Available online at www.sciencedirect.com

ScienceDirect

IFAC Papers Online
CONFERENCE PAPER ARCHIVE

IFAC-PapersOnLine 49-15 (2016) 120–125

Autonomous Apartment Exploration, Modelling and Segmentation for Service Robotics

Guido Manfredi * **Sandra Devin** * **Michel Devy** *
Daniel Sidobre **

* *LAAS-CNRS, 7 avenue du colonel Roche, F-31400 Toulouse, France (e-mail: michel@laas.fr).*
** *LAAS, UPS, Universite de Toulouse, F-31400 Toulouse, France.*

Abstract: This work proposes a full pipeline for a robot to explore, model and segment an apartment from a 2-D map. Viewpoints are found offline and then visited by the robot to create a 3-D model of the environment. This model is segmented in order to find the various rooms and how they are linked (windows, doors, walls) yielding a topological map. Moreover areas of interest are also segmented, in this case furniture's planar surfaces. The method is validated on a realistic three rooms apartment. Results show that, despite occlusion, autonomous exploration and modeling covers 95% of the apartment. For the segmentation part, 1 link out of 14 is wrongly classified while all the existing areas of interest are found.

Keywords: Autonomous, Exploration, Modeling, Segmentation, Service Robotics.

Fig. 1. The ADREAM apartment is used to illustrate the different algorithms and to perform the experiments described later. Note that the furniture visible in this view is slightly different from the real setup.

1. INTRODUCTION

Let us consider a robot designed to help humans at home. It faces an unknown environment, regularly evolving and always unpredictable. In order to operate in such an environment, the robot has to first explore and model it. This model can be a two dimensional map, e.g. for navigation, or a three dimensional map, e.g. for full body motion planning or scene understanding. When a model of the environment is available, given by the user or created by the robot, it must be segmented to extract places and areas of interest. Indeed, typical requests from the user will be "fetch OBJECT from PLACE in/on AREA", where PLACE and AREA are room and furniture names. This implies segmenting the environment model in various parts and associating spacial volumes with semantic information. On top of that, humans may lack precision, so a user could just ask "fetch OBJECT from AREA", omitting the place. In order to handle the lack of information, the environment model segmentation has to provide areas-places links.

The present work focuses on robots using directional sensors in indoor scenarios. The next section shows that in this case most existing works propose online exploration and modeling methods. To the best of our knowledge, there are no offline methods. In section 3, we propose an offline method to find the best 3-D viewpoints to explore a site. Then, a segmentation method based on local and global cues is presented to build a topological map of the site with the different places and how they interconnect. Areas of interest are also extracted and associated to the different places. Finally, section 4 describes an experiment in a real environment to demonstrate the possibility of autonomously discovering areas, places and how they are linked in an initially unknown environment. The results are discussed and error sources identified. The algorithms described hereafter are illustrated on the ADREAM apartment (Figure 1).

2. EXISTING WORK

As analysed in Aouina et al. (2014), an environment can be modeled at various levels: the geometric level, based on features; the topological level, based on views; the semantic level, based on objects and places. Though different, all these problems can be expressed using the SLAM formalism. In Rusu et al. (2009) the authors show good localization performances by using a SLAM approach where landmarks are known fixed objects and triangulated surfaces dynamically acquired. The authors of Nüchter and Hertzberg (2008) recommend the use of a semantic map to enable the robot knowledge to be reviewable and communicable. They use a 3-D laser scanner to acquire a 3-D map through 6-D SLAM. The SLAM considers coarse features, like walls, and finer features, like objects. The work of Aouina et al. (2014) shows how to decouple the construction of a localization model and of a dense 3-

Peer review under responsibility of International Federation of Automatic Control.
10.1016/j.ifacol.2016.07.719

D map. The localization is performed with a 2-D laser range finder while the 3-D modeling is done with a tilting laser. In the present work, two problems are tackled: site exploration and modeling, i.e. automated construction of a volumetric map, and segmentation of the site model into meaningful parts, i.e construction of a semantic map. We provide a quick review of the state of the art for these problems.

Though this task can appear similar to the modelling of an unknown object by a mobile sensor, like in Amigoni and Caglioti (2010), it is different. In this problem the site may be cluttered by objects preventing the robot from accessing some viewpoints. This problem is also different from the art gallery problem (Blaer and Allen (2009)). In the art gallery problem, guards have infinite view range and field of view; this is not the case for a robot. It is also different from a dense reconstruction (Newcombe et al. (2011); Arbeiter et al. (2012)) as we try to limit the robot motion by minimizing the number of viewpoints needed. According to Amigoni and Caglioti (2010) exploration strategies are divided among three types: fixed trajectories, random movements and observation positions. The first approach uses precomputed trajectories to explore any kind of site (Taylor and Kriegman (1993)). Though easy to implement, these methods do not adapt to the site's specificities and can fail for some geometrical configurations. In the random movements approach of Freda and Oriolo (2005), random points or trajectories are chosen and the robot explores them. This kind of approach has already been rather successfully applied to vacuum cleaner robots, see Tribelhorn and Dodds (2007). However, it suffers from the trade off between number of random draw, i.e. time spent, and quality of the coverage. Finally, the last type of methods determine the best viewpoints to visit depending on some constraints. Because they are adaptive and robust to the geometry of the site, we focus on these approaches. In Reed and Allen (2000), the authors start by acquiring manually a rough estimate of the environment. Then, the model is completed automatically. This method relies on a geometrical approach with volumes representing visibility, mobility and occlusion constraints. In Blaer and Allen (2009), the three dimensional exploration is initialized with the information of a two dimensional map. A voxel grid is filled with *empty*, *occupied* and *unknown* cells. Then, a greedy algorithm is used to select from a set of random viewpoint the one seeing the maximum number of *unknown* voxels. Ray casting is used to test visibility of every *unknown* voxel. A similar approach has been demonstrated at LAAS in Albalate et al. (2002) in the framework of the European project CAMERA. A voxel grid is progressively filled with voxels being classified as *unknown*, *empty*, *occupied*, *occluded*, *occplane* (occluded but adjacent to an empty voxel) and *border* (on the border of the line of sight). The next best view is selected based on the number of visible *unknown* voxels and on an estimation of the number of *occluded*, *occplane* and *border* voxels discovered. Finally, in Amigoni and Caglioti (2010), the authors develop a probabilistic framework based on information theory to choose the next best view according to given constraints like travel distance and expected information. In these four works, the aim is to find the best view, acquire it and search for the next best view, and so on.

In the context of this work, the robot has already many critical processes running (motion planning, human perception, actuators control, task planning, etc.) and limited computational power. Site exploration is not a critical activity, so instead of taking processing power, we advocate in favor of doing this offline, e.g. when the robot is idle. Thus the process finds all the good viewpoints at once and they can be explored when the apartment is quiet and not being modified, for example at night. Such method is proposed in the first part of section 3.

Once a 3-D map of the environment is available, the robot should segment it into meaningful parts. This segmentation depends on various criteria. In Wurm et al. (2008), the authors propose a segmentation algorithm for multi-agent exploration. They segment a 2-D map geometrically with a Voronoi Graph (VG) (Choset and Burdick (2000)). The graph is then partitioned by separating clusters at junction nodes which are: local minima, at least of degree 2 (two edges), with at least a neighbor of degree 3 and that lead from unknown to known areas. The work from Holz et al. (2010) elaborates on this method by changing the conditions to choose a critical node. The node must be: close to a Voronoi site, of degree 2, adjacent to a junction node or adjacent to a node adjacent to a junction node (2^{nd} degree adjacency). These modifications provide a better representation of locations such as doorways. In both cases, the segmentation is based on a geometrical criterion. A more comprehensive survey on the subject is available in Bormann et al. (2016)

However, the present work aims at a human representation of the site. In particular, we want to discover the location of rooms and how they are connected by doors and windows. This allows building a graph representation of the site where rooms are vertices and windows and doors are represented by edges. The segmentation method is shown in the second part of Section 3. The resulting graph with the rooms and their connectivity is of crucial importance for tasks like object search (see Rogers and Christensen (2013)) or learning areas-place context.

3. EXPLORATION, MODELLING AND SEGMENTATION

To deal with an unknown environment, the first step is to explore it and represent it as a model. In this work, the entities of interest are rooms and areas, so the second step is finding the different rooms and areas from the model. These two steps are described hereafter.

3.1 Exploration and Modelling

The exploration step searches for a set of observation points from which the environment can be modeled. Then, the modeling process aggregates the point clouds found at each observation point and handles occlusions. In the following, it is assumed that the robot has a 2-D map of the environment for localization purposes and is able to localize itself in the world. It also assumed that the environment is not heavily cluttered so that the robot can move around and observe most of the room.

For the exploration step, the world is considered flat, so there is no line of sight occlusion. Starting from the 2-