

A Survey on Ontologies for Human Behavior Recognition

NATALIA DÍAZ RODRÍGUEZ, Turku Centre for Computer Science (TUCS), Department of Information Technologies, Åbo Akademi University, Turku, Finland

M. P. CUÉLLAR, Department of Computer Science and Artificial Intelligence, University of Granada, Spain

JOHAN LILIUS, Turku Centre for Computer Science (TUCS), Department of Information Technologies, Åbo Akademi University, Turku, Finland

MIGUEL DELGADO CALVO-FLORES, Department of Computer Science and Artificial Intelligence, University of Granada, Spain

Describing user activity plays an essential role in ambient intelligence. In this work, we review different methods for human activity recognition, classified as data-driven and knowledge-based techniques. We focus on context ontologies whose ultimate goal is the tracking of human behavior. After studying upper and domain ontologies, both useful for human activity representation and inference, we establish an evaluation criterion to assess the suitability of the different candidate ontologies for this purpose. As a result, any missing features, which are relevant for modeling daily human behaviors, are identified as future challenges.

Categories and Subject Descriptors: I.2.4 [Artificial Intelligence]: Knowledge Representation Formalisms and Methods—*Semantic networks*; I.2.6 [Artificial Intelligence]: Learning—*Knowledge acquisition*; D.3.1 [Programming Languages]: Formal Definitions and Theory—*Semantics*

General Terms: Ontologies

Additional Key Words and Phrases: Human behavior recognition, activity recognition, context awareness

ACM Reference Format:

Natalia Díaz Rodríguez, M. P. Cuéllar, Johan Lilius, and Miguel Delgado Calvo-Flores. 2014. A survey on ontologies for human behavior recognition. *ACM Comput. Surv.* 46, 4, Article 43 (March 2014), 33 pages. DOI: <http://dx.doi.org/10.1145/2523819>

1. INTRODUCTION

Semantics is the branch of linguistics and logic concerned with meaning. Its two main areas are (a) *logical semantics*, concerned with matters such as sense, reference, pre-assumption, and implication, and (b) *lexical semantics*, dealing with the analysis of word meanings and relations between them.¹ In systems equipped with semantic tools, information is given well-defined meaning so that it enables computers and people to work in cooperation.

¹Oxford Dictionaries, <http://oxforddictionaries.com/definition/english/semantics>.

Authors' addresses: Natalia Díaz-Rodríguez and Johan Lilius, Turku Centre for Computer Science (TUCS), Department of Information Technologies, Åbo Akademi University, ICT-Talo, Joukahaisenkatu 3-5, FIN-20520 Turku (Finland); M. P. Cuéllar and Miguel Delgado Calvo-Flores, Department of Computer Science and Artificial Intelligence, University of Granada, Escuela Técnica Superior de Ingenierías Informática y de Telecomunicación. C/Daniel Saucedo Aranda, s/n 18071 Granada (Spain); email: {ndiaz, jolilius}@abo.fi, {manupc, mdelgado}@decsai.ugr.es.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2014 ACM 0360-0300/2014/03-ART43 \$15.00

DOI: <http://dx.doi.org/10.1145/2523819>

In this work, we make a survey and comparison of semantic tools for the recognition of human behavior. We focus on a more detailed analysis of ontologies, since properties like flexibility, reasoning, information sharing, and knowledge representation make these models one of the most promising tools for this purpose [Chen and Biswas 2009]. Both the environment and the user can be provided with semantics to help in the context definition process. Moreover, including semantics aids in the management of user information and the interaction with the system, facilitates the definition and comprehension of human behavior, and, thereafter, helps to develop new learning and recognition models.

In this introduction, we make first a contextualization of the problem and provide a general perspective on human behavior recognition. After that, Section 1.2 discusses context awareness tools to motivate the deeper analysis of ontologies in this survey. The aim of this section, in order to make this review more complete, is also to show different tools to model context and identify system requirements that serve to evaluate ontologies on behavior recognition. Finally, we provide the structure of this article.

1.1. Ambient Intelligence

The term *ambient intelligence* (AmI) was introduced by the European Commission in 2001 [Ducatel et al. 2001] as a response to new user needs in ubiquitous environments. One of the latest definitions [Augusto 2007] describes AmI as *digital and proactive environments, with capacity to sense the environment and assist users in their daily lives*. The major difference between *ubiquitous computing* and AmI is the introduction of AI in the latter. Thanks to that, user friendliness, user empowerment, human assistance, and easy interaction with efficient services are improved. Examples of AmI scenarios can be seen in Augusto [2007], Remagnino et al. [2005], and Cook et al. [2009].

AmI systems are usually composed by at least (a) a *perception* mechanism to gather information from both the user and the environment [Remagnino et al. 2005], (b) a set of *actuators* to modify the environment and communicate with users, and (c) a *reasoning/decision-making* module able to recognize what is happening to users in the environment, what they do, and what their aims are and to make decisions to assist them. These three abstract components often expand when designing AmI applications, since some scenarios also require the design of a sensor network, data fusion techniques, or a real-time response, among others [Remagnino et al. 2005; Ramos et al. 2008].

In AmI systems in general, the user occupies a central part. Thus, it becomes necessary to develop techniques to model, learn, recognize, and predict *what users are doing* in the environment, so that the system is able to make decisions about how to assist them. Usually, the literature calls *what users are doing human behavior* or *human activity* interchangeably [Remagnino et al. 2005; Cook et al. 2009; Rashidi and Cook 2009; Ros et al. 2013]. These terms usually mean *a sequence of human actions that can be tagged with a label*, that is, the corresponding activity/behavior. However, most of these authors agree to define *human action* as the simplest unit in the human activity, and it is usually associated with a sensor event. From our point of view, as long as new semantic approaches are being developed [Chen and Nugent 2009], new abstraction levels appear in the system. For this reason, in our opinion, a difference should be made between the terms *human activity* and *human behavior* to separate the concepts of what the user is really doing in the environment (activity), which is inferred from sensor data and machine-learning techniques, and the purpose or meaning it could have (behavior). However, since this work is a survey, we will not consider any difference between *human activity* and *behavior* in the rest of the article to preserve consistency with the existing literature.

Besides perception, actuation, and AI techniques, another important part included in the design of AmI systems is the analysis of the most suitable task model required

to achieve a well-designed approach. In Activity Theory [Nardi 1996], the fundamental unit of analysis is human activity. Activities' aim is to accomplish a goal. Sampling, analyzing, and modeling are examples of a user modeling process [Casas et al. 2008]. As the level of decomposition in task modeling depends on its purpose, task models should be rich in information and flexible to capture all the main activities that should be performed to reach the desired goals, as well as the different ways to accomplish them [W3C 2010]. A summary of existing task models and their limitations can be found in Gharsellaoui et al. [2012]. *HTA (Hierarchical Task Analysis)*, *GTA (Groupware Task Analysis)*, *CTT (Concur Task Trees)*, *UAN (User Action Notation)*, *TKS (Task Knowledge Structure)*, *DIANE+*, and *TOOD (Task Object-Oriented Description)* are examples of hierarchical, tree-based, and taxonomic structures. Common requirements often included in task models consider tasks, goals, activities, devices, time operators, and so forth. Despite the great advances produced in the last decade, the complexity and the quantity of possible complex activities [Naeem and Bigham 2007], the temporal interdependences among actions [Ros et al. 2013], the relevance of the semantics associated with a behavior [Chen and Nugent 2009], and the existence and interaction of several actors in the same environment [Singla et al. 2010; Cook et al. 2009], among others, make learning and recognition of human behavior nontrivial and bring up clear challenges in AmI research.

As new approaches for task modeling include semantics to represent the meaning of human activity, context awareness techniques become a more central part of AmI systems. To make this review more complete, the next subsection provides an introduction to the most widely used frameworks in context modeling.

1.2. Context Awareness

Context consists of any information that can be used to characterize the state of an entity [Dey and Abowd 2000]. Entities can include a person, an object, an environment, an application, or a device that interacts with the user. Context awareness is one of the drivers of the ubiquitous computing paradigm, and a well-designed model is a key accessor to the context in any context-aware system [Strang and Linnhoff-Popien 2004]. Proposals to model context can be integrated with human activity models provided with semantics. This subsection details models and frameworks to deal with context information as well as some ontologies used in these. The ontologies will be further studied in Section 3.

As regards data integration [Winograd 2001], we may distinguish among widgets, networked services, and blackboard models. Widgets may hide low-level details of sensing and ease application development. On the other hand, networked services can be less efficient than widgets, but at the same time, they may form a more robust and flexible approach (e.g., by using a widget manager discovery in a context server architecture). Finally, blackboard models have a data-centric view with event subscription capabilities, which provides simplicity for the addition of new data sources. A drawback is that they can have low efficiency in communication. Concerning the way data is captured, sensors can be classified as physical sensors (i.e., hardware), virtual sensors (context data from software applications or services), or logical sensors (combining physical and virtual sensors). Different physical sensors are available for diverse types of contexts, for example, light, visual context, audio, motion, acceleration, location, touch, temperature, or other physical attributes such as biosensors [Baldauf et al. 2007] or nonvisual tracking systems in general [Zhou and Hu 2008].

There are different approaches [Strang and Linnhoff-Popien 2004] to conceptually model context. Key-value models [Baldauf et al. 2007] are one of the simplest approaches that serve to describe service capabilities in service discovery matching. Markup scheme models use hierarchical data structures with tags to define, for example, profile instances. Graphical model tools [Strang and Linnhoff-Popien 2004] such

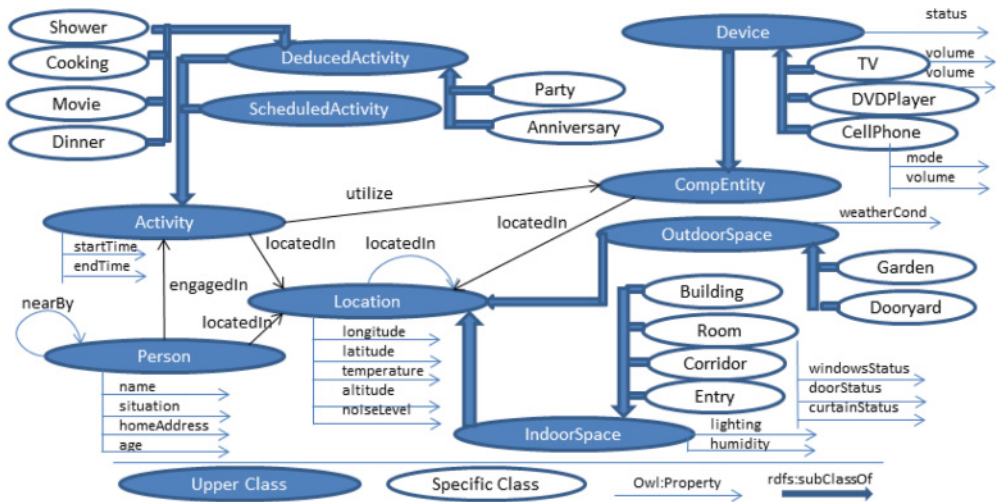


Fig. 1. OWL-encoded context ontology (CONON) [Wang et al. 2004] (partial).

as UML (Unified Modeling Language) or extensions to ORM (Object-Role Modeling) are useful to appropriately and easily model context by means of graphical interfaces. Object-oriented models take advantage of object-oriented features to encapsulate context processing and representation through well-defined interfaces. Logic-based models are formal ways to represent facts, expressions, and rules that allow an inference process to derive new facts based on existing rules. Finally, ontology-based models describe concepts and relationships in a high and formal expressiveness level [Baldauf et al. 2007]. Ontology-based context modeling overcomes the limitations of other models regarding simplicity, flexibility, extensibility, generality, expressiveness, and automatic code generation [Saleemi et al. 2011]. Interoperability solutions based on the ontology model can benefit from ontology reasoning, since ontologies are the most promising and expressive models fulfilling requirements for modeling context information [Baldauf et al. 2007].

The survey [Baldauf et al. 2007] shows advanced context models in a good compendium of design architectures with their respective advantages and disadvantages. Georgia Tech aware house [Kientz et al. 2008] is an example of a project that develops applications to support seniors living independently at home by using time-independent heterogeneous context sources. Regarding context-aware frameworks, examples of OWL-based approaches for context modeling are *CoBrA* and *SOCAM*. *CoBrA* [Chen et al. 2003] is an agent-based infrastructure for context modeling, reasoning, and knowledge sharing using context acquisition components. *SOUPA* and *CoBrA-Ont* ontologies are some of the related tools. User privacy control is also included. *SOCAM* (Service-Oriented Context Aware Middleware) [Gu et al. 2004] introduces a server-based architecture for building context-aware services focused on information sensing and context providers. Another project providing an OWL-encoded context ontology (*CONON*) is Wang et al. [2004]. As an example, a reduced part of *CONON*, for the home domain, can be seen in Figure 1. *CoDAMos* (Context-Driven Adaptation of Mobile Services) [Preuveneers et al. 2004] is another ontology, and it contains four main concepts: User, Environment, Platform, and Service. An excerpt of the *CoDAMos* ontology, with different environment conditions and locations, can be seen in Figure 2.

In *Gaia* [Roman et al. 2002], a metaoperative system is extended to include context awareness. Instead of using RDF Triples, *Gaia* uses 4-ary predicates (the fourth one is

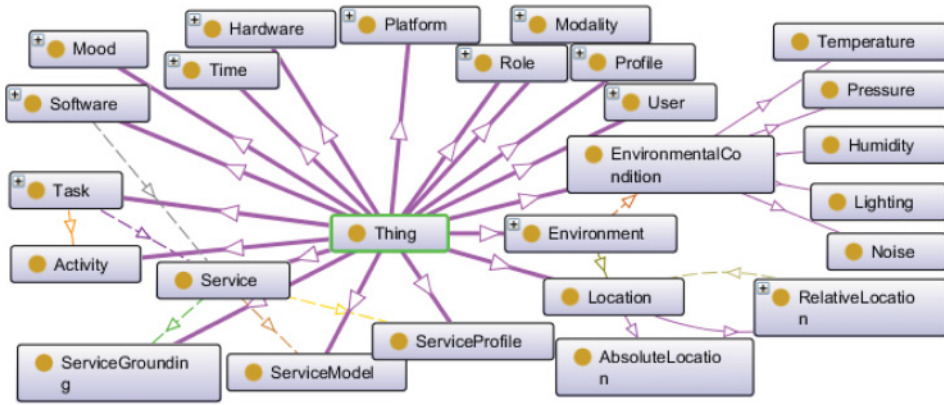


Fig. 2. Ontology concepts within the CoDAMos ontology [Preuveneers et al. 2004].

context type), first-order logic, and DAML+OIL. Other architectures for smartphone context-aware frameworks can be found in Kovacs et al. [2009], Yun Her [2010], and van Sinderen et al. [2006].

Context Toolkit [Dey and Abowd 2000] presents an approach to enable application development through reusable components. Situations are modeled on a system level, but there is no language-level situation modeling. However, a restriction in meaningfulness exists due to its attribute-value tuples, in contrast with RDF. Many systems use SQL rather than the semantic standard SPARQL. Another example is *HIPPIE* [Oppermann and Specht 2000], which utilizes existing users' information to distribute context information to their devices. *NESSIE* [Prinz 1999] focuses, on the other hand, on event-based awareness. For compensating its lack of handling interaction, *HIPPIE* was combined with *NESSIE* [Prinz 1999], but the result still lacked a semantic information description. In Bettini et al. [2008] and Agostini et al. [2009], the Context Aggregation and REasoning (*CARE*) middleware interacts with the *COSAR* [Riboni and Bettini 2011a] system to recognize human activities through hybrid ontological/statistical reasoners executed on personal mobile devices.

In summary, we can observe different frameworks to facilitate the creation of context-aware services. They use distinct context representation models and different sensors and infrastructure. However, we may notice that there is a predominance of OWL languages. In addition, we observed critical research issues such as the type of context modeling and reasoning, knowledge sharing, and user privacy [Chen et al. 2003]. Ontologies, as context-aware representation methods for human activities, will be analyzed in Section 3.

1.3. Article Organization

At this point, we have discussed some architectures for the design of AmI applications, with special emphasis on user activity modeling. We have shown different points of view for the development of human behavior recognition systems. We distinguish between data-driven approaches, mainly focused on the branch of pattern recognition and machine learning, and knowledge-based approaches, which include tools to model semantics and require a more accurate and refined activity knowledge representation. Section 1.2 provided us with an overview on the techniques and frameworks to model context and semantics in this application field.

Both types of approaches are studied in Section 2. Specifically, Section 2.1 focuses on machine-learning proposals used for human activity modeling and recognition, and

Section 2.2 focuses on knowledge-based approaches. Finally, Section 2.3 provides a simple taxonomy that summarizes the evaluation aspects considered for this literature review and identifies and discusses open challenges.

In this survey, we focus our efforts on the evaluation of ontology-based techniques for semantic modeling in human behavior recognition, since these techniques have shown better performance and possibilities than others in the management of context awareness for this application study, as shown in the previous subsection. Section 3 is aimed to this purpose. First, it introduces the concept of ontology reasoning and shows its main advantages in human behavior modeling. After that, in Section 3.1, we select a set of relevant ontologies for human activity modeling, and for context and environment representation in Section 3.2. Section 4 starts with a brief discussion about ontology evaluation methodologies, and Section 4.1 provides an assessment of each proposal. Then, the benefits and disadvantages of the presented ontologies and methodologies are discussed in Section 4.2. Finally, Section 5 concludes the review with a summary and promising research directions to be addressed.

2. TRENDS IN ACTIVITY RECOGNITION AND MONITORING

Activity recognition techniques can be divided into data-driven methods and knowledge-based approaches. Traditionally, techniques for activity recognition have focused on the branch of pattern recognition and machine learning and belong to the first group. These techniques have been extensively studied in the last decade; however, since they have been tackled from the pattern recognition perspective, they have not integrated mechanisms for semantic treatment or management. Despite this fact, they have meant an important step forward in the AmI discipline, and therefore, we devote Section 2.1 to their study. After that, a review on knowledge-driven and hybrid approaches is provided in Section 2.2, to describe those methods for activity inference and recognition from the knowledge engineering point of view. Finally, Section 2.3 ends with a taxonomy to evaluate the existing approaches and open challenges.

2.1. Data-Driven Approaches for Human Behavior Recognition

Approaches for the recognition of human activities and the detection of anomalies during their performance use the information provided by sensors to build, infer, or calibrate a behavior model. Machine-learning techniques have been extensively used with this purpose, and, more specially, probabilistic models, data mining, and inductive learning.

Probabilistic systems provide great flexibility when controlling different alternatives in the performance of behaviors and may be easily adapted to different environments. The work [Kim et al. 2010] collects a study on advantages and drawbacks on the use of stochastic techniques for human activity recognition. Although Bayesian networks [van Kasteren and Krose 2007; Rashidi and Cook 2009], Naïve Bayes classifiers [Tapia et al. 2004], or nonparametric Bayesian clustering methods [Tominaga et al. 2012] have been used in several cases, one of the most common approaches is Hidden Markov Models (HMMs) [Boger et al. 2006; Meng et al. 2006; Zhang et al. 2006]. The systems implementing these approaches usually build a model for each activity to be recognized. These activities are checked in parallel as the sensors cast events in order to find the most likely model that fits the current human behavior. These approaches have the power to be noise tolerant for sensor data and are capable to model sensor failure probability, but they have two main limitations: first, HMMs suffer from the first-order assumption; second, if a behavior may be performed in several ways, it is also necessary to calibrate multiple models, one for each. Multiple Behavioral HMM (MBHMM) [Naeem and Bigham 2007] and conditional random fields [Kim et al. 2010] overcome this last limitation. The determination of which tasks are currently active,

even if the activity has not been concluded by the user, is possible with these models. This is very helpful when there are several ways to finish an activity correctly or when actions within an activity can be executed in any order. Another work dealing with the loss of sensor event data is Naeem et al. [2007], where they use a size and relevance-based hierarchy of activities of daily living.

Traditionally, probabilistic methods have been the most used models, but data mining techniques are also found for behavior mining [Rashidi and Cook 2009; Wilson et al. 2005]. Another proposed machine-learning technique is inductive learning [Maurer et al. 2006; Delgado et al. 2009] to mine the most frequent actions of a behavior and their temporal relationship and to build decision trees that represent the ways in which an activity could be performed. These approaches provide a simple representation and a fast detection of the human behavior; however, their limitation arises when activities with cyclic actions should be modeled. Regarding other models, in Ermes et al. [2008], a neural network is designed to receive data from active sensors (acceleration, temperature, etc.), which are used to infer if the user is rowing, biking, playing football, walking, running, sitting, or hiking. In this case, the noise tolerance and low computation requirements to detect the activity in real time are key points of the approach. However, due to the inherent features of neural networks, their main limitations are difficulty in training the network with no local optima, its later adaptation to changes in the behavior, the validation of the results, and the interpretation of the network performance.

All previous approaches use passive sensors located in the environment or active wearable sensors to acquire data. One recent trend in human activity analysis is to use computer vision to avoid body markers. An approach-based taxonomy was designed in Aggarwal and Ryoo [2011] to categorize works in this area, distinguishing among nonhierarchical approaches, developed for the recognition of gestures and actions, and hierarchical approaches for high-level interaction analysis. A survey about these techniques may be found in Poppe [2010] and Chaaoui et al. [2012]. The main potential of these methods is that the information gathered by video sensors may provide much more information than passive sensors about the state of the user, its position, and its movements. In contrast to wearable sensors, these have the advantage to be much more comfortable and invisible to the user. However, the cost of these systems, the privacy loss, and the complexity of techniques to identify every action are limitations to be solved by these approaches nowadays.

Due to the growing interest in designing smart environments that reason about residents [Cook and Das 2005], intensive data collection is becoming more common. A great number of applications are tested on single user settings (naive Bayes classifiers [van Kasteren and Krose 2007], decision trees [Maurer et al. 2006], or conditional random fields [Vail et al. 2007]). Nowadays, open challenges are how to tackle multiple resident settings, interleaved activities, or social interaction, although we can find preliminary approaches in the literature [Singla et al. 2010; Cook et al. 2009; Crandall and Cook 2010; Li and Dustdar 2011]. For example, Singla et al. [2010] focuses on real-time recognition of interrupted and interleaved activities among multiple residents in smart environments. Manually labeled data from 40 residents were used to identify the most likely sequence of actions corresponding to a series of sensor events. The average accuracy was 60.60%. Parallel activities and those in cooperation have not been considered, since a strong constraint is the assumption of knowing the person's ID for each sensor event. With respect to handling social interaction, in Cook et al. [2009], an unsupervised learning algorithm is applied to detect social interaction and monitor ADLs in the CASAS smart environment [Rashidi and Cook 2009]. Activity and event density maps visualize sensor events for 15 days in a two-resident apartment. They applied a supervised learning algorithm with two HMMs (resident identifier and

activity identifier models) with an accuracy of 90%. However, since not all interactions can be determined by physical proximity, they suggest fusing the resident identification and the activity identifier models into a multilayer hierarchical model to improve the activity recognition task in multiresident settings.

Considering the adaptation of behaviors to changes in user routines, HMM and similar stochastic models have the limitation of being static; therefore, they cannot be applied to dynamic environments. A possible solution to overcome this limitation while maintaining the probabilistic nature of the system is Ros et al. [2013], where human activity modeling is enhanced with the online adaptation to habit changes, environmental variations, and temporal information. In this approach, actions are ordered into temporal execution levels, and learning automata are in charge of changing the temporal execution level associated to an action as the user performs changes in his or her routine. However, the main limitation of this technique is that it is not applicable to activities with an unknown number of cyclic execution times.

Other important issues are system scalability and tackling the presence of multiple users. Approaches that address these problems can be found in Crandall and Cook [2010] and Cook et al. [2009]. Likewise, it is important to consider the migration of learned behaviors to other spaces and their adaptation to other users [Rashidi and Cook 2010]. All these works assume probabilistic models such as HMM or Bayesian networks: a model is in charge of the inference of the user that is performing a behavior while a secondary one performs the activity recognition. In spite of these advances, there is no general solution that integrates all the aspects desired in an SS (i.e. scalability, multiuser support, adaptation to other users or routine changes, and support for interleaved activities and social behaviors). In addition, there is no consensus on what kind of sensors should be used and what activity models or what information treatment methodologies should be employed in each case. This is due to the fact that it mostly depends on the type of problem domain and task to be solved.

However, despite the current open challenges and the limitations of the existing data-driven approaches, the biggest strength of these models is that they are able to handle noise, uncertainty, and incomplete sensor data [Chen and Nugent 2009]. In addition, they have proven to be accurate in different domains where semantics are not key, and according to different problem constraints.

2.2. Knowledge-Driven Approaches for Human Behavior Recognition

Among knowledge-driven approaches to recognize human behavior, we can find event calculus [Kowalski and Sergot 1986]—describing agents and actions with durations and temporal relationships—and situation calculus [Levesque et al. 1998]—a logic-based framework for defining actions and changes in the representation of the world. Other variants [Chen et al. 2008] model temporal characterization of activities and causality relationships between activities and events. The problems found in these approaches to dynamic pervasive computing environments are related with interoperability and adaptation to different scenarios, since context data sources are dynamic and not known in advance. To solve this limitation, DAML+OIL and OWL 1 ontology languages are used to formally specify context data semantics and share it among heterogeneous entities. An approach based exclusively on ontology reasoning [Chen and Nugent 2009] uses ontologies to represent activities as well as each data source that can be used to recognize them, from sensors to actors. Coarse-grained activities are recognized by ontological reasoning based on the available data and refined as new information becomes available. However, OWL 1 operators have not enough expressiveness to define complex relationships, and tight integrations of OWL with expressive rule-based languages, such as SWRL (Semantic Web Rule Language), lead to undecidability [Riboni and Bettini 2011b]. OWL 2 tackles these problems by

allowing rule-based activity definitions with ontological axioms and preserving decidability, with the added advantage of providing unique semantics [Riboni and Bettini 2011b].

In Ejigu et al. [2007], upper- and lower-level ontologies are presented for a modeling context. The lower subclassification indicates domain-dependant views of context (hospital, home, car, campus), while generic context entities include, for example, person, device, network, physical environment, activity, location, and service. The RDF/OWL reification principle² is used to represent additional context attributes to the basic context triple. By means of the separation of static and dynamic contexts, context semantics, and (ontology-based and user-defined) rules, they focus on limited resource devices. Their query engine platform, however, relies on RDQL, an RDF Data Query Language prior to the standard SPARQL. In general, a wide expressiveness model is required to deal with the description of all possible features of a user and the functionalities provided by devices and systems [Mocholí et al. 2010]. For instance, in Mocholí et al. [2010], an ontology models ambient assisted living (AAL) services for the elderly in a domotic domain. Their particular aim is to facilitate the validation of accessibility (i.e., disability constraints) for users. A more in-depth analysis of these methods is provided in Sections 3 and 4.

Other models, based on fuzzy logic [Acampora and Loia 2005], have been also developed. For instance, in Hagrais et al. [2004], embedded agents are connected to sensors and effectors and use fuzzy logic-based *incremental synchronous learning* to define particularized rules instead of seeking to extract generalized rules. Based on rule bases built from previous (e.g., user or environment) occupiers, the learning time is minimized, since the system starts from a similar rule base to later refine the rules. Another approach in the same way uses agents to discretely control the Essex intelligent dormitory iDorm [Doctor et al. 2005] after an adaptive learning of human behavior. Agents use appliance parameters as input to a fuzzy logic controller acting over effectors. Through a fuzzy markup language (FML), a detailed structure of fuzzy control can be defined independently of its legacy representation, allowing agents to capture user habits and to apply an adaptive rule-based strategy.

Hybrid approaches (i.e., systems that combine data-driven and knowledge-driven approaches for activity recognition) are found in works such as evidential network-based activity inference [Hong et al. 2009] or COSAR [Riboni and Bettini 2011a]. Ontological reasoning with OWL 2 is used to recognize complex activities based on elementary sensor data and simple activities recognized through data-driven methods (in this case, statistical reasoning) [Riboni and Bettini 2011b]. COSAR is used together with the PalSPOT ontology within the context aggregation middleware CARE [Agostini et al. 2009]. The COSAR system retrieves information on simple human activities using hybrid ontological/statistical reasoners. They show how the recognition rate improves considerably as well as how the error rate is reduced by 45.43% with respect to the uniquely statistical technique. COSAR avoids misclassifications between activities characterized by similar body movements but different contexts. One disadvantage of this approach, however, is that it uses location association to infer activities and, therefore, this can be a problem when recognizing fine-grained activities that occur in a unique, confined, or small-size space.

Another example of a hybrid approach combines ontology-based context reasoning with computer vision research and integrates a scene tracking system with an ontological layer to avoid limitations that make classical object tracking procedures fail in complex scenarios [Gómez-Romero et al. 2011]. Abductive and deductive reasoning is

²A reified RDF data contains each original statement as a resource and the other additional statements made about it.

used to build different abstraction levels. The scene interpretation serves to generate feedback to the quantitative tracking procedure. In this way, DL reasoning reduces the complexity in the implementations for scene interpretation. One of the challenges, however, is not only the definition of suitable ontologies but also the detailed creation of abduction rules.

2.3. Evaluation and Open Challenges

Finally, to conclude the state of the art on human behavior monitoring and recognition, we create a taxonomy to summarize the issues considered in this section. Taking into account the approaches studied, the taxonomy in Figures 3 and 4 establishes the assessment criteria to evaluate human activity modeling. Each criterion and its correspondent key aspects are specified, together with a list of publications that tackle those criteria and key aspects. These publications are examples selected from previous sections as the most representative ones in their respective category.

This classification allows us to summarize the existing methods in the literature for human activity inference. In addition, it enables the classification of a human activity recognition system in a multimodal and taxonomic way, according to the methodology, special needs to model, and application domain considered. Next, we provide an evaluation of the mentioned techniques. We mainly focus on data-driven approaches, since knowledge-based ones, and more specifically ontologies, are analyzed in the following sections.

The first criterion, *learning procedure*, distinguishes among data-driven, knowledge-based, and hybrid approaches. Regarding data-driven methods, we notice that most of the proposals are based on supervised learning. On the other hand, ontological models are the most frequent methods used in knowledge-based techniques. The main difference between them, as stated in previous sections, is the inclusion of context awareness tools to include semantics in knowledge-based approaches. However, data-driven cases have provided very promising results for the inference and recognition of human activities in the environments where they have been tested. Unsupervised data-driven methods should be highlighted here, because their nature about not needing any previous training makes them suitable for the adaptability of systems to other environments, and they may contribute to the invisibility of the ubiquitous system to the user and to reduce the interaction between humans and the environment. However, in our opinion, these approaches are still in their first stages and much more work could be done to achieve this goal. In addition, self-adaptive supervised techniques are also under development and could provide suitable solutions for this problem.

The second criterion, *technique*, looks at the specific algorithm or model used for the recognition phase. We distinguish graphical models as those that provide a graphical representation of a behavior. This is an important feature to achieve modularity and ensure a consistent way to build user interfaces for data accessibility [Stengel 2003]. However, the complexity of these models for a nonexpert person makes the provision of this system feature difficult. Moreover, the model should be accurate enough to minimize the behavior detection failures, and nongraphical methods could provide better performance in some applications. In Ros et al. [2013], the authors suggest the separation of the graphical representation of the human activity from the learning procedure with this aim, although their approach is difficult to extend to all possible human behaviors.

Another criterion is the support to model *social activities* or human interaction, where multiuser settings or shared activity features are considered. Here, the current sensor technologies have the challenge to determine which user is performing each detected action, to model interleaved, social, and interchangeable behaviors. The approaches

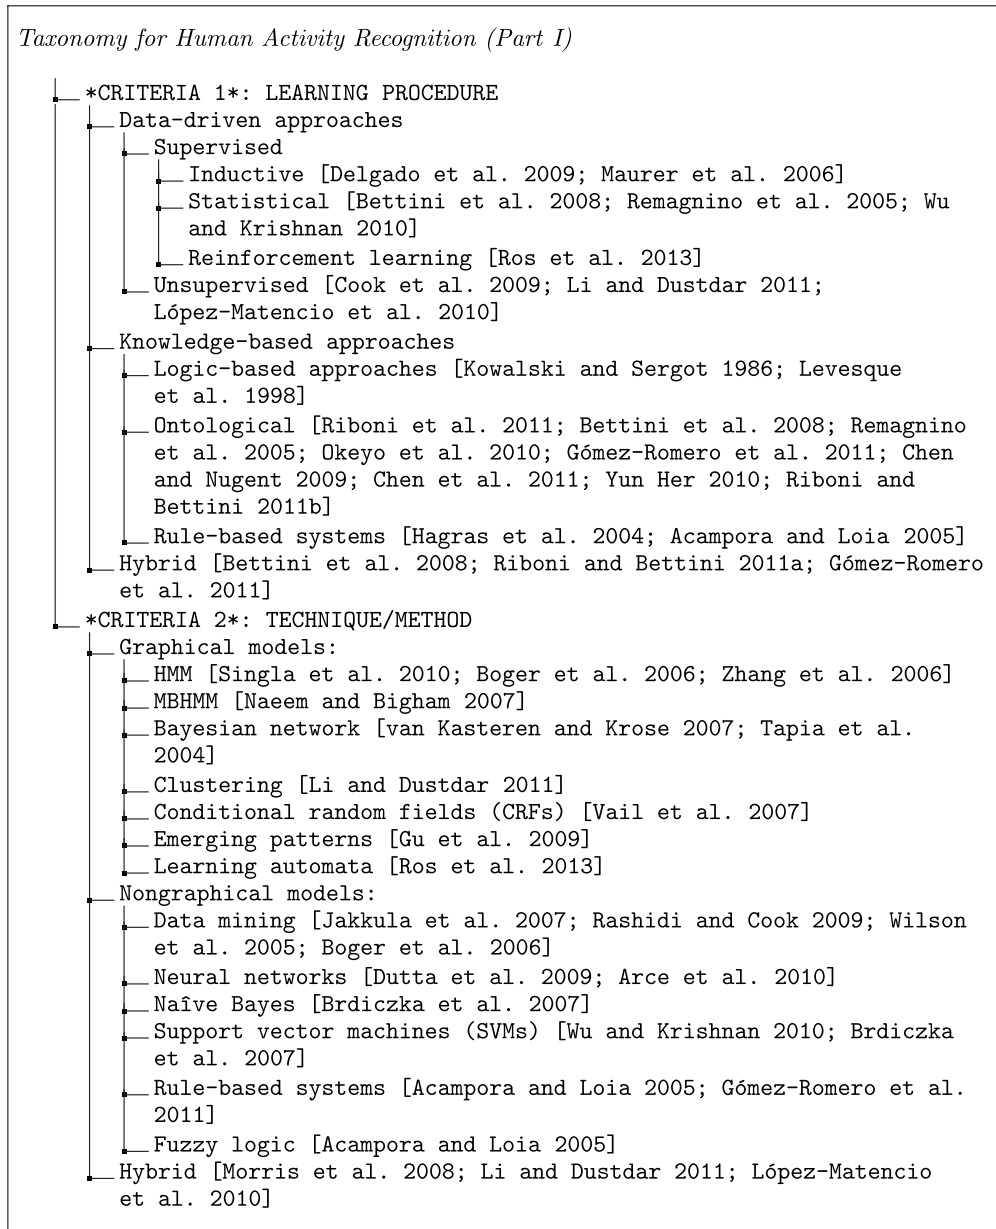


Fig. 3. Taxonomy for human activity recognition (Part I).

provided with this aim are promising, but the underlying sensor architecture of the environment is a key aspect for these techniques to succeed. In this way, the fourth criterion, *sensor infrastructure*, classifies the approaches according to the sensors used (i.e. video, passive, or wearable sensors). Recently, researchers have made a great effort to develop nonintrusive and accurate video sensors to overcome the previous limitations. However, the complexity and cost of these solutions are problematic in both hardware and activity modeling fields.

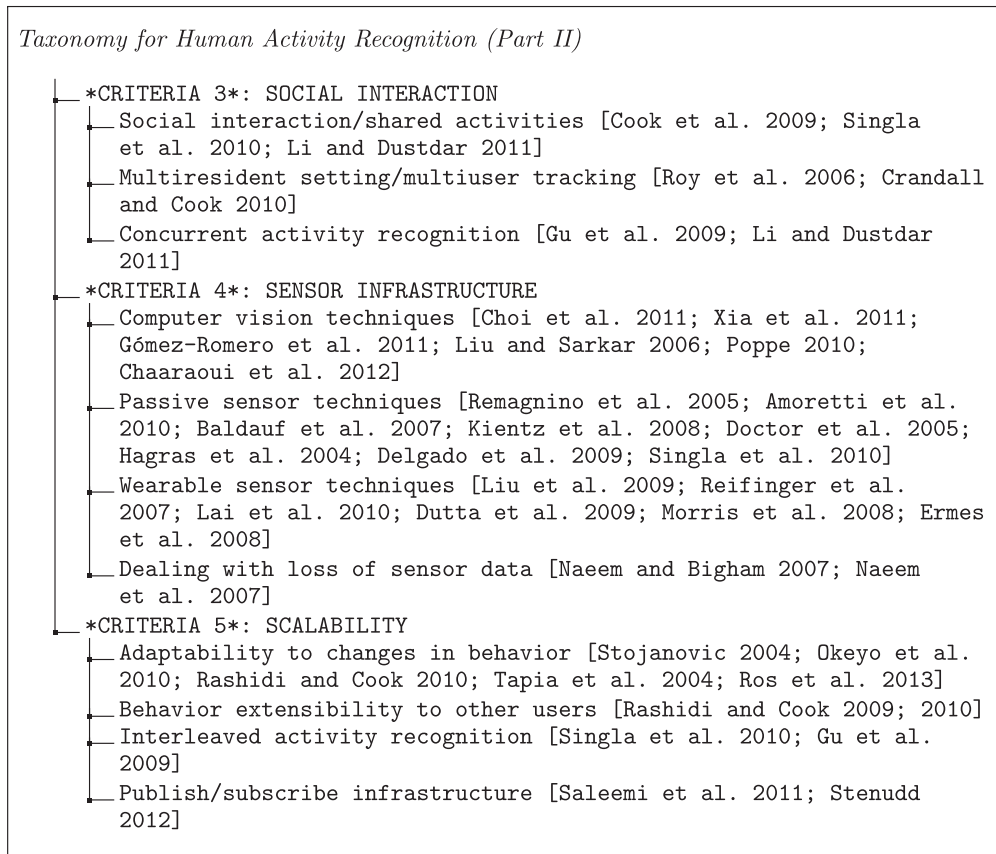


Fig. 4. Taxonomy for human activity recognition (Part II).

At last, *scalability* is a criterion that considers adaptation to changes in a routine, the accommodation of a given behavior to other users, and the scalability of the learned activities to other environments. In our opinion, this issue is essential to make commercial applications out of ubiquitous spaces. However, the goal of the environment, the sensor technologies required for each application, and the best inference and recognition model in each case are limitations that are difficult to address. The integration of semantics within these systems as a new abstraction layer could help overcome this challenge [Riboni and Bettini 2011b; Gómez-Romero et al. 2011].

3. ONTOLOGIES FOR HUMAN BEHAVIOR RECOGNITION

Data-driven approaches for activity recognition suffer from *ad hoc* static models, data scarcity, and scalability [Chen and Nugent 2009]. Semantic models can fulfill the needs of context-aware personalized activity-based environments where multimodal sensor technologies are constantly being introduced. Simple data systems can be modeled through key-value and markup models such as CC/PP [Klyne et al. 2004], while more complex domains require more sophisticated formalisms such as object-role-based models, spatial models of context, or ontologies (see Section 1.2). In general, expressiveness requirements in human behavior and environment representation include the ability to represent hierarchical structures, complex relationships among context instances,

and complex definitions based on simpler ones, usually using restrictions that may be spatial or temporal. Ontologies have shown, in the literature, to be one of the most promising tools to achieve these goals.

An ontology is a “formal specification of a shared conceptualization” [Borst 1997]. One of the main advantages of using ontologies is their way to represent and share knowledge by using common vocabulary. As providers of a format for exchanging knowledge, they promote interoperability, knowledge reuse, and information integration with automatic validation. They separate declarative and procedural knowledge, making the modularity of the knowledge base (KB) [Bobillo and Straccia 2011] easier. Ontologies allow information to become not only human but also machine readable and processable.

To be more specific, let us show an example of an ontological activity modeling using the expressiveness of OWL 2 language. The activity *WalpurgisParty* as a *friendly meeting in which all the participants are wearing a white hat* can be written as an OWL 2 axiom as follows:

Example 1: $WalpurgisParty \sqsubseteq FriendlyGathering \sqcap \forall hasActor.(Person \sqcap \exists isWearing.WhiteHat.)$

$FriendlyGathering \sqsubseteq Activity \sqcap \geq 2hasActor.Friend.,$

where a *FriendlyGathering* is an activity having at least two actors who are friends.

Not only the power of representation is key in ontologies. Reasoning capability is another important requirement for a knowledge-based activity recognition system. It is used to derive implicit information from explicit context data. For instance, the user’s current activity can be inferred based on his or her current location, posture, used objects, and surrounding people. Reasoning can also automatically detect inconsistencies in the KB [Riboni and Bettini 2011b]. Logical reasoning consists of deduction, abduction, and subsumption to extract a minimal set of covering models of interpretation from the activity model KB based on a set of observed actions, which could explain the observations [Chen and Nugent 2009]. For example, the simple statement “Two people working in the same project and institution are colleagues” may be formulated as follows, to infer which people are colleagues in the system:

Example 2: $Person(?x) \wedge Person(?y) \wedge worksInProject(?x, ?p) \wedge worksInProject(?y, ?p) \wedge worksForInstitution(?x, ?i) \wedge worksForInstitution(?y, ?i) \rightarrow isColleagueWith(?x, ?y).$

Description logics (DL) are the most used languages to model formal ontologies. DL reasoning supports tasks such as subsumption, satisfiability, equivalence and disjointness of classes, consistency, classification, instance retrieval, and realization [Riboni and Bettini 2011b]. DL reasoning can support incremental progressive activity recognition and assistance as the activity unfolds. Ontology-based activity recognition provides a number of advantages [Chen and Biswas 2009]. We can make further (fine-grained) inferences and it is easier to draw decisions to disclose other higher-level behaviors. It provides state-based modeling and a more robust reasoning. Since sometimes a mathematical description of a behavior (e.g., *morning routine*) cannot be trivially provided, ontology-based reasoning allows extra pieces of data to be used for behavior disclosure. Through the detection of low-level events reflecting the state of each individual entity, ontology-based reasoning can assert when a task or entity is different or the same as another one. However, to achieve this, the domain-specific knowledge needs to be unambiguously defined.

Other benefits are the ability to discriminate importance and urgency of activities through semantic descriptions [Chen and Biswas 2009], support for coarse-grained and fine-grained activity assistance, and the possibility for data fusion and semantic

reasoning [Chen and Nugent 2009], including activity learning, activity recognition, and activity assistance. When the conceptualization of activities and their interrelationships on ontological activity models encodes rich domain knowledge and heuristics in a machine-understandable way, a higher degree of automation is enabled for knowledge-based intelligent processing [Chen and Nugent 2009]. In addition, as new context sources are constantly being introduced in ubiquitous environments, data-driven approaches (e.g., supervised learning) require to retrain the complete model again, before being able to adapt to changes in the modeled activity. Furthermore, to rebuild the updated model, data collection with the new context data sources is required. However, knowledge-based methods allow previous activity recognition models to be used; updating only the affected context rules is enough to recognize the adapted activity.

Despite their advantages, a major limitation is the lack of support for imperfect information, since it is not inherent to OWL 2. Previous experience when defining complex activities has also shown some limitations (e.g., on the tree model property [Riboni and Bettini 2011b]). Adhering to this OWL 2 property guarantees decidability on reasoning problems but also limits the expressiveness, requiring every predicate (in an object property) to contain a quantified variable.³ On the other hand, most rule-based languages do not impose such forced restrictions [Riboni and Bettini 2011b].

3.1. Ontologies for Human Activity Representation

There is a broad variety of ontologies and vocabularies to model context in smart environments. Users are the central part, as well as what happens in their surroundings. The following ontologies show general user-centered approaches to model human activities:

- The *CoBrA-Ont* [Chen et al. 2003] ontology is an extension from SOUPA (Standard Ontologies for Ubiquitous and Pervasive Applications). It defines people, places, and activities. It was designed to reason within the CoBrA (Context Broker Architecture) infrastructure and defines key ontology categories such as action, agent, time (instant and interval), space, device, and so forth. One application scenario is the eBiquity group meeting ontology that models video presentations, recorded discussions, and other media material from meetings and their coordination by agents and actors. CoBrA specially takes care of places, distinguishing between atomic and compound places, depending on their containment property, defined as a spatial capability of subsuming other physical locations. It also distinguishes (person and software) agents, with their respective home pages, email addresses, and so forth. Each agent has his or her respective role (e.g., speaker or audience role, activity context, and location). CoBrA integrates a privacy policy language for user privacy protection that extends the Rei policy language.⁴ User privacy is considered by restricting the sharing of contextual information acquired by hidden sensors or agents.
- The *CoDAMoS* [Preuveneers et al. 2004] ontology defines four main core entities: user, environment, platform, and service. The aim behind this ontology design was to provide application adaptation, automatic code generation, code mobility, and generation of device-specific user interfaces. Resources are especially modeled (memory, network, power, storage resources), as well as service profiles, groundings, and

³For example, “an internal meeting is a meeting in which all actors are colleagues among themselves” is impossible to express in OWL 2 without giving up decidability. To solve this, restricting the activity definition to a specific domain is needed, for example, “an internal meeting of company X is a meeting in which all actors are employees of X” [Riboni and Bettini 2011b].

⁴Rei defines a set of ontology concepts for modeling rights, prohibitions, obligations, and dispensations in the domain of security.

different kinds of software (middleware, OS, rendering engine or virtual machine). Two levels of granularity offered are tasks and activities. Users can have a mood or be located in absolute or relative locations with different environmental variables (Figure 2).

- The *Delivery Context Ontology* [Cantera-Fonseca and Lewis 2009] (W3C) provides a formal model of environment characteristics in which different devices interact with concrete services. This ontology includes device characteristics, the software used to access the service, and the network (and network bearer) providing the connection, among others. Other entities modeled in this ontology are environment, hardware (battery, memory. . .), tactile input or text input types, cameras, aspect ratio, software (web browsers, script language, page markup restrictions), character set, Bluetooth profiles, location, unit conversions, and physical quantities (measures from Coulomb to inches).
- SOUPA* ontology [Chen et al. 2005] (Standard Ontology for Ubiquitous and Pervasive Applications) is divided into two main blocks called *SOUPA-Core* and *SOUPA-Extensions*. They are used in the CoBra architecture. *SOUPA-Core* defines general concepts to appear in different scenarios, for example, person, agent, policy (right, prohibition, obligation, dispensation, each of them with an associated actor and action), actions (preconditions and effects), events, geo-spatiality, space (locations' longitude, latitude, and altitude), time, and MoGATU BDI ontology (belief, desire and intention, goals, plans for agents). A policy in *SOUPA* is a set of rules defined by a policy creator, which is to be enforced by some policy enforcer. *SOUPA-Extensions* support particular concepts in narrower domains (e.g., home, office, entertainment). Extension ontologies demonstrate how to expand *SOUPA-Core* ontologies to define a set of vocabularies that support peer-to-peer data management in pervasive computing environments. Some examples of these ontologies consider instant time and intervals, spaces ((subsumed) spatial regions, geopolitical entities to which policies apply), a region connection calculus ontology, meeting ontology with event (re)schedules and cancellations, and user contact preferences. The *EasyMeeting* infrastructure facilitates typical user activities in meetings, such as setting up presentations, controlling services via speech, and adjusting lighting (light action control ontology) and background music, based on the state of the meeting. They also offer the priority ontology, which is established for a set of desires and intentions of an agent, and the ontology to describe conditional beliefs.
- The *mIO!* ontology [Villalon et al. 2010] is a network ontology, developed using the NeOn methodology, that represents user context to configure, discover, execute, and enhance different services in which the user may be interested. The NeOn methodology basically considers the reuse, merge, matching, and re-engineering of ontological resources. Eleven modular ontologies define the *mIO!* core: user (groups, organizations, their employment status, skills, mobility pattern, and online identities; reuses *FOAF* ontology as a whole), role (knowledge about profiles, preferences; reuses ontologies such as *Reco* for user preferences), environment (environmental conditions; reused from CoDAMoS), location (spatial entities and area and distance units from *SOUPA*, location coordinates, buildings, countries), time (temporal units and entities, instants, intervals, reuses W3C Time ontology), service (from business to *mIO!* services, with digital signature, input and output parameters, its components and functionalities), provider (wide categorization of simple and aggregated service providers, from business to software services), device (taxonomy categorization and “componency” pattern, charging mode, compatibility with standards, e.g., glucose meter, pulse oximeter, anemometer, etc.), interface (types, I/O modalities and characteristics), source (aggregated or not: user, device, service, etc.), and network (communication networks, network topologies, operators and administrators,

accessibility, price, coverage, etc.; network types and modes reused from *Delivery Context* ontology).

- The human activity recognition ontology [Riboni et al. 2011] from the *PalSPOT* project models individual and social activities. The types of interaction are modeled as acknowledgment, asking for opinion, comment, negative/positive opinion, proposal, or request of information. Activity granularity is slightly shown (basically, only one activity level); however, an extensive taxonomy is available for personal, physical, and professional activities, traveling activities, postures, and activities using different kinds of devices or artifacts. An interval-based representation of activities models the overlapping of these in time. Other entities are indoor (corridor, arcade, elevator, etc.) and outdoor (promenade, sidewalk, rails, etc.) communication routes. Symbolic locations (indoor, outdoor, pedestrian or not) and time granularity are provided. The PalSPOT ontology is used within the context aggregation middleware CARE [Agostini et al. 2009], which maps context data to ontological classes and properties and interacts with the COSAR system [Riboni and Bettini 2011a], which retrieves information about simple human activities using hybrid ontological/statistical reasoners.
- CONON* (CONtext ONtology) [Wang et al. 2004] defines general concepts in an upper ontology such as location (indoor, outdoor, with different environmental features and variables, weather conditions), activity, person, or computational entity (such as devices with a status—phone, TV, DVD player, etc.). CONON allows extensions in a hierarchical way by adding domain-specific concepts, where different intelligent environments are modeled (home, office, vehicle, etc.). Activities (with start and end time) are divided into deduced (dinner, movie) and scheduled (party, anniversary) activities. The status of indoor spaces entities (e.g., curtain, door, window) is also represented. Some domain-specific ontologies are the home domain (e.g., sleeping, showering, cooking, watching TV, having dinner) and office domain ontologies. Reasoning categories employed are DL ontology reasoning and user-defined reasoning using first-order logic (through customized rules).
- The Pervasive Information Visualization Ontology (PiVOn) is a formal context model [Hervás et al. 2010] composed by four independent ontologies (users, environment, devices, and services), which describes intelligent environments. Some properties of the main elements in the user ontology are location, identity, activity, and time. The context is analyzed from the perspective of the 5 Ws Theory⁵ to design context-aware systems. The result can be summarized in a two-dimensional taxonomy of context elements: the first one defined by the four main categories of the context (user, environment, device, and service) and the second by the 5 Ws. Events in the ontology have reminders, schedule, and are part of a user agenda. They involve contacts (FOAF) from the user, who can be in user situations, and possibly accompanied by some other user(s). Users perform tasks (which can have subtasks) that have a goal and use some services. Tasks have types, roles, significance levels, progress, time and space issues, and attention levels. User situations also play a role and belong to the user availability, a state of mind, or a task. The device ontology determines the types of devices (autonomous, dependent, sensor, actuator), the services provided by a device, the owner of the device, status, location, its hardware and software profiles, communication hardware profile, its use and compatibilities, and so forth. The environment ontology represents the colocation of objects, *near* (*inFrontOf*, *on*, *under*, *behind*), *includedIn*, *associated*, and so forth. Spaces are modeled with the area they are located in (building, wing, floor) and its purpose, structure, and capacity. The visualization service ontology includes, for each service, an associated visualization

⁵A journalism principle regarded as basic in information gathering (What, Who, Where, When, Why).

service (displayed by devices), which contains content. The content has a visual form and is transformed into data. The visual form has different types of scalability parameters (filtering, pagination, complexity, latency). A prototype scenario is modeled on an academic conference.

- The Situation Ontology [Yau and Liu 2006] is divided into situation and context layers. Situation is defined as a set of contexts in the application over a period of time that affects future system behavior. A context is any instantaneous, detectable, and relevant property of the environment, system, or users, such as location or available bandwidth. The ontology considers atomic and composite situations. The latter are composed by temporal, conjunction, disjunction, and negation situations. Composite situations can integrate atomic situations through Boolean context operators that act over context value domains. Similarly, a temporal situation has temporal operators over time periods. Regarding context, the Situation Ontology is classified into device, user, and environment context. An entity can satisfy a situation by having related context data, with a certain context value (e.g., float temperature value), within a context domain value (e.g., available-memory context). Context value domains are provided with data context operations. An example of a smart conference scenario could specify the situation *ReadyForMeeting* as the conjunction of two atomic situations: *InConferenceRoom* and *LightOn*, where the *InConferenceRoom* situation *location-context* value is the same as *crLocation*, and the *LightOn* situation is represented as “the *lightContext* value is true.”

3.2. Ontologies for Context and Environment Representation

We described, in the previous section, a set of ontologies that focus on the user. In this section, we describe a series of very helpful concrete domain ontologies to describe the context and the environment where human activities occur:

- Location: *PlaceTime.com*⁶ (URIs for places and times) contains instants and intervals in the Gregorian calendar and points in the WGS 84 datum, utilizing the RDFIG Geo vocabulary. Other vocabularies useful in object and human location tracking are *WGS84 Geo Positioning*⁷ or *GeoNames*.⁸
- Time ontology,⁹ developed by the W3C Semantic Web Best Practices and Deployment Working Group (SWBPD), describes temporal content and properties of web pages and web services. It also provides topological relations among instants and intervals, durations, datetimes, and world time zones.
- User profile and preferences: People can be modeled with the *FOAF* ontology.¹⁰ User Agent Profile (UAProf)¹¹ specification relates capabilities and preference information for wireless devices. The *CC/PP* (Composite Capabilities/Preference Profile) model [Klyne et al. 2004] is a W3C initiative that suggests an infrastructure (and vocabulary) to describe device capabilities and user preferences. The representation model can guide the adaptation of the content presented to the device, considering software terminals, hardware terminals, applications such as a browser, data types, protocols, and specification conformance of products (documents, producers, and consumers on the web). The hierarchical structure of components is divided into three areas: hardware, software, and application.¹² Also, a W3C Delivery Context Ontology and

⁶<http://placetime.com>.

⁷<http://schemapedia.com/schemas/geo>.

⁸<http://www.geonames.org/>.

⁹<http://www.w3.org/TR/owl-time/>.

¹⁰Friend of a Friend Vocabulary: <http://xmlns.com/foaf/spec/>.

¹¹<http://www.w3.org/wiki/UAProfIndex>.

¹²UAProf and CC/PP are encoded in RDF/S. <http://www.w3.org/TR/CCPP-struct-vocab/>.

a glossary of terms for device independence exists (with discontinued maintenance). The ontologies in the semantic desktop *Gnowsis* project focus on use cases such as tagging a file, email, or a website.¹³ The User Action Ontology in *Nepomuk* (Social Semantic Desktop) describes desktop events and their Calendar Ontology (NCAL) adapts the W3C ICALTZD ontology.¹⁴

- Online behavior, such as in social communities: The *SIOC* (Semantically Inter-linked Online Communities)¹⁵ ontology describes information from online communities (e.g., message boards, wikis, weblogs, etc.) on the semantic web. A natural extension to SIOC for online behavior is the *Open University Behavior Ontology* (OUBO), which allows user behavior to be captured over time and facilitates role inference in a way that a user's role in a given context can be derived through semantic rules [Rowe et al. 2013].
- Content extraction: The Image Annotation W3C¹⁶ ontology for semantic image annotation and retrieval can be used for deep multimedia analysis (e.g., image-based context recognition). The *Nepomuk Multimedia Ontology* (NMM) also defines metadata for multimedia files, and the *EXIF* ontology describes digital camera images and image management software.¹⁷
- The ASC (Aspect-Scale-Context) model [Strang et al. 2003] includes concepts such as aspects, scales, and context information, each aggregating one or more scales. Although useful to describe measurement units, it cannot describe more abstract context information, like user activities. The DAML-based Context Ontology Language (*CoOL*), derived from the model, can be used to enable context awareness and contextual interoperability.

4. DOMAIN-BASED CLASSIFICATION FOR ONTOLOGY EVALUATION

Now that the main ontologies for human behavior representation have been described in Section 3, we provide an evaluation of all relevant design aspects to select and develop new tools to improve the construction of competitive human activity recognition systems. There exist a large number of different scenarios deployed within heterogeneous ubiquitous spaces, as well as different technologies and approaches. Thus, the diversity of systems makes reaching a consensus in evaluation tools quite difficult. A lack of evaluation tools prevents the use of a well-formed hierarchical system classification, since there is no widely accepted model to be reused in different domain applications [Villalon et al. 2010].

In Baumgartner and Retschitzegger [2006], an evaluation framework for upper ontologies on situation and context awareness is provided. They analyze four ontologies: SAW (situation awareness ontology within SAWA project [Matheus et al. 2005]), situation ontology [Yau and Liu 2006], SOUPA [Chen et al. 2005], and CONON [Wang et al. 2004]. The evaluation framework focuses on top-level concepts, SAW concepts, and modeling characteristics of upper ontologies, such as *universality* or *articulation*. Top-level concepts evaluated are object, attribute, relation and role, event, and situation, while SAW-specific concepts consider space and time, thematic roles, situation types, and situations as objects. Even though none of the ontologies fulfills all or at least most of the overall criteria, we can affirm that SAWA satisfies the modeling of most concepts, followed by situation ontology, SOUPA, and CONON.

¹³<http://gnowsis.opendfki.de/>.

¹⁴<http://oscaf.sourceforge.net/ncal.html>.

¹⁵<http://rdfs.org/sioc/spec>.

¹⁶<http://www.w3.org/2005/Incubator/mmsem/XGR-image-annotation/>.

¹⁷*Shared-Desktop-Ontologies* (SDO): <http://oscaf.sourceforge.net/>.

Other examinations evaluate different ontologies focusing on licensing (open or free to use), maturity (showing evidence of use), structure (modularity, design), granularity of time and space, vagueness, and security [Semy et al. 2004], or availability, existence of management tools, expressiveness, and so forth [Feliu and Cabré 2002; Baldauf et al. 2007].

The taxonomy proposed in Abdualrazak et al. [2010] to evaluate a wide variety of pervasive computing systems presents some evaluation criteria that consists of the concepts *Architecture*, *Application Purpose*, *Autonomy*, *Integration*, *Interaction*, *Intelligence*, and *Service Availability*. The infrastructure defines the architecture and design. The application purpose describes, among other features, the quality of context, its reliability, fault tolerance, security, privacy, and effectiveness. Autonomicity considers how a pervasive computer system is initialized, how it evolves, how it automatically tackles failures adjusting to users, how it integrates new resources, and how it fends off attacks. The interaction criterion, on the other hand, identifies human-to-machine and machine-to-machine presentation capabilities. Intelligence criteria measure the proactiveness of the AmI technique as well as the quality of context and adaptability to changes. At last, Service Availability categorizes the pervasive service based on the ability to be “anywhere anytime” (e.g., discovery, deployment, mobility, etc.).

4.1. Evaluation of Ontologies for Human Activity Recognition

Despite the fact that general ontology evaluation methods exist, none of them is focused on the field of human activity recognition. In this work, we focus on specific features related with human behavior modeling in accordance with their environment. Our starting point is the previous study in Villalon et al. [2010], where different ontologies are compared with respect to their support for modeling different domains. Once we have made an in-depth analysis of the most relevant ontologies in Section 3, we have detected essential subdomains to properly model everyday human activities at different levels. These evaluation criteria are available in Tables I and II. When a subdomain is modeled in any way in the ontology, X is marked. However, if that domain is highlighted or specially treated in the ontology, this is marked with an increasing number of X s (XX , XXX).

When considering human activity representation, basic and obvious variables such as *user*, *role*, *location*, *environment*, *time*, *context sources*, and proper *behavior granularity* levels are needed. Indoor and outdoor spaces need to be taken into account within the environment. The *role* of a user determines one facet of his or her objectives, which means he or she can have different roles at different times of the day. This is another reason to consider *time* as an essential entity; for example, each event can have a unique timestamp associated to it. Also, the origin of the *context information source* can be crucial to determine its origin or further inferences. Since activity recognition is incremental, some kind of basic atomic *actions*, as well as more generic *activities* and/or *behaviors*, should be specified. It is out of atomic events, happening in a certain order, that a given activity can be specified and, therefore, recognized. It is also in this way that higher-level context can be inferred out of single sensor events, and so on.

Regarding the supply of assistance, as well as the interaction with the environment, taxonomies on *services*, *devices*, *interfaces*, *networks*, and *providers* must be taken into account, not only for modeling how to provide users help or support after a human activity has been recognized, but also for proper service grounding and context-aware user adaptation. In this domain, the *mIO!* ontology is the best candidate to completely support *service*, *network*, *device*, and *interface* supply.

Another element required to model human activity is *social interaction*, including *message exchange* among people, but also messages from the system to the user, and vice versa. Also, the *availability* of the ontology is important, as is its maintenance,

Table I. Subdomains Addressed by Available Context Ontologies (Ontology Set I)

Ontology/ Subdomain	CC/PP [Klyne et al. 2004]	CoBrA-Ont [Chen et al. 2003]	CoDAMoS [Preuveneers et al. 2004]	Delivery Context [Cantera- Fonseca and Lewis 2009]	SOUPA [Chen et al. 2005]
Device	X	X	X	X	
Environment		X	X	X	
Interface				XX (I. Restrictions)	
Location		XXX	X	X	X
Network	X			XX	
Provider					
Role		X	X		
Service			X	X	
Context Source					
Time		X	X		XX
User	X	X	X		X
Imprecision/ Uncertainty Management					
Message					
Behavior Granularity		Action	X (Task, Activity)		
Behavior Model					
Social Interaction					
Implementation Available	X	X	X		X
Other Specific Domains Modeled	Capabilities, profiles	Privacy policy, agents, compound places, Easy- Meeting (eBiquity Group Meeting Ontology)	Platform (Hardware, OS), hardware resources and software types, service (profile, model, and grounding)	Hardware and software (APIs, Bluetooth), physical quantities, unit conversions	Policy (also in privacy), Easy Meeting

scalability, and so forth. The more available and visible an ontology is, the more testing and usage will follow, encouraging also its evolution and adaptation to more concrete and real-life domains. In our evaluation, extra features modeled in the analyzed ontologies are also highlighted (e.g., OS features, hardware/software platforms, and other more specific domains).

We can observe that features such as modeling *uncertainty*, *imprecision*, and *vagueness*, typical of everyday life (and everyday language), are missing in all human context ontologies in Tables I and II. Formalization of messages as a way of interaction, as well as getting feedback from the user, need to be modeled, as they are samples of social networking or common behavior. In the next section, we will discuss the studied ontologies' response to each mentioned factor when modeling human behavior. Common specializations, as well as flaws of the existent ontologies, are to be discussed.

Table II. Subdomains Addressed by Available Context Ontologies (Ontology Set II)

Ontology/ Subdomain	mIO! [Villalon et al. 2010]	PalSPOT [Riboni et al. 2011]	CONON [Wang et al. 2004]	PiVOn [Hervás et al. 2010]	Situation Ontology [Yau and Liu 2006]
Device	X	X	X	XX	X
Environment	X	X	X	X	X
Interface	X				
Location	X	X	X	XX	X
Network	X		X		
Provider	X			X (device service provider)	
Role	X			X	
Service	X		X	XX	
Context Source	X	X		X (in device taxonomy)	X (context value domain)
Time	X	X	X	X	XX
User	X	X	X	X	X
Imprecision/ Uncertainty Management					
Message				X (schedule reminder)	
Behavior Granularity		X	X	X subtasks	X (Boolean & data context operator, atomic/ composite situation)
Behavior Model					
Social Interaction		X		X (companion)	
Implementa- tion Available	X	X			
Other Specific Domains Modeled	Games (Paddle) as service	Indoor/outdoor communica- tion routes, vehicles, traveling, device/ artifact usage, Snapshot ontology for ADLs	Home & office domain, scheduled/ deduced activities, furniture status	Spaces, visualiza- tion services, agenda, schedule, reminder, task progress	Context value domain, time period & operators, Boolean context operator, atomic/ composite situation, academic conference

4.2. Discussion

There are very heterogeneous methods for analyzing human activity. They specialize on a wide range of scenarios to be tracked, set of activities monitored, used methodology and algorithms, and further specific features such as interaction with other users or scalability of the method. These scenarios are usually surveillance, monitoring of public facilities (train stations, subways, airports), UAV¹⁸ surveillance, patient monitoring, or smart homes.

¹⁸Unmanned aerial vehicles.

The sources of context information are varied and the evaluation of activity recognition systems is complex because there is no standard taxonomy of human activities providing, for example, confusion matrices for each classifier's activities and their respective precision and recall. The way in which the methods execute the data collection and labeling processes is also relevant when it comes to the assessment of different approaches [Brush et al. 2010]. Additionally, it is not easy to model every context category precisely and generally enough. This is due to the reality that ontology-based activity recognition has also some drawbacks (e.g., it requires good knowledge engineering skills to model the domain). Besides, expressiveness limitations are found in OWL DL, mainly related to the lack of support for temporal reasoning. Directly with OWL, it is not possible to perform interval-based (i.e., overlapping) temporal reasoning, which is crucial for capturing complex temporal relations between activities. Even if ontology-based reasoning has a set of added advantages (discussed in Section 3 and Chen and Biswas [2009]), ontological reasoning can be computationally expensive [Riboni et al. 2011]. Thus, the design of an appropriate and complete ontology is crucial for ontology-based human activity modeling. As an example, we can mention that different data-driven approaches excel at different concrete aims (e.g., dealing with missing sensor readings). However, this works well when task models are small and manageable; otherwise, an ontology approach seems more efficient [Naeem and Bigham 2007]. This is just an example that shows that the size of the problem, specific domain, and concrete task are decisive issues to consider when selecting an ontology.

Some kind of different levels in action, activity, or behavior granularity can be seen only in CoDAMoS, PalSPOT, CONON, PiVON, and Situation Ontology. While ontologies like CoBrA-Ont only consider atomic actions, CoDAMoS includes a distinction between tasks and activities. PalSPOT also considers two levels of events with actions (e.g., moving arm) and activities (traveling, walking, bathing). However, activity granularity in the CONON ontology occurs only at a unique activity level, differentiating among scheduled and deduced activities. At last, PiVON includes a recursive subtask structure, while Situation Ontology allows a different approach of granularity based on Boolean and data context operators as well as atomic and composite situations. We could argue that PiVON's approach is perhaps one of the most flexible solutions to accommodate an infinite number of behavior granularity levels. However, having a hierarchical categorization of activities helps in modularizing and inferring different abstraction levels. This is due to the fact that, as we mentioned earlier, it is important to know the tasks happening but also the intention, behavior, or meaning associated to these events.

Looking at the human–computer interaction side, devices, interfaces, network, and services are represented in a large amount of ontologies. However, modeling *messages* in human interaction is only seen in a reduced form, with agenda schedule reminders, in PiVON. There are no clear ways of modeling the communication back from the system to the user and vice versa. That is, the kinds of interfaces that should be used in each moment and the time to communicate with the user are aspects that should also be context aware. Social interaction (i.e., human–human) is only modeled in PalSPOT and consists of a varied taxonomic distinction among acknowledgment, asking and giving opinion, comment, proposal, or request for information. For this matter, PiVON looks at the social interaction aspect by considering the companion of a user while executing a given task.

If we were to highlight what each analyzed ontology stands out for, we could say that mIO! stresses its modeling on device interaction, PalSPOT on the user activity in the environment and with others, CONON on activity planning and services, PiVON on location and device-based services, and Situation Ontology on temporal context operations. CC/PP enhances device and network capabilities, while CoBrA masters




Fig. 5. Overall ranking for the most representative human activity ontologies for the proposed evaluation criteria.

Table III. Ontology Pitfall Evaluation

Ontology	#Triples	#Classes	Pitfalls	Pitfall Rate
CoDAMoS	1291	106	143 <i>P8</i> — <i>P10</i> — 4 <i>P11</i> — 33 <i>P13</i> — <i>P22</i>	0.141
CoBrA	4144	88	5 <i>P4</i> — 186 <i>P8</i> — 22 <i>P11</i> — 3 <i>P12</i> — 67 <i>P13</i> — <i>P22</i> — 4 <i>P24</i> — 5 <i>Suggestions</i>	0.071
PalSPOT locont-2.0	5302	199	<i>P4</i> — <i>P5</i> — 251 <i>P8</i> — <i>P11</i> — 22 <i>P13</i> — 2 <i>P19</i> — <i>P22</i> — 10 <i>Suggestions</i>	0.055
SOUPA_policy	1304	30	12 <i>P8</i> — 3 <i>P11</i> — 7 <i>P13</i> — 1 <i>Warning</i>	0.025
Delivery-ContextAll	22573	134	<i>P4</i> — 37 <i>P11</i> — 97 <i>P13</i> — <i>P22</i> — <i>P24</i> — 4 <i>Sugg</i>	0.006
CCPPschema-20030226	19	134	Free of bad practice detectable by OOPS! Pitfall Scanner	0

locations. CoDAMoS brings out roles and (hardware and software) services; Delivery Context treats (hardware and software) interfaces and networks; and SOUPA, time, locations, and policies.

A summary of the human behavior ontology review and its respective coverage in subdomain modeling can be seen with an estimated evaluation score in Figure 5. The ranking is constructed in a straightforward way, assuming each evaluation criteria in Tables I and II as equally important. The histogram, based on a naive ranking scale, gives an idea of the number of (preselected) subdomains tackled by each context ontology. Similarly, the score shows a degree of specialization in a given subdomain. Eighteen categories were used for ontology evaluation (Tables I and II), and each ontology gets extra points (one per extra *X*) if a given domain is specially modeled or remarked. The overall grading is normalized between 0 and 10, according to how many subdomains are satisfied by the ontology and in what strength.

We are aware that a ranking number does not capture the many dimensions of an ontology, how it can be improved, or the problems it has [Vrandečić 2010]. Thus, ontology evaluation technologies rather meet their goals by pointing out if an ontology is bad, and in what way, instead of telling how good it is. From the structural point of view, the ontologies could be evaluated as regards their number of pitfalls (i.e., the number of features that could give problems in ontology-based reasoning). Table III shows the ontologies, their size in triples, and the recommended pitfalls cases to fix,

according to OOPS! (Ontology Pitfall Scanner) [Poveda-Villalón et al. 2012]. Due to spatial constraints, we refer the reader to the pitfall catalog¹⁹ for an accurate key pitfall description and suggested solution. Some of the most commonly identified pitfalls are about creating unconnected ontology elements (*P4*), missing annotations (*P8*), missing domain or range in properties (*P11*), using different naming criteria in the ontology (*P22*), or recursive definitions (*P24*). We define the Pitfall Rate evaluation parameter as:

$$\frac{\sum_{i=1}^n \#P_i}{\#T},$$

where $\#P_i$ is the number of pitfall cases occurring for pitfall type P_i and $\#T$ is the ontology size in number of triples. In this way, the pitfall rate symbolizes the average number of pitfalls per triple. A higher pitfall rate implies the appearance of a larger number of anomalies or errors in the ontology.

We are also aware that defining ontology quality is difficult, since it depends on different approaches. As a result, we can also appreciate that the semantic evaluation ranking in Figure 5 is independent of the structural evaluation in Table III. The majority of methods for ontology evaluation concentrate on structural evaluations, in a similar manner to what validators do. This is because semantic evaluation is subjective and application and domain dependent. Other methodologies that can help in choosing the right ontology include ONTOMETRIC [Lozano-Tello and Gómez-Pérez 2004] and an adaptation of the AHP method,²⁰ to help knowledge engineers choose the appropriate ontology for a given domain. Identified characteristics for the evaluation include tools, language, content, methodology, and costs [Lozano-Tello and Gómez-Pérez 2004]. Other evaluation methods can be found in Seyed [2012] or COAT²¹ (Cognitive Ontology AssessmentT), but most of the methods do not provide a versatile software interface such as OOPS'.

Following a methodology to guide ontology development has proven to be useful [Villalón et al. 2010], just as reusing knowledge resources and attending to good practices in the ontology development are. However, when searching for an existing ontology for a given domain, selecting and reusing context ontologies can become difficult due to the different purposes and requirements for which the ontologies are designed. For instance, in our domain, a hierarchical classification (of human actions and activities) for a granular behavior disclosure is required to incrementally infer new information from a collection of temporally evolving and atomic context data. However, specific requirements are not always, nor often, supported, and each application scenario will normally impose different requirements.

Elements such as location and time are essential when managing historical context data, in order to provide intelligent learning algorithms that can offer services after recognizing an activity. We believe that the possibility of associating machine-learning behavioral models to each behavior in the ontology can provide modular and proactive capabilities without depending on specific implementations of context-aware systems, but rather having a formulation in the ontological model itself. Some approaches close to this paradigm are hybrid systems such as *COSAR* context aggregation and reasoning middleware and its ontological/statistical hybrid reasoners. Their *Derivation of Possible Activities* algorithm, executed by an offline ontological reasoning module, takes an

¹⁹OOPS! Pitfall Scanner Catalog: <http://oeg-lia3.dia.fi.upm.es/oops/catalogue.jsp>.

²⁰Analytical Hierarchy Process, a measurement method based on preferential ordering.

²¹COAT (Cognitive Ontology AssessmentT) tool: <https://code.google.com/p/ontoeval>.

empty ABox and TBox as input to output a matrix of symbolic locations and their correspondent inferred human activities [Riboni and Bettini 2011a].

Future challenges in human activity ontologies need to further consider the modeling of imprecise and uncertain information for more accurate representation of everyday human tasks and human language. This is an aspect that no ontology tackles at the moment. Another issue to deal with is modeling social interaction, both virtual (for human-to-computer messages) and physical (for human-to-human messages). As we mentioned, having an adequate level of action granularity (e.g., actions, tasks, activities, behaviors, etc.) is crucial for a specialized and incremental discovery. Besides, a standardized common representation of universal entities, such as time and geographical indoor and outdoor locations, as well as environmental conditions would greatly help in these processes. At the same time, relations such as ownership of objects, rights, services, or privacy and service access are common issues of interest, usually modeled in different ways. Unfortunately, this causes the reuse and mapping of several heterogeneous ontologies to require a long time spent in curating and matching ontology concepts.

As a concluding remark, we can clearly point out that the integration of different methodologies (i.e. data-driven and knowledge-based ones) could help overcome current limitations in scenarios with several actors, providing semantics to social activities, user identification according to behavior semantics, and so forth. Current hybrid approaches such as Gómez-Romero et al. [2011] and Riboni and Bettini [2011a] have shown that these types of combinations can enhance the response of data-driven approaches as the environment complexity and the context awareness needs increase. In addition, knowledge-based approaches also could take advantage of features such as noise tolerance and uncertainty handling inherent to most used data-driven activity recognition models.

5. CONCLUSIONS AND FUTURE WORK

Current trends show that tracking and monitoring people is becoming an integral part of everyday life. Data-based approaches (HMM, Bayesian networks, decision trees, etc.) appear to stand out in contrast to the newly emergent knowledge-based techniques. The latter include, among others, information indexing and retrieval, hierarchical knowledge sources (taxonomies, ontologies, encyclopedias, dictionaries), representation languages, distributed knowledge, and logical or KB tools. In this article, focused on knowledge-based techniques, a set of ontologies representing human activity was described and analyzed from different perspectives.

The contribution of this survey consists in a study on available techniques for human behavior recognition. We have proposed an evaluation taxonomy for learning procedures, methods, models, and modeling capabilities. We also presented a set of upper ontologies designed to represent human activity, as well as domain ontologies that can serve the same aim in context-aware intelligent environments. A complete set of evaluation criteria was introduced to assess the current ontologies, having as the main focus the different subdomains required for human behavior representation, learning, and inference. The evaluation was performed by analyzing different useful domains and was concluded by giving an overall score to each ontology. Furthermore, as the semantic quality of the ontology ultimately depends on the specific domain to be modeled among other multiple aspects, we analyzed the structural problems, or *pitfalls*, found in each ontology. We can confirm that the broader an ontology is, the more situations will be possible to be modeled, in order to assist the users in their daily activities, and the less usable the ontology will be in order to achieve a particular goal [Villalon et al. 2010]. However, the more specific the ontology is, the fewer possibilities exist for reuse, but the more usable the ontology is.

As indicated earlier, the size of the problem, specific domain, and concrete task are key elements when selecting an ontology. However, what can be appreciated from this survey is that most of the works require a data-intensive-driven first approach to robustly identify the most basic-level actions or activities. Based on these sets of actions and activities, recognized first through precisely specified and robust models, further context assumptions can be integrated into the models to better handle the uncertainty inherent to the environment context. As in ubiquitous computing applications, it is not possible to specify all possible cases to recognize human behaviors; the implicit reasoning capabilities of knowledge-driven methods allow for more flexible and context-aware models (i.e., more knowledge can be expressed without explicit specification or knowledge redundancy [Yau and Liu 2006]). The latest research shows the benefits of introducing hybrid approaches to take advantage of each technique's best strengths. Combining ontology-based context reasoning with data-driven algorithms has shown to be a promising path to be explored. An example is combining ontological reasoning with computer vision research [Gómez-Romero et al. 2011]. Other works following these lines, with different (statistical, ontological) reasoning modules [Riboni and Bettini 2011a; Riboni et al. 2011], show that the usage of hybrid approaches with a preliminary phase on data-intensive methods can ease the way later, when inferring higher-level activities through knowledge-driven approaches. As time is not a feature inherently treated in knowledge-driven approaches such as logic-based systems, having hybrid methods with a first data-driven preprocessing stage appears to be the right direction to benefit from both data- and knowledge-driven computing paradigms. As ontological reasoning can be computationally expensive, this type of combination would achieve the best performance and efficiency from (time-dependent) data-driven methods and obtain the best adaptation for context awareness in each case.

When considering the real expressive power and usability of the reviewed ontologies, it is important to note that OWL 2 language is powerful for expressing knowledge, context information, and relations among entities. However, OWL 2 is insufficient to model context relations and rules with the form of cyclic relations [Riboni and Bettini 2011b] (e.g., relations such as `isColleagueWith` in the rule in Section 3, Example 2). Therefore, the ontologies discussed require an integration with a rule language (such as SWRL or SPIN²²) in order to express more complex and real-life context rules. The combination of DL with rule-based systems improves the reasoning capabilities. Rule-based languages enable definition of consistency rules, reducing ambiguity in the context information and thus maintaining and improving the information quality [Hervás et al. 2010]. For instance, SWRL (used, e.g., in PiVOn ontology [Hervás et al. 2010]) extends the semantics of OWL and defines antecedent-consequent rules and built-in (math, comparisons, string, and time) operators. Another example of rule-based inference support over OWL graphs is the Jena Semantic Web Toolkit,²³ used, for instance, in CONON ontology [Wang et al. 2004]. Another example is Flora-2 object-oriented knowledge base language,²⁴ the inference engine in CoBrA ontology. It can be concluded that the final expressivity of the ontology-based application will be a result of the combination of the elected ontology and its coupling with the chosen rule language, as an extension of OWL axioms, to express context rules.

A future challenge to be tackled is achieving adaptive processes of ontology evolution for adapting to natural changes in human behavior and/or environments. Adaptive activity modeling, such as the semiautomatic model in Chen et al. [2011], evolves from initial *seed* activity models through a continuous activity discovery, by using pattern

²²SPIN Rules: <http://spinrdf.org/>.

²³Jena Semantic Web toolkit: <http://jena.apache.org/>.

²⁴FLORA-2: <http://flora.sourceforge.net/>.

recognition and ontology-based learning and reasoning. Incremental activity recognition and the ability to recognize activities with incomplete sensor data are also key issues to be considered in order to speed the development of assistive technologies. The importance of embedding machine-learning techniques into the emerging research on knowledge-driven approaches will be crucial toward activity modeling and recognition.

More attention should be put into modeling uncertainty and fuzziness within OWL 2 [Bobillo and Straccia 2011]. Support for fuzzy linguistic labels would allow a more natural language-based modeling and, consequently, would improve the model's usability. In addition, this new feature included in human behavior ontologies would ease the integration with certain data-driven approaches as, for instance, statistical methods for activity recognition. The improvement in modeling of social interaction, both virtual (with human–computer messages) and physical (with human–human messages) could come from the integration of semantic tools within data-driven approaches to overcome their current limitations in this problem.

Equally important is having an adequate level of action granularity (e.g., tasks, activities, behaviors, etc.) for a specialized and incremental discovery. A formalized and common representation of universal entities, such as time, geographical indoor and outdoor locations, and environmental conditions, would greatly help in these processes. In our opinion, future works should make further efforts on a human behavior ontology that aims at filling in the identified missing gaps in the currently available proposals. Promising directions on behavior ontology development should consider the representation of imprecise and uncertain information for more accurate modeling of everyday human tasks and human language.

Another promising research area that should be further studied is change discovery, that is, (a) changes in a frequent behavior or environment migration and (b) when a predefined model of behavior is performed by a different user than the one it was designed for. Works such as Okeyo et al. [2010], focused on the ontology evolution process, could be a starting point. The change in an ontology consists here of six phases: change capturing, change representation, semantics of change, change implementation, change propagation, and change validation. This strategy allows activity learning and model evolution through ontology-based activity traces. It also makes it evident that activity models are key to support reusable, adaptive, and personalized activity recognition, as well as to improve scalability and applicability of any assistive system.

ACKNOWLEDGMENTS

We acknowledge the support of the TUCS (Turku Centre for Computer Science) graduate school, Hans Bang Foundation, the University of Granada CEI-BioTIC initiative, the MACMMEN project from the Spanish Ministry of Economy and Competitiveness, and the project *Development of an Intelligent System for Behaviour Detection and Control in a Tagged World* (TIN2009-14538-C02-01).

REFERENCES

- Bessam Abdulrazak, Yasir Malik, and Hen-I Yang. 2010. A taxonomy driven approach towards evaluating pervasive computing system. In *Proceedings of the Aging Friendly Technology for Health and Independence, and 8th International Conference on Smart Homes and Health Telematics (ICOST'10)*. Springer-Verlag, Berlin, 32–42.
- G. Acampora and V. Loia. 2005. Fuzzy control interoperability and scalability for adaptive domotic framework. *IEEE Transactions on Industrial Informatics* 1, 2 (2005), 97–111.
- J. K. Aggarwal and M. S. Ryoo. 2011. Human activity analysis: A review. *ACM Computing Surveys*. 43, 3, Article 16 (April 2011), 43 pages. DOI: <http://dx.doi.org/10.1145/1922649.1922653>
- Alessandra Agostini, Claudio Bettini, and Daniele Riboni. 2009. Hybrid reasoning in the CARE middleware for context awareness. *International Journal of Web Engineering and Technology* 5, 1 (2009), 3–23. DOI: <http://dx.doi.org/10.1504/IJWET.2009.025011>

- M. Amoretti, F. Wientapper, F. Furfari, S. Lenzi, and S. Chessa. 2010. Sensor data fusion for activity monitoring in ambient assisted living environments. In *Sensor Systems and Software*, Stephen Hailes, Sabrina Sicari, and George Roussos (Eds.). Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering, Vol. 24. Springer, Berlin, 206–221. DOI: http://dx.doi.org/10.1007/978-3-642-11528-8_15
- Francisco Arce, José Mario, and García Valdez. 2010. Accelerometer-based hand gesture recognition using artificial neural networks. *Soft Computing* 318 (2010), 67–77.
- J. Augusto. 2007. Ambient intelligence: The confluence of pervasive computing and artificial intelligence. In *Intelligent Computing Everywhere*, A. Schuster (Ed.). Springer, Berlin, 213–234.
- Matthias Baldauf, Schahram Dustdar, and Florian Rosenberg. 2007. A survey on context-aware systems. *International Journal of Ad Hoc and Ubiquitous Computing* 2, 4, 263–277. DOI: <http://dx.doi.org/10.1504/IJAHUC.2007.014070>
- N. Baumgartner and W. Retschitzegger. 2006. A survey of upper ontologies for situation awareness. In *Proceedings of the 4th IASTED International Conference on Knowledge Sharing and Collaborative Engineering*, 1–9.
- Claudio Bettini, Linda Pareschi, and Daniele Riboni. 2008. Efficient profile aggregation and policy evaluation in a middleware for adaptive mobile applications. *Pervasive and Mobile Computing* 4, 5 (2008), 697–718.
- Fernando Bobillo and Umberto Straccia. 2011. Fuzzy ontology representation using OWL 2. *International Journal of Approximate Reasoning* 52, 7 (2011), 1073–1094. DOI: <http://dx.doi.org/10.1016/j.ijar.2011.05.003>
- J. Boger, J. Hoey, P. Poupart, C. Boutilier, G. Fernie, and A. Mihailidis. 2006. A planning system based on Markov decision processes to guide people with dementia through activities of daily living. *IEEE Transactions on Information Technology in Biomedicine* 10, 2, 323–333.
- Willem Nico Borst. 1997. *Construction of Engineering Ontologies for Knowledge Sharing and Reuse*. Ph.D. Dissertation. Institute for Telematica and Information Technology, University of Twente, Enschede, The Netherlands.
- O. Brdiczka, P. Reignier, and J. Crowley. 2007. Detecting individual activities from video in a smart home. In *Proceedings of the International Conference on Knowledge-Based and Intelligent Information and Engineering Systems*, 363–370.
- A. J. Brush, J. Scott, and J. Krumm. 2010. Activity recognition research: The good, the bad, and the future. In *Proceedings of the 2010 Pervasive Workshop*.
- J. M. Cantera-Fonseca and R. Lewis. 2009. Delivery context ontology. W3C Working Draft.
- Roberto Casas, Rubén Blasco Marín, Alexia Robinet, Armando Roy Delgado, Armando Roy Yarza, John McGinn, Richard Picking, and Vic Grout. 2008. User modelling in ambient intelligence for elderly and disabled people. In *Proceedings of the 11th International Conference on Computers Helping People with Special Needs (ICCHP'08)*. Springer-Verlag, Berlin, 114–122. DOI: http://dx.doi.org/10.1007/978-3-540-70540-6_15
- Alexandros André Chaaraoui, Pau Climent-Pérez, and Francisco Flórez-Revuelta. 2012. A review on vision techniques applied to human behaviour analysis for ambient-assisted living. *Expert Syst. Appl.* 39, 12 (2012), 10873–10888. DOI: <http://dx.doi.org/10.1016/j.eswa.2012.03.005>
- Harry Chen, Tim Finin, and Anupam Joshi. 2003. An ontology for context-aware pervasive computing environments. In *Proceedings of the Workshop on Ontologies in Agent Systems*.
- Harry Chen, Tim Finin, and Anupam Joshi. 2005. The SOUPA ontology for pervasive computing. In *Ontologies for Agents: Theory and Experiences*, Valentina Tamma, Stephen Cranefield, Timothy W. Finin, and Steven Willmott (Eds.). Birkhäuser Basel, 233–258. DOI: http://dx.doi.org/10.1007/3-7643-7361-X_10
- Liming Chen, Chris Nugent, Maurice Mulvenna, Dewar Finlay, Xin Hong, and Michael Poland. 2008. Using event calculus for behaviour reasoning and assistance in a smart home. In *Smart Homes and Health Telematics*, Sumi Helal, Simanta Mitra, Johnny Wong, CarlK. Chang, and Mounir Mokhtari (Eds.). Lecture Notes in Computer Science, Vol. 5120. Springer, Berlin, 81–89. DOI: http://dx.doi.org/10.1007/978-3-540-69916-3_10
- Liming Chen and Chris D. Nugent. 2009. Ontology-based activity recognition in intelligent pervasive environments. *International Journal of Web Information Systems (IJWIS)* 5, 4 (2009), 410–430.
- Liming Chen, George Okeyo, Hui Wang, Roy Sterritt, and Chris D. Nugent. 2011. A systematic approach to adaptive activity modeling and discovery in smart homes. In *Proceedings of the 4th International Conference on Biomedical Engineering and Informatics (BMEI)*, 2192–2196.
- Luke L. Chen and Jit Biswas. 2009. Tutorial: an introduction to ontology-based activity recognition. In *Proceedings of the 7th International Conference on Mobile Computing and Multimedia (MoMM'09)*.

- Wongun Choi, C. Pantofaru, and S. Savarese. 2011. Detecting and tracking people using an RGB-D camera via multiple detector fusion. In *Proceedings of the 2011 IEEE International Conference on Computer Vision Workshops*. 1076–1083. DOI: <http://dx.doi.org/10.1109/ICCVW.2011.6130370>
- Diane J. Cook, Juan C. Augusto, and Vikramaditya R. Jakkula. 2009. Review: Ambient intelligence: Technologies, applications, and opportunities. *Pervasive Mobile Computing* 5, 4 (2009), 277–298. DOI: <http://dx.doi.org/10.1016/j.pmcj.2009.04.001>
- Diane J. Cook and Sajal K. Das. 2005. *Smart Environments: Technology, Protocols and Applications*. John Wiley.
- Aaron Crandall and Diane J. Cook. 2010. Learning activity models for multiple agents in a smart space. In *Handbook of Ambient Intelligence and Smart Environments*, Hideyuki Nakashima, Hamid Aghajan, and Juan Carlos Augusto (Eds.). Springer, New York, 751–769. DOI: http://dx.doi.org/10.1007/978-0-387-93808-0_28
- Miguel Delgado, Maria Ros, and Amparo Vila. 2009. Correct behavior identification system in a tagged world. *Expert Systems with Applications* 36, 6 (2009), 9899–9906. DOI: <http://dx.doi.org/10.1016/j.eswa.2009.01.077>
- Anind K. Dey and Gregory D. Abowd. 2000. The Context Toolkit: Aiding the development of context-aware applications. In *Proceedings of the Workshop on Software Engineering for Wearable and Pervasive Computing*.
- F. Doctor, H. Hagrais, and V. Callaghan. 2005. A fuzzy embedded agent-based approach for realizing ambient intelligence in intelligent inhabited environments. *IEEE Trans. Syst. Man Cybernet. Part A* 35, 1 (2005), 55–65.
- K. Ducatel, M. Bogdanowicz, F. Scapolo, J. Leijten, and J.-C. Burgelman. 2001. *Scenarios for Ambient Intelligence in 2010*. European Commission Community Research.
- Saibal Dutta, Amitava Chatterjee, and Sugata Munshi. 2009. An automated hierarchical gait pattern identification tool employing cross-correlation-based feature extraction and recurrent neural network based classification. *Expert Systems* 26, 2 (2009), 202–217. DOI: <http://dx.doi.org/10.1111/j.1468-0394.2009.00479.x>
- Dejene Ejigu, Marian Scuturici, and Lionel Brunie. 2007. An ontology-based approach to context modeling and reasoning in pervasive computing. In *Proceedings of the 5th Annual IEEE International Conference on Pervasive Computing and Communications. Workshops*. 14–19. DOI: <http://dx.doi.org/10.1109/PERCOMW.2007.22>
- M. Ermes, J. Parkka, J. Mantyjarvi, and I. Korhonen. 2008. Detection of daily activities and sports with wearable sensors in controlled and uncontrolled conditions. *IEEE Transactions on Information Technology in Biomedicine* 12, 1 (2008), 20–26. DOI: <http://dx.doi.org/10.1109/TITB.2007.899496>
- Asma Gharsellou, Yacine Bellik, and Christophe Jacquet. 2012. Requirements of task modeling in ambient intelligent environments. In *Proceedings of the 2nd International Conference on Ambient Computing, Applications, Services and Technologies*. IARIA, 71–78.
- Tao Gu, Hung Keng Pung, and Da Qing Zhang. 2004. A middleware for building context-aware mobile services. In *Proceedings of IEEE Vehicular Technology Conference*.
- Tao Gu, Zhanqing Wu, Xianping Tao, Hung Keng Pung, and Jian Lu. 2009. epSICAR: An emerging patterns based approach to sequential, interleaved and concurrent activity recognition. In *Proceedings of the 2009 IEEE International Conference on Pervasive Computing and Communications (PerCom'09)*. IEEE Computer Society, Washington, DC, 1–9. DOI: <http://dx.doi.org/10.1109/PERCOM.2009.4912776>
- Juan Gómez-Romero, Miguel A. Patricio, Jesús García, and José M. Molina. 2011. Ontology-based context representation and reasoning for object tracking and scene interpretation in video. *Expert Systems with Applications* 38, 6 (2011), 7494–7510. DOI: <http://dx.doi.org/10.1016/j.eswa.2010.12.118>
- H. Hagrais, V. Callaghan, M. Colley, G. Clarke, A. Pounds-Cornish, and H. Duman. 2004. Creating an ambient-intelligence environment using embedded agents. *IEEE Intelligent Systems* 19, 6 (2004), 12–20.
- Ramón Hervás, José Bravo, and Jesús Fontecha. 2010. A context model based on ontological languages: A proposal for information visualization. *Journal of Universal Computer Science* 16, 12 (2010), 1539–1555.
- Xin Hong, Chris Nugent, Maurice Mulvenna, Sally McClean, Bryan Scotney, and Steven Devlin. 2009. Evidential fusion of sensor data for activity recognition in smart homes. *Pervasive and Mobile Computing* 5, 3 (2009), 236–252. DOI: <http://dx.doi.org/10.1016/j.pmcj.2008.05.002>
- V. Jakkula, A. Crandall, and D. J. Cook. 2007. Knowledge discovery in entity-based smart environment resident data using temporal relations-based data mining. In *Proceedings of the ICDM Workshop on Spatial and Spatio-Temporal Data Mining*. 531–579.
- Julie A. Kientz, Shwetak N. Patel, Brian Jones, Ed Price, Elizabeth D. Mynatt, and Gregory D. Abowd. 2008. The Georgia Tech aware home. In *CHI'08 Extended Abstracts on Human Factors in Computing Systems*. ACM, New York, NY, 3675–3680. DOI: <http://dx.doi.org/10.1145/1358628.1358911>

- Eunju Kim, Sumi Helal, and D. Cook. 2010. Human activity recognition and pattern discovery. *IEEE Pervasive Computing* 9, 1 (2010), 48–53. DOI: <http://dx.doi.org/10.1109/MPRV.2010.7>
- G. Klyne, F. Reynolds, C. Woodrow, H. Ohto, J. Hjelm, M. H. Butler, and L. Tran. 2004. *Composite Capability/Preference Profiles (CC/PP): Structure and Vocabularies 1.0. W3C Recommendation*. Retrieved from <http://www.w3.org/TR/CCPP-struct-vocab/>.
- Laszlo Kovacs, Peter Matetelki, and Balazs Pataki. 2009. Service-oriented context-aware framework. In *Proceedings of the Young Researchers Workshop on Service-Oriented Computing*.
- R. Kowalski and M. Sergot. 1986. A logic-based calculus of events. *New Generation Computing* 4, 1 (1986), 67–95. DOI: <http://dx.doi.org/10.1007/BF03037383>
- Chin-Feng Lai, Yueh-Min Huang, Jong Hyuk Park, and Han-Chieh Chao. 2010. Adaptive body posture analysis for elderly-falling detection with multisensors. *IEEE Intelligent Systems* 25, 2 (2010), 20–30. DOI: <http://dx.doi.org/10.1109/MIS.2010.39>
- Hector J. Levesque, Fiora Pirri, and Ray Reiter. 1998. Foundations for the situation calculus. *Electronic Transactions on Artificial Intelligence* 2 (1998), 159–178.
- Fei Li and Schahram Dustdar. 2011. Incorporating unsupervised learning in activity recognition. In *Activity Context Representation (AAAI Workshops)*, Vol. WS-11-04. AAAI.
- Jiayang Liu, Zhen Wang, Lin Zhong, J. Wickramasuriya, and V. Vasudevan. 2009. uWave: Accelerometer-based personalized gesture recognition and its applications. In *Proceedings of the IEEE International Conference on Pervasive Computing and Communications (PerCom'09)*. 1–9. DOI: <http://dx.doi.org/10.1109/PERCOM.2009.4912759>
- Zongyi Liu and Sudeep Sarkar. 2006. Improved gait recognition by gait dynamics normalization. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 28 (2006), 2006.
- A Lozano-Tello and A Gómez-Pérez. 2004. ONTOMETRIC: A method to choose the appropriate ontology. *Journal of Database Management* 15, 2 (2004), 1–18.
- P. López-Matencio, J. Vales Alonso, F. J. González-Castaño, J. L. Sieiro, and J. J. Alcaraz. 2010. Ambient intelligence assistant for running sports based on k-NN classifiers. In *Proceedings of the 3rd Conference on Human System Interactions (HSI'10)*. 605–611. DOI: <http://dx.doi.org/10.1109/HSI.2010.5514507>
- Christopher J. Matheus, Mieczyslaw M. Kokar, Kenneth Baclawski, Jerzy Letkowski, Catherine Call, Michael Hinman, John Salerno, and Douglas Boulware. 2005. SAWA: An assistant for higher-level fusion and situation awareness. In *Proceedings of the 8th International Conference on Information Fusion*.
- U. Maurer, A. Smailagic, D. P. Siewiorek, and M. Deisher. 2006. Activity recognition and monitoring using multiple sensors on different body positions. In *Proceedings of the International Workshop on Wearable and Implantable Body Sensor Networks (BSN'06)*. 4pp.–116. DOI: <http://dx.doi.org/10.1109/BSN.2006.6>
- Xiaoning Meng, Ka Keung Lee, and Yangsheng Xu. 2006. Human Driving Behavior Recognition Based on Hidden Markov Models. In *Proceedings of the 2006 IEEE International Conference on Robotics and Biomimetics (ROBIO'06)*. 274–279. DOI: <http://dx.doi.org/10.1109/ROBIO.2006.340166>
- J. B. Mocholí, P. Sala, C. Fernández-Llatas, and J. C. Naranjo. 2010. Ontology for modeling interaction in ambient assisted living environments. In *Proceedings of the XII Mediterranean Conference on Medical and Biological Engineering and Computing 2010*, Panagiotis D. Bamidis, Nicolas Pallikarakis, and Ratko Magjarevic (Eds.). IFMBE Proceedings, Vol. 29. Springer, Berlin, 655–658. DOI: http://dx.doi.org/10.1007/978-3-642-13039-7_165
- D. Morris, B. Schazmann, Yangzhe Wu, S. Coyle, S. Brady, J. Hayes, C. Slater, C. Fay, King Tong Lau, G. Wallace, and D. Diamond. 2008. Wearable sensors for monitoring sports performance and training. In *Proceedings of the 5th International Summer School and Symposium on Medical Devices and Biosensors (ISSS-MDBS'08)*. 121–124. DOI: <http://dx.doi.org/10.1109/ISSMDBS.2008.4575033>
- U. Naeem and J. Bigham. 2007. A comparison of two hidden Markov approaches to task identification in the home environment. In *Proceedings of the 2nd International Conference on Pervasive Computing and Applications (ICPCA'07)*. 383–388. DOI: <http://dx.doi.org/10.1109/ICPCA.2007.4365473>
- Usman Naeem, John Bigham, and Jinfu Wang. 2007. Recognising activities of daily life using hierarchical plans. In *Proceedings of the 2nd European Conference on Smart Sensing and Context (EuroSSC'07)*. Springer-Verlag, Berlin, 175–189.
- B.A. Nardi. 1996. *Context and Consciousness: Activity Theory and Human-Computer Interaction*. MIT Press.
- George Okeyo, Liming Chen, Hui Wang, and Roy Sterritt. 2010. Ontology-enabled activity learning and model evolution in smart homes. In *Ubiquitous Intelligence and Computing*, Zhiwen Yu, Ramiro Liscano, Guanling Chen, Daqing Zhang, and Xingshe Zhou (Eds.). Lecture Notes in Computer Science, Vol. 6406. Springer, Berlin, 67–82. DOI: http://dx.doi.org/10.1007/978-3-642-16355-5_8
- R. Oppermann and M. Specht. 2000. A Context-Sensitive Nomadic Exhibition Guide. In *Proceedings of the 2nd Symposium on Handheld and Ubiquitous Computing*. Springer, 127–142.

- Ronald Poppe. 2010. A survey on vision-based human action recognition. *Image and Vision Computing* 28, 6 (June 2010), 976–990. DOI: <http://dx.doi.org/10.1016/j.imavis.2009.11.014>
- María Poveda-Villalón, MariCarmen Suárez-Figueroa, and Asunción Gómez-Pérez. 2012. Validating ontologies with OOPS! In *Knowledge Engineering and Knowledge Management*, Annette Teije, Johanna Völker, Siegfried Handschuh, Heiner Stuckenschmidt, Mathieu d’Acquin, Andriy Nikolov, Nathalie Aussenac-Gilles, and Nathalie Hernandez (Eds.). Lecture Notes in Computer Science, Vol. 7603. Springer, Berlin, 267–281. DOI: http://dx.doi.org/10.1007/978-3-642-33876-2_24
- Davy Preuveneers, Jan Van Den Bergh, Dennis Wagelaar, Andy Georges, Peter Rigole, Tim Clerckx, E. Berbers, Karin Coninx, and Koen De Bosschere. 2004. Towards an extensible context ontology for ambient intelligence. In *Proceedings of the 2nd European Symposium on Ambient Intelligence*. Springer-Verlag, 148–159.
- W. Prinz. 1999. NESSIE: An awareness environment for cooperative settings. In *Proceedings of 6th Conference on Computer-Supported Cooperative Work*.
- Carlos Ramos, Juan Carlos Augusto, and Daniel Shapiro. 2008. Ambient intelligence—the next step for artificial intelligence. *IEEE Intelligent Systems* 23, 2 (2008), 15–18. DOI: <http://dx.doi.org/10.1109/MIS.2008.19>
- P. Rashidi and D. J. Cook. 2009. Keeping the resident in the loop: Adapting the smart home to the user. *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans* 39, 5 (2009), 949–959.
- P. Rashidi and D. Cook. 2010. Multi home transfer learning for resident activity discovery and recognition. In *Proceedings of the International Workshop on Knowledge Discovery from Sensor Data*.
- Stefan Reifinger, Frank Wallhoff, Markus Ablassmeier, Tony Poitschke, and Gerhard Rigoll. 2007. Static and dynamic hand-gesture recognition for augmented reality applications. In *Proceedings of the 12th International Conference on Human-Computer Interaction: Intelligent Multimodal Interaction Environments (HCI’07)*. Springer-Verlag, Berlin, 728–737.
- P. Remagnino, H. Hagrais, N. Monekosso, and S. Velastin. 2005. Ambient intelligence. In *Ambient Intelligence*, Paolo Remagnino, Gian Foresti, and Tim Ellis (Eds.). Springer, New York, 1–14.
- Daniele Riboni and Claudio Bettini. 2011a. COSAR: Hybrid reasoning for context-aware activity recognition. *Personal and Ubiquitous Computing* 15, 3 (2011), 271–289.
- Daniele Riboni and Claudio Bettini. 2011b. OWL 2 modeling and reasoning with complex human activities. *Pervasive Mobile Computing* 7, 3 (2011), 379–395. DOI: <http://dx.doi.org/10.1016/j.pmcj.2011.02.001>
- Daniele Riboni, Linda Pareschi, Laura Radaelli, and Claudio Bettini. 2011. Is ontology-based activity recognition really effective? In *PerCom Workshops*. 427–431.
- M. Roman, C. Hess, R. Cerqueira, and A. Ranganathan. 2002. A middleware infrastructure for active spaces. In *IEEE Pervasive Computing*.
- M. Ros, M. P. Cuéllar, M. Delgado, and A. Vila. 2013. Online recognition of human activities and adaptation to habit changes by means of learning automata and fuzzy temporal windows. *Information Sciences* 220, 20 (2013), 86–101. DOI: <http://dx.doi.org/10.1016/j.ins.2011.10.005>
- Matthew Rowe, Miriam Fernandez, Sofia Angetou, and Harith Alani. 2013. Community analysis through semantic rules and role composition derivation. *Web Semantics: Science, Services and Agents on the World Wide Web* 18, 1 (2013). Retrieved from <http://www.websemanticsjournal.org/index.php/ps/article/view/293>.
- Nirmalya Roy, Abhishek Roy, and Sajal K. Das. 2006. Context-aware resource management in multi-inhabitant smart homes: A Nash H-learning based approach. In *Proceedings of the 4th Annual IEEE International Conference on Pervasive Computing and Communications (PERCOM’06)*. IEEE Computer Society, Washington, DC, 148–158. DOI: <http://dx.doi.org/10.1109/PERCOM.2006.18>
- Mohsin Saleemi, Natalia Díaz Rodríguez, Johan Lilius, and Ivan Porres. 2011. A framework for context-aware applications for smart spaces. In *Proceedings of the 4th Conference on Smart Spaces, (ruSMART’11)*, Sergey Balandin, Yevgeni Koucheryavi, and Honglin Hu (Eds.). LNCS, St. Petersburg, 14–25. Retrieved from <http://www.springerlink.com/content/d8618k217710th32/>.
- Salim K. Semy, Mary K. Pulvermacher, Leo J. Obrst, and Mary K. Pulvermacher. 2004. *Toward the Use of an Upper Ontology for U.S. Government and U.S. Military Domains: An Evaluation*. Technical Report. Workshop on Information Integration on the Web (IIWeb-04), in conjunction with VLDB-2004.
- Patrice Seyed. 2012. A method for evaluating ontologies—introducing the BFO-rigidity decision tree wizard. In *Proceedings of the 7th International Conference on Formal Ontology in Information Systems (FOIS’12)*. 191–204.
- Geetika Singla, Diane J. Cook, and Maureen Schmitter-Edgecombe. 2010. Recognizing independent and joint activities among multiple residents in smart environments. *Journal of Ambient Intelligence and Humanized Computing* 1, 1 (2010), 57–63. DOI: <http://dx.doi.org/10.1007/s12652-009-0007-1>

- Markus Stengel. 2003. *Introduction to Graphical Models, Hidden Markov Models and Bayesian Networks*. Tutorial, Toyohashi University of Technology, Japan.
- Sakari Stenudd. 2012. A model for using machine learning in smart environments. In *Proceedings of the 6th International Conference on Grid and Pervasive Computing (GPC'11)*. Springer-Verlag, Berlin, 24–33. DOI : http://dx.doi.org/10.1007/978-3-642-27916-4_4
- Ljiljana Stojanovic. 2004. *Methods and Tools for Ontology Evolution*. Ph.D. Dissertation. Karlsruhe Institute of Technology. Retrieved from <http://digbib.ubka.uni-karlsruhe.de/volltexte/1000003270>.
- Thomas Strang and Claudia Linnhoff-Popien. 2004. A context modeling survey. In *Proceedings of the Workshop on Advanced Context Modelling, Reasoning and Management, UbiComp 2004 - the 6th International Conference on Ubiquitous Computing*.
- T. Strang, C. Linnhoff-Popien, and K. Frank. Nov 2003. CoOL: A context ontology language to enable contextual interoperability. In *Proceedings of the 4th International Conference on Distributed Applications and Interoperable Systems*. IEEE Computer Society, 236–247.
- E. Tapia, S. Intille, and K. Larson. 2004. Activity recognition in the home using simple and ubiquitous sensors. *Pervasive Computing* 3001 (2004), 158–175.
- Shoji Tominaga, Masamichi Shimosaka, Rui Fukui, and Tomomasa Sato. 2012. A unified framework for modeling and predicting going-out behavior. In *Pervasive Computing*, Judy Kay, Paul Lukowicz, Hideyuki Tokuda, Patrick Olivier, and Antonio Krüger (Eds.). Lecture Notes in Computer Science, Vol. 7319. Springer, Berlin, 73–90. DOI : http://dx.doi.org/10.1007/978-3-642-31205-2_5
- Douglas L. Vail, Manuela M. Veloso, and John D. Lafferty. 2007. Conditional random fields for activity recognition. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'07)*. ACM, New York, NY. DOI : <http://dx.doi.org/10.1145/1329125.1329409>
- T. van Kasteren and B. Krose. 2007. Bayesian activity recognition in residence for elders. *IET Conference Publications* 2007, CP531 (2007), 209–212. DOI : <http://dx.doi.org/10.1049/cp:20070370>
- M. J. van Sinderen, A. T. van Halteren, M. Wegdam, H. B. Meeuwissen, and E. H. Eertink. 2006. Supporting context-aware mobile applications: an infrastructure approach. *IEEE Communications Magazine* 44, 9 (Sept. 2006), 96–104. DOI : <http://dx.doi.org/10.1109/MCOM.2006.1705985>
- Maria Poveda Villalon, Mari Carmen Suárez-Figueroa, R. García-Castro, and A. Gómez-Pérez. 2010. A context ontology for mobile environments. In *Proceedings of Workshop on Context, Information and Ontologies—CIAO 2010 Co-located with EKAW 2010*, Vol. 626. CEUR-WS, Germany.
- J. Vivaldi, J. Feliu, and M. T. Cabré. 2002. *Ontologies: A Review*, Technical Report eport IULA/INF034/02, University Pompeu Fabra, Institute for Applied Linguistics, Barcelona, Spain.
- Denny Vrandečić. 2010. *Ontology Evaluation*. Ph.D. Dissertation. KIT Karlsruhe Institute of Technology, Karlsruhe, Germany.
- D. H. Wilson, D. Wyaat, and M. Philipose. 2005. Using context history for data collection in the home. In *Proceedings of the 3rd International Conference on Pervasive Computing (PERVASIVE)*, Vol. 3468.
- Terry Winograd. 2001. Architectures for context. *Human Computer Interaction* 16, 2 (2001), 401–419. DOI : http://dx.doi.org/10.1207/S15327051HCI16234_18
- Yunfeng Wu and S. Krishnan. 2010. Statistical analysis of gait rhythm in patients with Parkinson's disease. *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 18, 2 (2010), 150–158. DOI : <http://dx.doi.org/10.1109/TNSRE.2009.2033062>
- X. H. Wang, D. Q. Zhang, T. Gu, and H. K. Pung. 2004. Ontology based context modeling and reasoning using OWL. In *Workshop Proceedings of the 2nd IEEE Conference on Pervasive Computing and Communications*.
- W3C Group. 2010. *Task Meta Models*. Retrieved from http://www.w3.org/2005/Incubator/model-based-ui/wiki/Task_Meta_Models.
- Lu Xia, Chia-Chih Chen, and J. K. Aggarwal. 2011. Human detection using depth information by Kinect. In *Proceedings of the 2011 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. 15–22.
- Stephen S. Yau and Junwei Liu. 2006. Hierarchical Situation Modeling and Reasoning for Pervasive Computing. In *Proceedings of the the 4th IEEE Workshop on Software Technologies for Future Embedded and Ubiquitous Systems, and the 2nd International Workshop on Collaborative Computing, Integration, and Assurance (SEUS-WCCIA'06)*. IEEE Computer Society, Washington, DC, 5–10. DOI : <http://dx.doi.org/10.1109/SEUS-WCCIA.2006.25>
- YoungTaek Jin Yun Her and Su-Kyoung Kim. 2010. A context-aware framework using ontology for smart phone platform. *International Journal of Digital Content Technology and Its Applications* 4, 5.

Dong Zhang, Dong Zhang, D. Gatica-Perez, S. Bengio, and I. McCowan. 2006. Modeling individual and group actions in meetings with layered HMMs. Retrieved from <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1632036>.

Huiyu Zhou and Huosheng Hu. 2008. Human motion tracking for rehabilitation, a survey. *Biomedical Signal Processing and Control* 3, 1 (2008), 1–18.

Received December 2012; revised March 2013; accepted August 2013