# Statistical Methods for Learning Maps with Mobile Robots

**Sebastian Thrun (PI), Dieter Fox, Wolfram Burgard, Steffen Gutmann, and Ben Kuipers**

**Problem:** Mapping is the problem of generating models of a mobile robot's environment based on sensor data. Based on a stream of measurements from the robot's sensors (shaft encoders, laser, sonar), the goal of mapping is to determine the location of various entities, such as landmarks or obstacles, in a global frame of reference. To generate consistent maps of large-scale environments, the robot also has to solve a concurrent localization problem, which arises because robot odometry is usually erroneous. Thus, the problem addressed in our research is a chicken-and-egg problem: Mapping is considerably simple if the robot's location is known; and localization is considerably simple if a map is readily available; but in combination the problem is hard.

**Impact:** Most successful navigation algorithms require the availability of maps, yet we lack general and scalable methods for automatically generating such maps. Such methods would advance the state of the art in two directions. First, they would facilitate the installation of mobile robots in indoor environments and enable robots to adapt to changes therein. At present, maps are usually generated manually, which involves significant time and money. Second, this research, if successful, can become a crucial step towards true robot autonomy. The methods investigated in this research can be combined with existing exploration and path planning strategies (in fact, we have done this in our lab), enabling robots to autonomously explore and build maps of large-scale environments.

**State of the Art:** Over the past two decades, the problem of building maps in indoor environments has received significant attention in the mobile robotics community. Previous approaches suffer from scaling limitations. Many of them make highly restrictive assumptions on the nature of the environment and the sensor input. With very few exceptions, previous approaches have not been able to map cyclic environments. They were often only tested in simulation.

**Approach:** Our approach is pervasively statistical in nature. It phrases the problem of concurrent mapping and localization as a *constrained maximum likelihood estimation problem*, where the objective is to find the most likely map given the data. With appropriate independence assumptions, the map likelihood function can be shown to be a function of two basic probabilistic models, a *motion model*, and a *perceptual model*. The motion model describes, using the language of probability theory, the effect of actions on a robot's *pose*, where pose comprises location and orientation. The perceptual model describes possible sensor readings for situations where both the map and the robot's pose are known. For mobile robots equipped with proximity sensors, both models are easily derived from the robot's physical properties.

Our approach applies the EM algorithm to find the most likely map under the data. EM is a well-known, efficient statistical algorithm for maximum likelihood estimation in high-dimensional spaces. In the context of our mapping problem, EM iterates two different steps, a *localization step* and a *mapping step*. In the localization step, the map is fixed, and our approach computes probabilities over the robot's various poses in the data set. Since the map is assumed to be known, this step is a generalized version of Markov localization, a well-known localization method. In the mapping step, our approach computes the most likely map based on the pose estimates generated in the localization step. Since here the robot's poses are fixed, generating the most likely map is mathematically straightforward.

Our approach, thus, tackles the concurrent mapping and localization problem by solving two simpler problems: the problem of localization with a fixed map, and the problem of mapping under knowledge of

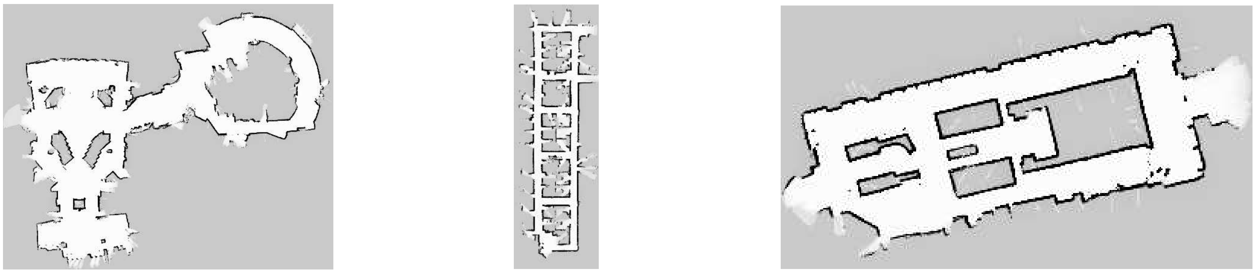

Figure 1: Maps generated at the Smithsonian Museum of American History (100 meters), a large office building (50 meters), and the Carnegie Museum of Natural History (40 meters).

the robot's poses. An interesting theoretical result is that the iteration of these two steps leads to a local maximum of the map likelihood function. This maximum usually coincides with the global maximum, and thus the most likely map.

This statistical approach has been implemented with two different representation; one in which all probabilities are approximated using piecewise constant functions, and one in which they are approximated using Normal distributions (Kalman filters). The first approximation is able to represent multi-modal distribution, hence can deal with global ambiguities that often arise in large-scale and cyclic environments. The second approximation can estimate the location of obstacles with floating-point accuracy, thus yields highly accurate maps. In combination, these methods have been found to generate maps whose size and accuracy are orders of magnitude better than maps generated with previous methods. Figure 1 shows examples of maps generated in environments of sizes up to 100 by 70 meters. In one case, the cumulative odometric error was larger than 70 meters.

**Future Work:** Our current approach is not able to generate and modify maps in real-time, since it is really meant as an off-line method. We envision to extend this method so that only a limited history of past sensor readings are memorized. We believe that the extension to multi-robot mapping is straightforward. It is desirable to address the fact that environment may change in a statistically sound way. We also intend to develop a statistical extension of the approach for autonomous exploration. We also plan to use the statistical framework for building large-scale visual maps from camera images.

**Selected References:**


[1] S. Thrun. Exploration and model building in mobile robot domains. In E. Ruspini, editor, *Proceedings of the ICNN-93*, pages 175–180, San Francisco, CA, March 1993. IEEE Neural Network Council.

[2] S. Thrun. Learning maps for indoor mobile robot navigation. *Artificial Intelligence*, 1998. to appear.

[3] S. Thrun, D. Fox, and W. Burgard. A probabilistic approach to concurrent mapping and localization for mobile robots. *Machine Learning and Autonomous Robots (joint issue)*, 1998. to appear.

[4] S. Thrun, D. Fox, and W. Burgard. Probabilistic mapping of an environment by a mobile robot. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 1998. to appear.

[5] S. Thrun, S. Gutmann, D. Fox, W. Burgard, and B. Kuipers. Integrating topological and metric maps for mobile robot navigation: A statistical approach. In *Proceedings of AAAI-98*. AAAI Press/The MIT Press, 1998. to appear.