itself, eliminating inaccuracies introduced when multiple people are present.

Another interesting approach can be found in this thesis proposal [11]. Although the research focuses on a single-person household, the topics are interesting in the context of actual activity determination (later on referred to as Activity Profile) as a pre-processing step before the problem of a multi-person household can be investigated. A historic overview on automatic health monitoring, people tracking and activity recognition is presented as well as an extensive list of completed work, including research on the use of binary sensors, the algorithms for Simultaneous Tracked And Recognition (STAR) through the adoption of Bayesian networks, classical data association methods and particle filter implementations. Furthermore, improvements on the current SOTA are presented as well as the results of experiments on real and simulated data.

The algorithm and framework presented and evaluated in [9] enables resident counting using simple (non vision-based) sensors. This initial work assumes a number of preconditions which limit its current usage, such as the pre-definition of the number of residents.

The evaluation of the algorithms presented in [8] is based on simulated data and limits its scope on the tracking of moving objects in a binary sensor network. This means that the aspect of activity recognition is not taken into account. Still, the research results tackle the issue of multiple moving objects (or in our case residents) in non-disjoint areas.

Although [10] in itself does not present technical SOTA in terms of platforms, software or algorithms, it does provide an argumentation of why the roll-out of sensors in the homes, together with the appropriate autonomous evaluation algorithms and tool-support, is beneficial for the overall well-being and living conditions of the residents. Specific official medical and clinical-trial-validated tests

³ <http://smart-pro.eu/>

are performed by medical professionals on elderly people to score their mobility and capability in daily living, etc. However, those tests are only performed every now and then by the doctor, in a professional care setting, and thus only take into account long-term evolution. This could be one of the major arguments for enabling more frequent assessments in domestic environments and as such more fine-grained monitoring of the patients evolution and this in the comfort of their own home.

The authors in [4] highlight two important barriers for wide-scale installation of sensor networks in a residential setting, namely the cost of installation and perhaps even more important, making sense of all that captured data. Results are presented from a data collection exercise, using three sensors at 17 households and adopting a black-box style algorithm to describe typical user behaviour, starting from 55 days of unannotated data.

3 Defining the Household / Activity Vocabulary

The heart of our proposed system is based on semantic technologies. An ontology is used to fix the vocabulary used throughout this system. This vocabulary is modelled into a taxonomy, in which relationships (i.e. properties) and constraints (i.e. axioms) link the terms of this taxonomy together. The generic ontology structure is graphically illustrated in Figure 1. To keep the ontology properly maintainable, adaptable and configurable, it has been split up in a number of re-usable ontology modules. All those ontology modules model a specific part of the domain. Should the scenario evolve in such a way that certain aspects are not needed anymore, then the import approach can be modified in such a way that the redundant modules are omitted.

The following subsections detail the purpose and the contents of the actual domain ontology modules.

3.1 SMARTpro BASE Taxonomy

All common concepts representing the minimal ontological commitment are included in the `SMARTpro_Base` ontology module, which is illustrated in Figure 2.

3.2 SMARTpro DOMAIN Taxonomy

Each taxonomy module in this subsection extends the base ontology module, and thus also imports the `http://users.atlantis.ugent.be/svrstich/SMARTpro/SMARTpro_Base.owl` ontology. One module for every room in the household or for every Activity Profile to be processed, has been created.

Finally, these are all brought together, at least those modules needed for the current demonstration scenario, in the overall SMARTpro ontology. A representative sample can be seen in Figure 2.

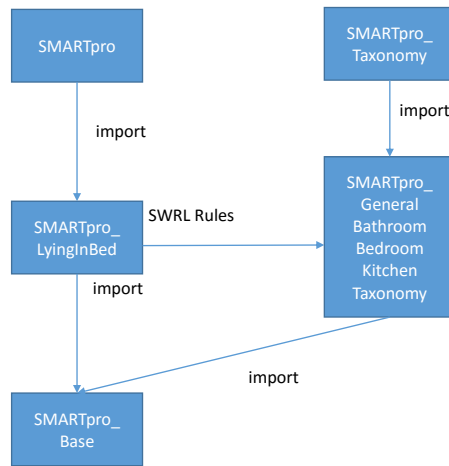


Figure 1. Ontology decomposition into separate re-usable independent modules, linked together by the import approach.

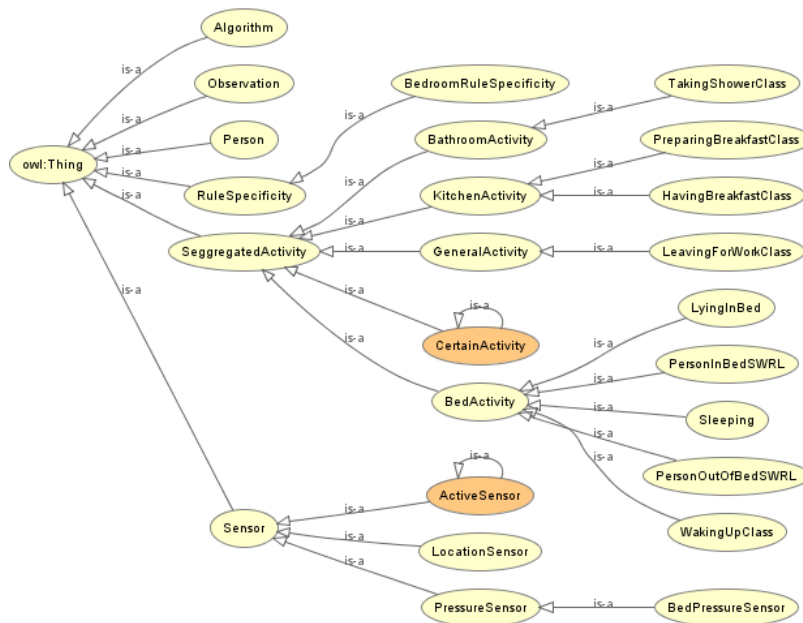


Figure 2. Overall asserted class hierarchy for the current SMARTpro scenario.

4 Modelling Behaviour using Finite State Machines

According to the Free Online Dictionary of Computing⁴, a FSM or also known as a Finite State Automaton (FSA) is an abstract machine consisting of a set of states (including the initial, intermediate and final states), a set of input events, a set of output events, and a state transition function.

4.1 FSM for the Activity Attribution Use-Case

We expect residents to behave in a more or less repetitive and structured manner. Note that this may not always be the case, and our approach will not be performing as desired, when the residents are rather unstructured in their daily active living. The states of the FSM are aligned with the SEGREGATEDACTIVITY concept from the domain ontology model, according to the scenario defined in Section 1.

As a base ontology model to represent the FSM, the results from the IST ELENA project [3] have been adopted, more specifically, the OWL ontology for state machines in the context of navigation and interaction modelling for web-based and hypermedia systems by Peter Dolog [2].

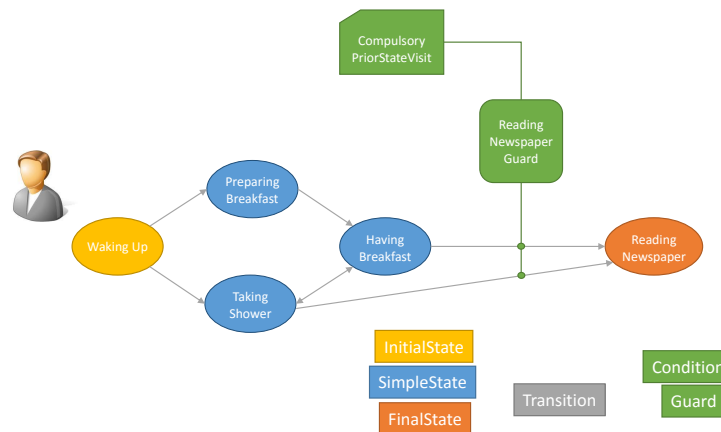


Figure 3. FSM for the multi-person household scenario.

In Figure 3, the actual FSM is presented. It has to be read from left to right. According to the defined scenario, for two residents: 1. The resident wakes up (Waking Up). 2. The resident takes either a shower (Taking Shower) or prepares breakfast (Preparing Breakfast). 3. After that, any of the two residents have breakfast (Having Breakfast) together. 4. One of the residents starts reading the newspaper (Reading Newspaper). 4a. However, before being allowed to do

⁴ <http://foldoc.org/>

that, there is a conditional guard (Reading Newspaper Guard) on the transition between having breakfast and reading the newspaper. 4b. In this situation this guard models that the resident should first have taken a shower (Compulsory Prior State Visit). 5. The other resident, which does not satisfy the condition defined in the guard, takes a shower (Taking Shower). 6. Now the second resident can also start reading the newspaper (Reading Newspaper).

In an M:N setting, i.e. more than one FSM belonging to more than one resident, multiple instances of such FSMs, each modelling slightly different situations, might be present. The goal of the detailed system is that, by analysing the incoming Activity Profiles from the low-level activity profile recognition algorithms – the actual detection of the activities is beyond the scope of this research, i.e. it is assumed that certain activities can be detected by an external system, but whom actually performed those is to be determined – it can find either one or more correct FSMs, still satisfying the sequence of already detected Activity Profiles. Moreover, as every FSM is linked in the system belonging to / representing the behaviour of a resident, we can conclude which resident has been responsible for the detection of those Activity Profiles. The system has been engineered in such a way that a true M:N situation can be supported. That is, every resident can have multiple instances of finite state machines. On the other hand, any given FSM can be linked to one or more residents.

4.2 Detecting the Appropriate FSM for the Resident, Using a Similarity Measure

The incoming sequence of detected Activity Profiles might not always correspond 100% with the pre-determined knowledge model. Perhaps, though, a similarity of 90% can be sufficient to make a well educated guess as to whom was actually executing those activities. To be able to calculate this similarity a certain similarity measure is needed.

The problem of measuring “similarity” of objects arises in many applications, and many domain-specific measures have been developed, e.g., matching text across documents or computing overlap among item-sets. Jeh and Widom [5] propose an approach, applicable in any domain with object-to-object relationships, which measures similarity of the structural context in which objects occur, based on their relationships with other objects. Given the graph-like triple nature of an FSM and indeed ontologies, this approach could be very effective. It calculates a measure that says “two objects are similar if they are related to similar objects.” This general similarity measure, called SIMRANK, is based on a simple and intuitive graph-theoretic model.

4.3 Finite State Machine Generator

In the previous subsections, the approach with using FSMs for the modelling and attribution of sequences of Activity Profiles to members of the household, has been detailed. Of course, before this methodology and corresponding algorithms can be used in a real-life setting, the actual FSMs corresponding to the typical

behaviour of the residents need to be instantiated. The approach adopted for this research has been to ask the residents to keep track of their typical behaviour during a number of days. The resulting logfiles were then fed into the system to generate the corresponding FSMs from. A small sample of such a logfile is given in the listing below.

```
29/06 7:00: Kate wakes up
29/06 7:02: Kate prepares breakfast
29/06 7:05: William wakes up
29/06 7:07: William takes a shower
29/06 7:20: Kate and William have breakfast together
29/06 8:00: William reads the newspaper
29/06 8:00: Kate takes a shower
```

5 Architecture

Figure 4 portrays the components and data flow involved in tracking multi-person households, expanding on previous research into our generic MASSIF [1] platform.

5.1 Data Producers (Module I)

Denoted with the title “DYAMAND”, reference is made towards the integration with the DYAMAND platform [7]. DYAMAND serves as the low-level sensor integration platform, used as a generic data provisioning component.

5.2 Conversion of Process Data into Activity Profiles (Module II)

The raw low-level data produced by the sensors can potentially generate high-frequency data samples. The aim of the services in “External Sources” module is to gather this data, execute proprietary algorithms on it in order to generate what has been referred to as “Activity Profiles” earlier on.

5.3 Preparing the Data (Module III)

Detected Activity Profiles are communicated to the MASSIF platform by means of the interface provided by the “Gateway” component.

Multiple sub-ontology modules have been created to support transparent personalisation and easy adaptation. The second component in Module III can be used to convert non-semantic data into semantic information, as well as to verify the correct formation of the incoming JSON fragment.

5.4 Data Broker (Module IV)

Module IV represents a data brokerage system. Its purpose is to streamline incoming data distribution. After all, the services responsible for the actual processing of the incoming data.

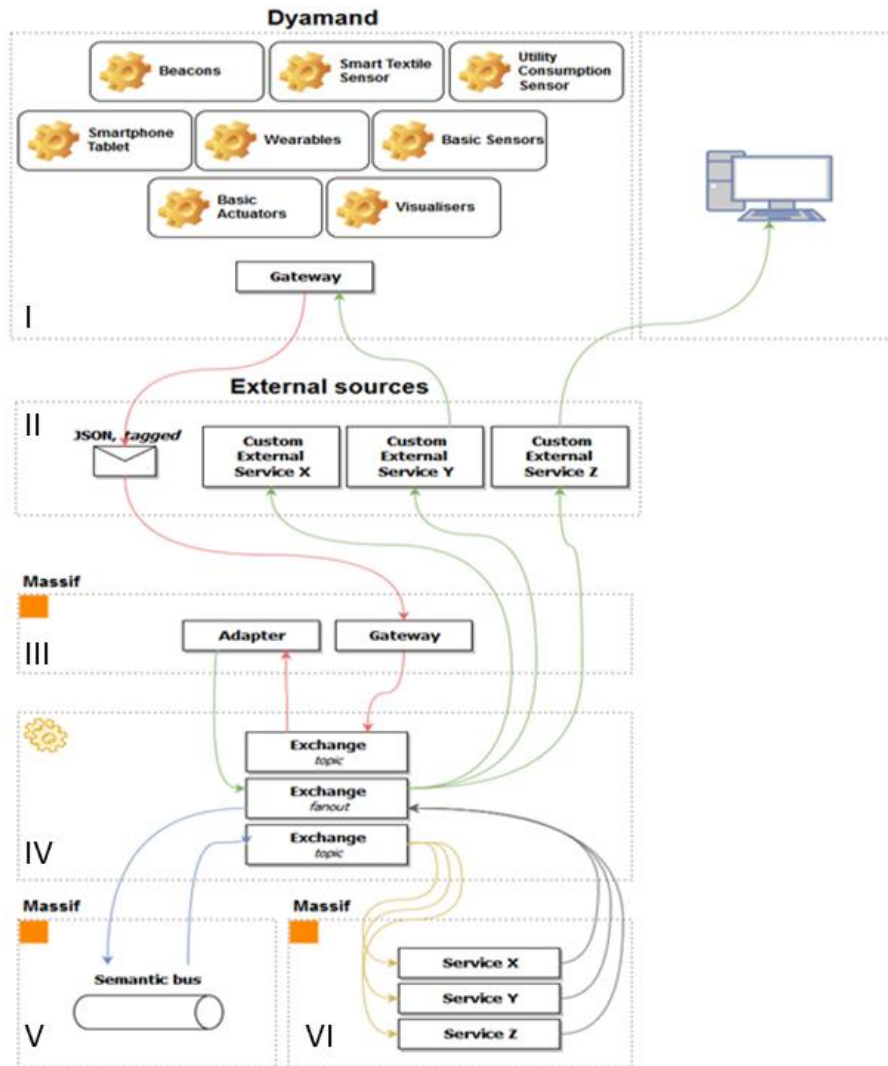


Figure 4. Overview of components required to demonstrate multi-person household activity attribution.

5.5 Intelligent Data Distribution (Module V)

The native message brokering system, as provided by the platform in Module IV, lacks support for intelligent data processing, e.g. by means of description-logics reasoning. The service subscribes for specific information, but cannot do this using first order logic. Therefore, the Semantic Communication Bus (SCB) enhances the data brokerage system by subscribing to all passing messages.

5.6 Activity Profile Clustering through FSM Analysis (Module VI)

In what has been presented so far, data engineering and communication mechanisms have been detailed.

The following paragraphs present in more detail how the algorithms are enabled in the MASSIF platform.

MCI Service In MASSIF a service that converts data into information, and processes this information to create knowledge is named a Meta-Context Information (MCI) Service.

Activity Profile Attribution Service This is the service responsible for attributing the detected Activity Profiles to a resident in the monitored household, and this based on what is specified in the knowledge model, i.e. in the FSMs. On a high level, the process of performing this attribution uses a combination of code-based programming and SPARQL queries (between brackets where applicable) and can be summarised as follows:

1. The FSM instance that contains a state representing the activity profile is queried. (=FIND_FSM_FOR_ACTIVITY_PROFILE)
2. For the given state from the FSM at hand, several specific cases need to be analysed and treated differently:
 - 2.a. Check whether the given state represents an initial state in the FSM. (=ASK_STATE_INITIAL)
 - 2.a.1. Execute the logic for an initial state in the FSM, i.e. check all clusters currently being found for all actors, until an empty cluster is found. After all, an initial state has to start from an empty cluster.
 - 2.a.2. If it is an initial state, add it to an empty cluster or create a new cluster.
 - 2.b. Check whether the given state represents a simple state in the FSM. (=ASK_STATE_SIMPLE)
 - 2.b.1. Execute the logic for a simple state in the FSM, i.e. check all clusters currently being found for all actors, until an appropriate cluster is found, i.e. a cluster where the last activity profile matches the previous FSM state of the current activity profile being analysed. (=QUERY_PREVIOUS_STATE)
 - 2.b.2. Iterate through all potential previous states according to the FSM ontology, until in the inner loop a cluster has been found where it can be added to.

2.b.3. Iterate through all clusters already present for one or more actors, and see whether the current activity profile being analysed can be added to that cluster.

2.b.4. Check whether a guard is attached to the transition between the two states, and if so, if the guard has been satisfied. (=QUERY_TRANSITION_GUARD_STATE)

2.c. Check whether the given state represents a final state in the FSM. (=ASK_STATE_FINAL)

2.c.1. Execute the logic for a Final State in the FSM, i.e. check all clusters currently being found for all actors, until an appropriate cluster is found, i.e. a cluster where the last activity profile matches the previous FSM state of the activity profile being analysed.

2.c.2. Finalise the current FSM.

Performance evaluation The performance in terms of Round Trip Time (RTT) has been evaluated. On a laptop, running Windows 10, with an Intel(R) Core(TM) i7-6600U CPU, running at 2.81 GHz and having 15.7 GB of RAM available, the following averaged results were achieved: *(i)* Platform startup: 29 024ms, *(ii)* Activity Profile Attribution service startup: 9ms, *(iii)* Gateway startup + adapter registration: 60 507ms, *(iv)* Activity Profile Attribution algorithm cycle: 426ms, *(v)* Shutdown: 64ms. Apart from *(iv)*, these timings represent a one-time execution. The algorithmic cycle can be subdivided into JSON fragment processing and the actual FSM matching aspect. On average 90% of the time is spent on the actual algorithm itself, with the remaining on the JSON processing.

6 Conclusions and Future Work

Attributing detected activities in a multi-person household is a difficult task. In this paper, a semantic approach to follow-up on daily activities in a multi-household setting has been presented. Both the software platform as well as the ontology model itself have been introduced. The algorithmic heart of the software is realised by means of ontology-based Finite State Machines. This system can help in allowing people to stay home in their habitual residences as long as possible.

The core concepts and technologies that have been used for this research are plenty, but appropriate for the tasks at hand. Ontologies have been used as central knowledge component. The SMARTpro project has allowed us to install sensors in residences, allowing us to collect data, together with a journal kept by the residents themselves. This data has been ingested through DYAMAND and processed by algorithms in the semantic MASSIF platform.

One aspect that needs more effort is the generation of the log files. It has been observed that keeping track of one's daily behaviour, in order to be able to generate a corresponding FSM, can be a very labour intensive task. More intelligent / automated approaches are definitely needed.

Acknowledgements

This research has been supported by the Flemish Innovation and Entrepreneurship Agency (VLAIO) in the context of the SMARTpro VIS-trajectory.

References

1. Pieter Bonte, Femke Ongenae, Femke Backere, Jeroen Schaballie, Dörthe Arndt, Stijn Verstichel, Erik Mannens, Rik Walle, and Filip Turck. The MASSIF platform: a modular and semantic platform for the development of flexible IoT services. *Knowledge and Information Systems*, 51(1):89–126, 2017.
2. Peter Dolog. The ontology for state machines. Online, 2017-07-04. <http://people.cs.aau.dk/~dolog/fsm/>.
3. Peter Dolog, Nicola Henze, Wolfgang Nejdl, and Michael Sintek. Personalization in distributed e-learning environments. In *Proceedings of the 13th international World Wide Web conference on Alternate track papers & posters*, pages 170–179. ACM, 2004.
4. Laura Fiorini, Filippo Cavallo, Paolo Dario, Alexandra Eavis, and Praminda Caleb-Solly. Unsupervised machine learning for developing personalised behaviour models using activity data. *Sensors*, 17(5):1034, 2017.
5. Glen Jeh and Jennifer Widom. SimRank: a measure of structural-context similarity. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 538–543. ACM, 2002.
6. Sebastian Müller. Ambient activity monitoring for medical applications in multi-person households. In *Proceedings of the 8th International Conference on Pervasive Computing Technologies for Healthcare*, pages 248–251. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2014.
7. Jelle Nelis, Tom Verschueren, Dieter Verslype, and Chris Develder. DYAMAND: Dynamic, adaptive management of networks and devices. In *37th Annual IEEE Conference on Local Computer Networks (LCN)*, pages 192–195. IEEE, 2012.
8. Songhwai Oh and Shankar Sastry. Tracking on a graph. In *Proceedings of the 4th international symposium on Information processing in sensor networks*, page 26. IEEE Press, 2005.
9. Jennifer Renoux, Marjan Alirezaie, Lars Karlsson, Uwe Köckemann, Federico Pecora, and Amy Loutfi. Context recognition in multiple occupants situations: Detecting the number of agents in a smart home environment with simple sensors. In *Knowledge-based techniques for problem solving and reasoning*, 2017.
10. Enno-E Steen, Thomas Frenken, Melina Frenken, and Andreas Hein. Functional assessment in elderlies homes: early results from a field trial. In *Ambient Assisted Living*, pages 3–17. Springer, 2014.
11. Daniel H Wilson and Chris Atkeson. Simultaneous tracking and activity recognition (STAR) using many anonymous, binary sensors. In *International Conference on Pervasive Computing*, pages 62–79. Springer, 2005.