Sebastian Thrun

PROBABILISTIC ROBOTICS

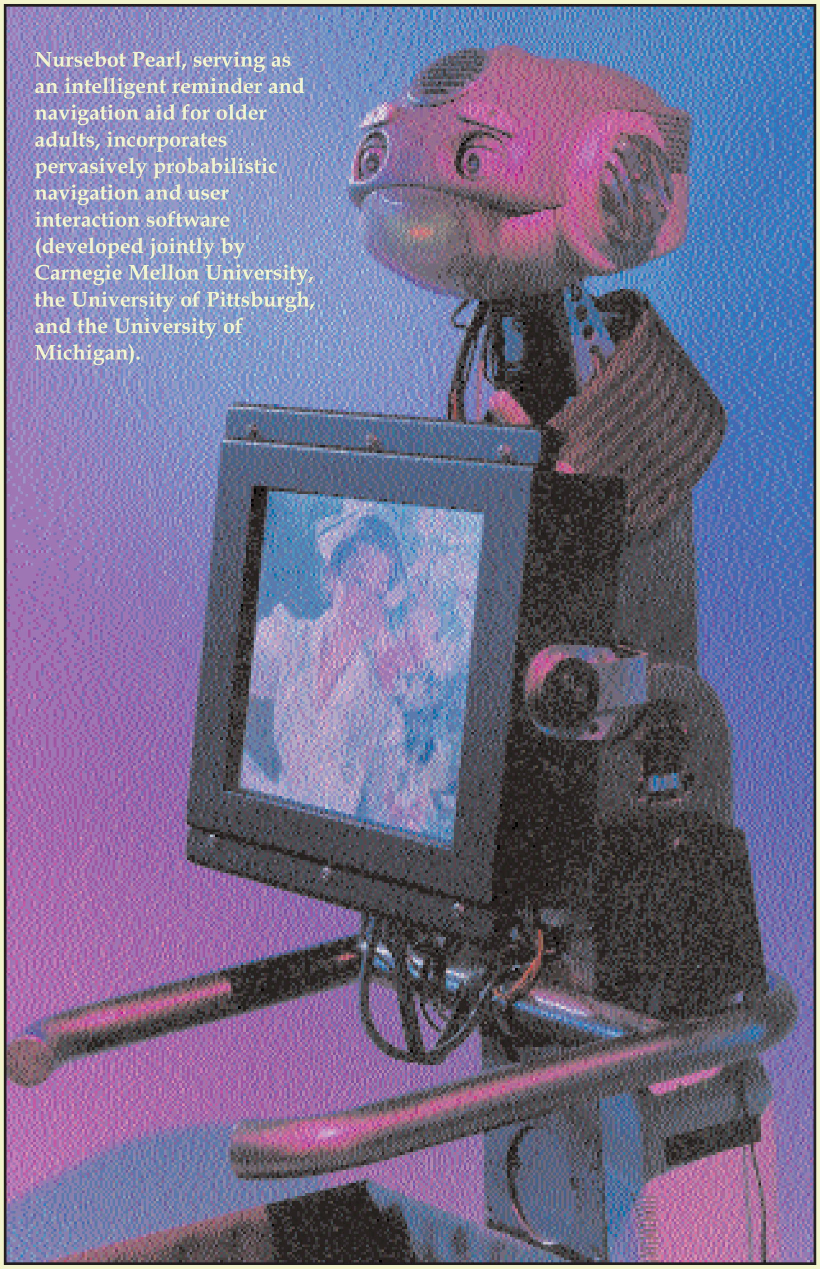

Nursebot Pearl, serving as an intelligent reminder and navigation aid for older adults, incorporates pervasively probabilistic navigation and user interaction software (developed jointly by Carnegie Mellon University, the University of Pittsburgh, and the University of Michigan).

CARNEGIE MELLON UNIVERSITY

Planning and navigation algorithms exploit statistics gleaned from uncertain, imperfect real-world environments to guide robots toward their goals and around obstacles.

In the past decade, the field of robotics has made notable progress in terms of flexibility, autonomy, and human interaction. Robots had been largely confined to factory floors and assembly lines, bound to perform the same narrow tasks over and over again. A recent series of successful robot systems has, however, demonstrated that robotics has advanced to where it is ready to conquer many new fields, including space exploration, medical applications, personal services, entertainment, and military support. Many of them are highly dynamic and uncertain. Uncertainty arises for many reasons, including the inherent limitations of a particular model of the world, the noise and perceptual limitations in a robot's sensor measurements, and the approximate nature of many algorithmic solutions. Finding mechanisms for handling this uncertainty is a primary challenge faced by robotics research today.

Figure 1 shows three examples of successful robot systems operating in uncertain environments: a commercially deployed autonomous straddle carrier [3]; an interactive museum tour-guide robot [7, 11]; and a prototype robotic assistant called Nursebot for the elderly. The straddle carrier, developed at the University of Sydney, is capable of transporting containers faster than trained human operators. The tour-guide robot, one in a series, safely guides visitors through densely crowded museums, including the Smithsonian Museum of Natural History in Washington, D.C. The Nursebot is being developed jointly by Carnegie Mellon University, the University of Pittsburgh, and the University of Michigan to interact with elderly people and assist them in their daily tasks. These robots all have to cope with uncertainty. For example, the straddle carrier faces intrinsic limitations when sensing its own location and that of the containers. A similar problem is faced by the museum tour guide, further aggravated by the presence of people. The Nursebot faces the added uncertainty of having to understand spoken language by elderly people and coping with their inability to clearly express their wishes. In all these application domains, the environments are highly unpredictable, and sensors are comparatively poor with regard to the performance tasks at hand.

As these examples suggest, the ability to accommodate uncertainty raises a number of key questions. For example, what type of internal world models should robots employ? How should sensor measurements be integrated into their internal states of information? How should they make decisions, even when uncertain about the world's most basic state variables?

The probabilistic approach addresses them through a single idea: how to represent information probabilisticly. In particular, world models in the probabilistic approach are conditional probability distributions describing the dependence of certain variables on other variables in probabilistic terms. A robot's state of knowledge is also represented by probability distributions derived by integrating sensor measurements into the probabilistic world models given to the robot. Probabilistic robot control anticipates various contingencies that might arise in uncertain worlds, thereby seamlessly blending information gathering (exploration) with robust performance-oriented control (exploitation).

The move to probabilistic techniques in robotics is paralleled in many other subfields of artificial intelligence, including computer vision, natural language understanding, and speech recognition. In the 1970s, most research in robotics presupposed the availability of exact models of robots and their environments.

Figure 1. Three robots controlled by probabilistic software: a robotic straddle carrier; a museum tour guide; and the Nursebot, a robotic assistant for the elderly.

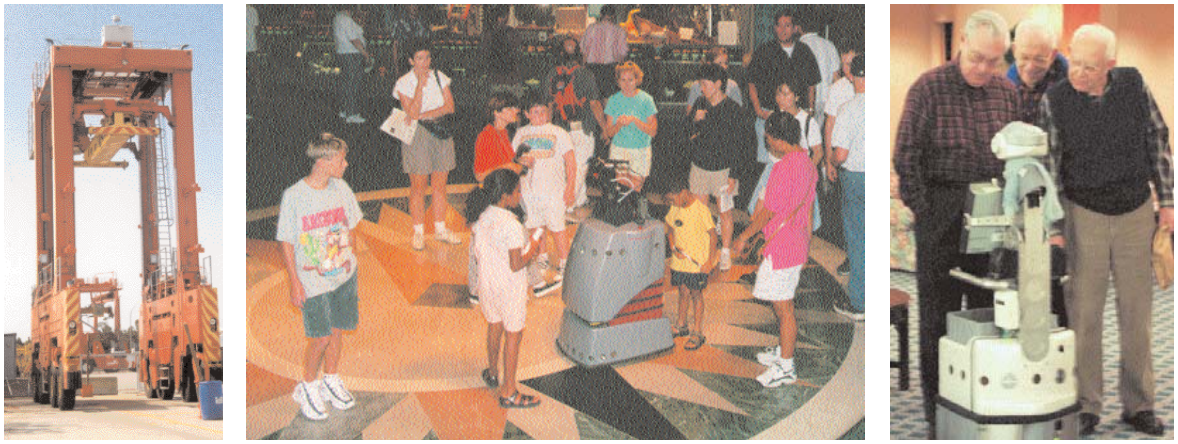

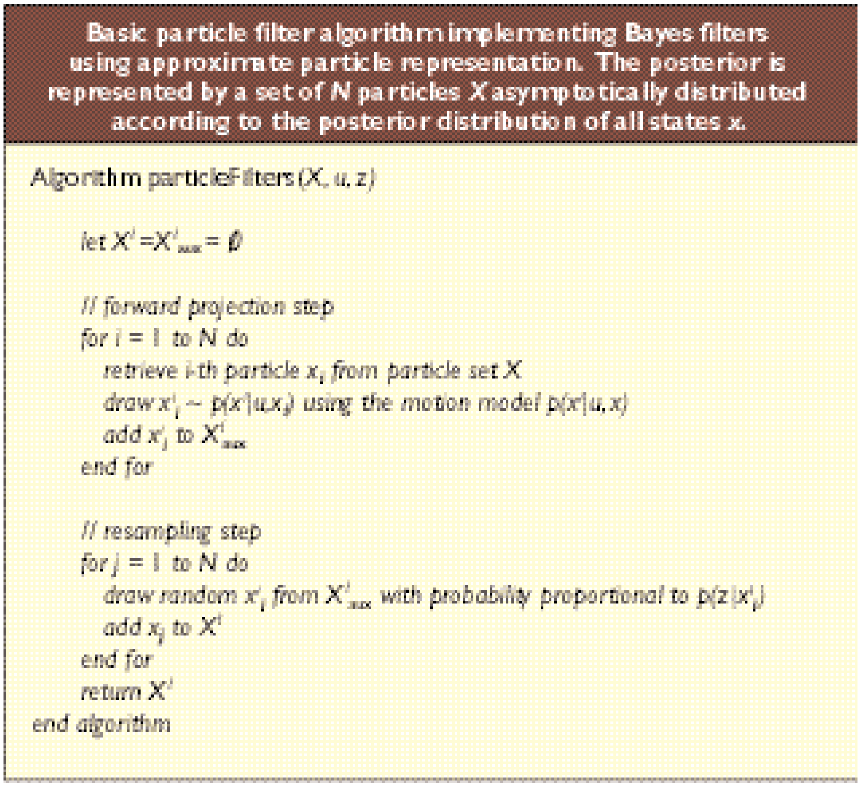

Basic particle filter algorithm implementing Bayes filters using approximate particle representation. The posterior is represented by a set of *N* particles *X* asymptotically distributed according to the posterior distribution of all states *x*.

```
Algorithm particleFilters($X, u, z$)

    let $X' = X'_{aux} = \emptyset$

    // forward projection step
    for $i$ = 1 to N do
        retrieve i-th particle $x_i$ from particle set $X$
        draw $x'_i \sim p(x'|u,x_i)$ using the motion model $p(x'|u,x)$
        add $x'_i$ to $X'_{aux}$
    end for

    // resampling step
    for $j$ = 1 to N do
        draw random $x'_j$ from $X'_{aux}$ with probability proportional to $p(z|x'_j)$
        add $x_j$ to $X'$
    end for
    return $X'$
end algorithm
```

Little emphasis was placed on sensing and the intrinsic limitations of modeling complex physical phenomena. This direction changed in the mid-1980s when the paradigm shifted toward reactive techniques. Reactive controllers rely on sensors to generate robot control. Probabilistic robotics has emerged since the mid-1990s, leveraging decades of research in probability theory, statistics, engineering, and operations research. More recently, probabilistic techniques have solved many outstanding robotics problems, leading to new theoretical insights into the structure of robotics problems and their solutions.

Models, Sensors, and the Physical World

Classical robotics textbooks usually describe at length the kinematics and dynamics of robotic devices, addressing the question of how controls affect the state of the robot and, more broadly, the world. However, the textbooks often suggest a deterministic relationship: The effect of applying control action u to the robot at state x is governed by the functional relationship $x'=f(u,x)$, for some (deterministic) function f. For example, x might be the configuration and velocity of a robotic arm, and u might be the motor currents asserted in a fixed time interval. Such an approach characterizes only idealized robots—free of wear and tear, inaccuracies, control noise, and the like. Operating in the real world, the outcomes of control actions are uncertain. For example, a robot that executes a control, leveraging, say, its position by one meter forward, might expect to be exactly one meter away from where it started but in reality will likely find itself in an unpredictable location nearby. The probabilistic approach accounts for this uncertainty by using conditional probability distributions to model robots. Such models, commonly denoted $p(x'|u,x)$, specify the posterior probability over states x' that might result when applying control u to a robot whose state is x. That is, instead of making a deterministic prediction, probabilistic techniques model the fact that the outcome of robot controls is uncertain by assigning a probability distribution over the space of all possible outcomes. As such, they generalize classical kinematics and dynamics to real-world robotics.

In the same vein, many traditional textbooks presuppose that the state of the robot x be known at all times. Usually, the state x comprises all necessary quantities relevant to robot prediction and control, including the robot's configuration, pose, and velocity, as well as the location of surrounding items, including physical obstacles and people. In an idealized world, the robot might incorporate sensors that can measure, without error, the state x. Such sensors may be characterized by a deterministic function g capable of recovering the full state from sensor measurements z, that is, $x=g(z)$. Real sensors are characterized by noise and, more important, by range limitations. For example, cameras cannot see through walls. The probabilistic approach generalizes this idealized view by modeling robot sensors by conditional probability distributions. Sensors may be characterized by forward models $p(z|x)$, reasoning from state to sensor measurements, or their inverse $p(x|z)$, depending on algorithmic details beyond the scope of this article.

As these ideas suggest, probabilistic models are indeed generalizations of their classical counterparts. But the explicit modeling of uncertainty raises fundamental questions as to what can be done with these world models. For instance, can we recover the state of the world? And can we still control robots so as to achieve set goals?

Probabilistic State Estimation

A first answer to these questions can be found in the rich literature on probabilistic state estimation, addressing the problem of recovering the state variables x from

Figure 2. Evolution of the conditional particle filter, from global uncertainty to successful localization and tracking.

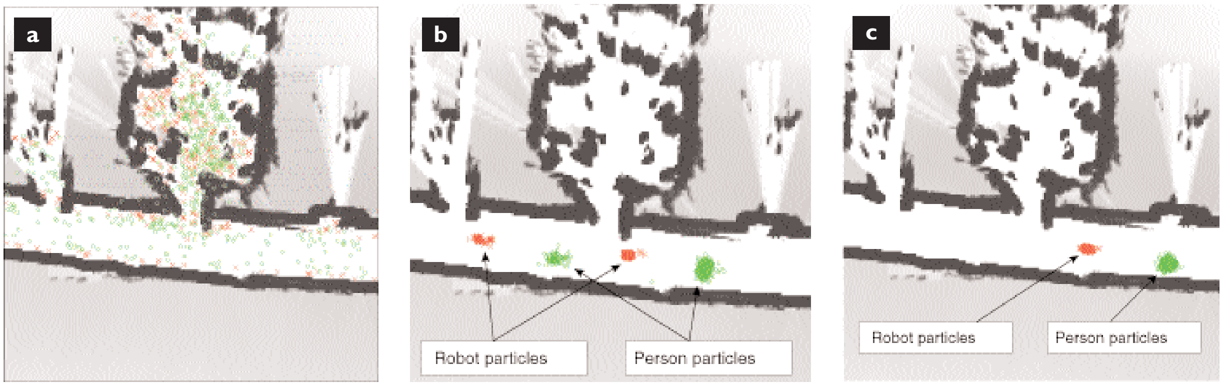

sensor data. Common state variables include:

- Parameters regarding the robot's configuration, such as its location relative to an external coordinate frame. The problem of estimating them is often referred to as localization;
- Parameters specifying the location of items in the environment, such as the location of walls, doors, and objects of interest. This problem, known as mapping, is regarded as one of the most difficult state estimations due to the high dimensionality of the parameter spaces [1]; and
- Parameters of objects whose positions change over time, including people, doors, and other robots. This problem is like the mapping problem, with the added difficulty of changing locations over time.

The predominant approach for state estimation in probabilistic robotics is Bayes filters, which offer a methodology for estimating a probability distribution over the state *x*, conditioned on all available data (controls and sensor measurements). They do so recursively, based on the most recent control *u* and measurement *z*, the previous probabilistic estimate of the state, and the probabilistic models $p(x'|x,u)$ and $p(z|x)$ discussed earlier. Thus, Bayes filters do not just "guess" the state *x*, they calculate the probability that any state *x* is correct. Popular examples of Bayes filters used for interpreting streams of data in applications, such as speech recognition, are hidden Markov models, Kalman filters, dynamic Bayes networks, and partially observable Markov decision processes [5, 10].

For low-dimensional state spaces, research in robotics and applied statistics has produced a wealth of literature on efficient probabilistic estimation. Remarkably popular is an algorithm known as particle filters (in computer vision also known as the condensation algorithm and in robotics as Monte Carlo localization) [2]. The particle filter algorithm approximates the desired posterior distribution through a set of particles. Particles are samples of states *x* distributed roughly according to the very posterior probability distribution specified by Bayes filters. The table here shows the basic particle filtering algorithm. Analogous to Bayes filters, it generates a particle set *X'* recursively, from the most recent control *u*, the most recent measurement *z*, and the particle set *X* representing the probabilistic estimate before incorporating *u* and *z*. It does so in two phases: First, it guesses states *x* based on particles drawn from *X* and the probabilistic motion model $p(x'|u,x)$. These guesses are subsequently resampled in proportion to the perceptual likelihood $p(z|x)$. The resulting sample set is approximately distributed according to the Bayesian posterior, taking *u* and *z* into account.

Figure 2 illustrates particle filters; a mobile robot, equipped with a laser range finder, simultaneously estimates its location relative to a 2D map of a corridor environment and the number and locations of nearby people. In (2a), the robot is globally uncertain of where it is. Consequently, the particles representing its location and that of the person are spread throughout the free space in the map. As the robot moves (2b), the particles representing the robot's location quickly converge on two distinct locations in the corridor, as do the particles representing the person's location. A few time steps later, the ambiguity is resolved, and both sets of particles focus on the correct positions in the map (2c). Algorithms based on particle filters are arguably the most powerful algorithms in existence for mobile robot localization. As the example illustrates, particle filters can represent a wide range of multimodal distributions. They are easily implemented as a resource-adaptive algorithm, capable of adapting the number of particles to the available computational resources. Finally, they converge for a large range of distributions, from globally uncertain to near-deterministic cases.

Figure 3. (a) 3D volumetric map, acquired in real time by a mobile robot; the lower part of the map is below the robot's sensors, hence is not modeled. (b) Map of underwater landmarks acquired by the submersible vehicle Oberon (courtesy Stefan Williams and Hugh Durrant-Whyte, University of Sydney).

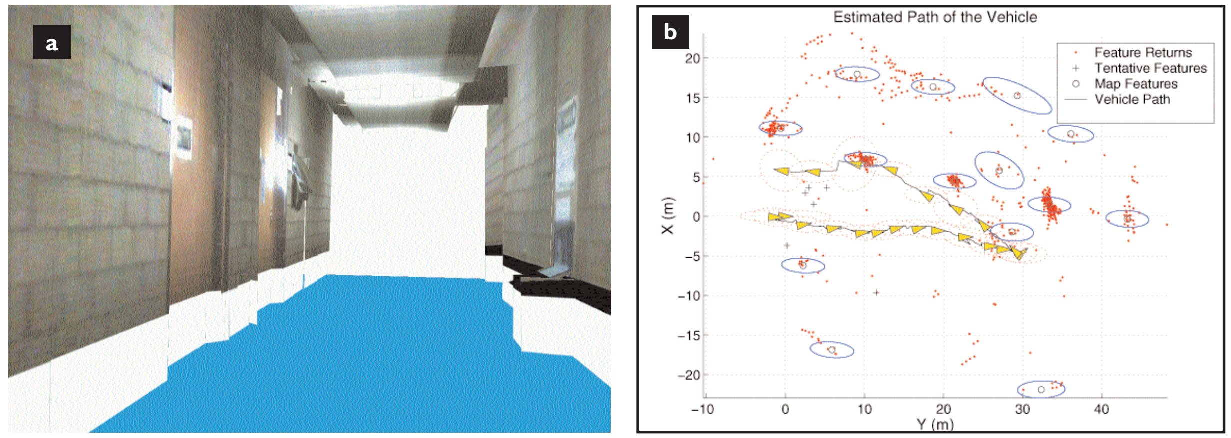

Toward Millions of Dimensions

In high-dimensional state spaces, computational considerations may pose serious obstacles when estimating state. Robot mapping, to cite a popular example of a high-dimensional problem, often involves thousands of dimensions, if not millions. For example, the volumetric map in Figure 3a consists of several million texture values, in addition to thousands of structural parameters. The computational considerations dealing with millions of dimensions raises the question of whether probabilistic techniques are equipped to perform state estimation in such high-dimensional spaces. The answer is intriguing. To date, virtually all state-of-the-art algorithms in localization, mapping, people tracking, and other areas are probabilistic.

Many probabilistic approaches estimate the mode of the posterior, or simply the most likely state x (there might be more than one). Some techniques, such as Kalman filters, also compute a covariance matrix to measure the curvature of the posterior at the mode. The specific techniques for estimating the mode and the covariance vary greatly, depending on the nature of the state estimation problem. In the robotic mapping problem, two of the most widely used algorithms are extended Kalman filters (EKFs) [5] and the expectation maximization (EM) algorithm [6]. Extended Kalman filters are applicable when the posterior is reasonably assumed to be Gaussian. The Gaussian posterior assumption is usually the case when mapping the locations of landmarks that can be uniquely identified. Kalman filter techniques have proved capable of mapping large-scale outdoor and underwater environments while simultaneously estimating the location of the robot relative to the map [1]. Figure 3b shows an example map of landmarks in an underwater environment obtained by researchers at the University of Sydney in Australia [12].

In the general mapping problem, the desired posterior may have exponentially many modes, not just one. Different modes commonly arise from uncertainty in calculating the correspondence between map items sensed at different points in time—a problem known as the data association problem. Many of today's best algorithms for state estimation with unknown data association are based on the EM algorithm [6], which performs a local hill-climbing search in the space of all states x (such as maps) with the aim of calculating the mode. The EM algorithm searches iteratively—by alternating a step that calculates expectations over the data association and related latent variables, followed by a step that computes a new mode under these fixed expectations. This search leads to a sequence of state estimates (such as maps) of increasing likelihood. In cases where both these steps can be calculated in closed form, EM can be highly effective for estimating the mode of complex posteriors. For example, the map in Figure 3a was generated through an online variant of EM, accommodating errors in the robot odometry and exploiting a Bayesian prior that biases the resulting maps toward planar surfaces [4]. In all these applications, probabilistic model selection techniques are employed for finding models of the "right" complexity.

Probabilistic Planning and Control

State estimation is only half the story. Clearly, the ultimate goal of any robotics software system is to control robotic devices. It is no surprise that probabilistic techniques specifically consider uncertainty when devising robot control. In doing so, they are robust to sensor noise and incomplete information. Probability theory provides a sound framework for active information gathering, smoothly blending exploration and exploitation for the control goals at hand.

Existing probabilistic control algorithms are mainly grouped into two categories: greedy and nongreedy. Each assumes the availability of a payoff function specifying the costs and benefits associated with the various control choices. Whereas greedy algorithms maximize the payoff for the immediate next time step, nongreedy algorithms consider entire sequences of controls, thereby maximizing the robot's (more appropriate) cumulative payoff. Clearly, nongreedy methods are more desirable from a performance point of view. However, the computational complexity of planning under uncertainty makes greedy algorithms welcome alternatives to their nongreedy counterparts, the reason they have found widespread application in practice.

The immediate next payoff is easily calculated by maximizing the conditional expectation of the payoff under the posterior probability over the state space. Thus, greedy techniques maximize a conditional expectation. For example, in the museum tour-guide project, this approach was employed to prevent the robot from falling down staircases. Similar techniques have been brought to bear in active-environment exploration with teams of robots [9] using payoff functions that measure the map's residual uncertainty.

Nongreedily optimizing robot control remains a challenging computational problem because the robot has to consider multiple contingencies during planning, paying tribute to the uncertainty in the world. Worse, the number of contingencies may increase exponentially with the planning horizon, making for a challenging planning problem [10].

Nevertheless, recent research in probabilistic robot control has led to a flurry of approximate algorithms that are computationally efficient. The coastal navigation algorithm described in [8] condenses the posterior belief to two quantities: the most likely state and the entropy of the posterior. This state space representation is exponentially more compact than the space of all posterior distributions. Still, it manages to capture the degree of uncertainty in the robot's posterior. Planning with this condensed state space has led to scalable robotic planning systems that cope with uncertainty. For example, in a mobile robot implementation reported in [8], the technique was found to navigate robots close to known landmarks to minimize the danger of getting lost, even though the strategy might also increase the overall path length. Experimentally, coastal navigation was shown to be superior to motion planners that disregard uncertainty in the planning process in environments densely populated with people and obstacles. This and many other examples in the literature illustrate how careful consideration of uncertainty often leads to superior control algorithms that explicitly consider uncertainty in planning and control.

Conclusion

This article offered an introduction to the vibrant field of probabilistic robotics. The key idea is a commitment to probability distribution as the basic representation of information. The related approaches provide sound solutions for the integration of inaccurate model information and noisy sensor data.

Probabilistic robotics is a rapidly growing subfield of robotics because of its ability to accommodate the sensor noise and uncertainty that arises naturally in robotics. While many research challenges remain, it has already led to fundamentally more scalable solutions to many difficult problems, specifically in the area of mobile robotics. Moreover, it has led to deep mathematical insights into the structure of the problems and their solutions, while probabilistic techniques have proved their value in practice.

References

1. Dissanayake, G., Newman, P., Clark, S., Durrant-Whyte, H., and Csorba, M. An experimental and theoretical investigation into simultaneous localization and map building (SLAM). In *Lecture Notes in Control and Information Sciences: Experimental Robotics VI,* P. Corke and J. Trevelyan, Eds. Springer Verlag, London, 2000, 265–274.
2. Doucet, A., de Freitas, J., and Gordon, N., Eds. *Sequential Monte Carlo Methods in Practice.* Springer Verlag, New York, 2001.
3. Durrant-Whyte, H. Autonomous guided vehicle for cargo handling applications. *Int. J. Robo. Res. 15,* 5 (1996).
4. Liu, Y., Emery, R., Chakrabarti, D., Burgard, W., and Thrun, S. Using EM to learn 3D models with mobile robots. In *Proceedings of the International Conference on Machine Learning* (Williams College, Williamstown, MA, June 28–July 1). Morgan Kaufmann Publishers, San Mateo, CA, 2001.
5. Maybeck, P. *Stochastic Models, Estimation, and Control, Vol. 1.* Academic Press, Inc., New York, 1979.
6. McLachlan, G. and Krishnan, T. *The EM Algorithm and Extensions.* Wiley Series in Probability and Statistics, John Wiley & Sons, Inc., New York, 1997.
7. Nourbakhsh, I., Bobenage, J., Grange, S., Lutz, R., Meyer, R., and Soto, A. An affective mobile robot with a full-time job. *Artif. Intelli. 114,* 1–2 (1999), 95–124.
8. Roy, N. and Thrun, S. Coastal navigation with mobile robot. In *Proceedings of the Conference on Neural Information Processing Systems* (Denver, Dec.). MIT Press, Cambridge, MA, 1999.
9. Simmons, R., Apfelbaum, D., Burgard, W., Fox, M., Moors, D., Thrun, S., and Younes, H. Coordination for multi-robot exploration and mapping. In *Proceedings of the AAAI National Conference on Artificial Intelligence* (Austin, TX, July). AAAI, Menlo Park, CA, 2000.
10. Sondik, E. *The Optimal Control of Partially Observable Markov Processes.* Ph.D. thesis, Stanford University, 1971.
11. Thrun, S., Beetz, M., Bennewitz, M., Burgard, W., Cremers, A., Dellaert, F., Fox, D., Haehnel, D., Rosenberg, C., Roy, N., Schulte, J., and Schulz, D. Probabilistic algorithms and the interactive museum tour-guide robot Minerva. *Int. J. Robo. Res. 19,* 11 (2000), 972–999.
12. Williams, S., Dissanayake, G., and Durrant-Whyte, H. Towards terrain-aided navigation for underwater robotics. *Adv. Robo. 15,* 5 (2001).

Sebastian Thrun (thrun@cs.cmu.edu) is a professor in the School of Computer Science of Carnegie Mellon University, Pittsburgh, PA.

The research described here is supported by the U.S. Defense Advanced Research Projects Agency (TMR, MARS, CoABS, and MICA programs) and the U.S. National Science Foundation (ITR, Robotics, and CAREER programs).