# Is Robotics Going Statistics? The Field of Probabilistic Robotics

Sebastian Thrun School of Computer Science Carnegie Mellon University http://www.cs.cmu.edu/~thrun

draft, please do not circulate

#### Abstract

In the 1970s, most research in robotics presupposed the availability of exact models, of robots and their environments. Little emphasis was placed on sensing and the intrinsic limitations of modeling complex physical phenomena. This changed in the mid-1980s, when the paradigm shifted towards reactive techniques. Reactive controllers rely on capable sensors to generate robot control. Rejections of models were typical for researchers in this field. Since the mid-1990s, a new approach has begun to emerge: probabilistic robotics. This approach relies on statistical techniques to seamlessly integrate imperfect models and imperfect sensing. The present article describes the basics of probabilistic robotics and highlights some of its recent successes.

#### 1 Introduction

In recent years, the field of robotics has made substantial progress. In the past, robots were mostly 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, however, has demonstrated that robotics has advanced to a level where it is ready to conquer many new fields, such as space, medical domains, personal services, entertainment, and military applications. Many of these new domains are highly dynamic and uncertain. Uncertainty arises for many different reasons: the inherent limitations to model the world, noise and perceptual limitations in a robot's sensor measurements, and the approximate nature of many algorithmic solutions. In this uncertainly lies one of the primary challenges faced by robotics research today.

Three examples of successful robot systems that operate in uncertain environments are shown in Figure 1: a commercially deployed autonomous straddle carrier [3], an interactive museum tourguide robot [7, 11], and a prototype robotic assistant for the elderly. The straddle carrier is capable of transporting containers faster than trained human operators. The tourguide robot—one in a series of many—can safely guide visitors through densely crowded museums. The Nursebot robot is presently being developed to interact with elderly people and assist them in various daily tasks. All of these robots have to cope with uncertainty. The straddle carrier faces intrinsic limitations when sensing its own location and that of the containers. A similar problem is faced by the museum tourguide robot, but here the problem is aggravated by the presence of people. The elderly companion robot faces the additional uncertainty of having to understand spoken language by elderly people, and coping with their inability to express their exact wishes. In all these application domains, the environments are highly unpredictable, and sensors are comparatively poor with regard to the performance tasks at hand.







Figure 1: Three robots controlled by probabilistic software: A robotic straddle carrier, a museum tourguide robot, and the Nursebot, a robotic assistant for nurses and the elderly.

As these examples suggest, the ability to accommodate uncertainty is a key requirement for contemporary robotic systems. This raises the question as to appropriate mechanisms for coping with uncertainty. What type of internal world models should robots employ? And how should sensor measurements be integrated into their internal states of information? How should robots make decisions even if they are uncertain about even the most basic state variables in the world?

The probabilistic approach to robotics addresses these questions through a single key idea: *representing information probabilistically*. In particular, world models in the probabilistic approach are conditional probability distributions, which describe the dependence of certain variables on others in probabilistic terms. A robot's state of knowledge is also represented by probability distributions, which are 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, such as computer vision, language, and speech. Probabilistic robotics leverages decades of research in probability theory, statistics, engineering and operations research. In recent years, probabilistic techniques have solved many outstanding robotics problems, and they have led to new theoretical insights into the structure of robotics problems and their solutions.

## 2 Models, Sensors, and The Physical World

Classical robotics textbooks often describe at length the kinematics and dynamics of robotic devices. These topics address the question of how controls affect the state of the robot and, more broadly, the world. However, 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 idealized robots only—free of wear and tear, inaccuracies, control noise, and the alike. In reality, the outcomes of control actions are uncertain. For example, a robot that executes a control leveraging 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. Put differently, instead of making a deterministic prediction, probabilistic techniques model the fact that the outcome of robot controls

#### Algorithm particleFilters(X, u, z)

```
\begin{split} & \text{let } X' = X'_{\text{aux}} = \emptyset \\ & \text{// forward projection step} \\ & \text{for } i = 1 \text{ to } N \text{ do} \\ & \text{retrieve } i\text{-th particle } x_i \text{ from particle set } X \\ & \text{draw } x'_i \sim p(x'|u,x_i) \text{ using the motion model } p(x'|u,x) \\ & \text{add } x'_i \text{ to } X'_{\text{aux}} \\ & \text{end for} \\ & \text{// resampling step} \\ & \text{for } j = 1 \text{ to } N \text{ do} \\ & \text{draw random } x'_j \text{ from } X'_{\text{aux}} \text{ with probability proportional to } p(z|x'_j) \\ & \text{add } x_j \text{ to } X' \\ & \text{end for} \\ & \text{return } X' \\ & \text{end algorithm} \end{split}
```

Table 1: Basic particle filter algorithm, which implements Bayes filters using approximate particle representation. The posterior is represented by a set of N particles X', which is roughly distributed according to the posterior distribution of all states x given the data that is commonly calculated by Bayes filters.

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, such as the robot's configuration, its pose and velocity, the location of surrounding items (obstacles, people, etc.). In idealized worlds, the robot might possess 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 importantly, 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), which reason from state to sensor measurements, or their inverse p(x|z)—depending on algorithmic details beyond the scope of this article.

As this discussion suggests, probabilistic models are indeed generalizations of their classical counterparts. The explicit modeling of uncertainty, however, raises fundamental questions as to what can be done with these world models. Can we recover the state of the world? Can we still control robots so as to achieve set goals?

#### 3 Probabilistic State Estimation

A first answer to these questions can be found in the rich literature on probabilistic state estimation. This literature addresses the problem of recovering the state variables x from sensor data. Common state variables include

• parameters regarding the robot's configuration, such as its location relative to an external coordinate




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


frame. The problem of estimating such parameters 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 one of the most difficult state estimation due to the high dimensionality of such parameter spaces [1], and
- parameters of objects whose position changes over time, such as people, doors, and other robots. This problem is similar to the mapping problem, with the added difficulty changing locations over time.

The predominant approach for state estimation in probabilistic robotics is known as *Bayes filters*. Bayes filters 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 in the previous section. Thus, Bayes filter do not just "guess" the state x. Rather, they calculate the probability that any state x is correct. Popular examples of Bayes filters 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, which in computer vision is known as condensation algorithm and in robotics as Monte Carlo localization [2]. This algorithm approximates the desired posterior distribution through a set of particles. Particles are samples of states x which are distributed roughly according to the very posterior probability distribution specified by Bayes filters. Table 1 states the basic particle filtering algorithm. In analogy to Bayes filters, the algorithm generates a particle set X' recursively, from the most recent control u, the most recent measurement z, and the particle set X that represents the probabilistic estimate before incorporating u and z. It does so in two phases: First, it "guesses" states  $x_i'$  based on particles drawn from X and the probabilistic motion model p(x'|u,x). Subsequently, these guesses are resampled in proportion to the perceptual likelihood,  $p(z|x_i')$ . The resulting sample set is approximately distributed according to the Bayesian posterior, taking u and z into account.

Figure 2 illustrates particle filters via an example. A mobile robot, equipped with a laser range finder, simultaneously estimates its location relative to a two-dimensional map of a corridor environment and the number and locations of nearby people. In the beginning (Panel 2(a)), the robot is globally uncertain as to






Figure 3: (a) 3D volumetric map, acquired by a mobile robot in real-time. 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 at the University of Sydney. Courtesy of Stefan Williams and Hugh Durrant-Whyte.

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 (Panel 2(b)), the particles representing the robot's location quickly converge to 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, as shown in Panel 2(c). Localization algorithms based on particle filters are arguable the most powerful algorithms in existence. As this example illustrates, particle filters can represent a wide range of multi-modal distributions. They are easily implemented as *resource-adaptive algorithm*, capable of adapting the number of particles to the available computational resources. And finally, they converge for a large range of distributions, from globally uncertain to near-deterministic cases.

### 4 Towards Millions of Dimensions

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

Many probabilistic approaches estimate the mode of the posterior, which is simply the most likely state x (there might be more than one). Some techniques, such as Kalman filters, also compute a covariance matrix, which measures the curvature of the posterior at the mode. The specific techniques for estimating the mode and the covariance vary widely, 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 can reasonably assumed to be Gaussian. This is usually the case when mapping the locations of landmarks that can be uniquely identified. Kalman filter techniques have proven to be capable of mapping

large-scale outdoor and underwater environments while simultaneously estimating the location of the robot relative to the map [1]. Figure 3(b) shows an example map of landmarks in an underwater environment, obtained by researchers at the University of Sydney [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 commonly known as data association problem. Many of today's best algorithms for state estimation with unknown data association are based on the EM algorithm [6]. This algorithm performs a local hill-climbing search in the space of all states x (e.g., maps), with the aim of calculating the mode. The "trick" of the EM algorithm is to search 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 leads to a sequence of state estimates (e.g., maps) of increasing likelihood. In cases where both these steps can be calculated in closed form, EM can be a highly effective algorithm for estimating the mode of complex posteriors. For example, the map shown in Figure 3(a) has been generated through an on-line variant of the EM algorithm, accommodating errors in the robot odometry and exploiting a Bayesian prior that biases the resulting maps towards planar surfaces [4]. In all these applications, probabilistic model selection techniques are employed for finding models of the "right" complexity.

## 5 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 should come at no surprise that probabilistic techniques specifically take uncertainty into consideration when devising robot control. By 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 as most beneficial for the control goals at hand.

Existing probabilistic control algorithms can mainly be grouped into two categories: *greedy* and *non-greedy*. Both families assume the availability of a payoff function, which specifies the costs and benefits associated with the various control choices. Whereas greedy algorithms maximize the payoff for the immediate next time step, non-greedy algorithm consider entire sequences of controls, thereby maximizing the (more appropriate) cumulative payoff of the robot. Clearly, non-greedy methods are more desirable from a performance point of view. The computational complexity of planning under uncertainty, however, make greedy algorithms welcome alternatives that have found widespread applications 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. In the museum tourguide project, such an approach was successfully employed to prevent the robot from falling down staircases. Similar techniques have been successfully brought to bear for active environment exploration with teams of robots [9], using payoff functions that measure the residual uncertainty in the map.

Non-greedily optimizing robot control—over multiple time steps—remains a challenging computational problem. This is because the robot has to consider multiple contingencies during planning, paying tribute to the uncertainty in the world. Worse so, the number of contingencies may increase exponentially with the planning horizon, which makes for a most challenging planning problem [10].

Nevertheless, recent research 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. It captures, however, still the degree of uncertainty in

the robot's posterior. Planning with this condensed state space has led to scalable robotic planning systems that can cope with uncertainty. For example, in a mobile robot implementation reported in [8], this technique has been found to navigate robots closed to known landmarks, in order to minimize the danger of getting lost—even though this might increase the overall path length. Experimentally, coastal navigation was shown to be superior to motion planners that do not regard uncertainty in the planning process, in densely populated environments. This and many other examples in the literature illustrate how a careful consideration of uncertainty often leads to superior control algorithms, which explicitly consider uncertainty in planning and control.

### 6 Conclusion

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

To date, probabilistic robotics is one of the most rapidly growing subfield of robotics. While many research challenges remain, the approach has already led to fundamentally more scalable solutions to many hard robotics problems, specifically in the area of mobile robotics. They have led to deep mathematical insights into the structure of robotics problems and solutions, And finally, probabilistic techniques have proven their value in practice. They are at the core of dozens of successful robotic systems to date.

This article was necessarily brief, and the interested reader is invited to consult the rich literature on this topic. Additional introductory material can be found at the author's Web site http://www.cs.cmu.edu/~thrun.

## Acknowledgment

The author acknowledges the invaluable contributions by various the members of CMU's Robot learning lab. Generous financial support by DARPA (TMR, MARS, CoABS and MICA programs) and NSF (ITR, Robotics, and CAREER programs) is also gratefully acknowledged.

#### References

- [1] G. Dissanayake, P. Newman, S. Clark, H.F. Durrant-Whyte, and M. Csorba. An experimental and theoretical investigation into simultaneous localisation and map building (SLAM). In P. Corke and J. Trevelyan, editors, *Lecture Notes in Control and Information Sciences: Experimental Robotics VI*, pages 265–274, London, 2000. Springer Verlag.
- [2] A. Doucet, J.F.G. de Freitas, and N.J. Gordon, editors. *Sequential Monte Carlo Methods In Practice*. Springer Verlag, New York, 2001.
- [3] H.F. Durrant-Whyte. Autonomous guided vehicle for cargo handling applications. *International Journal of Robotics Research*, 15(5), 1996.
- [4] Y. Liu, R. Emery, D. Chakrabarti, W. Burgard, and S. Thrun. Using EM to learn 3D models with mobile robots. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2001.
- [5] P. Maybeck. Stochastic Models, Estimation, and Control, Volume 1. Academic Press, Inc, 1979.

- [6] G.J. McLachlan and T. Krishnan. *The EM Algorithm and Extensions*. Wiley Series in Probability and Statistics, New York, 1997.
- [7] I. Nourbakhsh, J. Bobenage, S. Grange, R. Lutz, R. Meyer, and A. Soto. An affective mobile robot with a full-time job. *Artificial Intelligence*, 114(1–2):95–124, 1999.
- [8] N. Roy and S. Thrun. Coastal navigation with mobile robot. In *Proceedings of Conference on Neural Information Processing Systems (NIPS)*, 1999. to appear.
- [9] R. Simmons, D. Apfelbaum, W. Burgard, M. Fox, D. an Moors, S. Thrun, and H. Younes. Coordination for multi-robot exploration and mapping. In *Proceedings of the AAAI National Conference on Artificial Intelligence*, Austin, TX, 2000. AAAI.
- [10] E. Sondik. *The Optimal Control of Partially Observable Markov Processes*. PhD thesis, Stanford University, 1971.
- [11] S. Thrun, M. Beetz, M. Bennewitz, W. Burgard, A.B. Cremers, F. Dellaert, D. Fox, D. Hähnel, C. Rosenberg, N. Roy, J. Schulte, and D. Schulz. Probabilistic algorithms and the interactive museum tour-guide robot minerva. *International Journal of Robotics Research*, 19(11):972–999, 2000.
- [12] S. Williams, G. Dissanayake, and H.F. Durrant-Whyte. Towards terrain-aided navigation for underwater robotics. *Advanced Robotics*, 15(5), 2001.