A Testbed for Autonomous Robot Surveillance (Demonstration)

Stefan J. Witwicki
EPFL, Switzerland
stefan.witwicki@epfl.ch

José Carlos Castillo
U. Carlos III de Madrid, Spain
jocastil@int.uc3m.es

Jesús Capitán
University of Seville, Spain
jcapitan@us.es

João V. Messias, João C. Reis, and Pedro U. Lima
Instituto Superior Técnico,
Universidade de Lisboa, Portugal
{jmessias,jreis,pal}@isr.ist.utl.pt

Francisco S. Melo
INESC-ID / Inst. Superior Técnico,
Universidade de Lisboa, Portugal
fmelo@inesc-id.pt

Matthijs T. J. Spaan
Delft University of Technology,
The Netherlands
m.t.j.spaan@tudelft.nl

ABSTRACT

This demo paper serves to describe a research platform for autonomous surveillance by mobile robots. System components from various technical disciplines (including robot navigation, computer vision, discrete event systems, and human-computer interaction) have been combined together into a cohesive, modular system architecture. We employ our architecture to experiment with state-of-the-art methods in planning and decision making under uncertainty, which we demonstrate enable robots to effectively patrol and respond to surveillance events in a large indoor environment.

Categories and Subject Descriptors

I.2.1 [**Artificial Intelligence**]: Applications

General Terms

Design, Experimentation

Keywords

Planning under Uncertainty, HRI, Discrete Event Systems

1. INTRODUCTION

The increasing prevalence of surveillance in our society means bigger and more sophisticated surveillance systems, monitoring larger environments and processing higher volumes of information. Human operators of conventional security systems inevitably face an increasing burden of tedious work that likely causes fatigue, lack of attention, and errors that might compromise safety and security in the surveilled environment. Here, we present a next-generation surveillance system that strives to ameliorate these negative effects. Our system still employs human operators, but eases their job by automating various aspects such as image processing and surveillance-event handling and visualization.

In contrast to other research advances in automated surveillance (e.g., [2]), a particular novelty of our system is the incorporation of autonomous mobile surveillance robots, with the potential to replace human patrols as well as to serve as mobile camera nodes. Each robot embodies an intelligent agent, facing a challenging decision problem of how to act given local and global percepts over the course of a dynamic series of random event occurrences.

We approach the problem of surveillance robot reasoning as one of sequential decision making under uncertainty [4], which we frame in Section 2. Our application of high-level reasoning to a real robotic system is supported by a robust surveillance infrastructure, outlined in Section 3. It leads us to demo the resulting behavior, shown in the following **video link**: `http://users.isr.ist.utl.pt/~jccastillo/aamas_demo/`.

2. EXAMPLE SURVEILLANCE SCENARIO

As an illustration of the capabilities of the system and the challenges of autonomous surveillance, consider a robot surveilling the environment depicted in Figure 1, which comprises several connected rooms and hallways of a building with static cameras mounted at strategic positions. A perpetual objective is to *patrol* the environment, ensuring that the key rooms (e.g, the soccer field, the coffee room, and the elevator hallway) are continually visited. Meanwhile, a *fire* could break out in the coffee room. The robot should respond by entering and using its sensor to assess the fire. Another event that the robot may face is an *assistance request*, wherein a visitor appears from the elevator and waves to the camera. To respond effectively, the robot should meet the visitor and offer guidance to a desired destination. These events constitute additional objectives that should be balanced with patrol according to significance and urgency.

For careful planning of its movements in anticipation of such randomly-occuring events, we equip the robot with a planner that uses Markov Decision Processes (MDPs) [4] to account for event likelihoods; from here it computes a policy that prescribes a high-level movement action for each possible state (i.e., the combination of robot location, patrol state, and state of event detections). Particular challenges that we face in applying decision-theoretic planning include (1) specifying accurate event probabilities, which we approximate by inferring a bounded-parameter model (as detailed our earlier paper [6]), and (2) accounting for false event detections, which we model with partial observations.

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*

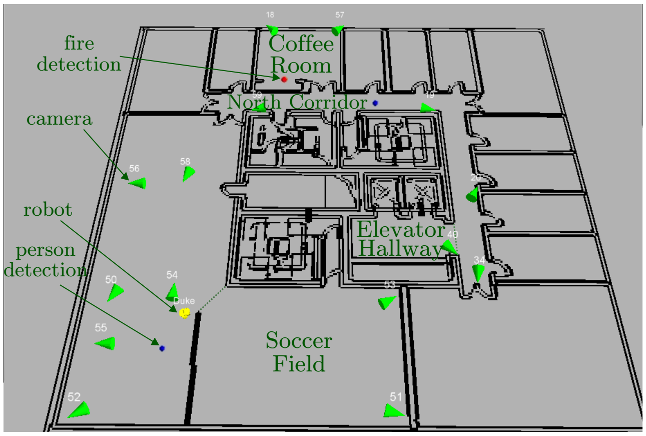

Figure 1: A real physical surveillance environment as visualized using the central command interface.

3. SYSTEM COMPONENTS

To support autonomous surveillance, we have designed our system as a collection of interdependent modules (tied together with ROS [`http://www.ros.org`]), thereby integrating the necessary technologies from multiple disciplines.

3.1 Robot Navigation

The Pioneer 3-AT robot, equipped with a SICK laser range finder, already has an off-the-shelf low-level controller provided by ROS libraries that navigates from waypoint to waypoint on an obstacle-free path from. However, most high-level decision-making methods do not work over a continuous space of locations. Thus, we prescribe a layer of abstraction that partitions our laser map of the environment into rooms and corridors, each with an associated waypoint. We also define high-level movement actions that navigate between spatially-adjacent areas (by invoking the lower-level path planner), as if the robot were transitioning between nodes of a graph. In addition to controlling real robots, our robot navigation controller may be plugged directly into a 2-D Stage simulation, allowing us to simulate robot surveillance movements for the purposes of testing and debugging.

3.2 Image Processing for Event Detection

In our surveillance environment, feeds from a dozen high-resolution network cameras are processed in real time. For the scenario described in this paper, there are two functions at play: **People detection** records $\langle position, timestamp\rangle$ pairs for each human found in the environment; and **waving detection** indicates whether or not he or she is waving [1], though our system could easily interface with additional detection modules that generate other events. Because the vision algorithms are not perfect (e.g., depending heavily on lighting conditions and calibration), and despite our additional development of an event filter, sometimes a waving event is falsely triggered, such that a robot believes there is a person awaiting assistance when really there is not.

3.3 Event Communications

A high volume of data is transmitted across our network over the course of normal operation, so some packet loss is expected. However, in order to ensure robust transmission of critical data, such as events sent to robots, we have developed a novel communication mechanism that employs two different modes of transmission. While most data (e.g., video frames) is transmitted using Reconfigurable and Adaptive Time Division Multiple Access, the critical data is transmitted with a special protocol that tolerates packet loss by employing redundancy [3].

3.4 Robot Planning and Execution

Putting together event detections with our robot location abstraction (Sec. 3.1), we have a discrete representation of state that is compatible with a wide array of planning methods. For their application, we have implemented in ROS a robot state estimator, that maintains the current values of all predicates (location and event-based). For instance, as soon as the robot crosses the doorway between the "north corridor" and the "coffee room", it updates its state. Likewise, we have implemented a policy controller that automatically looks to a plan or policy (in the case of an MDP-based-planner) for the appropriate action, whenever any predicate changes value. These flexible components serve as the interface into which various planning algorithms can be plugged. In particular, we have integrated a suite of existing planners called the *Multiagent Decision Process* (*MADP*) *Toolkit* [5].

3.5 Interface for Human Security Personnel

Operating on a dedicated computer, the control interface, whose map view is shown in Figure 1, provides real-time status of robot positions, robot states, and events. It can be seen as an evolution of conventional security displays in that event positions are visualized directly on a map of the environment. This enables human security personnel to monitor the robots' behavior as well as to interact with the system, for instance by triggering an event that was not otherwise detected or by canceling a falsely-detected event. The interface supports both physical- and simulated-robot interaction.

4. ACKNOWLEDGMENTS

This work was partially supported by the Portuguese Fundação para a Ciência e a Tecnologia (FCT), through strategic funding for ISR/LARSyS under project PEst-OE/EEI/LA0021/2013 and through the Carnegie Mellon-Portugal Program under project CMU-PT/SIA/0023/2009.

5. REFERENCES

[1] P. Moreno, A. Bernardino, and J. Santos-Victor. Waving detection using the local temporal consistency of flow-based features for real-time applications. In *Image Analysis and Recognition*, pages 886–895. Springer, 2009.

[2] V. Onut, D. Aldridge, M. Mindel, and S. Perelgut. Smart surveillance system applications. In *Proceedings of the 2010 Conference of the Center for Advanced Studies on Collaborative Research*, pages 430–432, 2010.

[3] J. C. G. Reis, P. U. Lima, and J. Garcia. Efficient distributed communications for multi-robot systems. In *Proceedings of the RoboCup Symposium*, LNAI. Springer, 2013.

[4] S. Seuken and S. Zilberstein. Formal models and algorithms for decentralized decision making under uncertainty. *JAAMAS*, 17(2):190–250, 2008.

[5] M. T. J. Spaan and F. A. Oliehoek. The MultiAgent Decision Process toolbox: software for decision-theoretic planning in multiagent systems. In *MSDM Workshop*, 2008.

[6] S. J. Witwicki, F. S. Melo, J. Capitán, and M. T. Spaan. A flexible approach to modeling unpredictable events in MDPs. In *Proceedings of ICAPS*, pages 260–268, Rome, Italy, 2013.