CMU's 3D Multi-Robot Modeling Project

Sebastian Thrun, Reid Simmons, and various other members of the Lab

Problem: CMU's 3D Modeling Project seeks new ways of generating highly accurate and complete models of indoor and outdoor environments with mobile robots. The problem of generating 3D models is challenging, since multiple sources of error all interfere: from robot odometry error to error in cameras and range sensors. The goal of this project is to devise a statistical framework that can correct for these different types of errors, and generate accurate, detailed 3D models of indoor environments.

The 3D Modeling Project also studies the problem on online multi-robot exploration and data fusion. When robots explore so as to build a model, they have to make decisions as to where to move and where to sense. This problem is significantly harder in 3D than it is in 2D, Building a complete and accurate model in 3D poses a challenging optimization problem.

Impact: The ability to generate 3D models would enable people to inspect and "fly through" an environment without actually entering it. Just imagine you could take a virtual, interactive walk through any major building in the world! Accurate 3D models may play a vital role in future tele-presence and tele-medicine projects, in the fields of real-estate and architecture, in video-games, and it might also increase the awareness of the military, who funds this research through DARPA.

State of the Art: Current mobile robot modeling algorithms almost exclusively operate in 2D. The problem of building models is usually referred to as the *concurrent mapping and localization* problem, to indicate its chicken-and-egg nature. Building a map with *noise-free* odometry is relatively easy, as is the opposite: estimating the location of a robot when an accurate map is available. In combination, however, the problem is hard.

The literature has devised a range of approaches. To of the most popular ones are SLAM and EM. SLAM algorithms [4, 12] use Kalman filters for estimating the robot pose and the map together. However, instead of generating a full map, mathematical limitations of Kalman filtering forces them to extract a small number of *uniquely identifiable* features, and determine the location of those. EM [10, 14] approaches can cope with more difficult modeling problems, but they are offline and computational very complex. 3D modeling algorihtms for robots can be found at [6, 8, 9, 13]. 3D scene reconstruction has been studied in the computer vision and graphics literature [1, 2, 5, 7], but often manual intervention is required for generating 3D models.

Approach: Our approach is probabilistic: Modeling is formulated as a posterior estimation problem, in which the map and the robot pose is estimated as a posterior conditioned on the available data. Since posterior estimation in the general case is computationally intractable, we are investigating a range of approximations that can reduce the computational complexity while maintaining robustness and accuracy in practice. Current results suggest that maps and poses can be estimated with high accuracy in real-time, under limiting conditions such as that the initial poses of robots relative to each other are known. See also [13].

Our exploration approaches are also probabilistic [3, 11]. Exploration is formalized as maximizing expected negative entropy (knowledge gain), and we have devised a greedy algorithm for partitioning the exploration task between multiple robots. In future research, we intend to develop non-greedy algorithms that can plan over sequence of exploratory actions.

We are equally interested in finding learning algorithms that map high-dimensional sensor data to lower dimensional 3D maps. Our existing 3D maps (see, e.g., Figure 1) are composed of tens of thousands of small

Figure 1: (a) 3D map in VRML format for interactive data display. (b) Team of robots that concurrently explore and map a building. (c) 2D map acquired by autonomous robot team, along with the paths of the exploring robots.

---

polygons. Our latest approach is to apply variants of the EM algorithm to extract models that are orders of magnitude less complex, composing maps of generic features such as walls, doors, cylindrical objects, and so on [8].

Future Work: Future work might include the development of modeling algorithms that can cope with changing environments, and the development of more efficient probabilistic control and navigation algorithms for mobile robots.

Selected References:

[1] P.K. Allen and Ioannis Stamos. Integration of range and image sensing for photorealistic 3D modeling. 2000.

[2] R. Bajcsy, G. Kamberova, and Lucien Nocera. 3D reconstruction of environments for virtual reconstruction. In *Proc. of the 4th IEEE Workshop on Applications of Computer Vision*, 2000.

[3] W. Burgard, D. Fox, M. Moors, R. Simmons, and S. Thrun. Collaborative multi-robot exploration. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, San Francisco, CA, 2000. IEEE.

[4] G. Dissanayake, H. Durrant-Whyte, and T. Bailey. A computationally efficient solution to the simultaneous localisation and map building (SLAM) problem. ICRA'2000 Workshop: Mobile Robot Navigation and Mapping, April 2000.

[5] S.E. Hakim and P. Boulanger. Sensor based creation of indoor virtual environment models. In *Proc. of the 4th Internetional Conference on Virtual Systems and Multimedia (VSMM)*, 1997.

[6] M. Bajracharya L. Iocchi, K. Konolige. Visually realistic mapping of a planar environment with stereo. In *Proceesings of the 2000 International Symposium on Experimental Robotics*, Waikiki, Hawaii, 2000.

[7] M. Levoy. The digital michelangelo project. In *Proc. of the Second Int'l Conference on 3D Imaging and Modeling*, 1999.

[8] Y. Liu, R. Emery, D. Chakrabarti, W. Burgard, and S. Thrun. Using em to learn 3d models with mobile robots. Submitted for publication, 2000.

[9] H.P. Moravec and M.C. Martin. Robot navigation by 3D spatial evidence grids. Mobile Robot Laboratory, Robotics Institute, Carnegie Mellon University, 1994.

[10] H Shatkay and L. Kaelbling. Learning topological maps with weak local odometric information. In *Proceedings of IJCAI-97*. IJCAI, Inc., 1997.

[11] R. Simmons, D. Apfelbaum, W. Burgard, M. Fox, D. an Moors, S. Thrun, and H. Younes. Coordination for multi-robot exploration and mapping. In *Proceedings of the AAAI National Conference on Artificial Intelligence*, Austin, TX, 2000. AAAI.

[12] R. Smith, M. Self, and P. Cheeseman. Estimating uncertain spatial relationships in robotics. In I.J. Cox and G.T. Wilfong, editors, *Autonomous Robot Vehnicles*, pages 167–193. Springer-Verlag, 1990.

[13] S. Thrun, W. Burgard, and D. Fox. A real-time algorithm for mobile robot mapping with applications to multi-robot and 3D mapping. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, San Francisco, CA, 2000. IEEE.

[14] S. Thrun, D. Fox, and W. Burgard. A probabilistic approach to concurrent mapping and localization for mobile robots. *Machine Learning*, 31:29–53, 1998. also appeared in Autonomous Robots 5, 253–271.