Meet HanS, the Health&Safety Autonomous Inspector

Emanuele Bastianelli[1], Gianluca Bardaro[2], Ilaria Tiddi[1,3], and Enrico Motta[1]

[1] Knowledge Media Institute, The Open University, UK
{name.surname}@open.ac.uk
[2] Dip. di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, Italy
gianluca.bardaro@polimi.it
[3] Department of Computer Science, VU University Amsterdam, NL
ilaria.tiddi@vu.nl

Abstract. We present here one of the demonstrators we implemented as part of our research on the integration of robots in smart cities, where an autonomous mobile platform is employed for monitoring and assessing the Health&Safety rules in a smart office environment in combination with a centralised infrastructure for data integration, processing and reasoning. Monitoring of the status of infrastructures and environments still presents a number of challenges in terms of knowledge acquisition and representation, as data need to constantly be re-evaluated due to their high dynamism. Common solutions, ranging from human monitoring to sensor deployment, fail in flexibility, costs and, in the case of large scale scenarios, scalability. We focus on the idea that autonomous mobile agents can be used as moving sensors deployed by a larger, knowledge-based infrastructure, where the central unit collects and reasons over the information produced by the agents. In particular, the paper presents HanS, the Health&Safety inspector, with the goal of showing that applications integrating robots as data consumers and collectors can be deployed thanks to a combination of state-of-the-art semantic and robotics technologies.

Keywords: Robots · Knowledge Acquisition · Mobile Sensors

1 Introduction

Monitoring the status of infrastructures and environments is a task common to many domains including, for instance, the smart cities scenarios. In these cases, a centralised infrastructure is in charge of collecting and integrating data from a variety of sources (energy consumption, transport and mobility data, citizens opinions etc.) to support decision making and resource optimisation [?]. In our research, we are looking into integrating autonomous mobile agents into this "system of systems", with the idea that they can provide valuable services in urban scenarios, but also that their capabilities can be improved by using the vast external knowledge provided within the city infrastructure.

Managing data in these scenarios is still a challenging task in terms of knowledge acquisition, as data can change and/or lose validity in time, resulting in the central infrastructure to become outdated relatively quickly, hence affecting reasoning and decision making. To cope with this, several solutions are possible, e.g. areas can be monitored either with recurrent human inspections or by deploying static sensors regularly streaming information. Besides being expensive in time, resource and maintenance, these solutions are neither flexible (e.g. sensors are bound to specific locations) nor scalable, especially in large scale scenarios such as urban environments.

Works in Robotics have demonstrated that using an autonomous mobile agent as a moving sensor is a valid alternative solution, but have rather focused on achieving robotics tasks (perception, planning, navigation, etc.) with high precision [**?**]. We focus instead on integrating robots as part of a larger knowledge-based system, whose role is to store and manage the information that robots collect, share and revise continuously for the achievement of their daily activities. As semantic technologies have proven to be successful in contexts where knowledge from heterogeneous sources needed to be integrated to enable a variety of light-weight applications [**?**], our question here becomes *can we use the available semantic technologies to deploy applications where robots act both as data collectors and data consumers for a centralised knowledge base?*

To answer this question, we implemented the HanS system, where a mobile robot is deployed as part of a central knowledge-based system to autonomously assess the correct compliance with the Health&Safety rules holding in the Knowledge Media Institute (KMi). We focus on rules concerning the use of appropriate signage for emergency appliances and runaways, e.g. "*Are fire extinguishers clearly labelled?*" (Rule01), or the presence of forbidden objects in restricted areas, e.g. "*Are electric heaters away from confined areas?*" (Rule02). This demo presents the implementation of the system and a simulation of HanS at work, showing users how robots can be integrated in simple applications through combining the available semantic and robotics technologies.

2 System Architecture

HanS is implemented as a modular architecture with the idea that each module can be easily replaced with more advanced implementations depending on platforms and tasks. As shown in Figure 1, the system is articulated in (i) the *Knowledge Component*, that manages the knowledge level, and (ii) the *Sensing Component*, performing monitoring and data collection.

The *Knowledge Component* is composed by three modules, i.e. the **Knowledge Base**, the **Triplestore**, and the **RESTful Server**. The Knowledge Base (KB) module contains all data necessary for the system to reason over H&S rule violations – namely, information about the physical environment, called **Semantic Map** [**?**], and the definition of the **H&S Rules**. The semantic map describes objects and areas in terms of their position w.r.t. the geometrical map used by the robot for its localisation and navigation, as well as other objects/areas' properties such as types and names. For the sake of simplification, rules are encoded as constraints or restrictions about specific objects in specific locations, following a ba-

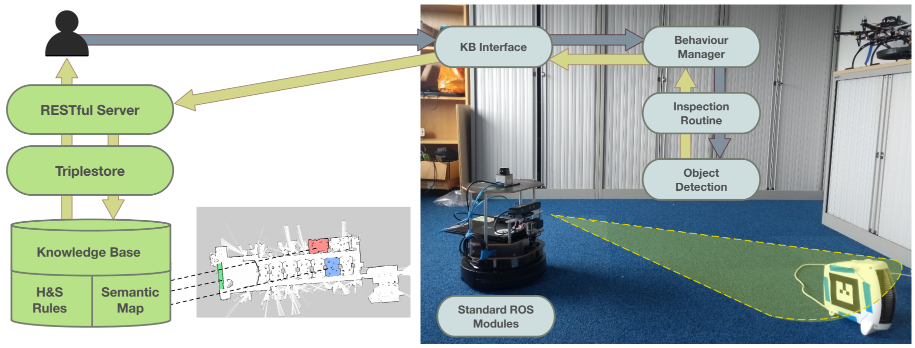

Fig. 1. Architecture of the HanS system. Modules belonging to the Knowledge Component are in green, while the ones belonging to the Sensing Component are in blue.

sic schema ⟨`:RestrictionRule,:hasForbiddenObject,:Object`⟩. ⟨`:RestrictionRule,:hasLocation,:Area`⟩. For example, a `:HeaterRestriction` establishes that a `:Heater` cannot be located in a space of type `:Activity`[1]. The knowledge base is managed by the Blazegraph triplestore[2], which natively supports both geospatial reasoning and inference checking. The RESTful server, implemented in Blazegraph as a simple Servlet, is the module in charge of communication between a user, the knowledge base and the robotics platform.

The *Sensing Component* is in charge of collecting information and includes the physical robot and all the modules to operate it. The **platform** we use is a Turtlebot 2 equipped with a Hokuyo Laser Rangefinder (for localisation and navigation), and an Orbbec Astra RGBD sensor for object detection. All the software modules, managing specific functionalities of the robot, are implemented through the Robot Operating System[3] (ROS) framework using a combination of custom and **standard ROS modules**. For example, we use `move_base` for navigation and `kobuki_node` for motor control. A custom **Object Detection** module is used for the task of detecting an object and its position when navigating an area. This is currently implemented as an ARTag[4] detection process, associating an ARTag with an object class (e.g. `:Heater`), but can be easily extended with more sophisticated systems. An **Inspection Routine** module then implements the logic of the exploration process necessary to check a specific rule: for example, the module will instruct the robot to perform a 360°-spin when reaching an area (Rule02), or to perform an inspection along a wall with a specific height (Rule01). Finally, the **Behaviour Manager** uses behavioural trees [**?**] to activate a specific robot behaviour depending on the type of rule to be checked (currently activated upon user requests, e.g "check for Rule02"), while the *KB interface* manages the communication with the server.

[1] Activities are portions of open space in KMi.
[2] https://www.blazegraph.com/
[3] http://www.ros.org/
[4] https://en.wikipedia.org/wiki/ARTag

3 Demonstration

The demonstration will show HanS in action using recorded videos and, if conditions allow, a simulation of the robot and environment. Users will be shown the whole system in-depth, and will be able to check live for a wider range of rules. Discussion on possible extensions (e.g. reasoning through SPIN[5] rules, object detection through neural network approaches etc.) will also be held. A demo of the process is currently available online[6], described as follows.

1. At first, the knowledge base includes only spatial information about the areas in the environment and the H&S rules. A rule can apply to several or no areas. The robot is idle waiting for commands.
2. A user instructs the robot to check for Rule02, and the corresponding behaviour tree is triggered by the Behaviour Manager.
3. The robot queries the triplestore to know all the areas where Rule02 applies. The qeury result set comprehends the `:Activity2`, the `:RearPodium` and the `:RoboLab`.
4. The robot starts the navigation going towards `:Activity2`. When the area is reached, the Inspection Routine module starts the inspection behaviour associated with Rule02 (i.e. spinning for 360° while seeking for heaters).
5. The robot carries on and repeats step 4 for the `:RearPodium`, where an ARTag is detected. The class corresponding to the ARTag and the position are sent to the server, which (1) updates the knowledge base with the new information collected and (2) queries the KB for any rule violation. Since Rule02 applies to the `:RearPodium`, a violation warning is fired to the user.
6. Finally, the robot navigates to the `:RoboLab`, where a new ARTag associated with the class `:Chair` is detected. Since no rule imposes restrictions on the presence of chairs in the robotics lab, the knowledge base is then updated with the newly collected observation.

References

1. Colledanchise, M., gren, P.: How Behavior Trees Modularize Hybrid Control Systems and Generalize Sequential Behavior Compositions, the Subsumption Architecture, and Decision Trees. IEEE Transactions on Robotics **33**(2), 372–389 (April 2017)
2. d'Aquin, M., Davies, J., Motta, E.: Smart cities' data: Challenges and opportunities for semantic technologies. IEEE Internet Computing **19**(6), 66–70 (2015)
3. Lécué, F., Tucker, R., Bicer, V., Tommasi, P., Tallevi-Diotallevi, S., Sbodio, M.: Predicting severity of road traffic congestion using semantic web technologies. In: European Semantic Web Conference. pp. 611–627. Springer (2014)
4. Nüchter, A., Hertzberg, J.: Towards semantic maps for mobile robots. Robot.is and Autonomous Systems **56**(11), 915–926 (2008)
5. Paola, D.D., Milella, A., Cicirelli, G., Distante, A.: An autonomous mobile robotic system for surveillance of indoor environments. International Journal of Advanced Robotic Systems **7**(1), 8 (2010)

[5] http://spinrdf.org/
[6] https://tinyurl.com/ycs7t9c2