Activity Context Representation: Techniques and Languages
AAAI Technical Report WS-12-05

Towards Activity Recognition Using Probabilistic Description Logics

Rim Helaoui[1], Daniele Riboni[2], Mathias Niepert[1], Claudio Bettini[2], Heiner Stuckenschmidt[1]
[1] KR&KM research group, University of Mannheim
[2] EveryWare Lab, Universita' degli Studi di Milano, D.I.Co.
Email: {rim,mathias,heiner}@informatik.uni-mannheim.de, {daniele.riboni,claudio.bettini}@unimi.it

Abstract

A major challenge of pervasive context-aware computing and intelligent environments resides in the acquisition and modelling of rich and heterogeneous context data. Decisive aspects of this information are the ongoing human activities at different degrees of granularity. We conjecture that ontology-based activity models are key to support interoperable multilevel activity representation and recognition. In this paper, we report on an initial investigation about the application of probabilistic description logics (DLs) to a framework for the recognition of multilevel activities in intelligent environments. In particular, being based on Log-linear DLs, our approach leverages the potential of highly expressive description logics with probabilistic reasoning in one unified framework. While we believe that this approach is very promising, our preliminary investigation suggests that challenging research issues remain open, including extensive support for temporal reasoning, and optimizations to reduce the computational cost.

Introduction

The spectacular progress of low-cost and low-power sensing has given rise to appealing computing fields. Designated as pervasive computing, this new paradigm is witnessing an increasing interest and a growing research community. Especially, the concept of emerging intelligent environments is progressively evolving into a commonplace. Such environments necessitate context awareness to support and assist the user with reactive and proactive services. Ongoing human activities are a decisive aspect of such contextual information. Hence, many researchers have been recently concerned with automatically recognizing human activities from lightweight dense sensing. Accordingly, several approaches have been proposed that can be classified as data-driven, knowledge based, or hybrid approaches. A multitude of works tried to apply diverse machine learning algorithms to recognize human activities. The majority proposed supervised learning methods such as Hidden Markov Models (Patterson et al. 2005) and conditional Random Fields (Buettner et al. 2009). The bottleneck of obtaining large amounts of training data has motivated some recent works to explore weakly supervised learning (Stikic and Schiele 2009) as well as unsupervised learning (Gu et al. 2010).

Despite being well suited for simple activities, data-driven techniques have a number of problems with the recognition of complex high-level activities. Their poor portability as well as the severe scalability problems they face, make them hardly deployable in other environments. Adding new activities without newly training and possibly designing the given model remains also out of reach. Furthermore, these techniques are generally doomed to a flat representation of human activities which contradicts their hierarchical nature. Finally, the lack of formal semantics prevents data-driven models from encoding the inherent common-sense knowledge underlying human activities.

As an alternative to data-driven approaches, some researchers adopted different logical modelling and reasoning algorithms to address human activity recognition. These include (Cirillo et al. 2009) where general predefined rules are used to recognize activities based on a constraint-based temporal reasoning framework (OMPS). In (Bouchard, Giroux, and Bouzouane 2006), the authors leverage action description logics and lattice theory for plan recognition to predict human behaviour. Despite their ability to cope with some of the limitations of data-driven approaches, rule-based systems suffer from many restrictions, including limited support for uncertainty and temporal reasoning.

Combining both paradigms has been the motivation of our previous works (Helaoui, Niepert, and Stuckenschmidt 2011) and (Riboni and Bettini 2011). In the former, we applied Markov Logic Networks to unite both logical statements as well as probabilistic ones in one single framework. This allows to address the uncertain aspect of human activities and integrate crucial background knowledge. Due to the specific rules and the learning phase, the same system can not be reused under different settings. The capability of thoroughly exploiting the semantic features between the context and the activities leaves much to be desired. In the latter, we proposed a loosely coupled technique to integrate supervised learning methods with DLs. We defined an ontology to formally model the semantics of activities, and exploited ontological reasoning to refine the prediction of the machine learning algorithm based on the current context. While useful for improving the recognition rates of simple actions, that method was not well suited to recognize complex activities, since the ontological language lacked support for uncertainty and temporal reasoning.

In this paper, we explore the convergence of our previous works by the use of probabilistic DLs to exploit semantic features of dense sensing for activity recognition. Compared to symbolic activity recognition based on manual specification of a set of rules, our approach aspires to facilitate portability and interoperability of the recognition system. As such, we propose to use ontologies as a commonly shared conceptual knowledge description standard. To leverage the full potential of DLs, we use the same unified language for both modelling multi-level activities and their context and probabilistically reasoning about them. This hybrid framework is achieved through adoption of log-linear DLs as proposed in (Niepert, Noessner, and Stuckenschmidt 2011).

The rest of the paper is structured as follows. In the next section we describe related works. Then, we present our overall recognition framework and techniques. Finally, we conclude the paper with a discussion of our approach and future work.

Related Work

In this section, we describe approaches closely related to our work. Especially, we discuss those that utilize ontologies and semantic information to enable and improve activity recognition. There are only few works that keep both semantic description of the activities and their recognition process tightly-coupled. Such an approach has been adopted in (Chen, Nugent, and Wang 2011) and in (Springer and Turhan 2009). Similarly to our work, the authors load the current contextual information to populate their generated ontology then employ inference reasoner to obtain the most specific equivalent activity or situation. Hence, activity recognition can be mapped to equivalency and subsumption reasoning. Especially, Chen $el\,al.$ proceed to an incrementally specific recognition of the activities through the progressive activation of the sensors. Neither works, however, address the inherent uncertainty aspect in human activities. Apart from some experiments with noisy sensor data, both systems do only reason with facts and implicitly assume a deterministic mapping from the context data to the activities' descriptions.

Bridging the gap between such a purely symbolic approach and supporting uncertainty, was the concern of several works recently. Following a lazy instance based approach, Knox $et\,al$ (Knox, Coyle, and Dobson 2010) use a vector of the sensors' values to define their cases. A semantically extended case base is created through extracting ontological relationships between sensors, locations and activities. This allows them to reduce the resulting number of cases. Further efforts to exploit semantic information to improve the recognition system are detected in (Yamada et al. 2007) and (Wang et al. 2007). Relying on the subsumption hierarchy, the former involves ontology to handle unlearned objects and map them into learned classes. At the recognition step, parametric mixture models are applied. In the latter, the subsumption hierarchy helps automatically inferring probability distributions over the current actions given the object in use. Thus, the integrated common-sense knowledge is used to learn a Dynamic Bayesian Network-based activity classifier. Other attempts to cope with uncertainty involve applying a hierarchy Bayesian networks based on the ontology's instances such as in (Latfi, Lefebvre, and Descheneaux 2007). All these works dissociate the inference step from the semantic model. This aspect limits the ability of incorporating rich background and common sense knowledge. It also strips the system from other advantages of symbolic reasoning such as consistency check.

To the best of our knowledge, (Hong et al. 2009) is the only ontology-based tightly-coupled human activity recognition approach from dense-sensing. The authors model the interrelationships between sensors, contexts and activities. They use the resulting hierarchical network of ontologies to generate evidential networks. Following Dempster-Shafer theory of evidence, they calculate and define the heuristic relationships between the network's nodes in form of evidential mappings. These mappings are used through seven steps of evidential operations as inference process. Obviously, their evidential network discloses limited expressiveness compared to our DLs language. To the limitations also belongs cardinality constraints for example. This, such as other rich background knowledge, can be flexibly described and respected through the reasoning and consistency check process. Moreover, using a wide-spread language such as OWL 2 (Grau et al. 2008) also offers the potential of using existing ontologies as well as mining techniques for automatic ontology learning.

Representing and Recognizing Multilevel Activities with Log-Linear DLs

In this section, we describe our technique for multilevel activity recognition.

System overview

As shown by the following example, an activity recognition system should be able to recognize activities at both coarse- and fine-grained level of detail.

Example 1. *Consider a healthcare system to remotely monitor the activities of an elderly person. Recognized activities are periodically communicated to a medical center, in order to evaluate the evolution of the patient's physical and cognitive capabilities. A similar system should be able to recognize not only high-level activities such as "preparing breakfast" and "cleaning house", but also the more specific actions that compose those activities. For instance, it should be possible, for the medical center, to inspect the simpler activities that were executed during breakfast preparation (e.g., washing fruit, cooking egg, fill a jug with milk, . . .), to better evaluate the patient's behavior.*

Figure 1 shows a multilevel decomposition of complex activities (Level 1) in more specific components; we adopt the structure proposed in (Lukowicz et al. 2010) for the Opportunity dataset. Level 2 represents simple activities, which can be described as temporal sequences of manipulative gestures and modes of locomotion (Level 3). Finally, Level 4

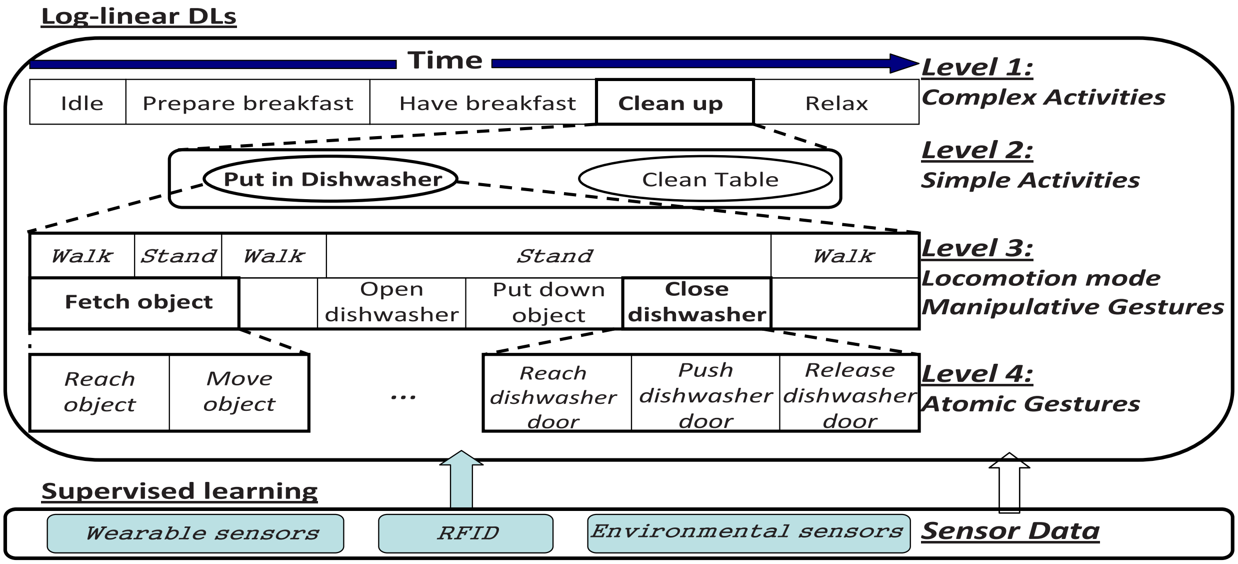

Figure 1: Overview of the system for multilevel activity recognition

describes the atomic gestures that characterize manipulative gestures.

As shown in Figure 1, atomic gestures may be recognized through the application of supervised learning methods, based on data acquired from body-worn sensors (e.g., accelerometers to track movements of the body), environmental sensors, and RFID tags to detect objects usage. In order to achieve the goal of higher-level activity recognition, in our technique we propose the use of DLs to express background knowledge about the domain, and to model the relationships among high-level activities and their simpler component actions. In order to cope with the uncertainty of both context facts and activity definitions, we adopt *log-linear description logics (DLs)*, a probabilistic DL that has been recently proposed in (Niepert, Noessner, and Stuckenschmidt 2011).

Description Logics and Log-linear Description Logics

Description logics are a commonly used representation for knowledge bases. There are numerous tools and standards for representing and reasoning with knowledge using DLs. The DLs framework allows one to represent both facts about individuals (concept and role assertions) as well as axioms expressing schema information. Log-linear DLs integrate description logics with probabilistic log-linear models (Niepert, Noessner, and Stuckenschmidt 2011). In particular, Log-linear DLs allow to model both probabilitic and deterministic dependencies between DL axioms through extending uncertain axioms with weights. Regarding the expressiveness of the language, Log-linear DL supports the same operators of the well-known OWL 2 language.

The syntax of log-linear DLs is equivalent to that of the underlying DL except that it is possible to assign weights to general concept inclusion axioms (GCIs), role inclusion axioms (RIs), and assertions. We explicitly allow assertions (facts about individuals) in the knowledge base and we will use the terms *constraint box* (CBox) and *knowledge base* (KB) interchangeably. Moreover, for ease of presentation, we will use the term *axiom* to denote GCIs, RIs, *and* concept and role assertions. A log-linear knowledge base $\mathcal{C} = (\mathcal{C}^{\mathsf{D}}, \mathcal{C}^{\mathsf{U}})$ is a pair consisting of a *deterministic* CBox $\mathcal{C}^{\mathsf{D}}$ and an *uncertain* CBox $\mathcal{C}^{\mathsf{U}} = \{(c, w_c)\}$ with each c being an axiom and w_c a real-valued weight assigned to c. The *deterministic* CBox contains axioms that are known to hold and the *uncertain* CBox contains axioms with weights. The greater the weight of an axiom the more evidence there is for it to hold. Each axiom can either be part of the deterministic or the uncertain CBox but not both. The deterministic CBox is assumed to be coherent and consistent.

The simple semantics of log-linear DLs is based on probability distributions over *coherent* and *consistent* knowledge bases. The weights of the axioms determine the log-linear probability distribution. For a log-linear knowledge base $\mathcal{C} = (\mathcal{C}^{\mathsf{D}}, \mathcal{C}^{\mathsf{U}})$ and a CBox $\mathcal{C}'$ with $\mathcal{C}^{\mathsf{D}} \subseteq \mathcal{C}' \subseteq \mathcal{C}^{\mathsf{D}} \cup \{c : (c, w_c) \in \mathcal{C}^{\mathsf{U}}\}$, we have that

$$\mathrm{Pr}_{\mathcal{C}}(\mathcal{C}') = \begin{cases} \frac{1}{Z} \exp\left(\sum_{\{c \in \mathcal{C}' \setminus \mathcal{C}^{\mathsf{D}}\}} w_c\right) & \text{if } \mathcal{C}' \text{ is } \textit{coherent} \text{ and } \textit{consistent}; \\ 0 & \text{otherwise} \end{cases}$$

where Z is the normalization constant of the log-linear distribution $\mathrm{Pr}_{\mathcal{C}}$. An axiom with weight 0 that is not in conflict with any other axiom has the marginal probability of 0.5. Hence, the semantics of log-linear DLs leads to a distribution compatible with the open-world assumption.

Example 2. *Let* S *and* P *be two concepts and let* $\mathcal{C} = (\mathcal{C}^{\mathsf{D}}, \mathcal{C}^{\mathsf{U}})$*, with* $\mathcal{C}^{\mathsf{D}} = \emptyset$ *and* $\mathcal{C}^{\mathsf{U}} = \{\langle \mathsf{S} \sqsubseteq \mathsf{P}, 0.8\rangle, \langle \mathsf{P} \sqsubseteq \mathsf{S}, -1.0\rangle, \langle \mathsf{S} \sqcap \mathsf{P} \sqsubseteq \bot, 0.4\rangle\}$*. Then,*

$$\begin{aligned}
\Pr_{\mathcal{C}}(\emptyset) &= Z^{-1}\exp(0) &&\approx 0.17\\
\Pr_{\mathcal{C}}(\{\mathsf{S} \sqsubseteq \mathsf{P}\}) &= Z^{-1}\exp(0.8) &&\approx 0.38\\
\Pr_{\mathcal{C}}(\{\mathsf{P} \sqsubseteq \mathsf{S}\}) &= Z^{-1}\exp(-1.0) &&\approx 0.06\\
\Pr_{\mathcal{C}}(\{\mathsf{S} \sqcap \mathsf{P} \sqsubseteq \bot\}) &= Z^{-1}\exp(0.4) &&\approx 0.25\\
\Pr_{\mathcal{C}}(\{\mathsf{S} \sqsubseteq \mathsf{P}, \mathsf{P} \sqsubseteq \mathsf{S}\}) &= Z^{-1}\exp(-0.2) &&\approx 0.14\\
\Pr_{\mathcal{C}}(\{\mathsf{S} \sqsubseteq \mathsf{P}, \mathsf{S} \sqcap \mathsf{P} \sqsubseteq \bot\}) &= 0\\
\Pr_{\mathcal{C}}(\{\mathsf{P} \sqsubseteq \mathsf{S}, \mathsf{S} \sqcap \mathsf{P} \sqsubseteq \bot\}) &= 0\\
\Pr_{\mathcal{C}}(\{\mathsf{P} \sqsubseteq \mathsf{S}, \mathsf{S} \sqsubseteq \mathsf{P}, \mathsf{S} \sqcap \mathsf{P} \sqsubseteq \bot\}) &= 0
\end{aligned}$$

with $Z = \exp(0.8) + \exp(0.4) + \exp(0) + \exp(-0.2) + \exp(-1.0) \approx 5.90$.

Maximum A-Posteriori Inference Under the given syntax and semantics the first central inference task is the maximum a-posteriori (MAP) query: *"Given a log-linear ontology, what is a most probable coherent ontology over the same concept and role names?"* In the context of probabilistic DLs, the MAP query is crucial as it infers a *most probable* classical ontology from a probabilistic one. The ELOG reasoner (Noessner and Niepert 2011) solves MAP queries by transforming the probabilistic ontology into an integer linear program. It iteratively queries a reasoner such as Pellet to derive explanations for incoherences or inconsistencies and adds those as constraints to the ILP, resolves, and so forth.

Context Representation

In our framework, we use log-linear DLs to represent context data, including human activities. It is well known that traditional ontological models of context have strong points in terms of representation of heterogeneous context data, interoperability, automatic reasoning, and representation of complex relationships. However, they fall short in representing uncertain context facts, since the OWL language does not natively support uncertain reasoning (Bettini et al. 2010). The use of log-linear DLs allows us to naturally extend OWL with uncertain reasoning, while retaining the expressiveness of the original language. By using a unified language for representing both activities and other context data, we exploit context information to recognize current activities by maximum a-posteriori inference on the probabilistic ontology instantiated with current context data.

Representation of Multilevel Activities

The use of a probabilistic DLs has two main advantages with respect to existing approaches in representing activities: i) use of ontological reasoning to derive higher-level activities based on their component simple actions and current context; ii) representation of the uncertainty about both activity axioms and context facts through weights, which determine their log-linear probability distribution. Multilevel activities, as those shown in Figure 1, are represented by concepts, which are associated to weighted axioms.

Example 3. *Suppose that we want to represent the complex activity "clean up" in terms of its component simple activities "put in dishwasher" and "clean table". Hence, we can add the following axioms to the knowledge base:*

$$\textsc{CleanUp} \sqsubseteq \textsc{ComplexActivity} \sqcap \forall\, \textsc{hasActor}.\Big(\textsc{Person} \sqcap \exists\, \textsc{hasSimpleActivity}.\textsc{PutInDishwasher}\Big), 1.8 \tag{1}$$

$$\textsc{CleanUp} \sqsubseteq \textsc{ComplexActivity} \sqcap \forall\, \textsc{hasActor}.\Big(\textsc{Person} \sqcap \exists\, \textsc{hasSimpleActivity}.\textsc{CleanTable}\Big), 1.5 \tag{2}$$

*Axiom (*1*) has weight* 1.8*, and essentially states that if a person is currently putting things in the dishwasher, she is probably performing the complex activity "clean up". Axiom (*2*) is similar, apart that it has a slightly lower weight, and considers activity "clean table" instead of "put in dishwasher". Note that the axioms weights can be manually defined based on background knowledge, or automatically learned from a training set of performed activities.*

A major issue with our approach is that the description logic underlying Log-linear DL (which has essentially the same expressiveness of OWL 2 (Grau et al. 2008)) does not natively support temporal reasoning. This is a serious limitation, since reasoning with temporal intervals is needed to recognize many human activities. Consider, for instance, the simple activity "put in dishwasher" depicted in Figure 1. A possible situation describing that activity is: *"the individual takes an object, walks to the dishwasher, opens its door, puts the object inside, and closes the dishwasher door"*. Other definitions for the same activity may exist, based on the individual's habits. For instance, many objects may be put inside the dishwasher while its door is left open. In order to express these temporal relationships among activities and actions, we should be able to reason with temporal intervals. Unfortunately, this kind of reasoning is not naturally supported by the description logic underlying Log-linear DL, due to some restrictions to the language operators that are necessary for preserving the decidability of ontological reasoning problems. This shortcoming is shared by any activity recognition system strictly based on OWL 2. On the contrary, DLs for interval-based temporal languages, like the one presented in (Artale and Franconi 1998), are well suited for reasoning with the temporal characteristics of activities and actions, but lack support for some expressive operators that are admitted in OWL 2, which are needed to model complex context data. An alternative approach we are considering is to exploit an external reasoner to perform temporal reasoning, while maintaining the whole expressiveness of our ontological language. We further discuss this issue in the last section of this paper.

Our Ontology for the Opportunity Dataset

As a testbed for our technique, we chose to use the dataset acquired within the *Opportunity* EU research project (Activity and Context Recognition with Opportunistic Sensor Configurations)[1]. The dataset has been acquired in a smart home

[1] http://www.opportunity-project.eu/

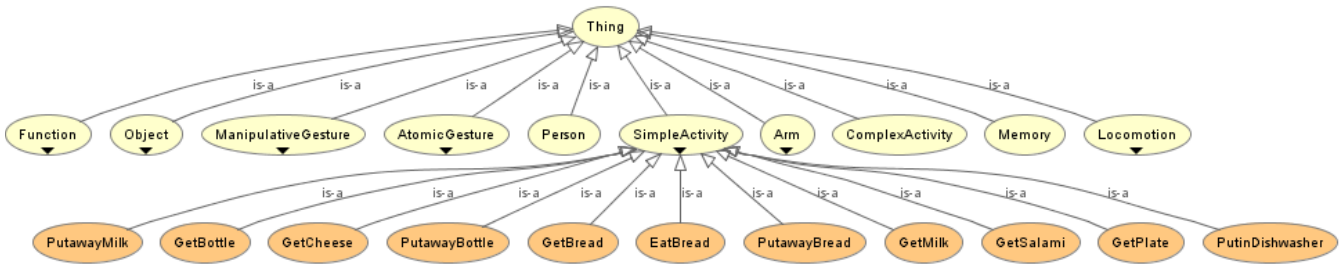

Figure 2: A snapshot of our ontology

with dense sensing equipments (72 sensors of 10 modalities, including accelerometers, RFID tags attached to objects, and environmental sensors for audio and video). A number of multilevel activities in an "early morning" setting has been acquired by 12 subjects. Gestures and activities are structured similarly as shown in Figure 1. The dataset is presented in detail in (Lukowicz et al. 2010).

The Opportunity dataset provides labels for atomic gestures (Level 4). The goal of our experiments will be to evaluate the feasibility of our technique to recognize gestures and activities at Levels 3, 2, and 1 starting from atomic gestures. For the sake of these experiments, we will assume that a state-of-the-art method is used to recognize activities at Level 4 based on raw context data. For our experiments, we have developed an ontology for representing the multilevel activites considered in the Opportunity dataset. The ontology has been developed using the Protégé OWL editor (Knublauch et al. 2004); axioms weights have been added as annotations to axioms definitions. Figure 2 shows a snapshot of our ontology.

Determining Axioms Weights

According to our log-linear model, the marginal probability of an axiom is the sum of the probabilities of each consistent and coherent ontology in which that axiom holds. We adopt the association rules learning principle to obtain a-priori probabilities of individual axioms: We first expand our training data with the corresponding multi-level activity labels. This is achieved based on the concept definitions provided in the $TBox$. Each of these axioms is considered as a rule. For example, considering the lowest activity level, the *Atomic Gesture* would be the antecedent and the *Manipulative Gesture* the consequent. After deleting the consecutive duplicate events, we determine the support value for the considered rules which actually coincide with the probability values of the corresponding axioms. In the example above, the obtained support value corresponds to the conditional probability of *Manipulative Gesture given an Atomic Gesture*. Since the semantics of log-linear DLs is based on log-linear probability distributions, we compute the a-priori weights by taking the logit of the a-priori probabilities.

The Recognition Framework

At the time of writing, we are developing a prototype system in Java to evaluate our activity recognition technique. At first, the Java program adds to the assertional part of our ontology an instance of PERSON, which represents the current individual. Then, the program parses the Opportunity dataset to acquire the atomic gestures executed by the individual. Each gesture is added as an instance of its subclass of ATOMICGESTURE, and a role HASATOMICGESTURE is added, to relate the gesture with the individual. At each new gesture, we use the ELOG reasoner to calculate the most probable consistent ontology, given the current gestures. Then, we use the Pellet[2] reasoner for realizing the ontology calculated by ELOG, in order to derive the specific manipulative gestures, simple, and complex activities performed by the individual.

Discussion and Future Work

In this paper, we propose a tightly-coupled hybrid system for human activity recognition. Our framework unites both symbolic and probabilistic reasoning. This is achieved through the adoption of highly expressive log-linear DLs to represent and reason about the current activity at different granularities and complexity levels simultaneously. This paper outlines an initial implementation of the proposed framework. It describes the recognition algorithm and illustrates the concepts of our work through several examples. While the full evaluation of the proposed approach awaits further steps and experiments, our ontology-based approach combines promising features to address human activity recognition. The benefits of the proposed approach are manifold. Unlike the majority of related works, it supports the inherent uncertain nature of human activities without sacrifying the advantages of ontological reasoning. These advantages include consistency checking, the ability of integrating rich background knowledge, and the simultaneous recognition of coarse and fine-grained activities. The use of a standard description formalism enhances the portability and reusability of the proposed system, and supports the representation of heterogeneous and uncertain context data. Moreover, the declarative nature of DLs reinforces the flexibility and intelligibility of the system.

As pointed out earlier, however, there are some challenging open issues that need to be addressed. Currently, our approach does not support temporal reasoning, which is a key requirement for human activity recognition. There are two alternative approaches for enabling temporal reasoning in

[2]http://clarkparsia.com/pellet/

a DL framework. The first one consists in the use of temporal description logics, in which the temporal and terminological domains are tightly coupled. A relevant instance of these languages was proposed in (Artale and Franconi 1998), in which a temporal DL is used for reasoning about actions and plans. Applied to human activity recognition, actions essentially represents instantaneous activities (like atomic and manipulative gestures), while plans represent more complex activities, which are defined as temporally-constrained sequences of actions. However, in order to preserve decidability, the expressiveness of non-temporal operators is limited, and there is no support for reasoning with uncertainty. The second approach consists in the use of a loosely-coupled technique, in which time is treated as a *concrete domain* (Lutz and Milicic 2007). With this approach, ontology instances are related to values of the temporal domain by functional properties, and an external reasoner is used to deal with qualitative and/or quantitative relationships among the time intervals corresponding to activities executions. We believe that the latter approach is a promising research direction to follow. A further issue is that the computational time of activity recognition with a large knowledge base (as the one for the dataset considered in this work) may not be compatible with the requirements of ambient intelligence systems. Computational costs could be reduced by pruning, given the context at run time, those ontology axioms that are not involved in the activity derivation. Future work also includes the validation of our approach with additional datasets that include richer context sources.

References

Artale, A., and Franconi, E. 1998. A temporal description logic for reasoning about actions and plans. *J Artif Intell Res* 9:463–506.

Bettini, C.; Brdiczka, O.; Henricksen, K.; Indulska, J.; Nicklas, D.; Ranganathan, A.; and Riboni, D. 2010. A survey of context modelling and reasoning techniques. *Pervasive and Mobile Computing* 6(2):161–180.

Bouchard, B.; Giroux, S.; and Bouzouane, A. 2006. A smart home agent for plan recognition of cognitively-impaired patients. *Journal of Computers* 1(5):53–62.

Buettner, M.; Prasad, R.; Philipose, M.; and Wetherall, D. 2009. Recognizing daily activities with rfid-based sensors. In *Proc. of Ubicomp'09*, 51–60. ACM.

Chen, L.; Nugent, C. D.; and Wang, H. 2011. A knowledge-driven approach to activity recognition in smart homes. *IEEE T. Knowl. Data En.* 99(PrePrints).

Cirillo, M.; Lazellotto, F.; Pecora, F.; and Saffiotti, A. 2009. Monitoring domestic activities with temporal constraints and components. In *Proc. of IE'09*, 117–124. IOS Press.

Grau, B. C.; Horrocks, I.; Motik, B.; Parsia, B.; Patel-Schneider, P. F.; and Sattler, U. 2008. OWL 2: The next step for OWL. *J Web Semant* 6(4):309–322.

Gu, T.; Chen, S.; Tao, X.; and Lu, J. 2010. An unsupervised approach to activity recognition and segmentation based on object-use fingerprints. *Data Knowl Eng* 69(6):533–544.

Helaoui, R.; Niepert, M.; and Stuckenschmidt, H. 2011. Recognizing interleaved and concurrent activities using qualitative and quantitative temporal relationships. *Pervasive Mob Comput* 7(6):660 – 670.

Hong, X.; Nugent, C.; Mulvenna, M.; McClean, S.; Scotney, B.; and Devlin, S. 2009. Evidential fusion of sensor data for activity recognition in smart homes. *Pervasive Mob Comput* 5(3):236–252.

Knox, S.; Coyle, L.; and Dobson, S. 2010. Using ontologies in case-based activity recognition. In *Proc. of FLAIRS'10*. AAAI Press.

Knublauch, H.; Fergerson, R. W.; Noy, N. F.; and Musen, M. A. 2004. The protégé owl plugin: An open development environment for semantic web applications. In *Proc. of ISWC'04*, volume 3298 of *LNCS*, 229–243. Springer.

Latfi, F.; Lefebvre, B.; and Descheneaux, C. 2007. Ontology-based management of the telehealth smart home, dedicated to elderly in loss of cognitive autonomy. In *Proc. of OWLED'07*, volume 258 of *CEUR Workshop Proceedings*. CEUR-WS.org.

Lukowicz, P.; Pirkl, G.; Bannach, D.; Wagner, F.; Calatroni, A.; Förster, K.; Holleczek, T.; Rossi, M.; Roggen, D.; Tröster, G.; Doppler, J.; Holzmann, C.; Riener, A.; Ferscha, A.; and Chavarriaga, R. 2010. Recording a complex, multi modal activity data set for context recognition. In *Proc. of ARCS '10*, 161–166. VDE Verlag.

Lutz, C., and Milicic, M. 2007. A tableau algorithm for description logics with concrete domains and general tboxes. *J. Autom. Reasoning* 38(1-3):227–259.

Niepert, M.; Noessner, J.; and Stuckenschmidt, H. 2011. Log-linear description logics. In *Proc. of IJCAI'11*, 2153–2158. IJCAI/AAAI.

Noessner, J., and Niepert, M. 2011. ELOG: a probabilistic reasoner for OWL EL. In *Proc. of RR'11*, 281–286. Springer.

Patterson, D. J.; Fox, D.; Kautz, H.; and Philipose, M. 2005. Fine-grained activity recognition by aggregating abstract object usage. In *Proc. of ISWC'05*, 44–51. IEEE Comp. Soc.

Riboni, D., and Bettini, C. 2011. Cosar: hybrid reasoning for context-aware activity recognition. *Personal Ubiquitous Comput.* 15(3):271–289.

Springer, T., and Turhan, A.-Y. 2009. Employing description logics in ambient intelligence for modeling and reasoning about complex situations. *J Ambient Intell Smart Environ* 1(3):235–259.

Stikic, M., and Schiele, B. 2009. Activity recognition from sparsely labeled data using multi-instance learning. In *Proc. of LoCA'09*, volume 5561 of *LNCS*. Springer. 156–173.

Wang, S.; Pentney, W.; Popescu, A.-M.; Choudhury, T.; and Philipose, M. 2007. Common sense based joint training of human activity recognizers. In *Proc. of IJCAI'07*, 2237–2242. Morgan Kaufmann.

Yamada, N.; Sakamoto, K.; Kunito, G.; Isoda, Y.; Yamazaki, K.; and Tanaka, S. 2007. Applying ontology and probabilistic model to human activity recognition from surrounding things. *IPSJ Digital Courier* 3:506–517.