

A Sample of Monte Carlo Methods in Robotics and Vision

Frank Dellaert
dellaert@cc.gatech.edu
College of Computing, Georgia Institute of Technology
801 Atlantic Drive, Atlanta, GA, 30332 (USA)

Abstract

Approximate inference by sampling from an appropriately constructed posterior has recently seen a dramatic increase in popularity in both the robotics and computer vision community. In this paper, I will describe a number of approaches in which my co-authors and I have used Sequential Monte Carlo methods and Markov chain Monte Carlo sampling to solve a variety of difficult and challenging inference problems. Very recently, we have also used sampling over variable dimension state spaces to perform automatic model selection. I will present two examples of this, one in the domain of computer vision, the other in mobile robotics. In both cases Rao-Blackwellization was used to integrate out the variable dimension-part of the state space, and hence the sampling was done purely over the (combinatorially large) space of different models.

This paper describes joint work with many collaborators over the past 5 years, both at Carnegie Mellon University and at the Georgia Institute of Technology, including Dieter Fox, Sebastian Thrun, Wolfram Burgard, Zia Khan, Tucker Balch, Michael Kaess, Rafal Zboinski, and Ananth Ranganathan.

1 Introduction

Two very popular uses of sampling-based approximate inference are mobile robot localization [14, 23, 49, 50] and visual target tracking [30, 31]. Their widespread use is indicative of a larger trend in the robotics and computer vision community, both of which are increasingly borrowing methods from applied probability and machine learning to take on difficult inference problems. In this paper, I describe some of the work that my collaborators and myself have done in the last 5 years, and which I believe to be representative of an emerging class of research, in which ever more sophisticated modeling techniques from statistics are being used. In particular, in computer vision the use of Markov chain Monte Carlo sampling is gaining in popularity [21, 53, 15, 18, 16, 2], and in robotics the use of Rao-Blackwellized particle filters is noteworthy [20, 42, 41, 40].

Very recently, we have also used sampling over variable dimension state spaces to perform automatic model selection. I will present two examples of this, one in the domain of computer vision, the other in mobile robotics. In both cases Rao-Blackwellization was used to integrate out the variable dimension-part of the state space, and hence the sampling was done purely over the (combinatorially large) space of different models.



Figure 1: Monte Carlo localization of a mobile robot: the posterior density over the three-dimensional robot pose (x , y , and orientation) is represented by a set of weighted particles which is sequentially updated over time (a particle filter). (a) Initially, the position of the robot is totally unknown, and particles are distributed uniformly over the environment. (b) After a number of sonar and odometry measurements, the uncertainty is reduced to essentially a bimodal density in this symmetric environment.

2 Monte Carlo Localization

Two key problems in mobile robotics are global position estimation and local position tracking. Global position estimation is the ability to determine the robot's position in an a priori or previously learned map. Once a robot has been localized in the map, local tracking is the problem of keeping track of that position over time. Both these capabilities are necessary to enable a robot to execute useful tasks, such as office delivery or providing tours to museum visitors.

Kalman-filter based techniques have proven to be robust and accurate for keeping track of the robot's position. However, a Kalman filter cannot represent ambiguities and lacks the ability to globally (re-)localize the robot in the case of localization failures. To deal with these shortcomings, Burgard et al. [6] introduced a histogram-based localization approach, which can represent arbitrarily complex probability densities at fine resolutions. However, the computational burden and memory requirements of this approach were considerable. In addition, the grid-size and thereby also the precision at which it can represent the state had to be fixed beforehand.

In [14, 23, 49, 50] we introduced the **Monte Carlo Localization** method (MCL) that takes a sequential Monte Carlo approach to the robot localization problem. Here, as illustrated graphically in Figure 1, a weighted sample approximation of the posterior over robot pose is maintained over time by means of a particle filter. Particle filters were invented in the seventies [28] but were deemed unpractical at the time. However, they were recently rediscovered independently in the target-tracking [25], statistical [33] and computer vision literature [30, 31], and have gained enormously in popularity since. Partly, this is due to the ease by which they can be implemented and understood (as foreshadowed in the seminal paper by Smith and Gelfand [46]).

In the context of robot localization, using a particle filter has several other key advantages with respect to earlier work: (1) in contrast to Kalman filtering based techniques, it is able to represent multi-modal distributions and thus can *globally* localize a robot; (2) it drastically reduces the amount of memory required compared to histogram-based localization methods, and it can integrate measurements




Figure 2: (a) 20 ants are being tracked by an MCMC-based particle filter. Targets do *not* behave independently: whenever one ant encounters another, some amount of interaction takes place, and the behavior of a given ant before and after an interaction can be quite different. (b) The motion of ants during an interaction can be modeled using a Markov Random field, which is built on the fly at each time-step.

at a considerably higher frequency; (3) it is more *accurate* than histogram-based methods with a fixed cell size, as the state represented in the samples is not discretized. As a result, Monte Carlo Localization is now *the* most popular mobile robot localization method in use throughout the robotics community.

3 An MCMC-Based Particle Filter for Multi-Target Tracking

Sequential Monte Carlo techniques can also be applied to the domain of visual target tracking [30, 31], and we have specifically looked at the problem of tracking multiple, *interacting* targets. The classical tracking literature approaches the multi-target tracking problem by performing a data-association step after a detection step. Most notably, the multiple hypothesis tracker [44] and the joint probabilistic data association filter (JPDAF) [1, 22] are influential algorithms in this class. *Interacting* targets cause problems for these traditional approaches. The basic assumption on which all established data-association methods rely is that targets maintain their behavior before and after the targets visually merge. However, consider the example in Figure 2a, which shows 20 ants being tracked in a small arena. In this case, the targets do *not* behave independently: whenever one ant encounters another, some amount of interaction takes place, and the behavior of a given ant before and after an interaction can be quite different.

A Markov random field motion prior (see Figure 2b), built on the fly at each time step, can adequately model these interactions. Again our approach will be based on the particle filter [25, 30, 7]. The standard particle filter weights particles based on a likelihood score, and then propagates these weighted particles according to a motion model. Incorporating an MRF to model interactions is equivalent to adding an additional *interaction factor* to the importance weights in a joint particle filter.

However, a joint particle filter suffers from exponential complexity in the number of tracked targets, n , and computational requirements render the joint filter unusable for more than three or four targets. As a solution, we replace the traditional importance sampling step in the particle filter with a Markov chain Monte Carlo (MCMC) [24, 37] sampling step, which samples from the joint motion model. This approach has the appealing property that *the filter behaves as a set of individual particle filters when the targets are not interacting, but efficiently deals with complicated interactions when*

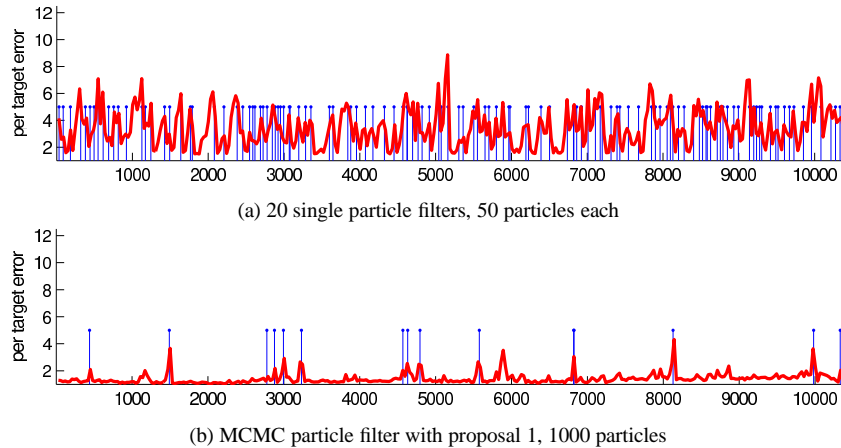


Figure 3: Qualitative comparison of (a) individual particle filters vs. (b) an MCMC-based joint filter, each tracking 20 ants using an equivalent sample size of 1000. Tick marks show when tracking failures occur throughout the sequence. The time series plot shows average distance from ground truth.

targets approach each other. Using MCMC in the sequential importance resampling (SIR) particle filter scheme has been explored before, in [4], and our approach can be consider a specialization of this work with an MRF-based joint posterior and an efficient proposal step to achieve reasonable performance.

The MCMC-based tracker performs significantly better than independent particle filters with a comparable number of samples, both in track quality and failures reported. Figure 3 shows the result graphically. An MCMC tracker with 1000 samples had 16 failures, compared to 125 for independent particle filters with 50 particles each. As a measure of trajectory quality, we recorded for each frame the average distance of the targets to their ground truth trajectories. This is shown in the figure as a time series.

4 Rao-Blackwellized EigenTracking



Figure 4: Subspace representations approximate the appearance of a target using a limited set of “eigen-images”, which are linearly combined according to normally distributed appearance coefficients. In a tracking context, this introduces a non-trivial continuous “appearance space” to be tracked as well.

Sequential Monte Carlo methods are traditionally based on importance sampling, and hence their performance degrades as the dimensionality of the state space increases. As discussed in the previous section, using MCMC sampling instead is one way to deal with the limited representational ability of SMC methods. Another approach is to analytically integrate out a part of the state space, a process which in this context is often referred to as Rao-Blackwellization [42, 8]

In our work, we have made use of this device to incorporate more complex appearance models in a particle filtering framework. In particular, subspace representations [52, 3, 13] have been a long-standing and popular way to model appearance and shape in computer vision. These methods model the

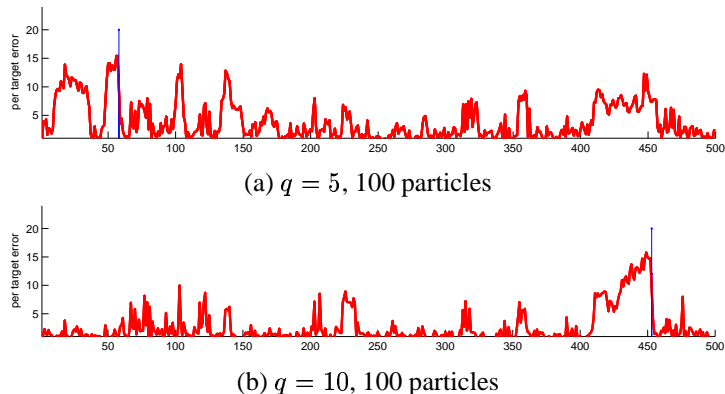


Figure 5: Qualitative comparison of trackers with PPCA measurement models of increasing complexity tracking 2 unmarked dancing bees in a 500 frame test sequence. q is the number of principal components in the PPCA probabilistic subspace model. Tick marks show when tracking failures occur.

density over a high-dimensional space of feature vectors using a generative model, where each vector is assumed to be a corrupted version of a linear combination of a small set of basis vectors. In figure 4, this is illustrated for a honey-bee tracking application. Subspace representations were also successfully used for tracking, e.g. in Jepson and Black’s influential paper on EigenTracking [5].

Incorporating subspace representations as the measurement model in a particle filter proves problematic, however. The number of samples in the particle filter needs to increase exponentially with the dimensionality of the state space, which now includes the subspace coefficients. We propose to resolve this problem by integrating out the appearance subspace coefficients of the state vector, leaving only the original target state. The advantage of this is that fewer samples are needed since part of the posterior over the state is analytically calculated, rather than being approximated using a more expensive sample representation. We use probabilistic principal component analysis (PPCA) method, a probabilistic subspace model for which the integral can be computed analytically [51, 45].

The resulting trackers work much better than using simpler appearance models, and some results are shown in Figure 5. As model complexity of the PPCA-based appearance model was increased, the number of tracker errors decreased and track quality, as measured by per target tracking error, increased.

5 Piecewise Continuous Curve-Fitting

An new avenue of research we are exploring concerns model selection, in which we bring together MCMC sampling and Rao-Blackwellization. In particular, we investigated the reconstruction of piecewise smooth 3D curves from multiple images. Among other applications, this is useful for the reconstruction of shards and other artifacts that are known to have “jagged edges”. Such objects frequently show up in large museum collections and archaeological digs, and hence a possible application is the automatic reconstruction of archaeological artifacts [32, 12]. The problem of representing and performing inference in the space of piecewise smooth curves is of interest in its own right.

To model piecewise smooth curves we use *tagged subdivision curves* as the representation. Subdivision curves [47] are simple to implement and provide a flexible way of representing curves of any type, including all kinds of B-splines and extending to functions without analytic representation. In [29], Hoppe introduces piecewise smooth subdivision surfaces, allowing to model sharp features such as creases and corners by tagging the corresponding control points. We apply the tagging concept to




Figure 6: A tagged 2D subdivision curve. Tagged points are drawn as circles. (a) the original control points, (b) the converged curve with non-smoothly interpolated points.

subdivision curves to represent piecewise smooth curves, as illustrated in Figure 6.

To infer the parameters of these curves from the data, *as well as automatically determine the number and the nature of the control points*, we use Markov chain Monte Carlo (MCMC) sampling [24] to obtain a posterior distribution over the discrete variables, while the continuous control point locations are integrated out after a non-linear optimization step. We sample over the number of control points, using the framework of reversible jump MCMC that was introduced by Green [26] and later described in a more easily accessible way as trans-dimensional MCMC [27]. In related work, Denison and Mallick [17, 38] propose fitting piecewise polynomials with an unknown number of knots using RJMCMC sampling. Punsakaya [43] extends this work to unknown models within each segment with applications in signal segmentation. DiMatteo [19] extends Denison’s work for the special case of natural cubic B-splines. With our method, we are working with a much reduced sample space, as we directly solve for optimal control point locations and hence only sample over the boolean product space of corner tags.

Some successfully recovered pot-shards are shown in Figure 7 along with a control image. Note that the curves are recovered in three dimensions, and that we have a distribution over them rather than a single point estimate. Hence, it is easy to obtain marginal statistics such as a histogram over the number of control points or the number of tagged points, etc.

6 Probabilistic Topological Maps

The use of Rao-Blackwellized sampling for model selection can also be used advantageously in mobile robotics, bringing us full circle in this paper. In particular, consider the problem of localizing a robot in an *unknown* environment. In that case, we are faced with a chicken and egg problem: if the robot knows where it is at any given moment, it can record in a map what it sees at that moment. On the other hand, the ability to localize itself on the map already presupposes a map. Algorithms that solve this conundrum are called “simultaneous localization and mapping” (SLAM) algorithms [36, 9, 48].

Probabilistic approaches have been very successful in dealing with the uncertainties associated with robot perception. However, almost all the work in the literature that applies probabilistic methods to




Figure 7: Piecewise continuous curve-fitting results for two different pot shards. Projections of the control points are drawn as yellow '+' for untagged and yellow 'x' for tagged ones, and the corresponding subdivision curve is drawn in white. Six views are used for the fitting process of the first shard, while only four are used for the second shard. In both cases, two of those are shown in columns (a) and (b). The third column (c) shows a view that is not used for fitting and is taken from a lower angle than all other images.

the map building problem deals with *metric maps*. While metric maps provide detailed information about the size and shape of obstacles and free space in the environment, they are expensive to build and maintain due to this very reason. Hence, another class of mapping algorithms construct *topological maps* [39, 35, 10, 11], which are typically graphs where the vertices denote rooms or other recognizable places, and the edges denote traversals between these places. Topological maps are quite useful for planning and symbolic manipulation and, unlike metric maps, do not require precise knowledge of the robot's environment. Unfortunately, there is no straightforward way to incorporate uncertainty into the topological map representation.

We use Markov chain Monte Carlo (MCMC) sampling [24] to extend the highly successful Bayesian framework to the space of topological maps. We define the *topology* of an environment as a partition of landmark observations into a set of equivalence classes, which denote that these observations stem from the same landmark in the environment. We begin our consideration by assuming that the robot observes K "special places" or landmarks during a run, i.e. we assume that the robot is equipped with a "landmark detector". Furthermore, we assume the existence of a the set of sensor measurements O recorded by the robot, which can include odometry as well as appearance measurements taken at the landmark locations. The problem then is to compute the discrete posterior probability distribution $P(T|O)$ over the space of topologies T given O .

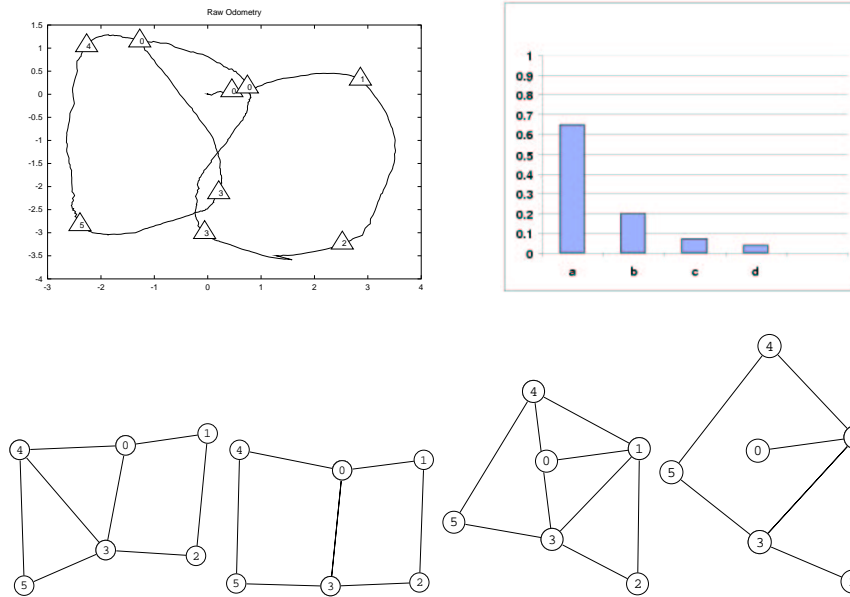


Figure 8: Probabilistic Topological Maps represent knowledge about the environment as an empirical histogram over a set of possible topological maps, obtained by approximate inference. Above, this is illustrated for a small environment in the shape of an “8”. The correct topology is present in the set of hypotheses but is not accorded the highest probability due to odometric error.

Note that this is not obvious, as there is a one-to-one mapping between topologies and a type of combinatorial objects known as *set partitions*. Indeed, each topology can be viewed as a disjoint set partition of the set of landmark observations, and the cardinality of the space of topologies over a set of K landmark observations is identical to the number of disjoint set partitions of the K -set. This number is called the *Bell number* b_K [34] and grows hyper-exponentially with K , for example $b_1 = 1$, $b_2 = 5$ but $b_{15} = 190899322$. Clearly, it is impossible to exhaustively evaluate this combinatorial space for any practical situation involving a significant number of landmarks observations.

We overcome the combinatorial nature of this problem by using MCMC sampling over the space of partitions, using as the target distribution the posterior probability of the topology T given the measurements O . A topology cannot be scored, however, without reference to the distortion in odometry measurements it induces. Since each topology induces a different, high-dimensional continuous space over the set of robot poses, it seems that we need to sample in a large, mixed-dimensionality space. However, once again Rao-Blackwellization provides a convenient exit: for each proposed topology in the MCMC sampling scheme, we optimize for the maximum likelihood robot trajectory, and analytically integrate out the probability mass in the continuous space by using Laplace’s approximation.

Thus, by sampling we introduced a novel concept: **Probabilistic Topological Maps** (PTM’s). As illustrated in Figure 8, a PTM represents knowledge about the environment as an empirical histogram over a set of possible topological maps, obtained by approximate inference through MCMC sampling. PTM’s can also be seen as a principled, probabilistic way of dealing with “closing the loop” in the context of SLAM, but on a much larger scale. Indeed, we not only consider simple loops, but consider the entire space of possible topologies at once. The key to making this work is assuming a simple but very effective prior on the density of “special places” in the environment. Given this prior the additional




Figure 9: While it is not our intended goal to produce highly accurate metric maps, it is indicative of the success of the approach that good-looking maps can be obtained using odometry only. The figure above shows two maps, one (a) obtained by plotting laser measurements from a laser range-finder using raw odometry, and (b) those same measurements using odometry derived from the inferred topology.

sensor information used can be very scant indeed. In fact, while our method is general and can deal with any type of sensor measurement (or, for that matter, prior knowledge) the results we present below were obtained using *only* the odometry measurements, and yet yield very nice maps of the environment.

In Figure 9, we show the results of an experiment that involved running a robot around a complete floor of the building containing our lab. Nine landmark observations were recorded during the run. The raw odometry with laser readings plotted on top is shown in Fig 9a. Shown in Figure 9b is the map obtained by optimizing over the odometry for the ground truth topology, which had the largest probability mass in the PTM built for this run. This figure demonstrates that probabilistic topological maps have the power to close the loop even in large environments.

7 Conclusion

The above sample of Monte Carlo methods, spanning 5 seemingly very different applications in robotics and computer vision, represent but a small part of the increasingly diverse array of approximate inference methods being deployed in these applied fields. These methods also include variational methods such as belief propagation and expectation-maximization, as well as graph-theory based approaches based on efficient max-flow/min-cut algorithms. However, to deal with challenging problems, especially those that include a model-selection component, it is to be expected that sequential Monte Carlo and Markov chain Monte Carlo approaches will steadily gain in popularity.

Acknowledgments

The work described in this paper was partially funded over the years by various government agencies and companies, including the Office of Naval Research, DARPA, Sony, Kirin Breweries, and the National Science Foundation, whose support is gratefully acknowledged.

References

- [1] Y. Bar-Shalom, T.E. Fortmann, and M. Scheffe. Joint probabilistic data association for multiple targets in clutter. In *Proc. Conf. on Information Sciences and Systems*, 1980.
- [2] A. Barbu and S.-C. Zhu. Graph partition by Swendsen-Wang cuts. In *Intl. Conf. on Computer Vision (ICCV)*, pages 320–327, Nice, France, 2003.
- [3] P. N. Belhumeur, J P Hespanha, and David J Kriegman. Eigenfaces vs. fisherfaces: Recognition using class specific linear projection. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 1996.
- [4] C. Berzuini, N. G. Best, W.R. Gilks, and C. Larizza. Dynamic conditional independence models and Markov chain Monte Carlo methods. *Journal of the American Statistical Association*, 92:1403–1412, 1996.
- [5] M.J. Black and A.D. Jepson. Eigentracking: Robust matching and tracking of articulated objects using a view-based representation. In *Eur. Conf. on Computer Vision (ECCV)*, 1996.
- [6] W. Burgard, D. Fox, D. Hennig, and T. Schmidt. Estimating the absolute position of a mobile robot using position probability grids. In *AAAI Nat. Conf. on Artificial Intelligence*, pages 896–901, 1996.
- [7] J. Carpenter, P. Clifford, and P. Fernhead. An improved particle filter for non-linear problems. Technical report, Department of Statistics, University of Oxford, 1997.
- [8] G. Casella and C.P. Robert. Rao-Blackwellisation of sampling schemes. *Biometrika*, 83(1):81–94, March 1996.
- [9] P. Cheeseman, B. Kanefsky, R. Kraft, J. Stutz, and R. Hanson. Super-resolved surface reconstruction from multiple images. In G. R. Heidbreder, editor, *Maximum Entropy and Bayesian Methods*, pages 293–308. Kluwer, the Netherlands, 1996.
- [10] H. Choset. *Sensor based motion planning: The hierarchical generalized voronoi graph*. PhD thesis, California Institute of Technology, 1996.
- [11] H. Choset and K Nagatani. Topological simultaneous localization and mapping (SLAM): toward exact localization without explicit localization. *IEEE Trans. on Robotics and Automation*, 17(2), April 2001.
- [12] D. Cooper, A. Willis, S. Andrews, J. Baker, Y. Cao, D. Han, K. Kang, W. Kong, F. Leymarie, X. Orriols, S. Velipasalar, E. Vote, M. Joukowsky, B. Kimia, D. Laidlaw, and D. Mumford. Bayesian virtual pot-assembly from fragments as problems in perceptual-grouping and geometric-learning. In *ICPR*, pages 11–15. IEEE Computer Society Press, August 2002.
- [13] T.F. Cootes, Edwards G, J, and C.J. Taylor. Active appearance models. In *Eur. Conf. on Computer Vision (ECCV)*, 1998.
- [14] F. Dellaert, D. Fox, W. Burgard, and S. Thrun. Monte Carlo Localization for mobile robots. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 1999.
- [15] F. Dellaert, S.M. Seitz, C.E. Thorpe, and S. Thrun. Structure from motion without correspondence. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2000.
- [16] F. Dellaert, S.M. Seitz, C.E. Thorpe, and S. Thrun. EM, MCMC, and chain flipping for structure from motion with unknown correspondence. *Machine learning*, 50(1-2):45–71, January - February 2003. Special issue on Markov chain Monte Carlo methods.

- [17] D. G. T. Denison, B. K. Mallick, and A. F. M. Smith. Automatic bayesian curve fitting. *Journal of the Royal Statistical Society, Series B*, 60(2):333–350, 1998.
- [18] A.R. Dick, P.H.S. Torr, and R. Cipolla. A Bayesian estimation of building shape using MCMC. In *Eur. Conf. on Computer Vision (ECCV)*, pages 852–866, 2002.
- [19] I. DiMatteo, C. R. Genovese, and R. E. Kass. Bayesian curve fitting with free-knot splines. In *Biometrika*, pages 1055–1071, 2001.
- [20] A. Doucet, J.F.G de Freitas, K. Murphy, and S. Russell. Rao-Blackwellised particle filtering for dynamic Bayesian networks. In *Proceedings of the 16th Annual Conference on Uncertainty in AI (UAI)*, 2000.
- [21] D.A. Forsyth, S. Ioffe, and J. Haddon. Bayesian structure from motion. In *Intl. Conf. on Computer Vision (ICCV)*, pages 660–665, 1999.
- [22] T.E. Fortmann, Y. Bar-Shalom, and M. Scheffe. Sonar tracking of multiple targets using joint probabilistic data association. *IEEE Journal of Oceanic Engineering*, 8, July 1983.
- [23] D. Fox, W. Burgard, F. Dellaert, and S. Thrun. Monte Carlo Localization – Efficient position estimation for mobile robots. In *AAAI Nat. Conf. on Artificial Intelligence*, 1999.
- [24] W.R. Gilks, S. Richardson, and D.J. Spiegelhalter, editors. *Markov chain Monte Carlo in practice*. Chapman and Hall, 1996.
- [25] N.J. Gordon, D.J. Salmond, and A.F.M. Smith. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proceedings F*, 140(2):107–113, 1993.
- [26] P. Green. Reversible jump Markov chain Monte Carlo computation and bayesian model determination. *Biometrika*, 82:711–732, 1995.
- [27] Peter J. Green. Trans-dimensional markov chain monte carlo. Chapter in *Highly Structured Stochastic Systems*, 2003.
- [28] J.E. Handschin. Monte Carlo techniques for prediction and filtering of non-linear stochastic processes. *Automatica*, 6:555–563, 1970.
- [29] Hugues Hoppe, Tony DeRose, Tom Duchamp, Mark Halstead, Hubert Jin, John McDonald, Jean Schweitzer, and Werner Stuetzle. Piecewise smooth surface reconstruction. *Computer Graphics*, 28(Annual Conference Series):295–302, 1994.
- [30] M. Isard and A. Blake. Contour tracking by stochastic propagation of conditional density. In *Eur. Conf. on Computer Vision (ECCV)*, pages 343–356, 1996.
- [31] M. Isard and A. Blake. Condensation – conditional density propagation for visual tracking. *Intl. J. of Computer Vision*, 29(1):5–28, 1998.
- [32] I. Kanaya, Q. Chen, Y. Kanemoto, and K. Chihara. Three-dimensional modeling for virtual relic restoration. In *MultiMedia IEEE*, volume 7, pages 42–44, 2000.
- [33] G. Kitagawa. Monte Carlo filter and smoother for non-gaussian nonlinear state space models. *J. of Computational and Graphical Statistics*, 5(1):1–25, 1996.
- [34] D.L. Kreher and D.R. Stinson. *Combinatorial Algorithms: Generation, Enumeration, and Search*. CRC Press, 1999.
- [35] B. J. 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.
- [36] J.J. Leonard and H.F. Durrant-Whyte. Simultaneous map building and localization for an autonomous mobile robot. In *IEEE Int. Workshop on Intelligent Robots and Systems*, pages 1442–1447, 1991.
- [37] D.J.C. MacKay. Introduction to Monte Carlo methods. In M.I. Jordan, editor, *Learning in graphical models*, pages 175–204. MIT Press, 1999.

- [38] B. K. Mallick. Bayesian curve estimation by polynomials of random order. *J. Statist. Plan. Inform.*, 70:91–109, 1997.
- [39] M. J. Matarić. A distributed model for mobile robot environment-learning and navigation. Master’s thesis, MIT, Artificial Intelligence Laboratory, Cambridge, January 1990. Also available as MIT AI Lab Tech Report AITR1228.
- [40] M. Montemerlo and S. Thrun. Simultaneous localization and mapping with unknown data association using FastSLAM. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2003.
- [41] M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit. FastSLAM: A factored solution to the simultaneous localization and mapping problem. In *AAAI Nat. Conf. on Artificial Intelligence*, 2002.
- [42] K. Murphy and S. Russell. Rao-Blackwellised particle filtering for dynamic bayesian networks. In A. Doucet, N. de Freitas, and N. Gordon, editors, *Sequential Monte Carlo Methods in Practice*. Springer-Verlag, New York, January 2001.
- [43] E. Punskeya, C. Andrieu, A. Doucet, and W.J. Fitzgerald. Bayesian curve fitting using MCMC with applications to signal segmentation. *IEEE Transactions on Signal Processing*, 50:747–758, March 2002.
- [44] D.B. Reid. An algorithm for tracking multiple targets. *IEEE Trans. on Automation and Control*, AC-24(6):84–90, December 1979.
- [45] S. Roweis. EM algorithms for PCA and SPCA. In Michael I. Jordan, Michael J. Kearns, and Sara A. Solla, editors, *Advances in Neural Information Processing Systems*, volume 10. The MIT Press, 1998.
- [46] A.F.M. Smith and A.E. Gelfand. Bayesian statistics without tears: A sampling-resampling perspective. *American Statistician*, 46(2):84–88, 1992.
- [47] E.J. Stollnitz, T.D. DeRose, and D.H. Salesin. *Wavelets for computer graphics: theory and applications*. Morgan Kaufmann, 1996.
- [48] S. Thrun. Robotic mapping: A survey. In G. Lakemeyer and B. Nebel, editors, *Exploring Artificial Intelligence in the New Millenium*. Morgan Kaufmann, 2002. To appear, also available as Carnegie Mellon TR CMU-CS-02-111.
- [49] S. Thrun, D. Fox, W. Burgard, and F. Dellaert. Robust Monte Carlo localization for mobile robots. *Artificial Intelligence*, 128(1-2):99–141, 2000.
- [50] S. Thrun, D. Fox, F. Dellaert, and W. Burgard. Particle filters for mobile robot localization. In Arnaud Doucet, Nando de Freitas, and Neil Gordon, editors, *Sequential Monte Carlo Methods in Practice*. Springer-Verlag, New York, January 2001.
- [51] M.E. Tipping and C.M. Bishop. Probabilistic principal component analysis. Technical Report NCRG/97/010, Neural Computing Research Group, Aston University, September, 1997.
- [52] M. Turk and A.P. Pentland. Face recognition using eigenfaces. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 586–591, 1991.
- [53] S.C. Zhu. Stochastic jump-diffusion process for computing medial axes in markov random fields. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 21(11), November 1999.