Department of
Information Engineering
and Computer Science
DISI

UNIVERSITY
OF TRENTO - Italy

DISI - Via Sommarive 14 - 38123 Povo - Trento (Italy)
http://www.disi.unitn.it

HUMAN-LIKE CONTEXT MODELLING FOR ROBOT SURVEILLANCE

F. Giunchiglia, E. Bignotti, M. Zeni

Technical Report # DISI-16-013

ISWER 2017: Proceedings of the First International Workshop on Semantics for Engineering and Robotics at ICSC 2017, IEEE Computer Society Press, San Diego, California, USA, 2017.

Human-like context modelling for robot surveillance

Fausto Giunchiglia
DISI - University of Trento
Via Sommarive, 9
I-38123 Povo (TN)
Email: fausto.giunchiglia@unitn.it

Enrico Bignotti
DISI - University of Trento
Via Sommarive, 9
I-38123 Povo (TN)
Email: enrico.bignotti@unitn.it

Mattia Zeni
DISI - University of Trento
Via Sommarive, 9
I-38123 Povo (TN)
Email: mattia.zeni@disi.unitn.it

***Abstract*—Robot surveillance requires, in addition to detection, to make sense of what is happening in a real world scenario, which is what humans do with contexts. This is critical when robots have to interact with people. Thus, the main issue is how to model human-like context to be mapped to robots, so that they can mirror human understanding. We propose a context model formalizing its relation with the robot, organized according to different dimensions of the environment. We then introduce the notions of *endurants* and *perdurants* to account for how space and time, respectively, aggregate context for humans, and in turn can help the surveillance from robots. To map real world data, i.e., sensory inputs, to our context model, our idea is a system capable of managing both the robots sensors and interacting with sensors from other devices. A possible use case is a robot, using the system fusing sensory inputs and the context model, patrolling an university building. Our contribution is a modelling of context that, while accounting for how human represent it, can be used in robots for surveillance purposes, coupled with a proposed system for exploiting sensors from both the robot and the environment.**

***Index Terms*—Ontologies, Context modelling, Robotics, Surveillance, Sensors**

I. Introduction

While robots can be well suited for surveillance operations, they still present many challenges. One of them is how they can understand what is happening while patrolling an area [12]. This requires not only to sense elements of the environment, but also to understand what they are and how and if they relate to each other. Humans do so innately via *context*, i.e., "a theory of the world which encodes an individual's subjective perspective about it" [5]. Context is an extremely powerful tool for humans [2], since it must be able to account for the complexity of the world [7]. To address this issue, a modelling human-like context to be mapped to robots, allowing them to mirror human understanding [16], is required.

We propose a model of context in accordance with [5], which is then organized according to the different dimensions of a robot surroundings. Furthermore, our model relies on the notions of *endurant* and *perdurant* [6]. In formal ontology [4], endurants are "individuals wholly present whenever they are present, and that persist in time while keeping their identity", e.g., people and objects, while perdurants are "individuals composed of temporal parts", e.g., processes and activities. These notions can be exploited for building ontologies representing the way humans use either space or time to guide the aggregation of their environment.

While humans map their context with the real world through their senses, robots need to use sensors instead. To do so, our idea is to adapt the previous work on the i-Log [24] system to the needs of robots. I-Log allows not only to map sensor data to the context model, but it also allows the robots to interact with sensors from other devices.

The use case for our approach is a robot, using our system fusing sensory inputs and the context model while patrolling a university floor. During its patrol, it switches between different context depending on its surveillance task needs, exploiting both its and the surrounding sensors. This scenario shows how, depending on which context is active for the robot, the elements to be recognized change and how the sensing strategies can change as well.

The reminder of this paper is organized as follows. Section II provides our definition of human-like context for robots, while Section III explains the notions of endurant and perdurant contexts. Then, Section V presents our system for sensing, expanding previous research. Section IV illustrates a scenario where our robot patrols a university building in an average day, and Section VI describes works similar to ours. Finally, Section VII concludes the paper.

II. Modelling Context

Consider a robot that enters an office where there are people doing some activities, e.g., having a meeting or moving in the office. This office is full with objects, some of which are devices capable of connection, e.g., thermostates, smartphones and so on.

Fig. 1 shows the office scenario as a knowledge graph, representing the context understood by a robot. It is a (human-like) context since it focuses only on certain elements of the real world and abstracts to a certain degree, e.g., being in a meeting and recognizing smart devices such as Enrico's smartphone and the room thermostate. Each node represents an entity, e.g., the office, with its respective attributes and their attribute values. For instance, attributes of Enrico in Fig. 1 are "Class", "Name", and "Role", and their corresponding values are "Person", "Enrico", and "PhD student", respectively. Edges represent relations between entities, e.g., "Enrico" has two relations: "Attend" for "Meeting" and "Own" for "Smartphone".

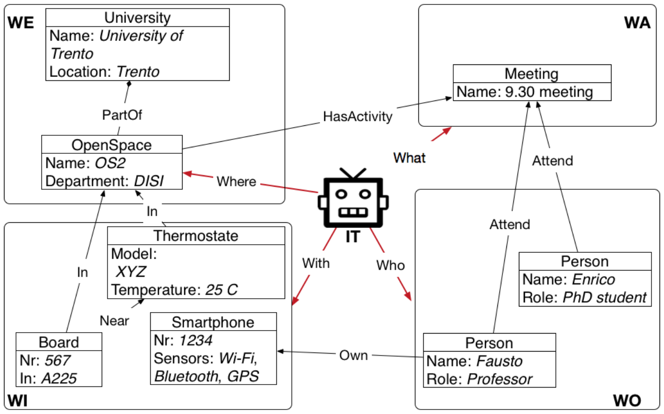

Fig. 1. The four dimensions of context, centered on the subject.

We represent the relation between the context as a partial representation of the real world and the robot it is centered on from Fig. 1 as $MyWorld =< it, Cxt >$ where:

- *it* is the robot on which the context is focused, represented as an entity with its attributes and relations.
- *Cxt* is the (real world) context, aggregating different elements surrounding the robot.

In Fig. 1, the red arrows represent this relation between *it* and *Cxt*, since they link elements of the context directly to the robot, e.g., where the robot is or the devices that it can communicate with.

Furthermore, we model Cxt as: $Cxt = WA \cup WE \cup WO \cup WI$, where:

- WA is the Temporal context, i.e., the context generated from the question "**WhAt** are you doing?".
  In Fig. 1, WA consists in the main activities taking place, i.e., the meeting.
- WE is the Spatial context, i.e., the context generated from the question "**WhE**re are you?".
  In Fig. 1, WE list the most relevant location, i.e., the office.
- WO is the Social context, i.e., the context generated from the question "**WhO** are you with?".
  In Fig. 1, WO focuses on the people in the office, i.e., Enrico and Fausto.
- WI is the Object context, i.e., the context generated from the question "**W**hat are you w**I**th?".
  In Fig. 1, WI covers different devices together with their sensors, e.g., the room thermostate or Enrico's smartphone.

III. Endurants and Perdurants

In addition to its components, humans tend to aggregate context according to two specific criteria: time and space. We account for this with the notions of endurants and perdurants. According to [4], endurants are "individuals wholly present whenever they are present, and that persist in time while keeping their identity", e.g., buildings and people, while perdurants are "individuals composed of temporal parts", e.g., events.

In an endurant context, one could say "I am in the office", thus focusing on the location; notice that this is still effective at different levels of abstraction, since it could also be understandable to say "I'm at my workplace". In a perdurant context, one could say "I'm in a meeting", while other activities may be going on, e.g., somebody leaving the office or people working and so on. In both cases the objective state of affairs of the world is the same, but the representation is different.

Depending on which element is more salient, the representation of the same real world scenario may change. For instance, from Fig. 1, Fig. 2 shows that the focus can be on:

- **Endurant context:** being in the office is more salient, so the activities to be performed are fixed: the possible activities are either studying or having a meeting.
- **Perdurant context:** meeting is the most salient event, so the possible locations, and their granularity, are less relevant and may be very different types of locations; for instance, a meeting may be virtual while a person is in his or her own house.

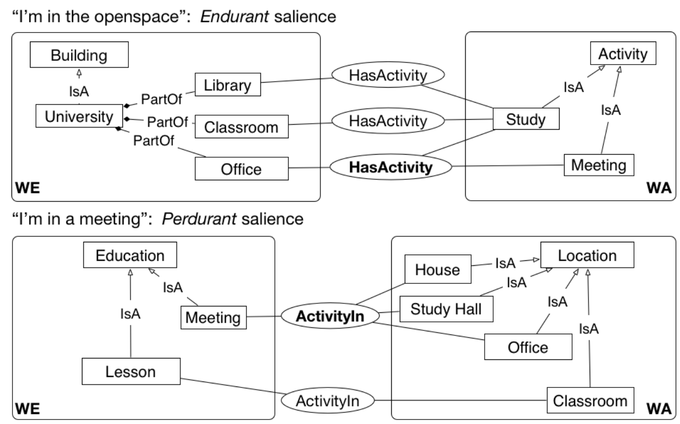

Fig. 2. The difference between the notions of endurant and perdurant when describing the context.

Notice that "HasActivity" and "ActivityIn", which map locations in the *WE* context and activities in the *WA* context and vice versa, are not inverse functions, but they link different elements depending on the saliency of the context. This means that $ActivityIn = HasActivity^{-1}$ is not necessarily true. In fact, in Fig. 2 in the case of endurant context "HasActivity" maps the location "Office" to both the activities "Study" and "Meeting", whereas in the perdurant context "ActivityIn" maps the activity "Meeting" to many more elements, i.e., "House", "Study Hall", "Classroom", and "Office". This shows that the structure changes depending on the viewpoint, since the relations do not map to the same elements.

As noted in [6], these notions affect the activity recognition process, since, depending on which context is active, the elements to be recognized and to be expected change. In addition, from a surveillance point of view, endurant contexts can be used during the patrolling activities where the environment around the robot is more relevant, e.g., checking whether some objects are missing or some analytics on a room like temperature. On the other hand, perdurant contexts can be used for situations where activities to be recognized are more relevant, e.g., understanding and focusing on specific activities being performed in a room.

IV. Perdurant- and endurant-based context surveillance

For our proposed use case, we assume an autonomous mobile robot, equipped with a system for aggregating sensory information from the robot and the other devices it encounters in the building. Since we focus on the sensing part of the patrolling in this scenario, we do not account for other issues in the area of autonomous mobile robot surveillance, e.g., self configuration or optimization and planning of the patrolling path [17], [12] and so on.

The idea is that this robot patrols the university by moving around in the building and providing analytics for any location modelled in our ontology, switching between perdurant and endurant contexts depending on the surveillance needs. The analytics depend on the properties to be recognized for the specific location, e.g., temperature. The robot is equipped with sensors, described in detail in Section V, and has access to the schedule of the activities taking place on the day of its patrol. Once it has checked a room, it produces a report specifying the environmental parameters of the room, together with additional notes on activities, people and objects in the room if needed.

Our scenario starts with the robot patrolling a floor of the university where there are different types of rooms, e.g., offices, classrooms, bathrooms, service rooms and hallways, where a set of activities can be performed, e.g., meetings or lessons. We will focus on five different points of of the scenario to highlight aspects of both context modelling and the interactions of our sensing system.

1) **Occupancy of classroom "A201":** The first stop for the robot consists in checking whether there are people in the classroom "A201", where a lesson is expected to start soon, as shown in Fig. 3. To do so, it switches to an edurant context and checks for people in the classroom. The robot can estimate that the room is not (too) crowded by checking different sensors in its *WI* context and comparing them to the maximum room occupancy. It initially checks for smartphones, which imply human presence, but since not every person may carry around or possess a smartphone, it relies on other environmental inputs. Thus, it checks for the level temperature by connecting to the classroom thermostate, since the presence of people generally leads to a rise in temperature. Indeed, the temperature is higher than its average recorded in the past days, when the room was empty. Also, it checks that the light is on through its own dedicated sensors, since people would not stay in a dark room. Once all sensory data are analyzed and the results shows (estimated) occupancy to be within the room limit, it goes out of the classroom and continues its patrol.
2) **Broken thermostate in "Lases" meeting room:** The robot comes into the "Lases" meeting room, where

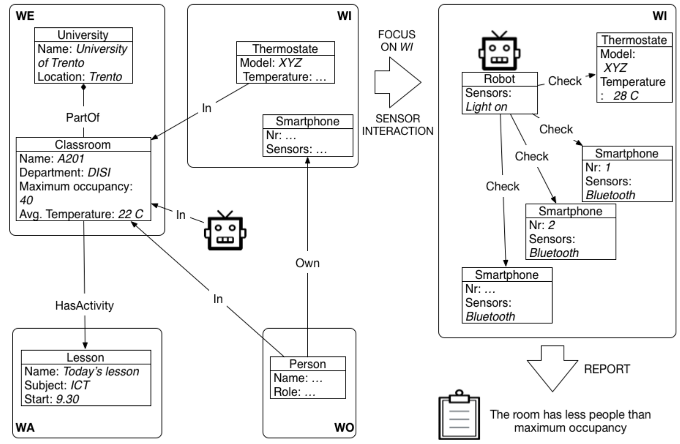

Fig. 3. The case of occupancy as a knowledge graph, showing our sensing system coupled with the endurant context model.

there are reports of issues with the heating. The robot checks to be in the correct room by using both Wi-Fi positioning and the RFID tag of the chairs in the room. While it understands that it is in the correct room, it notices that one chair is missing from the room. Then, its temperature sensor checks for the temperature in the room, and detects that it is lower than the average recorded for the past day. Thus, it tries to contact the thermostate; however, the connection fails and so the robot sends a report stating that the thermostate is not working and that a chair is missing from meeting room "Lases". Then, the robot resumes its patrol and leaves the room
3) **Lesson in classroom "A201:** Now the robot comes back to the classroom "A201" from 1) and checks whether the scheduled lesson is actually taking place. To do so, it switches to a perdurant context to check for the activity "Lesson". This leads to detecting additional devices for the lesson, e.g., the projector, in addition to the sensors used in 1). By ensuring that the projector is turned on and that the previous sensory inputs are aligned with the current state of the room, the robot can report that the lesson is taking place according to schedule. It then continues its patrol.
4) **Meeting in "Garda" meeting room:** According to schedule, there should be a meeting in room "Garda". However, when entering the room, the same sensor strategy in 3) shows that the projector is off, and that no voice can be heard. The robot checks for any device that may belong to people, e.g., smartphones, but to no avail. It infers that no meeting is going on, since no humans are in the room, it then turns off the light, either by actuating or notifying that it should be turned off in its report, and moves on with its patrol.
5) **Night time patrolling** The last case of our scenario is at night time. While being an endurant context, the environmental parameters change from day time, since now nobody must be present and all lights must be off. Thus, the robot patrols by ensuring that these parameters

hold for every room. During its patrol, the robot faces two particular cases: *i)* like in 4), a room is empty with the lights on, so it proceeds to switch the light off but no alert is sent; *ii*) in a hallway, the light is off but a human noise is heard, so an alert is sent, similarly to [9], so that a security guard is sent to check.

V. WORLD SENSING

Although the representation of context presented in this work is in line with human understanding, it must still connect with the real world. This means that the data collected from the sensors involved, i.e., the sensors on the robot and in the environment, must be mapped with the context. The generation of these abstractions is the key element of our approach. To do this, we envision a system that collects sensory data and process them, mapping the results to the corresponding elements of the context.

A. System Architecture

For this work, we rely on the i-Log [24] system conveniently modified. i-Log is an infrastructure composed by different elements seamlessly interacting with each other to move from sensor data to the level of the contex model. i-Log is built according to a client-server architecture. The client side consists in a mobile application for Android smartphones that collects personal data from the users in a transparent and energy efficient way [23]. The server side is dedicated to memorize the data, make them available for future analysis, and it generally makes the mapping between sensor data and the context model possible. Its components are the following:

- **API system:** Enables the communication between the client and the server in both ways, in order to upload the raw sensor data and download the context information.
- **Sensor data storage system (STB):** It stores the streams of data coming from the sensors in a NOSQL database. The NOSQL paradigm is best suited for needs of scalability. In fact, the client side can generate a considerable amount of information, up to 1GB per day per client instance and both writes and subsequent reads should be fast. Moreover, Cassandra[1] is the best candidate as the storage engine because of the linear scalability it provides. In fact, it allows for linear scalability by increasing the number of nodes in the database cluster. This allows us to increase the number of users without the need of modifying the server side implementation but simply by adding nodes to the cluster.
- **An entity storage (EB/KB):** This part of the system stores the contextual information as entities, their attributes and their corresponding values. This structured way of storing the data allows for future queries.
- **Service module (SM):** It performs the mapping between the sensor data and the services. It can be implemented using a distributed computational framework (Apache Spark) as for the database. This is also due to the need to be transparent to the system users. Just like the database, adding workers to the cluster increases the computational power and reduces the mapping time between streams of sensor data and the context model. This components analyzes the streams stored in the STB and updates with the results of the values in the EB/KB.

For this work, the idea is to design an ad-hoc solution for the client side in addition to the original i-Log smartphone implementation. In fact, since humans tend to carry the smartphone with them, it is the perfect piece of technology to collect personal data. In the case of robots, smartphones are less suited for the task. In fact, a smartphone is something that requires interaction with a touch screen, which of course a robot cannot do. Moreover, in this scenario we could use much simpler devices other than a smartphone, thus improving battery life, if necessary. A good candidate is an embedded device, e.g., a Raspberry Pi[2], specifically designed to collect data from the sensors attached on it on the robot. Our implementation accounts for a communication part in charge of connecting wirelessly with the server to exchange the collected streams of sensor data. Additionally, the same wireless interface can also connect to objects and other sensors and actuators in the environment to read sensor values from them and to send actuation signals, e.g, turn off the light (see Section IV).

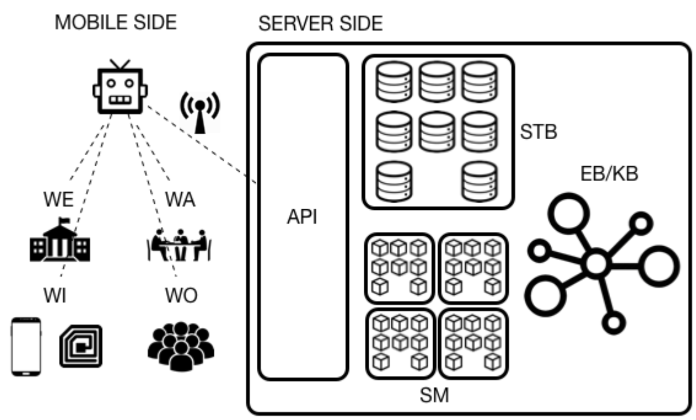

Fig. 4. System architecture showing client and server.

At the same time, the idea is to adapt also the server side to deal with the new robot configuration at the client side. The system should be designed so that it can work with any stream of data conveniently formatted in the appropriate way, and also IoT devices. For this specific use case, we propose to extend the original API part so that it could correctly parse the information coming from the new client, since they would be formatted differently from the original implementation involving smartphones. Finally, this will require uploading the entities, with their respective attributes and attribute values in our Entity Base database, while the Knowledge Base stores the knowledge about the structure of the entities, i.e., their attributes and relations.

The whole proposed system, with both the client and the server sides is shown in Figure 4.

[1]http://cassandra.apache.org/

[2]https://www.raspberrypi.org/

TABLE I
TABLE SHOWING THE MAPPING BETWEEN SENSORS AND THE ROBOT CONTEXT MODEL FOR THE 5 SCENARIOS AS ILLUSTRATED IN SECTION IV

Scenario	Sensors	Rules	Context Element
1	Wi-Fi	Wi-Fi can be detected here	**WE:** Room A201
	Thermometer, Light	People increase room temperature, stay with lights on and talk. People own smartphones.	**WO:** People
	Microphone	There is at least a vocal track	**WA:** Speaking
	Bluetooth	Objects have a Bluetooth	**WI:** Smartphone, Thermostate
2	Wi-Fi	Wi-Fi can be detected here	**WE:** Lases meeting room
	Bluetooth, RFID	Objects have a Bluetooth or RFID	**WI:** Thermostat, Chair
3	Wi-Fi	Wi-Fi can be detected here	**WE:** Room A201
	Thermometer, Light	People increase room temperature and stay with lights on. People own smartphones	**WO:** People
	Microphone	There is a major vocal track. Lessons require the projector to be on	**WA:** Lesson
	Bluetooth	Objects have a Bluetooth	**WI:** Smartphone, Thermostat, Projector
4	Wi-Fi	Wi-Fi can be detected here	**WE:** Garda meeting room
	Microphone	There is a major vocal track. Meetings require the projector to be on	**WA:** Meeting
	Bluetooth	Objects have a Bluetooth	**WI:** Smartphone, Projector
5	Wi-Fi	Wi-Fi can be detected here	**WE:** Hallway, Room
	Light	People stay with lights on and make noise	**WO:** Person
	Microphone	In a quiet environment, a noise stands out	**WA:** Movement

B. Robot and Environment Sensor Mapping

By combining together sensor data and semantic representations, we propose to account for the process of the robot for understanding the environment, allowing the system to elaborate the sensor information and map it to the elements composing the context, with the final goal of understanding the elements of the scenario in Section IV, e.g., "There is lesson in classroom A201". The algorithms for the mapping process are implemented in the modules of the SM component of the server. Each algorithm uses a combination of sensor data and semantic rules to generate the output. For example, consider the first element from the scenario in Section IV, i.e., "There is lesson in classroom A201". First of all, the robot could understand its location using, among others, the Wi-Fi triangulation thanks to multiple routers in the building. The next step is to detect how many people are in the room using the Bluetooth interface to detect the users' smartphones. Additionally, this element can be inferred by checking the temperature of the room using the embedded thermometer on the robot or read the temperature value from the thermostat in the room. Finally, as an additional check, the light sensor can be used to infer the presence of people. In fact, one can argue that if a person is in the room, the light would be on. This last element shows how sensory data and higher level semantic information can be combined together to match the context elements. Table I shows for every scenario described in Section IV what are the sensors involved and the semantic rule generating the corresponding context element.

VI. RELATED WORK

As shown in the surveys by [17] and [12], robot surveillance is gaining a lot of momentum, especially in the case of autonomous mobile robots. A direct comparison with our proposed robot is the iBotGuard [9], a security system using autonomous mobile robots on wheels for detecting intruders. This system alerts for any unusual activity while patrolling via an image feed to a monitoring station. It relies on an invariant face recognition approach for recognizing intruders. Our main differences with respect to these types of robots is that we focus on sensing rather than other research areas of robotics (as mentioned in Section IV). Furthermore, we do not assume video cameras on our robot, although it is common for surveillance robots [17], [12]. This is because we expect our robot to be able to interact with humans controllers to use cameras in the building.

We are not the first proposing to adopt ontologies for (mobile) autonomous robots. The main work concerning ontologies in robotics is done by the IEEE RAS Ontology for Robotics and Automation Working Group (ORA WG) [15], which is divided in several subgroups to study different domains of robotics and standardize the knowledge about them. Among these groups, the Autonomous Robots (AuR) has been working on developing an ontology for these types of robots. [11] introduces the fundamental notions and some use case for the application of this ontology, e.g., mine hunting and space explorations. [1] discusses the requirements for the development of this ontology, which is still ongoing, in terms of robotic platforms, e.g., types of unmanned robots, subsystems and components, e.g., sensors, system modelling, i.e., state estimations of autonomous robots, and aspects of autonomy such as path planning and robot control.

Ontologies are also an important tool for context modelling, especially in the area of ubiquitous and pervasive computing. One of the earliest works is CoBrA [3], an agent-based infrastructure capable of performing several context operations such as modeling, reasoning, and knowledge sharing. CONON [19] focuses on the modelling of locations via an upper ontology and domain-specific ontologies connected hierarchically. Similarly, CaCONT [20] concentrates on locations, by providing different levels of abstraction to describe the location of entities, e.g., GPS and location hierarchies. PiVOn [8] consists of four ontologies. The first one describes users while the second one models services for the users. The third one accounts for devices and their specifications, and the last one represents the locations and objects in them. Our main novelty with respect to these works is the notions of endurants and perdurants. However, the Mining Minds Context Ontology [18] models context, consisting in locations, activities and emotions, by organizing it according to a scenario, e.g., amusement, housework, commuting and so on, which could be treated either as a perdurant or endurant context. In addition, several works in ontology based activity recognition use ontologies to model contexts or elements such as activities of daily living (ADLs) within smart homes (SHs) [14]. Context is used mainly for informing the activity recognition phase, allowing for reasoning and reducing the amount of activities to be recognized [22], [13]. From a more theoretical point of view, [21] proposes a top level ontology for smart environments, and [10] describes an ontological model for human activity in smart homes.

VII. Conclusions

In this work we addressed the modelling context for robots, in the area of robot surveillance, in human-like fashion, in order to help them making sense of the real world. We proposed a context model organized according to different dimensions of the environment. We also adopted the notions of *endurants* and *perdurants* to account for how space or time, respectively, are used by humans as the main criterion for structuring their context. To account for sensor data, we proposed to extend the i-Log [24] system for managing the sensors of both robots and other available devices. We presented a possible scenario of a robot patrolling a university building, switching between perdurant and endurant contexts. Our contribution is a modelling of context that, being human-like, could be used in robots for surveillance tasks, coupled with a proposed system able to exploit sensors from both the robot and its surroundings.

Acknowledgment

This work has been supported by the SmartSociety (http://www.smart-society-project.eu) EU FP7 project, Grant Agreement No. 600854.

References

[1] B. Bayat, B. Bayat, J. Bermejo-Alonso, J. Bermejo-Alonso, J. Carbonera, J. Carbonera, T. Facchinetti, T. Facchinetti, S. Fiorini, S. Fiorini *et al.*, “Requirements for building an ontology for autonomous robots,” *Industrial Robot: An International Journal*, vol. 43, no. 5, pp. 469–480, 2016.

[2] P. Bouquet and F. Giunchiglia, “Reasoning about theory adequacy. a new solution to the qualification problem,” *Fundamenta Informaticae*, vol. 23, no. 2, 3, 4, pp. 247–262, 1995.

[3] H. Chen, T. Finin, and A. Joshi, “An intelligent broker architecture for context-aware systems,” *PhD proposal in computer science, University of Maryland, Baltimore, USA*, 2003.

[4] A. Gangemi, N. Guarino, C. Masolo, A. Oltramari, and L. Schneider, “Sweetening ontologies with dolce,” in *International Conference on Knowledge Engineering and Knowledge Management*. Springer, 2002, pp. 166–181.

[5] F. Giunchiglia, “Contextual reasoning,” *Epistemologia, special issue on I Linguaggi e le Macchine*, vol. 16, pp. 345–364, 1993.

[6] F. Giunchiglia, E. Bignotti, and M. Zeni, “Personal context modelling and annotation,” submitted to *1st International Workshop on Annotation of useR Data for UbiquitOUs Systems (ARDUOUS)*, 2017.

[7] F. Giunchiglia, E. Giunchiglia, T. Costello, and P. Bouquet, “Dealing with expected and unexpected obstacles,” *Journal of Experimental & Theoretical Artificial Intelligence*, vol. 8, no. 2, pp. 173–190, 1996.

[8] R. Hervás, J. Bravo, and J. Fontecha, “A context model based on ontological languages: a proposal for information visualization.” *J. UCS*, vol. 16, no. 12, pp. 1539–1555, 2010.

[9] J. N. Liu, M. Wang, and B. Feng, “iBotGuard: an internet-based intelligent robot security system using invariant face recognition against intruder,” *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, vol. 35, no. 1, pp. 97–105, 2005.

[10] Q. Ni, I. Pau de la Cruz, and A. B. García Hernando, “A foundational ontology-based model for human activity representation in smart homes,” *Journal of Ambient Intelligence and Smart Environments*, vol. 8, no. 1, pp. 47–61, 2016.

[11] L. Paull, G. Severac, G. V. Raffo, J. M. Angel, H. Boley, P. J. Durst, W. Gray, M. Habib, B. Nguyen, S. V. Ragavan *et al.*, “Towards an ontology for autonomous robots,” in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2012, pp. 1359–1364.

[12] T. D. Räty, “Survey on contemporary remote surveillance systems for public safety,” *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, vol. 40, no. 5, pp. 493–515, 2010.

[13] D. Riboni and C. Bettini, “Cosar: hybrid reasoning for context-aware activity recognition,” *Personal and Ubiquitous Computing*, vol. 15, no. 3, pp. 271–289, 2011.

[14] N. D. Rodríguez, M. P. Cuéllar, J. Lilius, and M. D. Calvo-Flores, “A survey on ontologies for human behavior recognition,” *ACM Computing Surveys (CSUR)*, vol. 46, no. 4, p. 43, 2014.

[15] C. Schlenoff, E. Prestes, R. Madhavan, P. Goncalves, H. Li, S. Balakirsky, T. Kramer, and E. Miguelanez, “An IEEE standard ontology for robotics and automation,” in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2012, pp. 1337–1342.

[16] J. M. Taylor, “Mapping human understanding to robotic perception,” *Procedia Computer Science*, vol. 56, pp. 514–519, 2015.

[17] T. Theodoridis and H. Hu, “Toward intelligent security robots: A survey,” *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, vol. 42, no. 6, pp. 1219–1230, 2012.

[18] C. Villalonga, O. Banos, W. A. Khan, T. Ali, M. A. Razzaq, S. Lee, H. Pomares, and I. Rojas, “High-level context inference for human behavior identification,” in *International Workshop on Ambient Assisted Living*. Springer, 2015, pp. 164–175.

[19] X. H. Wang, D. Q. Zhang, T. Gu, and H. K. Pung, “Ontology based context modeling and reasoning using owl,” in *Pervasive Computing and Communications Workshops, 2004. Proceedings of the Second IEEE Annual Conference on*. Ieee, 2004, pp. 18–22.

[20] N. Xu, W. S. Zhang, H. D. Yang, X. G. Zhang, and X. Xing, “Cacont: a ontology-based model for context modeling and reasoning,” in *Applied Mechanics and Materials*, vol. 347. Trans Tech Publ, 2013, pp. 2304–2310.

[21] J. Ye, G. Stevenson, and S. Dobson, “A top-level ontology for smart environments,” *Pervasive and mobile computing*, vol. 7, no. 3, pp. 359–378, 2011.

[22] ——, “Kcar: A knowledge-driven approach for concurrent activity recognition,” *Pervasive and Mobile Computing*, vol. 19, pp. 47–70, 2015.

[23] M. Zeni, F. Giunchiglia, and Z. Melkamu Mogesse, “Energy efficient service-based context recognition,” submitted to *IEEE International Conference on Pervasive Computing and Communications (PerCom 2017)*, 2017.

[24] M. Zeni, I. Zaihrayeu, and F. Giunchiglia, “Multi-device activity logging,” in *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication*. ACM, 2014, pp. 299–302.