Information Fusion 27 (2016) 170–188

Contents lists available at ScienceDirect

Information Fusion

journal homepage: www.elsevier.com/locate/inffus

Particle filter robot localisation through robust fusion of laser, WiFi, compass, and a network of external cameras

A. Canedo-Rodríguez [a,*], V. Álvarez-Santos [a], C.V. Regueiro [b], R. Iglesias [a], S. Barro [a], J. Presedo [a]

[a] CITIUS, University of Santiago de Compostela, R/Jenaro de la Fuente Dominguez s/n, 15782 Santiago de Compostela, Spain
[b] Department of Electronics and Systems, University of A Coruña, Campus de Elviña s/n, 15071 A Coruña, Spain

ARTICLE INFO

Article history:
Received 4 September 2014
Received in revised form 5 February 2015
Accepted 28 March 2015
Available online 12 April 2015

Keywords:
Robot localisation
WiFi localisation
Multi-camera network
Particle filter
Sensor fusion

ABSTRACT

In this paper, we propose a multi-sensor fusion algorithm based on particle filters for mobile robot localisation in crowded environments. Our system is able to fuse the information provided by sensors placed on-board, and sensors external to the robot (off-board). We also propose a methodology for fast system deployment, map construction, and sensor calibration with a limited number of training samples. We validated our proposal experimentally with a laser range-finder, a WiFi card, a magnetic compass, and an external multi-camera network. We have carried out experiments that validate our deployment and calibration methodology. Moreover, we performed localisation experiments in controlled situations and real robot operation in social events. We obtained the best results from the fusion of all the sensors available: the precision and stability was sufficient for mobile robot localisation. No single sensor is reliable in every situation, but nevertheless our algorithm works with any subset of sensors: if a sensor is not available, the performance just degrades gracefully.

1. Introduction

We want to build robots that are able to operate in environments where we live and work, such as hospitals or museums. These environments are rather static, in the sense that layout changes are not frequent (e.g. new walls, furniture movement, etc.). However, the environment conditions are inherently dynamic: there will always be people moving around the robots, the illumination conditions will change depending on the time of the day, etc. Robots must execute their tasks correctly under these conditions, and they must do it on continuous basis. To do this, they must be able to determine their position accurately and robustly at all times (a task known as mobile robot localisation [1]). In order to achieve this, we have designed a localisation system that is able to work robustly in these environments.

In our case, this localisation system will be used by a tour-guide robot that we have developed in the past. This robot can work in different environments, such as museums, large buildings, events or conferences [2]. It can follow humans around the environment [3] and interact with them (via voice and gestures) [4]. This robot can also record routes of interest and verbal explanations from instructors and reproduce them for visitors. We have also designed a network of camera agents (Fig. 1) that detect and inform the robot about situations that require its presence (e.g. people that might need information about the event) [5,6]. This way, our tour guide robot is aware of situations that happen outside of its range of perception. Our purpose is to use this system in events where our robots interact with attendees and offer them information or entertainment.

Taking into account the kind of environments that we are considering, our robot will work in two stages: deployment and operation. During the deployment stage, we will place the camera agents, construct the map of the environment, and calibrate all the sensors. To create the map, we will use Simultaneous Localisation and Mapping (SLAM) techniques [1]. In practice, a functional approach is to assume that we can create the map when there is no people in the environment (e.g. non-working hours), and that this map will not be modified once created. Of course, we could update the map continuously, but taking into account that the working environment does not change significantly, this adds little value while it consumes scarce computational resources. Moreover, continuous SLAM may eventually integrate people around the robot as part of the map, which would lead to failures. For instance, in the experiment of Fig. 2 our robot was following a group of people, and the Gmapping SLAM [7] algorithm integrated their legs as part of the free environment (dots on the corridors).

* Corresponding author. Tel.: +34 881816397.
E-mail address: adrian.canedo@usc.es (A. Canedo-Rodríguez).

http://dx.doi.org/10.1016/j.inffus.2015.03.006
1566-2535/

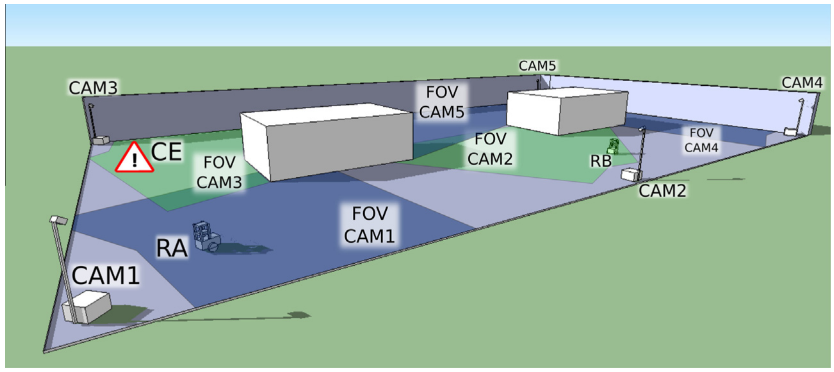

Fig. 1. Example of operation of our system. We can observe 5 cameras (CAM1–CAM5), their Fields of View (FOV CAM1–FOV CAM5), and 2 robots (RA and RB). Camera CAM3 is detecting an event that requires the presence of the robot (call event, CE).

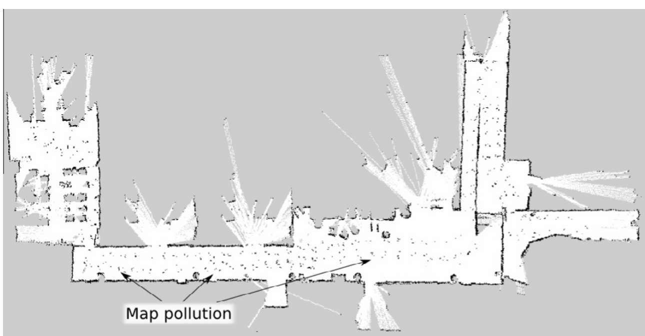

Fig. 2. Laser occupancy map created with the Gmapping algorithm. The map was polluted due to the presence of people around the robot.

These situations are still very challenging for state-of-the-art SLAM [8].

During the operation stage, the robot must navigate towards where the users require it and offer them information, assistance or entertainment. Therefore, the robot must determine its pose (position and orientation) relative to the map that we have created. Most localisation systems rely on the information provided by one sensor, such as sonar sensors [9], 2D lasers [1], 3D lasers [10], cameras mounted on the robot [11], etc. Nevertheless, there is not such a thing as a perfect sensor: each sensor has its limitations, and no sensor is applicable to all situations. Therefore, we cannot expect these systems to respond robustly in every single situation that may happen in the real world. Localisation works have tackled this problem mostly by exploring the use of new sensors and by designing increasingly complex algorithms. In a complementary way, we want to explore the use of multi-sensor fusion techniques [12]: we believe that we can increase the robustness and redundancy of localisation systems by combining the information of sensors of different nature, that will fail in different situations.

Conceptually, we can categorise information sources depending on whether they are better at providing global or local estimates of the robot position.

1. Global estimates. Sources such as wireless localisation or external cameras can provide rough global estimates of the robot position. For instance, when a camera detects a robot, the robot knows that it must be within its FOV. Similarly, when the robot receives a certain signal power from three or more wireless transmitters, the area where this reception is possible can be roughly delimited. In both cases it is unlikely to receive the same measurement in a totally different area, therefore the position estimate will be of high confidence (although it might not be highly accurate).
2. Local estimates. Sources such as 2D/3D laser range finders or on-board cameras might not be able to provide global estimates with great confidence at all times. This is because they might provide similar information on distant areas. For instance, for a camera all the corridors of a hospital might look alike, therefore distinguishing among them might not be a trivial task. On the other hand, these sources usually can achieve highly accurate estimates, specifically when previous estimates of the robot pose are available.

Sensors from the first group can provide a rough estimate and ensure that the robot is not completely lost, while the sensors of the second group can help on achieving higher accuracy. In addition, we can combine both on-board sensors (e.g. lasers or wireless positioning systems) and off-board sensors (e.g. external cameras). Most works combine only the information of sensors on-board the robot. These works assume self-contained, stand-alone robots that do all the sensing, deliberation and action selection on board, based only on their own perceptions. However, on-board and off-board sensors can provide the robot position from a different perspective (e.g. a group of people might oclude the vision of an on-board laser but not the vision of an external camera, and vice versa). This links with recent advances in Ubiquituous Robotics [13] and Networked Robots [14], which have demonstrated the benefits of coordinating robots and sensor networks. Moreover, since on-board and off-board sensors are independent, their fusion increases the redundancy of the system.

In this paper, we propose a multi-sensor fusion algorithm for mobile robot localisation based on particle filters [1]. Our algorithm combines sensors that provide coarse global position estimates and sensors that provide accurate local position estimates. Moreover, we will combine both on-board and off-board sensors. With these requirements, we have chosen a reasonably low-cost set of information sources consisting on a 2D laser range finder, a WiFi positioning system designed by us, a magnetic compass and a network of USB webcams. We expect the WiFi positioning system and the cameras to provide a coarse estimation of the robot position, and to ensure that the robot never gets completely lost. Meanwhile, we expect the laser and the compass to refine these estimates and provide sufficient accuracy. We would like to remark that our goal is not to propose yet another algorithm for mobile robot localisation, but to demonstrate the benefits of multi-sensor fusion in a real context. Our main contributions are: