STATISTICAL METHODS FOR COOPERATIVE MULTI-ROBOT MAPPING

Principal Investigator: Sebastian Thrun

Center for Automated Learning and Discovery
Carnegie Mellon University
Pittsburgh, PA 15213-3891

Phone: (412) 268-8077
FAX: (412) 268-5576
Email: thrun@cs.cmu.edu

PROJECT SUMMARY

We propose research on a new family of statistical algorithms for mobile robot mapping of indoor environments. The problem of cooperative multi-robot mapping is the problem of generating consistent maps of unknown environments, based on sensor data acquired distributedly on multiple mobile robot platforms. Since robot sensors and robot odometry are erroneous, and since the robots might be unaware of their initial position relative to each other, the mapping problem implies a localization problem, that is, a problem of determining each robots' pose (location and heading direction) relative to each other and to the map. In addition, the mapping problem also implies a distributed exploration problem.

Cooperative mobile robot mapping is an important problem in the area of robotics. This is because a large number of successful mobile robot systems (including several commercial systems) navigate using maps; yet, we still lack methods that enable teams of robots to acquire maps in reasonably large environments and to adapt to changes therein.

The project focuses on a new, statistical framework for concurrent mapping and localization with teams of robots. The framework poses the problem of building maps as a maximum likelihood problem, of finding the most likely map given the data acquired by a team of robots. Exploration is posed as a maximum information gain problem, of moving and sensing so as to maximally reduce the robots' uncertainty. Fast statistical methods, such as the well-known EM algorithm, are employed to solve the various estimation and maximization problems while the robots are in motion. The project places additional emphasis on extensions of the basic framework, that will

- minimize computational and communication requirements (memory, time, bandwidth), by using fast statistical estimation techniques (in particular: EM and sampling-based methods), and by utilizing compact, object-centered representations of space.
- efficiently coordinate teams of robots when exploring unknown environments, by applying information-theoretic measures of information gain, and by developing a communication scheme to communicate "intent."
- rapidly adapt to changes in the environment, by invoking exponential filters that decay information over time, and by focusing exploration on parts of the environment where the map is inconsistent.
- accommodate communication bottlenecks, by developing flexible inter-robot communication protocols that minimize the amount of information transmitted between the robots, and that are robust to temporary radio link failures.

We have already implemented a limited (single-robot) proto-type of our new approach. Initial results, presented in the body of this proposal, demonstrate unprecedented scalability to large-scale environments. The research proposed here, if successful, will enable teams of robots to explore unknown environments in a coordinated fashion. They will generate maps that are at last an order of magnitude larger, more accurate, more compact, and more up-to-date than what previous methods were able to accomplish. Albeit from developing the basic framework and methods, this project seeks to materialize our claims and document them through thorough systematic empirical verification.

The results of this research will be disseminated thorough publications, software sharing, and educational means.

PROJECT DESCRIPTION

1 Introduction

1.1 Problem

This proposal addresses the problem of building maps of indoor environments with teams of mobile robots. The problem of mapping is the problem of determining, from sensor data, the locations of *quantities-of-interest* (obstacles, landmarks, walls, objects, etc.) in a global frame of reference (e.g., a Cartesian coordinate system). A variety of factors make this problem challenging:

1. **Perceptual limitations.** Sensor measurements are typically corrupted by noise, making it imperative to integrate and resolve possible conflicts between multiple sensor measurements when building a map. The perceptual range of most sensors is limited to a small range close to the robot. To acquire maps of large-scale environments, robots have to explore their environments collaboratively.
2. **Drift and slippage.** Robot motion is inaccurate. Odometric errors accumulate over time. Thus, the problem of mapping implies a localization problem, that is, the problem of determining a robot's location relative to past locations.
3. **Multi-robot alignment.** If multiple robots collaboratively map an environment, there exists an additional localization problem, which is the problem of determining the location of each robot relative to all other robots. In the most difficult version of cooperative multi-robot mapping (which this proposal addresses), the robots may not know their initial location relative to each other, requiring them to localize themselves under *global uncertainty* while acquiring a map.
4. **Dynamic environments.** Environments change constantly; yet, virtually all existing mapping methods assume static environments.

The problem of concurrent mapping and localization is generally acknowledged as one of the most significant "open" problems in mobile robotics [BEF96, Ren93, Cox91]. A large number of successful mobile robot architectures are based on maps (see various papers in [CW90, Sim95, KBM98]), yet we currently lack methods for automatically building maps that scale up to large environments. Thus, the ability to generate large-scale maps will almost certainly lead to more capable autonomous robotic systems, and reduce the effort required for installing a mobile robot in a new environment.

1.2 Deliverables

The central deliverable is a family of statistical algorithms that enables a team of robots to collaboratively explore an unknown environment and acquire a single, unified geometric map. Assuming that this research is successful, these algorithms will advance the state-of-the-art by enabling teams of robots

- *to build maps of environments that are an order of magnitude larger than maps built by previous methods*, through using novel, statistical methods for mapping and coordinated exploration,
- *to generate maps whose resolution will an order of magnitude higher than what current methods are able to generate*, by replacing previous grid-based representation through object-centered vector representations, and
- *to adapt to changes (dynamics) in the environment*, by using efficient exponential filters and entropy-based exploration methods that permit rapid adaptation to changes in the environment,

thereby overcoming problems with perceptual limitations, sensor noise, drift and slippage, and environment dynamics.

1.3 Practical Impact

We envision that this research will have an impact on a variety of upcoming mobile robot applications, such as:

- **Search and rescue:** Robots that assess the situation and find survivors in buildings damaged by catastrophes such as earthquakes.

- **Service applications:** Robots that, for example, clean large buildings such as department stores.
- **Health-care:** Robots that assist elderly and disabled people in assisted living facilities, where they perform functions such as guidance, find-and-fetch, surveillance, mobile manipulation, etc.

The research proposed here, if successful, will help reduce the costs of deploying mobile robots, by generating maps at almost no costs and without the need for human intervention. At present, installation costs are a significant fraction of the overall costs of a mobile robot system (Helpmate, Cybermotion, etc.). In addition, existing methods cannot adapt well to changes in the environment, which occur regularly. The research proposed here seeks to facilitate long-term operation of mobile robots in dynamic environments.

The PI closely cooperates with a leading mobile robotics company, which commercializes robotic solutions for the application domains listed above. An endorsement letter has been enclosed.

2 Approach

The focus of this project is on a new approach to concurrent mapping, localization and exploration for teams of mobile robots. Our approach employs efficient statistical estimation techniques to estimate both, the map and the robots' locations (current and past). All knowledge will be represented in form of probability density functions. For example, a robot's pose (a pose comprises the x-y-location of a robot and its heading direction, θ) is expressed by a probability density function that assigns to every possible pose a probability that reflects its plausibility under the sensor data. Maps, too, are represented probabilistically. In the most simple case (for which we have an implemented single-robot proto-type), a map assigns to each x-y location probabilities for the various possibilities (e.g., free, occupied, part-of-landmark, etc.). Alternatively, a map is a probability distribution over objects in the world, their size, shape, orientation, location, etc. Due to the pervasive use of probabilistic representations, our approach is never certain about the state of the world. Instead, it maintains multiple beliefs, weighted by a numerical probability factor. The advantage of probabilistic representations over conventional ones, which are used in almost all existing approaches, lies in their increased robustness to odometric errors, sensor/control noise, and changes in the environment. In addition, these methods are well-suited to solve global localization problems that arise if initially, the robots do not know their location relative to each other. They provide effective ways to coordinate multi-robot exploration.

2.1 Statistical Foundations

Our approach formulates the problem of mapping as a statistical *maximum likelihood estimation problem* [CB90, TFB98a]. Suppose N robots explore an environment collaboratively. To generate a map, we assume that each robot is given a stream of data, denoted

$$d_n \quad = \quad \{o_n^{(1)}, u_n^{(1)}, o_n^{(2)}, u_n^{(2)}, \ldots o_n^{(T-1)}, u_n^{(T)}, o_n^{(T)}\}, \tag{1}$$

where $o_n^{(t)}$ stands for an *observation* that the n-th robot made at time t, and $u_n^{(t)}$ for an action command that this robot executed between time t to time $t+1$. T denotes the total number of time steps in the data, i.e., $t = 1, \ldots, T$, and $n = 1, \ldots, N$. Without loss of generality, we assume that each robot's data set d_n is an alternated sequence of actions and observations. Let $d := \{d_1, \ldots, d_n\}$ denote the set of all data.

In statistical terms, the problem of mapping is the problem of finding the most likely map given the data. *Maps* will be denoted by $m = \{m_{x,y}\}_{x,y}$. A map is an assignment of "properties" $m_{x,y}$ to each x-y-locations in the world.

The literature actually distinguishes two paradigms for mobile robot mapping: *metric* [BK91, CL85, Elf87, Elf89, Mor88] and *topological* [EM92, KW94, KB91, Mat90, PK94, Tor94, YB96, Zim96]. In topological approaches to mapping, the properties-of-interest are usually locations of landmarks [CKK95] or, alternatively, locations of significant places [KB91, Cho96]. Metric approaches, on the other hand, usually use the location of obstacles as properties-of-interest [CL85, Mor88, LM97, Thr98b]. Our statistical framework encompasses both (see [TGF+98]).

Our approach assumes that each robot is given two basic, probabilistic models, one that describes the robot kinematics (robot motion), and one that models its perception.

- **The motion model**, denoted $P(\xi'|u, \xi)$, describes the probability that the robot's pose is ξ', if it previously executed action u at pose ξ. Here ξ is used to denote a pose, where *pose* refers to the x-y-location of a robot together with

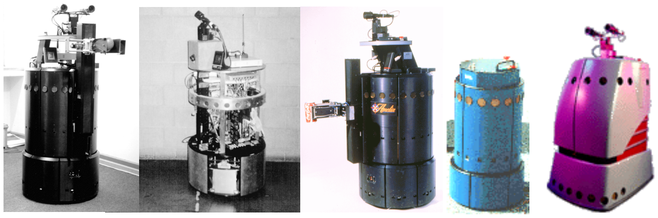

Figure 1: The robots used in our research: RHINO, XAVIER, and AMELIA, RWI B14 robot, and MINERVA.

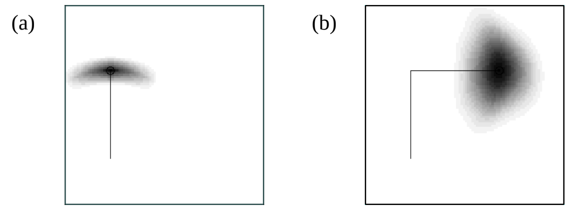

Figure 2: Probabilistic model of robot motion: Accumulated uncertainty after moving (a) 40 meter, (b) 80 meter.

its heading direction. Figure 2 illustrates a probabilistic motion model, by showing the probability distribution for ξ' upon moving as shown (40 and 80 meter). Notice that in these and other diagrams, poses are projected into x-y-space (the heading direction is omitted).

- **The perceptual model**, denoted $P(o|m,\xi)$, models the likelihood of observing o in situations where both the world m and the robot's pose ξ are known. For low-dimensional sensors such as proximity sensors, a perceptual models can readily be found in the literature [BFHS96, Mor88]. Figure 3 illustrates a probabilistic model of robot perception for a planar 2D laser range finder. Figure 3a shows a laser scan o and a map m, and Figure 3b shows the likelihood $P(o|m,\xi)$ for different poses ξ. The darker a pose, the more likely it is under this observation. As can be seen there, the laser scan determines that with high probability the robot is somewhere in the main corridor. Other poses are less likely. The exact pose is, of course, not known.

These three quantities—the data d, the motion model $P(\xi'|u,\xi)$, and the perceptual model $P(o|m,\xi)$—form the statistical basis of our approach, from which everything else follows. Without loss of generality, we assume that all robots use the same motion model and perceptual model (the extension to different motion models and perceptual models is straightforward, but complicates the notation).

2.2 The Map Likelihood Function

In statistical terms, the problem of mapping is the problem of finding the most likely map given the data

$$m^* \quad = \quad \operatorname*{argmax}_m P(m|d). \tag{2}$$

As shown in detail in [TFB98a, TFB98b], under appropriate assumptions the likelihood function is equivalent to

$$P(m|d) \quad = \quad \alpha \int \cdots \int \prod_{n=1}^{N} \prod_{t=1}^{T} P(o_n^{(t)}|m,\xi_n^{(t)}) \prod_{t=1}^{T-1} P(\xi_n^{(t+1)}|u_n^{(t)},\xi_n^{(t)})\ d\xi^{(1)}, d\xi^{(2)}, \ldots, d\xi^{(T)}. \tag{3}$$

where α is a constant (normalizer) which can safely be ignored when maximizing $P(m|d)$. The variable $\xi^{(t)}$ denotes the robot's pose at time t. Formally, Equation (3) makes a *Markov* assumption [Chu60, Put94], which states that *noise* in perception and motion are *independent* random variables. This Markov assumption is made throughout the literature on mapping, localization, and exploration (it is usually implicit).

Equation (3) demonstrates that the maximization problem is well-defined, since $P(m|d)$ is exclusively a function of the data d, the perceptual model $P(o|m,\xi)$, and the motion model $P(\xi'|u,\xi)$. Unfortunately, maximizing (3) is computationally challenging. This is because finding the most likely map involves search in the space of all maps. For the size environments considered here, this space often has 10^6 dimensions or more. To make matters worse, the evaluation of a single map would require integrating over all possible poses of all robots at all points in time, which can easily require integration over more than 10^6 variables.

2.3 Efficient Estimation

Fortunately, there exists an efficient technique for hill-climbing in likelihood space: the *EM algorithm* [DLR77], which in the context of Hidden Markov Models is often referred to as *Baum-Welch* or *alpha-beta algorithm* [RJ86]. EM is a hill-climbing routine in likelihood space, which alternates two steps, an *expectation step* (E-step) and a *maximization*

(a) laser scan and map (b) probability distribution for different poses

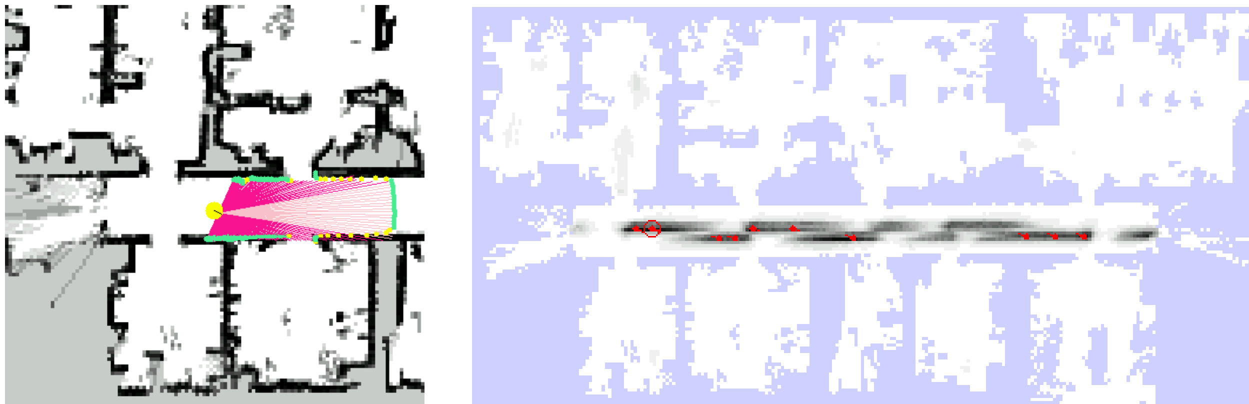

Figure 3: Probabilistic model of perception: (a) Laser range finder scan, projected into a map built previously by the same robot. (b) The sensor scan constraints robot poses probabilistically. Shown here is the probabilistic model of laser scans, $Prob(scan \mid pose, map)$, projected into 2D. The darker a pose, the more likely it is. Based on a single sensor scan, the robot assigns high likelihood for being somewhere in the main corridor.

---

step (M-step). In the context of robot mapping, these steps correspond roughly to a localization step and a mapping step (see also [KS96, SK97, OHD97]):

1. In the E-step, each robot computes (locally) probabilities $P(\xi_n|m,d)$ for its poses ξ_n at the various points in times, based on the currently best available map m (in the first iteration, there will be no map).
2. In the M-step, the society of robots determine the most likely map by maximizing argmax$_m$ $P(m|\xi,d)$, using the location estimates computed in the E-step. To do so, they first compute maps locally, which are then combined using a simple multiplicative rule.

The E-step corresponds to a localization step with a fixed map, whereas the M-step implements a mapping step which operates under the assumption that the robots' locations (or, more precisely, probabilistic estimates thereof) are known. Iterative application of both rules leads to a refinement of both, the location estimates and the map. As a side effect, the robots localize themselves relative to each other, and they also identify errors in their odometry. Thus, the multi-robot localization problem (*where is robot A relative to robot B?*) is solved indirectly, by statistically "comparing" data collected by the different robots when constructing the most likely map (and localizing each robot therein).

EM is fast! In all our experiments, EM converged to the "right" solution in 3-4 iterations—which is much faster than a comparable gradient descent scheme would be. Like any approach based on EM, however, our approach is a hill-climbing procedure that does not provide a guarantee of global optimality; given the complexity of the problem, however, it is unclear whether a computationally feasible and globally optimal routine exists at all.

2.3.1 The E-Step

In the E-step, each robot uses the current-best map m along with its data d_n to compute probabilities $P(\xi_n^{(t)}|d,m)$ for poses at times $t = 1,\ldots,T$. With appropriate assumptions, $P(\xi_n^{(t)}|d,m)$ can be expressed as the normalized product of two terms

$$P(\xi_n^{(t)}|d,m) \quad = \quad \eta \; \underbrace{P(\xi_n^{(t)}|o_n^{(1)},\ldots,o_n^{(t)},m)}_{:=\alpha_n^{(t)}} \; \underbrace{P(\xi_n^{(t)}|u_n^{(t+1)},\ldots,o_n^{(T)},m)}_{:=\beta_n^{(t)}} \tag{4}$$

Here η is a normalizer that ensures that the left-hand side of Equation (4) sums up to one (see [TFB98a] for a mathematical derivation). Both terms, $\alpha_n^{(t)}$ and $\beta_n^{(t)}$, are computed separately. Computing the first term amounts to *localization forwards in time*, whereas computing the second corresponds to *localization backwards in time*.

The reader should notice that the computation of the α-values is a version of *Markov localization*, which has recently been used with great success for robot localization in *known* environments by various researchers [BFHS96, KCK96, KS96, NPB95, SK95]. The β-values add additional knowledge to the robot's pose, typically not captured in Markov-localization. They are, however, essential for revising past belief based on sensor data that was received later in time, which is a necessary prerequisite for building large-scale maps.

Computing the α-Values: Initially, one of the robots (the robot with index 1) is assumed to be at the center of the global reference frame, $\alpha_1^{(1)}$ is given by a Dirac distribution centered at $(0,0,0)$:

$$\alpha_1^{(1)} \;=\; P(\xi_1^{(1)}|o_1^{(1)},m) \;=\; \begin{cases} 1, \text{ if } \xi_1^{(1)} = (0,0,0) \\ 0, \text{ if } \xi_1^{(1)} \neq (0,0,0) \end{cases} \tag{5}$$

If the starting poses of the other robots relative to the first are known, their $\alpha_n^{(1)}$ are initialized correspondingly. In the most general case, however, where the individual robots do not know their relative starting pose, $\alpha_n^{(1)}$ is initialized by a uniform distribution for all other robots $n = 2,\ldots,N$. In this case, these robots can localize themselves relative to the first right after the first iteration of EM (in which a first map is constructed).

The values $\alpha_n^{(t)}$ for $t = 2,\ldots,T$ are computed recursively:

$$\alpha_n^{(t)} \;=\; \eta\; P(o_n^{(t)}|\xi_n^{(t)},m)\; P(\xi_n^{(t)}|o_n^{(1)},\ldots,u_n^{(t-1)},m) \tag{6}$$

where η is again a probabilistic normalizer. The rightmost term of (6) can be transformed to

$$P(\xi_n^{(t)}|o_n^{(1)},\ldots,u_n^{(t-1)},m) \;=\; \int P(\xi_n^{(t)}|u_n^{(t-1)},\xi_n^{(t-1)})\; \alpha_n^{(t-1)}\; d\xi_n^{(t-1)} \tag{7}$$

Substituting (7) into (6) yields a recursive rule for the computation of all $\alpha_n^{(t)}$. See [TFB98a, TFB98b] for a more detailed derivation. Notice that $\alpha_n^{(t)}$ can be computed locally on each robot.

Computing the β-Values: The computation of $\beta_n^{(t)}$ is completely analogous but takes place backwards in time. The "initial" $\beta_n^{(T)}$, which expresses the probability that the n-th robot's final pose is $\xi_n^{(T)}$, is uniformly distributed, since $\beta_n^{(T)}$ does not depend on data. All other β-values are computed in backwards order:

$$\beta_n^{(t)} \;=\; \eta \int P(\xi_n^{(t+1)}|u_n^{(t)},\xi_n^{(t)}) P(o_n^{(t+1)}|\xi_n^{(t+1)},m)\; \beta_n^{(t+1)}\; d\xi_n^{(t+1)} \tag{8}$$

The derivation of the equations are analogous to that of the computation rule for α-values and can be found in [TFB98a]. The result of the E-step, the product $\alpha_n^{(t)}\beta_n^{(t)}$, is an estimate of the poses of each robot's pose for each point in time.

2.3.2 The M-Step

The M-step maximizes $P(m|\xi,d)$, that is, in the M-step the robots compute the most likely map based on the pose probabilities computed in the E-step. Generating maps with *known* robot poses, which is basically what the M-step amounts to, has been studied extensively in the literature on mobile robot mapping (see e.g., [BK91, Elf89, Mor88]).

By applying Bayes rule and with the appropriate assumptions, the estimation problem can be decomposed into

$$P(m|\xi,d) \;=\; \alpha \prod_{n=1}^{N}\prod_{t=1}^{T} P(o_n^{(t)}|\xi_n^{(t)},m) \tag{9}$$

where α is a normalizer that can safely be ignored in the maximization. It is common practice to decompose the problem spatially, by solving the optimization problem independently for different x-y-locations:

$$\operatorname*{argmax}_{m} P(m|\xi,d) \;=\; \left\{ \operatorname*{argmax}_{m_{x,y}} \prod_{n=1}^{N}\prod_{t=1}^{T} P(o_n^{(t)}|\xi_n^{(t)},m_{x,y}) \right\}_{x,y} \tag{10}$$

These local maximum likelihood estimations problems are highly tractable, since each of them involves only a single, discrete random variable. In fact, the M-step (mapping) can be broken down into two steps, one in which each robot computes a map based on its own data d_n, and one where the local maps are integrated into a single, global map. Thus, robot need only to communicate maps (which are much more compact than the pose estimates ξ_n).

(a) (b)

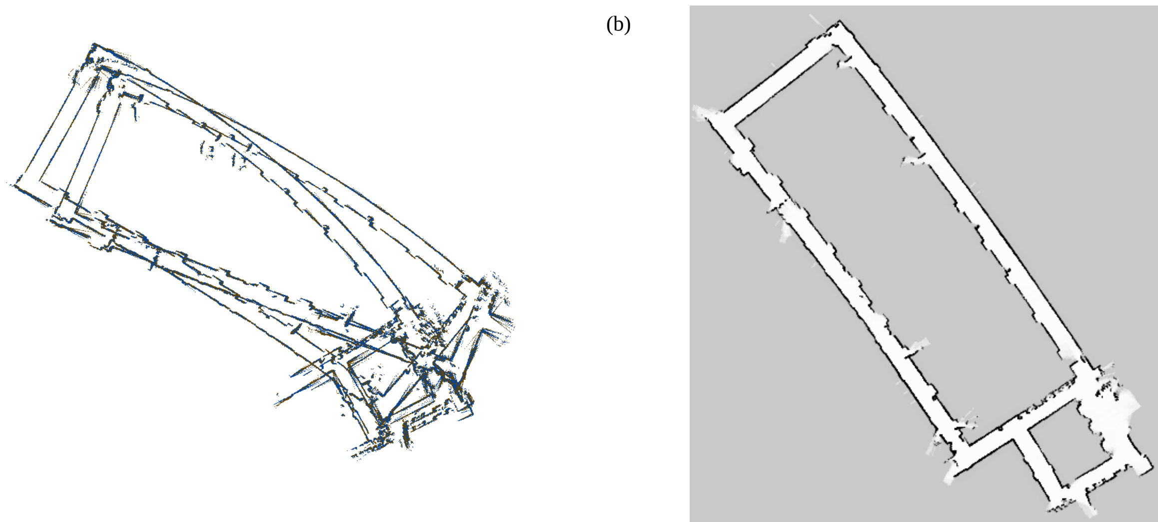

Figure 4: (a) Raw data, obtained in an environment of size 80 by 25 meters. The challenge here arises from the circular nature of the environment: notice that the final odometric error is 24.9 meter. (b) Occupancy grid map, generated using our proto-type implementation. To our knowledge, no other approach has ever generated maps of similar size.

2.3.3 Distributed EM

For a team of robots to collaboratively build a map, EM must be implemented in a distributed fashion. Fortunately, the basic EM algorithm lends itself nicely to distributed mapping with low-bandwidth communication.

The E-step will carried out locally, using a map built cooperatively by the team of robots. Thus, the data d and the distributions α and β, which consume the major bulk of memory, are kept locally on each robot. The M-step will be implemented by two steps, one in which each robot determines a map locally (leaving all probabilistic information intact), and one where the robots communicate their local maps and build a single, global map. This decomposition exploits the fact that Equation (10) can be decomposed into a product of maps generated by individual robots. Only the latter step involves communication. Using commercially available wireless communication hardware (e.g, 2 Mbit/sec Breezecom), transmission of maps of the size and resolution obtained using our existing prototype requires less than 0.1 seconds per robot. By moving from grid-representation to object-centered vector representation, as outlined below, we expect to reduce the communication overhead by at least one order of magnitude, since the number of objects (walls, desks, ...) are much smaller than the number of grid cells.

We will develop a flexible communication protocol, which will take the asynchronous nature of the information into account. In particular, we will investigate whether information exchange is needed for *every* M-step, or whether it is possible to interleave multiple iterations of local EM with a single map exchange. If such a scheme works in practice, the communication between different robots could accommodate unreliable radio links.

2.4 Initial Results

We have already applied the approach successfully to the problem of single-robot mapping. As described in more detail elsewhere [TGF+98], we used a mixture of grid-based methods and Kalman-filter methods for representing the various probability distributions (maps, pose estimates). Grid-based methods suffers from excessive memory and processing time requirements, since the underlying densities are fairly complex. Kalman filters can only model uni-modal distributions; which is insufficient for mapping due to the large number of ambiguities. While the initial results are presented to illustrate the potential of the statistical approach proposed here, the reader should not dismiss that we intend to move away entirely from such representations, and replace them by sampling techniques (for poses ξ) and object-centered vector representations (for maps m), as described further below.

One of our benchmark data sets is shown in Figures 4. This data set has been collected in our university building. As Figure 4a illustrates, the final odometric error is quite significant: 24.9 meter. What makes this dataset particularly challenging is the large circular hallway (60 by 25 meter). When traversing the circle for the first time, the robot

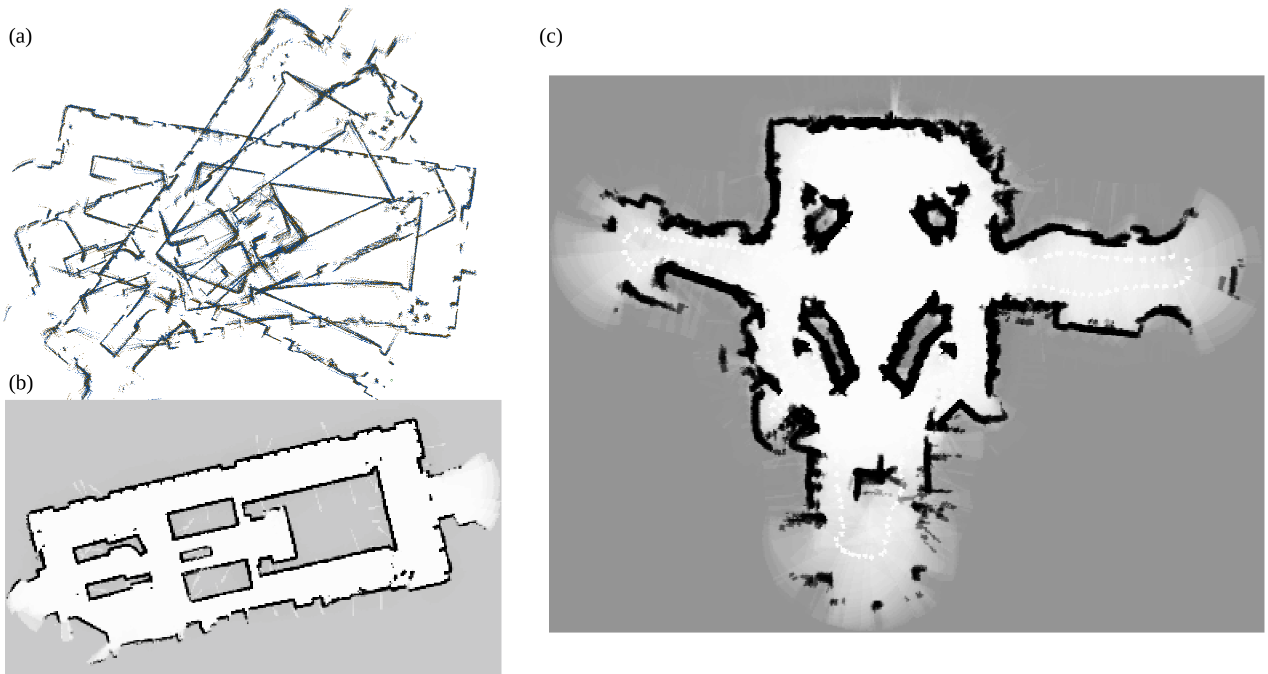

Figure 5: Raw dataset has been recorded in the Dinosaur Hall of the Carnegie Museum of Natural History. The final odometric error is more than 35 meters (translation) and more than 80 degrees (rotation). (b) The final map was constructed in less than 60 minutes, including the data collection. (c) Map generated for the Smithsonian's National Museum of American History (60 by 40 meter). Based on this map, our robot MINERVA reliably navigated for 2 weeks (44 km, approx. 50,000 visitors), using the localization methods described here (the α values).

---

cannot exploit landmarks to improve its pose estimates; thus, it accumulates odometric error. Since significant places are indistinguishable, it is difficult to determine the robot's position when the circle is closed for the first time (here the odometric error is larger than 14 meter). Only as the robot proceeds through known territory it can use its perceptual clues to estimate where it is (and was), in order to build a consistent map. Figure 4b shows the result of mapping, represented as a occupancy grid map [BK91, Elf89, Mor88]. This map is well-suited for our current navigation routines [GN97, TBB+98].

Other examples are shown in Figures 5. The map shown in Figures 5c (right diagram) was constructed in the Smithsonian's National Museum of American History (NMAH), where one of our robots (Minerva) was used as a tour-guide for visitors. Based on this map, Minerva localized itself successfully despite the fact that there were crowds of people blocking its sensors. For localization, we used the very statistical approach outlined above: The values α are pose estimates of the robot, and computing α is known as *Markov localization* [BFHS96, KCK96, KS96, NPB95, SK95]. The robot tracked its position successfully during almost 100 hours of operation, during which it traversed more than 44 km, demonstrating an unprecedented level of robustness, which we attribute to the statistical nature of our approach.

We also performed preliminary experiments to check the feasibility of applying our approach to multi-robot exploration. One of the difficult problems is that the robots have to find out where they are relative to each other—see the E step above. Thus, for our approach to be successful, robots must be able to *globally* localize themselves in the learned map. In all experiments conducted thus far, we found that our statistical approach is well-suited for global localization. For example, in the NMAH we repeatedly initialized our robot under global uncertainty, and it always localized itself successfully. An example is shown in Figure 6. Here Minerva localizes itself in a previously learned map—strictly speaking, this data has been collected by one robot only, but the results can be directly transferred to a multi-robot scenario where one robot localizes itself in a map learned by one or more other robots. Figure 6a depicts the pose estimate $P(\xi)$ after incorporating a single sensor scan (laser range scan), Figure 6b depicts the pose estimate after integrating a second. As is easy to be seen there, after only integrating two range scans, the robot knows with high certainty (and accuracy) its position, even though initially it didn't. These results, which can be viewed as a restricted version of the EM scheme proposed here (it is restricted because there is no propagation of β values), demonstrate the feasibility of our proposed approach.

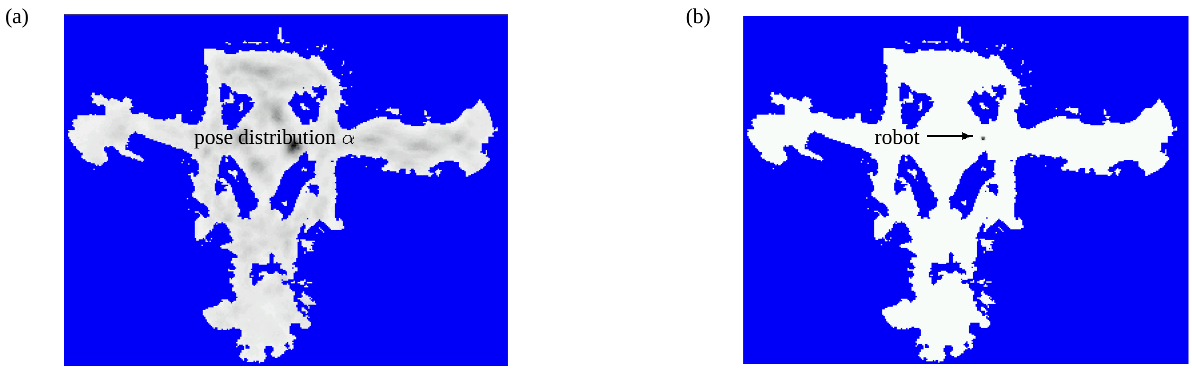

Figure 6: A second dataset is integrated into the first, both collected in the Smithsonian NMAH. (a) Pose probability distribution after integrating a first laser scan (projected into 2D), (b) after integrating a few more, the robot knows its pose with high certainty/accuracy. Global localization is essential for cooperative multi-robot mapping, as robots have to determine their pose relative to each other.

To our knowledge, the maps shown here and elsewhere [TFB98a, TFB98b] are an order of magnitude larger than the maps produced by any other algorithm. We attribute this difference to the use of statistical techniques, which can handle ambiguities and uncertainty in a mathematically elegant way, and which provide natural means for belief revision. All maps shown in this proposal were generated in less than 2 hours computational time; with the bulk of time required because we used grid-based representations. With our new representations described below, we hope to reduce the computational time by at least an order of magnitude.

3 Research Issues

We seek to further develop the basic mathematical framework with the goal of devising a new set of statistical algorithms for efficient multi-robot mapping. Apart from the research issues mentioned above (which will include the transition from single-robot to multi-robot algorithms), this project will pursue research on the following issues:

3.1 Vector-Based and Sampling-Based Representations

A primary goal of this research is to move away from grid-based representations, since their memory- and time-complexity imposes inherent scaling limitations. To this end, our research will develop new representations for the two probability distributions involved: Map and pose distributions.

The proposed approach, if successful, will be an improvement over previous methods in that is scales much better to complex environments. It will also yield qualitatively different maps. Instead of breaking environments into equally-sized grid cells, it will generate maps that are composed of "typical" objects in indoor environments. These representations are more compact than previous representations, and hence more scalable, and they are also better suited for tracking changes that inevitably occur.

- **Object-centered representations for maps.** We will investigate the utility of *object-centered vector representations* for representing maps. These representations will compose maps as a superposition of generic prototype objects. Each object is annotated with a *parameter vector*, specifying shape, size, and location of the object within the map's global coordinate frame. We believe that the same statistical methods can be applied to such object-centered maps. Instead of estimating distributions for the occupancy of grid cells $\langle x, y \rangle$, our new approach will estimate distributions over object-specific parameter vectors (Equation (10) is modified accordingly). To determine the appropriate number of objects, Occam's razor will be used for map regularization. Occam's razor will penalize complex maps (many objects), to avoid fragmentation of larger objects (such as long walls).

  Initially, our library of atomic objects will contain walls, rectangular objects (chairs, tables), and curved objects. Each object will possess a generic set of parameters (width, length, orientation, diameter, etc.). If the set of atomic objects is insufficient, we will augment it by polygons or, possibly, hierarchical compositions of atomic objects. While or current grid-based maps often contain as many as 10^6 entities, we expect typical environments to be decomposable into a much smaller number of atomic objects (10^3 or less). Thus, the resulting maps will reduce the

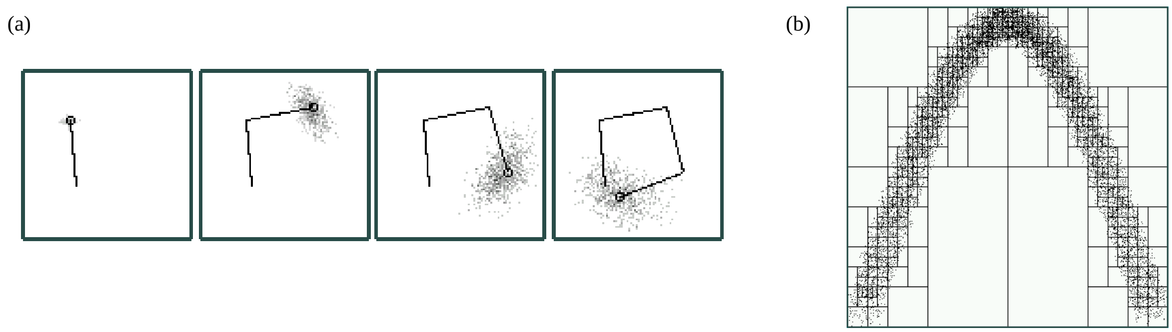

Figure 7: (a) Sampling-based methods represent probability distributions by samples drawn from this distribution. The example here is taken from a proto-type implementation, which led to a hundred-fold increase in speed over a previous grid-based implementation. (b) Sample sets are approximated using oct-trees. Oct-trees are extremely fast to access, and they adaptively place resolution where needed. Both diagrams have been obtained using a proto-type implementation.

---

computational and memory complexity of our approach, and they will also reduce the communication overhead when teams of robot collaboratively build maps.

- **Sampling-based methods for pose estimation.** We will investigate the utility of *sampling-based* methods for representing pose densities—in fact, we have already implemented a first proto-type using sampling-based methods. Currently, pose distributions (e.g., $P(\xi), \alpha, \beta$) are represented by evenly spaced grids. Their enormous memory and time complexity has forced us to work with extremely coarse representations (1 meter spatial resolution, 5 degree angular resolution).

  Sampling-based methods place computation and memory where needed. The idea is to represent probability distributions through samples, where the *density of samples corresponds to their likelihood*. Figure 7a shows an example: Instead of representing $P(\xi)$ as a grid, it is represented through a set of samples. The idea of sampling to represent distributions is commonly used in statistics: likelihood-weighted sampling is the basis of the Metropolis algorithm for stochastic integration [Win95], and similar methods have recently been applied for visual tracking of deformable objects with remarkable success [IB98]. However, to our knowledge the use of sampling-based methods in mapping is new.

  To apply sampling-based methods to localization (i.e., when computing α and β), the initial pose $\alpha_1^{(1)}$ (the only "known" pose) is initialized using a single sample (0,0,0). All other pose estimates $\alpha_n^{(1)}$ ($n \neq 1$) and $\beta_n^{(t)}$ are initialized using a set of samples generated according to a uniform distribution (to reflect uniform uncertainty, c.f., Section 2.3.1). When computing $\alpha_n^{(t)}$ (with $t > 1$), our approach will first sample from $\alpha_n^{(t-1)}$ according to the conditional probability density $P(\xi'|u, \xi)$ (the motion model). This is done by drawing samples from $\alpha_n^{(t-1)}$ randomly, and then applying the motion model by randomly "guessing" specific error parameters. Next, the likelihood of each resulting sample is re-assessed using the perceptual model (the conditional likelihood $P(o|m, \xi)$). This assessment phase will weigh the likelihood of each sample according to the perceptual consistency with the map. Since samples are generated according to likelihood, the computation will predominately focus on high-likelihood poses. (The process for sampling β-values is fully analogous.)

  The only complication arises when combining $\alpha_n^{(t)}$ and $\beta_n^{(t)}$ to obtain $P(\xi_n^{(t)}|d, m)$, as prescribed by Equation (4). Here we have to *multiply* two probability densities represented by two different sets of samples. Our proposal is to transform each sample set into an *oct-tree*, as previously proposed in a different context by several other researchers [BDFC98, KF98, Sam89b, Sam89a, Moo90, Moo98]. Figure 7b shows an example tree. Oct-trees can be generated so that each leaf covers approximately the same number of samples, thereby placing higher density at more likely poses. When computing $P(\xi_n^{(t)}|d, m)$, samples from one set ($\alpha_n^{(t)}$ or $\beta_n^{(t)}$) will be weighted according to the likelihood obtained through a tree constructed from the respective other distribution.

  In preparation of developing this research proposal, we recently implemented a proto-type algorithm using the sampling idea. In initial tests using small data sets, we managed to reduce the computational complexity consistently by two orders of magnitude when compared with our highly-optimized grid-based methods, while slightly increasing the spatial resolution when compared to the grid-based mapper. These results require systematic verification; yet, they demonstrate the advantage of sampling-based methods.

3.2 Cooperative Multi-Robot Exploration

Extending our statistical framework, we will pursue research on algorithms for cooperative exploration. Our previous research has led to several efficient algorithms for single-robot exploration [CNT98, Thr98b]. This project will address the question: *How can a team of robots maximize the knowledge gained when collaboratively exploring and mapping an unknown environment?*

The key idea of our approach will be to use the *entropy* in a robot's internal belief to drive exploration (see also [BFT97, FBT98a]). The entropy is a statistical measure for the uncertainty of an estimator (such as a map). By maximizing the rate at which entropy is diminished, robots can explore their environment efficiently.

More specifically, let $u_{\Delta x,\Delta y}$ denote the action of moving $\langle \Delta x, \Delta y\rangle$ relative to a robot's current location. Each possible action has an effect on the expected entropy:

$$H[m] \quad = \quad -\sum_{x,y}\sum_{\tau} P(m_{x,y}=\tau)\log P(m_{x,y}=\tau) \tag{11}$$

is the entropy of map m, where τ sums over all possible values a map can take at each location (e.g., occupied, unoccupied, part-of-object-i). In fact, the entropy can be computed separately for each map location $\langle \Delta x, \Delta y\rangle$. $H[m]$ is the *cumulative entropy of the map* m, that is, it is a quantitative measure for the uncertainty in m. The *expected entropy upon executing action* u *in pose* ξ, denoted $H[m|u,\xi]$, is defined as the entropy one expects to result if the robot executes action u at position ξ (see [BFT97, FBT98a] for a precise mathematical definition). Since the robot does not know where it is, the *expected entropy upon executing action* u is given by

$$H[m|u] \quad = \quad \int H[m|u,\xi]\ P(\xi)\ d\xi \tag{12}$$

The *rate of entropy loss*, thus, is the quotient of $H[m]-H[M|u]$, divided by the time required to execute action u. In our previous work [BFT97, FBT98a, Thr98b], we devised an algorithm for estimating this time. More specifically, we have already devised a planning algorithm for executing actions under uncertainty in the robot's position. This algorithm, which uses dynamic programming for motion planning [Bel57], computes the expected time it takes to execute an action based on a map m and a pose belief $P(\xi)$. By putting both things together, the expected entropy loss and the results of this planner, the resulting action selection mechanism will select actions that maximize the rate of knowledge gain, thus, lead to efficient exploration.

To extend this algorithm to multi-robot exploration, we will devise a communication scheme for communicating robot *intent*. Notice that the general exploration problem is NP-hard (in the number of robots) [Rei79, CD88]; thus, any practical strategy will only approximate the theoretically optimal solution. Our approach will rely on the observation that the entropy—and with it, the expected change in entropy—can be evaluated for each map location separately. To coordinate the exploration, each robot determines locally the action that best suits the knowledge gain, as described above. The action is then projected into a map that describes the expected reduction in uncertainty *upon* executing this action, evaluated separately for each location $\langle x, y\rangle$. This map (called the *expected entropy reduction map*) is then communicated to all other robots, which will then reassess their optimal exploration actions accordingly. Thus, our approach will enable multiple robots to coordinate their exploration efforts, while still performing the necessary computation locally (on-board), and in real-time.

The approach proposed here, if successful, will be an improvement over previous mobile robot exploration techniques developed by the PI and others [CNT98, Thr98b, YB96]. Previous mobile robot exploration techniques were well-suited for exploring unknown terrain. However, they are incapable of taking the robot's uncertainty into account, i.e., they all assume that places that were once visited by the robot are mapped *correctly*. This is an important deficiency. If the robots are not quite certain about the shape of the map (a common situation in our experiment), they sometimes have to traverse previously explored terrain to gather more information. The proposed method will automatically balance these two exploration goals—moving to unexplored terrain to extend the map and traversing known terrain to improve the map—through a single objective: minimizing expected entropy.

3.3 Continual Adaptation to Dynamic Environments

Finally, we will extend our approach to one that enables robot teams to map their environment continuously, so that they can operate over long periods of time. Virtually all indoor environments change at some point, yet the vast majority of existing methods for mobile robot mapping (with the exception of [BCF+98a, Thr98b, YLS+98]) can only cope with static environments. We seek to extend the basic statistical framework in a way that it can accommodate changes in the environments. This is important for virtually all practical applications, as virtually all indoor environments tend to change over time.

To cope with environment dynamics, we propose to extend our approach by an *exponential filter*, which will be applied to sensor evidence. More specifically, in the M-step sensor evidence $o_n^{(t)}$ will be weighted by a factor γ^{T-t} where T is the current time, t is the time $o_n^{(t)}$ was recorded, and γ is an exponential decay factor that is smaller than 1. Thus, the modified M-step becomes

$$\operatorname*{argmax}_{m} P(m|\xi, d) \;=\; \left\{ \operatorname*{argmax}_{m_{x,y}} \prod_{n=1}^{N} \prod_{t=1}^{T} \gamma^{T-t} P(o_n^{(t)}|\xi_n^{(t)}, m_{x,y}) \right\}_{x,y} \tag{13}$$

This modified M-step (cf. Equation (10)) weighs more recent sensor readings exponentially stronger than more distant ones. The underlying probabilistic assumption here is that environments change randomly, with a fixed probability $(1-\gamma)$ per time interval.

Different objects change at different rates. For example, the location of walls change much less frequently than that of chairs. Conventional grid-based and other iconic representations cannot differentiate between different objects. Thus, a key advantage of our proposed object-centered representations (Section 3) is that different time constants γ can be used for different type objects. We will generally assume that objects such as chairs and tables have a faster decay rate than walls and other large objects. Our research will characterize empirically the ability of the proposed approach to construct maps in changing environments.

Since our current statistical method re-estimates poses backwards in time, the time required for each update increases linearly with T, and the basic approach is not applicable for continuous "lifelong" mapping. To remedy this problem, we propose to develop a method that considers a fixed history window when estimation poses backwards in time. For poses beyond the history window, the Viterbi algorithm [RJ86] will be applied to calculate the *most likely* poses, which will then be compiled incrementally into a fixed map. As a result, only a few, past poses will be estimated incrementally, whereas the bulk of poses will be assumed to be known. The resulting approach should inherit the advantages of the approach proposed here, while enabling robots to map their environment continuously, over extensive periods of time.

Our statistical approach has an interesting side effect on exploration: When the environment changes, the certainty in the map is reduced; hence the robots will direct their exploration towards such areas (trying to exploit the opportunity to reduce entropy). As a result, we project that once a change has been spotted, robots will move there to rapidly adjust their map. We will run experiments to characterize this effect and evaluate the benefits of entropy-based exploration methods in dynamic environments.

3.4 Evaluation

The approach will be developed and evaluated using a range of mobile robot platforms, including those shown in Figure 1. If funded, we will perform systematic evaluations in several environments, including office buildings, museums, private homes, and possibly environments such as assisted living facilities (see the enclosed endorsement letter).

Our evaluations will characterize the scalability of the basic algorithms in terms of size of the environment, structural regularity/complexity, dynamics, amount of sensor noise, and so on. It is our goal to integrate the results of this work into RWI's BeeSoft navigation package (see below), to facilitate its dissemination into the scientific community at large. However, the requested funds will exclusively be used to support basic research activities; funds for integrating the practical results of this research into BeeSoft will be covered through other sources.

4 Related Work

Over the last decade, there has been a flurry of work on map building for mobile robots (see e.g., [CL85, LDWC92, Ren93, Thr98b]). A detailed survey of recent literature on map building can be found in [Thr98b]. Our approach differs from other work in the field (see the surveys in [Thr98b, LM97]) in four important technical aspects.

1. First, robot poses are revised forward *and* backwards in time—as pointed out by Lu and Milios [LM97], most existing approaches do not revise pose estimations backwards in time. The ability to revise pose estimates backwards in time is essential to build large-scale maps, specifically in cyclic environments.
2. Second, by using probabilistic representations, the approach considers multiple hypotheses as to where a robot might have been, which facilitates the recovery from errors.
3. Third, the probabilistic nature of the approach makes possible that multiple robots collaboratively map unknown environments. Existing approaches are unable to localize robots relative to each other; thus fail to integrate information collected on multiple platform.
4. Fourth, our proposed object-centered representations will generate more natural and more compact maps than previous methods, which also facilitate adaptation to changes over time.

Recently, several groups have proposed algorithms that revise estimates backwards in time, thereby overcoming the first limitation listed above. Koenig and Simmons investigated the problem under the assumption that a topologically correct sketch of the environment is available, which simplifies the problem somewhat [KS96]. They proposed a probabilistic framework similar to the one described here (in fact, it is a special case), but by relying on a sketch of the environment, their approach is unable to generate maps from scratch. Shatkay and Kaelbling [SK97] generalized this approach for mapping in the absence of prior information. Their approach consults local geometric information to disambiguate different locations. Both approaches differ from ours in that they build only topological maps. They do not explicitly estimate global geometric information (i.e., x-y-θ positions). As acknowledged in [SK97], the latter approach fails to take the cumulative nature of rotational odometric error into account. It also violates a basic "additivity property" of geometry (see [SK97]). Even in the absence of odometric error, it is still unclear to us if the approach will always produce the correct map.

Lu and Milios [LM97] have proposed a method that matches laser scans into partially built maps, using Kalman filters for positioning. Together with Gutmann [Gut96], they have demonstrated the appropriateness of this algorithm for mapping environments with cycles. Their approach is incapable of representing ambiguities and multi-modal densities. It can only compensate a limited amount of odometric error in x-y-θ-space, due to the requirement of a "sufficient overlap between scans" [LM97]. In all cases studied in [Gut96, LM97], the odometric error was an order of magnitude smaller than even the one reported in out initial experiments. In addition, the approach is largely specific to robots equipped with laser range finders. It is unclear to us if the approach can cope with less accurate sensors such as sonars.

To the best of our knowledge, the topic of collaborative multi-robot mapping has previously only been studied by Lopez and colleagues [LLdMS97]. Like ours, their approach models the uncertainty of a robot's location explicitly, and it also takes this uncertainty into account when building maps. However, their approach lacks a method for localization. As the uncertainty grows larger than a prespecified threshold, mapping is simply terminated, thereby imposing tight, intrinsic bounds on the size of environments that can be mapped (they are *very* small!). Due to the lack of a localization component, robots cannot localize themselves in another robot's map. Thus, the robots must know their position relative to each other.

While the issue of efficient robot exploration has been studied by the PI and others [YB96, Thr98b], the issue of collaborative multi-robot exploration has never been investigated. Most approaches to multi-robot collaboration only invoke local behaviors; thus, do not necessarily lead to coordinated mapping behavior [Par96, Mat97, LLdMS97].

The approach proposed here also relates to work in the field of Markov localization. Markov localization addresses the problem of localizing a robot under the assumption that a map to be given. According to Cox [Cox91], "Using sensory information to locate the robot in its environment is the most fundamental problem to providing a mobile robot with autonomous capabilities." Recently, Markov localization has been employed by various groups with remarkable success [BFHS96, KCK96, KS96, NPB95, SK95, TBB+98, Thr98a]. In our own work, Markov localization played a key role in a recent installation in the Deutsches Museum Bonn, where one of our robots provided interactive tours to visitors [BAC+98, BCF+98b]. In more than 18.5km of autonomous robot navigation in a densely crowded environment (top speed 80 cm/sec, average speed 36 cm/sec), Markov localization was absolutely essential for the robot's safety and success [FBT98b].

Technically, Markov localization is equivalent to the computation of the α-values (cf. Section 2.3.1). Thus, the statistical framework proposed here directly subsumes Markov localization, and extends it by a mapping component.

If this research is successful, maps do not have to be crafted manually but can now be generated automatically for future installations of the tour-guide robot. This will effectively reduce the installation time of the robot's navigation component from several days to only a few hours.

5 Long-Term Goals

The proposed research is part of a larger research effort carried out by the PI. Our long-term objective is to contribute to the basic scientific foundations of an upcoming generation of mobile service robots, such as mobile assistants in industrial settings and in the private or medical sector, or personal assistants for elderly and disabled people.

The societal need for service robots is enormous (cf. [VN94, FKAM84]). For example, according to the U.S. Senate Special Committee on Aging (1982), more than 80% of the population over 65 have at least one chronic illness, the most frequent being arthritis. Arthritis may cause limitations of functions in one or more activities of daily living. When fine hand movement and mobility become severely impaired, independent self-care becomes painfully difficult or even impossible. Currently 12.5% of US population is over 65 and this fraction is expected to double by 2050. Nursing homes cost $30-60,000 annually, part-time home care is about $10,000/year, while full-time home care is over $100,000/year. The current nursing home population is 1,800,000, and home health care costs have grown by 25% over the last three years and is expected to grow at this rate into the next century. The Alliance for Aging Research estimates that for every month that we postpone entry into a nursing home for those who will go, the economic benefit is $3 billion, not to mention impact on the quality of life. Even with careful estimates, reducing the need for nursing homes by as little as 1% with robots that cost up to $10,000 would save $2 to $4 billion annually. It is important that the U.S. will not loose its leading position to its major competitors (Japan in particular), as partially happened in other aspects of robotics. While robots that provide assistance to elderly and handicapped people are still out of reach, we believe that the ability to acquire maps and localize robots using such maps will become an essential enabling component for such applications. Map-based navigation is the most successful paradigm in mobile robot navigation to date; yet without methods for acquiring and maintaining maps, it is not possible to deploy such robots at low costs in environments such as the ones targeted here.

As part of our effort towards this long-term goal, we have developed an integrated multi purpose mobile robot control architecture (called "BeeSoft"). BeeSoft is a generic software package for autonomous robot navigation and human robot interaction (see [TBB+98, BBC+95]). It currently consists of the following components:

1. a fast, reactive module for collision avoidance [FBT96, FBT97, FBT98b],
2. a Markov localization method that can reliably localize the robot provided that a map is given [BFHS96, BFH97, BFT97, FBT98a]—in fact, this method can be viewed as a special case of the topic proposed here,
3. a point-to-point path path planner (based on an any-time algorithm) [Thr93, Thr98b],
4. a fast segmentation-based vision system for finding obstacles and identifying specific objects [BBC+95, MT98],
5. a system for tracking people and faces in real-time [WTRM98, WTR98],
6. a gesture-based control interface for mobile robots [WTRM98, WTR98],
7. a mission-control programming interface [BFT97, HBL98],
8. a real-time stereo system for obstacle avoidance and map building [FB96b, FB96a],
9. a graphical control and tele-operation interface, and
10. a graphical, interactive simulator.

Mobile robots controlled by our software can navigate through unknown populated environments with a speed of up to 1.6 meter/sec while avoiding collisions with obstacles. Several of our robots have successfully been employed for finding and fetching small free-standing objects such as soda cans, instructed by humans through gestures.

To make the public aware of the potential impact of mobile robotics, and to further test the robustness of our approach in "realistic" settings, we recently installed mobile robots in the Deutsches Museum Bonn, and in the Smithsonian's National Museum of American History [BAC+98, BCF+98a, Thr97a, Thr98c]. Our first robot, "RHINO," guided more than 2,000 visitors through the museum during a six-day installation period, at a top speed of 80 cm/sec and for a total distance of over 18.5 km. "Minerva," which operated at the National Museum of American History for 2 weeks, led approximately 50,000 people, at top speeds of 163 cm/sec and for a total distance of 44.0 km. Both robots sparked enthusiasm among people in all age groups, but especially in the younger generation. Both robots navigated using maps—the museums were *not* modified in any way. While Rhino's map was constructed manually, Minerva's map was learned using the prototype described in this proposal.

Other robots controlled by the BeeSoft architecture won first, and second, price at the AAAI autonomous mobile robot competitions in 1996, and 1994, respectively [BBC+95, TBB+98, Thr97b]. The BeeSoft architecture is also distributed by Real World Interface Inc., as the only navigation software distributed with their top-of-the-line mobile robots. It is already being used at 20 (or more) academic sites worldwide, and this number seems to be increasing rapidly.

Building maps of indoor environments with teams of robots is a key open problem. It is one of the major technical obstacles in deploying robots in health care facilities and private homes. Private homes in particular pose unique challenges on the flexibility and robustness of the approach that go beyond the difficulties encountered in structured laboratory environments. On the one hand, the architectural diversity of private homes is enormous. On the other hand, it is generally undesirable to generate maps of private homes by hand. Methods that enable a human operator to teach-in the system with a limited amount of time and technical understanding are clearly needed, but not yet available. The proposed research aims at filling this gap. However, the funds requested here will not be used to develop robot applications (such as the service robots described here). Instead, they will be used to pursue the basic scientific research necessary for the development of more scalable mapping technology.

6 Work Plan

Year 1. We will develop and implement the basic multi-robot mapping and localization approach as outlined above. We will develop an efficient communication protocol for teams of robots, and a synchronization scheme for coordinating the local EM on multiple robots. We will begin research on multi-robot exploration, object-centered representations and sampling methods for efficient density propagation.

Year 2. We will develop and implement the distributed exploration as outlined above. We will further develop object-centered representations, sampling methods for efficient density propagation, and initiate or work on and extending our approach to dynamic environments.

Year 3. We will fully develop our approach on all aspects of object-centered representations and sampling-based methods, and fully develop our methods for dynamic environments. We will thoroughly test and document this research, using testbeds similar in nature to those we used in the past (e.g., a museum).

The results of this research will be disseminated through (1) conference and journal publications, (2) code sharing (using BeeSoft as vehicle, and using the infrastructure provided by DARPA's "Tactical Mobile Robotics" program, which partially supports the PI (the scope of this research is beyond CMU's TMR contract), (3) by creating up a Web site that documents the latest problems, solutions, and results, and (4) by using some of the results of this research as course material, both at the undergraduate and the graduate level.

7 Budget Justification

The funds will be used to support two enthusiastic graduate students, Nicholas Roy and Frank Dellaert, who have contributed to the current research and who, if this project is funded, are very likely to write Ph.D. theses on the topic proposed here. No funds are requested for faculty support or robotic hardware, even though the PI expects to dedicate significant time to this project.

BIBLIOGRAPHY

[BAC+98] W. Burgard, A.B., Cremers, D. Fox, D. Hähnel, G. Lakemeyer, D. Schulz, W. Steiner, and S. Thrun. The interactive museum tour-guide robot. In *Proceedings of the AAAI Thirteenth National Conference on Artificial Intelligence*, 1998.

[BBC+95] J. Buhmann, W. Burgard, A.B. Cremers, D. Fox, T. Hofmann, F. Schneider, J. Strikos, and S. Thrun. The mobile robot Rhino. *AI Magazine*, 16(1), 1995.

[BCF+98a] W. Burgard, A.B. Cremers, D. Fox, D. Hähnel, G. Lakemeyer, D. Schulz, W. Steiner, and S. Thrun. Experiences with an interactive museum tour-guide robot. Submitted *Artificial Intelligence*, 1998.

[BCF+98b] W. Burgard, A.B. Cremers, D. Fox, D. Hähnel, G. Lakemeyer, D. Schulz, W. Steiner, and S. Thrun. Experiences with an interactive museum tour-guide robot. Technical Report CMU-CS-98-139, Carnegie Mellon University, Computer Science Department, Pittsburgh, PA, 1998.

[BDFC98] W. Burgard, A. Derr, D. Fox, and A.B. Cremers. Integrating global position estimation and position tracking for mobile robots: The dynamic markov localization approach. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'98)*, 1998. To appear.

[BEF96] J. Borenstein, B. Everett, and L. Feng. *Navigating Mobile Robots: Systems and Techniques*. A. K. Peters, Ltd., Wellesley, MA, 1996.

[Bel57] R. E. Bellman. *Dynamic Programming*. Princeton University Press, Princeton, NJ, 1957.

[BFH97] W. Burgard, D. Fox, and D. Hennig. Fast grid-based position tracking for mobile robots. In *Proceedings of the 21th German Conference on Artificial Intelligence*, pages 289–300, Berlin, 1997. Springer Verlag.

[BFHS96] W. Burgard, D. Fox, D. Hennig, and T. Schmidt. Estimating the absolute position of a mobile robot using position probability grids. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, Menlo Park, August 1996. AAAI, AAAI Press/MIT Press.

[BFT97] W. Burgard, D. Fox, and S. Thrun. Active mobile robot localization. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI)*, San Mateo, CA, 1997. Morgan Kaufmann.

[BK91] J. Borenstein and Y. Koren. The vector field histogram – fast obstacle avoidance for mobile robots. *IEEE Journal of Robotics and Automation*, 7(3):278–288, June 1991.

[CB90] G.C. Casella and R.L. Berger. *Statistical Inference.* Wadsworth & Brooks, Pacific Grove, CA, 1990.

[CD88] J.F. Canny and B.R. Donald. Simplified voronoi diagrams. *Discrete and Computational Geometry*, 3:219–236, 1988.

[Cho96] H. Choset. *Sensor Based Motion Planning: The Hierarchical Generalized Voronoi Graph*. PhD thesis, California Institute of Technology, 1996.

[Chu60] K.L. Chung. *Markov chains with stationary transition probabilities*. Springer Publisher, Berlin, 1960.

[CKK95] E. Chown, S. Kaplan, and D. Kortenkamp. Prototypes, location, and associative networks (plan): Towards a unified theory of cognitive mapping. *Cognitive Science*, 19:1–51, 1995.

[CL85] R. Chatila and J.-P. Laumond. Position referencing and consistent world modeling for mobile robots. In *Proceedings of the 1985 IEEE International Conference on Robotics and Automation*, 1985.

[CNT98] H. Choset, K. Nagatani, and S. Thrun. Towards exact localization without explicit localization: The topological voronoi graph. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 1998.

[Cox91] I.J. Cox. Blanche—an experiment in guidance and navigation of an autonomous robot vehicle. *IEEE Transactions on Robotics and Automation*, 7(2):193–204, 1991.

[CW90] I.J. Cox and G.T. Wilfong, editors. *Autonomous Robot Vehicles.* Springer Verlag, 1990.

[DLR77] A.P. Dempster, A.N. Laird, and D.B. Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society, Series B*, 39(1):1–38, 1977.

[Elf87] A. Elfes. Sonar-based real-world mapping and navigation. *IEEE Journal of Robotics and Automation*, RA-3(3):249–265, June 1987.

[Elf89] A. Elfes. *Occupancy Grids: A Probabilistic Framework for Robot Perception and Navigation*. PhD thesis, Department of Electrical and Computer Engineering, Carnegie Mellon University, 1989.

[EM92] S. Engelson and D. McDermott. Error correction in mobile robot map learning. In *Proceedings of the 1992 IEEE International Conference on Robotics and Automation*, pages 2555–2560, Nice, France, May 1992.

[FB96a] T. Fröhlinghaus and J.M. Buhmann. Real-time phase-based stereo for a mobile robot. In *Proceedings of the 1st Euromicro Workshop on Advanced Mobile Robots*. IEEE Computer Society Press, 1996.

[FB96b] T. Fröhlinghaus and J.M. Buhmann. Regularizing phase-based stereo. In *Proceedings of the 13th International Conference on Pattern Recognition*, Vienna, Austria, 1996.

[FBT96] D. Fox, W. Burgard, and S. Thrun. Controlling synchro-drive robots with the dynamic window approach to collision avoidance. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'96)*, 1996.

[FBT97] D. Fox, W. Burgard, and S. Thrun. The dynamic window approach to collision avoidance. *IEEE Robotics and Automation*, 4(1), 1997.

[FBT98a] D. Fox, W. Burgard, and S. Thrun. Active markov localization for mobile robots. *Robotics and Autonomous Systems*, 1998. To appear.

[FBT98b] D. Fox, W. Burgard, and S. Thrun. A hybrid collision avoidance method for mobile robots. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 1998.

[FKAM84] T. Fukuda, K. Kosuge, F. Arai, and H. Matsuura. *Reasearch Activities -1993-*. Laboratory of Robotics and Mechatronics, Department of Mechano-Informatics and Systems, Nagoya University, Japan, 1984.

[GN97] J.-S. Gutmann and B. Nebel. Navigation mobiler roboter mit laserscans. In *Autonome Mobile Systeme.* Springer Verlag, Berlin, 1997.

[Gut96] J.-S. Gutmann. Vergleich von algorithmen zur selbstlokalisierung eines mobilen roboters. Master's thesis, University of Ulm, Ulm, Germany, 1996. (in German).

[HBL98] D. Haehnel, L. Burgard, and G. Lakemeyer. GOLEX: Bridging the gap between logic (GOLOG) and a real robot. In *Proceedings of the 22st German Conference on Artificial Intelligence (KI 98)*, Bremen, Germany, 1998.

[IB98] M. Isard and A. Blake. Condensation: conditional density propagation for visual tracking. *International Journal of Computer Vision*, in press 1998.

[KB91] B. Kuipers and Y.-T. Byun. A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations. *Journal of Robotics and Autonomous Systems*, 8:47–63, 1991.

[KBM98] D. Kortenkamp, R.P. Bonassi, and R. Murphy, editors. *AI-based Mobile Robots: Case studies of successful robot systems*, Cambridge, MA, 1998. MIT Press.

[KCK96] L.P. Kaelbling, A.R. Cassandra, and J.A. Kurien. Acting under uncertainty: Discrete bayesian models for mobile-robot navigation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 1996.

[KF98] D. Koller and R. Fratkina. Using learning for approximation in stochastic processes. In *Proceedings of the International Conference on Machine Learning (ICML)*, 1998.

[KS96] S. Koenig and R. Simmons. Passive distance learning for robot navigation. In L. Saitta, editor, *Proceedings of the Thirteenth International Conference on Machine Learning*, 1996.

[KW94] D. Kortenkamp and T. Weymouth. Topological mapping for mobile robots using a combination of sonar and vision sensing. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*, pages 979–984, Menlo Park, July 1994. AAAI, AAAI Press/MIT Press.

[LDWC92] J.J. Leonard, H.F. Durrant-Whyte, and I.J. Cox. Dynamic map building for an autonomous mobile robot. *International Journal of Robotics Research*, 11(4):89–96, 1992.

[LLdMS97] M. Lopez, R. Lopez de Mantaras, and C. Sierra. Incremental map generation by low cost robots based on possibility/necessity grids. In *Proceedings of the Thirteenth International Conference on Uncertainty in AI*, pages 351–357, Providence, RI, August 1997.

[LM97] F. Lu and E. Milios. Globally consistent range scan alignment for environment mapping. *Autonomous Robots*, 4:333–349, 1997.

[Mat90] M. J. Matarić. A distributed model for mobile robot environment-learning and navigation. Master's thesis, MIT, Cambridge, MA, January 1990. also available as MIT AI Lab Tech Report AITR-1228.

[Mat97] M. J Matarić. Reinforcement learning in the multi-robot domain. *Autonomous Robots*, 4(1):73–83, January 1997.

[Moo90] A. W. Moore. *Efficient Memory-based Learning for Robot Control*. PhD thesis, Trinity Hall, University of Cambridge, England, 1990.

[Moo98] A.W. Moore. Very fast EM-based mixture model clustering using multiresolution kd-trees. In *Advances in Neural Information Processing Systems (NIPS)*, Cambridge, MA, 1998. MIT Press.

[Mor88] H. P. Moravec. Sensor fusion in certainty grids for mobile robots. *AI Magazine*, pages 61–74, Summer 1988.

[MT98] D. Margaritis and S. Thrun. Learning to locate an object in 3d space from a sequence of camera images. In *Proceedings of the International Conference on Machine Learning (ICML)*, 1998.

[NPB95] I. Nourbakhsh, R. Powers, and S. Birchfield. DERVISH an office-navigating robot. *AI Magazine*, 16(2):53–60, Summer 1995.

[OHD97] S. Oore, G.E. Hinton, and G. Dudek. A mobile robot that learns its place. *Neural Computation*, 9:683–699, 1997.

[Par96] L. E. Parker. On the design of behavior-based multi-robot teams. *Journal of Advanced Robotics*, 10(6), 1996.

[PK94] D. Pierce and B. Kuipers. Learning to explore and build maps. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*, pages 1264–1271, Menlo Park, July 1994. AAAI, AAAI Press/MIT Press.

[Put94] M.L. Puterman. *Markov Decision Processes.* John Wiley & Sons, New York, 1994.

[Rei79] J.H. Reif. Complexity of the mover's problem and generalizations. In *Proceesings of the 20th IEEE Symposium on Foundations of Computer Science*, pages 421–427, 1979.

[Ren93] W.D. Rencken. Concurrent localisation and map building for mobile robots using ultrasonic sensors. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2129–2197, Yokohama, Japan, July 1993.

[RJ86] L.R. Rabiner and B.H. Juang. An introduction to hidden markov models. In *IEEE ASSP Magazine*, 1986.

[Sam89a] H. Samet. *Applications of Spatial Data Structures.* Addison-Wesley Publishing Inc., 1989.

[Sam89b] H. Samet. *The Design and Analysis of Spatial Data Structures.* Addison-Wesley Publishing Inc., 1989.

[Sim95] R. Simmons. The 1994 AAAI robot competition and exhibition. *AI Magazine*, 16(1), Spring 1995.

[SK95] R. Simmons and S. Koenig. Probabilistic robot navigation in partially observable environments. In *Proceedings of IJCAI-95*, pages 1080–1087, Montreal, Canada, August 1995. IJCAI, Inc.

[SK97] H Shatkay and L. Kaelbling. Learning topological maps with weak local odometric information. In *Proceedings of IJCAI-97*. IJCAI, Inc., 1997.

[TBB+98] S. Thrun, A. Bücken, W. Burgard, D. Fox, T. Fröhlinghaus, D. Henning, T. Hofmann, M. Krell, and T. Schmidt. Map learning and high-speed navigation in RHINO. In D. Kortenkamp, R.P. Bonasso, and R Murphy, editors, *AI-based Mobile Robots: Case Studies of Successful Robot Systems.* MIT Press, 1998.

[TFB98a] 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.

[TFB98b] 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.

[TGF+98] S. Thrun, J.-S. Gutmann, D. Fox, W. Burgard, and B. Kuipers. Integrating topological and metric maps for mobile robot navigation: A statistical approach. In *Proceedings of the AAAI Thirteenth National Conference on Artificial Intelligence*, 1998.

[Thr93] S. Thrun. Exploration and model building in mobile robot domains. In E. Ruspini, editor, *Proceedings of the IEEE International Conference on Neural Networks*, pages 175–180, San Francisco, CA, 1993. IEEE Neural Network Council.

[Thr97a] S. Thrun. The museum tourguide project: Experiences with a deployed service robot. In *Proceedings of the IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA)*, Monterey, CA, 1997.

[Thr97b] S. Thrun. To know or not to know: On the utility of models in mobile robotics. *AI Magazine*, 18(1):47–54, 1997.

[Thr98a] S. Thrun. Bayesian landmark learning for mobile robot localization. *Machine Learning*, 33(1), 1998.

[Thr98b] S. Thrun. Learning metric-topological maps for indoor mobile robot navigation. *Artificial Intelligence*, 99(1):21–71, 1998.

[Thr98c] S. Thrun. When robots meet people: Research directions in mobile robotics. *IEEE Intelligent Systems*, 1998.

[Tor94] M. C. Torrance. Natural communication with robots. Master's thesis, MIT Department of Electrical Engineering and Computer Science, Cambridge, MA, January 1994.

[VN94] VDI-N. Serviceroboter verlassen die Fabrik. *VDI nachrichten, Nr. 15, S11, 15.April*, 1994.

[Win95] G. Winkler. *Image Analysis, Random Fields, and Dynamic Monte Carlo Methods*. Springer Verlag, Berlin, 1995.

[WTR98] S. Waldherr, S. Thrun, and R. Romero. A gesture-based interface for human-robot interaction. Submitted *Autonomous Robots*, 1998.

[WTRM98] S. Waldherr, S. Thrun, R. Romero, and D. Margaritis. Template-based recognition of pose and motion gestures on a mobile robot. In *Proceedings of the AAAI Thirteenth National Conference on Artificial Intelligence*, 1998.

[YB96] B. Yamauchi and R. Beer. Spatial learning for navigation in dynamic environments. *IEEE Transactions on Systems, Man, and Cybernetics - Part B: Cybernetics*, Special Issue on Learning Autonomous Robots, 1996. also located at http://www.aic.nrl.navy.mil/∼yamauchi/.

[YLS+98] B. Yamauchi, P. Langley, A.C. Schultz, J. Grefenstette, and W. Adams. Magellan: An integrated adaptive architecture for mobile robots. Technical Report 98-2, Institute for the Study of Learning and Expertise (ISLE), Palo Alto, CA, May 1998.

[Zim96] U.R. Zimmer. Robust world-modeling and navigation in a real world. *Neurocomputing*, 13(2–4), 1996.

BIBLIOGRAPHICAL SKETCH

Sebastian Thrun
Assistant Professor
Department of Computer Science, Robotics Institute, and
Center for Automated Learning and Discovery
Carnegie Mellon University
Pittsburgh, PA 15213-3891

Phone: (412) 268-8077
FAX: (412) 268-5576
E-mail: thrun@cs.cmu.edu

Education

- 1995: Ph.D., Computer Science and Statistics, University of Bonn, Germany.
- 1993: M.Sc., Computer Science and Mathematics, University of Bonn, Germany.
- 1988: B.Sc., Computer Science, Economics, Medicine, University of Hildesheim, Germany.

Professional Affiliations

- 1998-today. Assistant Professor, Computer Science Department and Center for Automated Learning and Discovery, Carnegie Mellon University.
- 1995-1998. Research Computer Scientist, Computer Science Department, Carnegie Mellon University.
- 1993-1995 Research Associate, University of Bonn, Germany.

Scientific Awards

- 1998: Outstanding Paper Award, National Conference on Artificial Intelligence (AAAI-98)
- 1996: First place award at the 5th AAAI autonomous mobile robot contest.
- 1994: Second place award at the 3rd AAAI autonomous mobile robot contest.
- 1992: Third place award at the 1st AAAI autonomous mobile robot contest.

Selected Professional Activities

- General Chair, Conference on Automated Learning and Discovery, CONALD'98.
- Editorial Board Member, Journal for Artificial Intelligence Research and Neural Computing Surveys. Guest editor, Machine Learning Journal, special issue on "Inductive Transfer," special issue on "Robot Learning."
- Editor, "Learning To Learn" and "Recent Advances in Robot Learning" both Kluwer Academic Publishers.
- Program committee member for various conferences and workshops.

Five Related Publications

- S. Thrun, D. Fox, and W. Burgard. A Probabilistic Approach to Concurrent Mapping and Localization for Mobile Robots. *Machine Learning* 31, and *Autonomous Robots* 5, (joint issue), 1998.
- S. Thrun. Bayesian Landmark Learning for Mobile Robot Localization. *Machine Learning* 33(1), 1998
- S. Thrun. Learning Metric-Topological Maps for Indoor Mobile Robot Navigation, *Artificial Intelligence* 99(1), 21–71, 1998.
- S. Thrun. When Robots Meet People: Research Directions In Mobile Robotics. *IEEE Intelligent Systems*, May/June 1998.
- 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 the AAAI National Conference on Artificial Intelligence, 1998.

Five Other Publications

- W. Burgard, A.B. Cremers, D. Fox, D. Haehnel, G. Lakemeyer, D. Schulz, W. Steiner, S. Thrun. The Interactive Museum Tour-Guide Robot. In Proceedings of the AAAI National Conference on Artificial Intelligence. Outstanding paper award, 1998.
- S. Waldherr, S. Thrun, R. Romero, and D. Margaritis. Template-Based Recognition of Pose and Motion Gestures On a Mobile Robot. In Proceedings of the AAAI National Conference on Artificial Intelligence, 1998.
- 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.
- D. Fox, W. Burgard, and S. Thrun. Active Markov Localization for Mobile Robots. To appear in *Robotics and Autonomous Systems*, 1998.
- K. Nigam, M. McCallum, S. Thrun, and T. Mitchell. Learning to Classify Text from Labeled and Unlabeled Documents. To appear in *Machine Learning*, 1998.

Recent Collaborators, advisors, and advisees

- Hans Berliner, Joseph O'Sullivan, Pradeep Khosla, Frank Pfenning, Han Kiliccote, David Touretzky, Reid Simmons, Howie Choset, Roseli Romero, Andrew McCallum, Benjamin Kuipers, Dieter Fox, Wolfram Burgard, Gerhard Lakemeyer, Dirk Hähnel, Dirk Schulz, Maren Bennewitz
- Armin B. Cremers, Joachim Buhmann, and Tom Mitchell (advisors)
- Frank Dellaert, John Langford, Joelle Pineau, Dimitris Magaritis, Jamie Schulte, Chuck Rosenberg, Nicholas Roy, Stefan Waldherr, Corinna Richter, Jörg Stachowiak, Frank Schneider, and Peter Wallossek (advisees)