Efficient Monitoring for Planetary Rovers

Vandi Verma
vandi@ri.cmu.edu

Geoff Gordon
ggordon@cs.cmu.edu

Reid Simmons
reids@cs.cmu.edu

Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

Planetary rovers operate in environments where human intervention is expensive, slow, unreliable, or impossible. It is therefore essential to monitor the behavior of these robots so that contingencies may be addressed before they result in catastrophic failures. This monitoring needs to be efficient since there is limited computational power available on rovers.

We propose an efficient particle filter for monitoring faults that combines the *Unscented Kalman Filter* (UKF) [7] and the *Variable Resolution Particle Filter* (VRPF) [16]. We begin by using the UKF to obtain an improved proposal distribution for a particle filter which tracks discrete fault variables as part of its state space. This requires computing an unscented transform for every particle and every possible discrete transition to a fault or nominal state at each instant in time. Since there are potentially a large number of faults that may occur at any instant, this approach does not scale well. We use the VRPF to address this concern. The VRPF tracks abstract states that may represent single states or sets of states. There are many fewer transitions between states when they are represented in abstraction. We show that the VRPF in conjunction with a UKF proposal improves performance and may potentially be used in large state spaces. Experimental results show a significant improvement in efficiency.

1 Introduction

A number of future space exploration missions include rovers. To prevent mission failure, it is essential to be able to detect and recover from faults. Here a *fault* is defined as a deviation from the expected behavior of the system, while a *failure* is a complete interruption of the system's ability to perform its required operations. The faults addressed here include mechanical component failures, such as broken motors and gears; faults due to environmental interactions, such as a wheel stuck against a rock; and sensor failures, such as broken encoders. Interpreting these faults requires context sensitive interpretation of sensor data that can be obtained by continuously monitoring the dynamics of the system, which tend to differ according to operating conditions. For example, for a rover, high power draw on flat ground may be a cause for concern, but high power draw on a slope might be perfectly acceptable. Sensors do not directly report these dynamics because they are noisy and limited, i.e., they do not have complete information about the state of the rover and the environment that it is operating in. In addition, there are a large number of components that can fail in various combinations at any instant in time and the computational resources are too limited to consider all possible combinations.

We focus here on the rover domain. The fault monitoring problem is formulated in terms of providing a distribution over the unobservable discrete fault and operational states of a rover from noisy measurements of continuous sensor readings. We address this problem using a probabilistic technique called Bayes filtering. Bayes filters estimate the distribution over the state space (the belief state) of a dynamic system conditioned on the data. Tracking an exact posterior is intractable for the fault monitoring problem because the state space is very large. We use a special case of a Bayes filter called the *particle filter* to approximate the distribution over the state space [11]. Particle filters [3, 6] are a Monte Carlo method for monitoring dynamic systems by approximating the belief state with a set of samples or "particles". The benefits of this approach are that it is non-parametric and that it can represent arbitrary distributions. Both discrete and continuous variables can be represented with a single particle filter. Particle filters are easily implemented based on a forward simulation.

The main drawback of particle filters is that the sampling process has a high variance, particularly in high-dimensional spaces. As a result, an extremely large number of particles are needed to obtain a reasonable approximation of the belief state. It is not feasible to use such a large number of particles, given the computational constraints of a planetary rover. In this paper we present an improved particle filter that reduces the variance of the particle filter estimate by taking into account the next measurement when generating particles. This is done by computing an approximately optimal proposal distribution for each transition using an *Unscented Kalman Filter* (UKF) [7]. Particles are then generated from this proposal distribution. To improve the scalability of this approach we use the *Variable Resolution Particle Filter (VRPF)*. The VRPF introduced the notion of abstract states that may represent individual states or sets of states [16]. The VRPF dynamically varies the resolution of the state space for computational efficiency. Where belief is strong, resolution is fine. Where belief is low, resolution is coarse, abstracting multiple similar states together. The VRPF reduces the number of next state transitions. Experimental results show a significant improvement in efficiency.

2 BAYESIAN MODEL FOR MONITORING FAULTS

Let D represent the finite set of discrete fault and operational modes of the rover, $d_t \in D$ the state of the rover at time t and $\{d_t\}$ the discrete, first order Markov chain representing the evolution of the state over time. The problem of monitoring the state of the rover consists of providing a belief (a distribution over the state set D) at each time step as it evolves based on the following transition model:

$$p(d_t = j \mid d_{t-1} = i), \quad (i, j \in D) \tag{1}$$

Each of the discrete fault and operational modes changes the dynamics of the rover. Let $x_t \in \mathbb{R}^{n_x}$ denote the multivariate continuous state of the rover at time t. The non-linear conditional state transition models are denoted by $p(x_t \mid x_{t-1}, d_t)$. The state of the rover is observed through a sequence of measurements, $\{z_t\}$, based on the measurement model $p(z_t \mid x_t, d_t), z_t \in \mathbb{R}^{n_z}$.

3 CLASSICAL PARTICLE FILTER

The monitoring problem consists of estimating the marginal $p(d_t \mid z_{1..t})$ of the posterior distribution $p(x_t, d_t \mid z_{1..t})$. A recursive estimate of this posterior distribution may be obtained using the Bayes filter:

$$p(x_t, d_t \mid z_{1...t}) = \eta_t \, p(z_t \mid x_t, d_t) \int \sum_{d_{t-1}} p(x_t, d_t \mid x_{t-1}, d_{t-1}) dx_{t-1} \tag{2}$$

There is no closed form solution to this recursion, hence we use a particle filter approximation. A particle filter (PF) [5, 8] is a Monte Carlo approximation of the posterior in a Bayes filter. PFs approximate the posterior with a set of N fully instantiated state samples or particles $\{(d_t^{[1]}, x_t^{[1]}) \ldots (d_t^{[N]}, x_t^{[N]})\}$ and importance weights $\{w_t^{[i]}\}$:

$$\hat{p}_N(x_t, d_t \mid z_{1\ldots t}) = \sum_{i=1}^{N} w_t^{[i]} \; \delta_{x_t^{[i]}, d_t^{[i]}}(x_t, d_t) \tag{3}$$

where $\delta(\cdot)$ denotes the Dirac delta function. It can be shown that as $N \to \infty$ the approximation in (3) approaches the true posterior density [14]. Because it is difficult to draw samples from the true posterior, we instead draw them from a more tractable distribution $q(\cdot)$, called the proposal (or importance) distribution. There are a large number of possible choices for the proposal distribution, the only condition being that its support must include that of the posterior. The importance weights are used [12, 13] to account for the discrepancy between the proposal distribution $q(\cdot)$ and the true distribution $p(x_t, d_t \mid z_{1:t})$, and for a given sample $x_t^{[i]}, d_t^{[i]}$ the importance weight is

$$w_t^{[i]} = p(x_t^{[i]}, d_t^{[i]} \mid x_{t-1}^{[i]}, d_{t-1}^{[i]}, z_t) \;/\; q(x_t^{[i]}, d_t^{[i]}) \tag{4}$$

The simplest choice for the proposal distribution is the transition probability

$$q(x_t, d_t) = p(x_t, d_t \mid x_{t-1}^{[i]}, d_{t-1}^{[i]}) = p(x_t \mid x_{t-1}^{[i]}, d_t) p(d_t \mid d_{t-1}^{[i]})$$

in which case the importance weight is equal to the likelihood

$$p(z_t \mid x_{t-1}^{[i]}, d_{t-1}^{[i]}, z_{1\ldots t-1})$$

This is the most widely used proposal distribution [1, 4, 9] and is simple to compute, but it can be inefficient since it ignores the most recent measurement z_t. Particularly in the fault diagnosis domain, where there are a large number of possible faults that may occur at any instant in time, the most recent measurement can be very informative.

4 METHODS FOR ENHANCING MONITORING EFFICIENCY

A well-known problem with particle filters is that a large number of particles are often needed to obtain a reasonable approximation of the posterior distribution. For real-time fault detection and identification, maintaining such a large number of particles is typically not practical. However, the variance of the particle-based estimate can be high with a limited number of samples, since a large number of faults may potentially occur at any instant. In addition, faults are typically not very likely and so some parts of the state space transition to other parts with very low probability. Consider the problem of diagnosing locomotion faults on a robot. The probability of a stalled motor is low and wheel faults on the same side generate similar observations. Motors on any of the wheels may stall at any time. A particle filter that produces an estimate with a high variance is likely to result in identifying some arbitrary wheel fault on the same side, rather than identifying the correct fault. Here we describe a method that enables efficient monitoring by producing low variance estimates even for small sample sizes. The approach combines the Unscented Kalman Filter with the Variable Resolution Particle Filter.

4.1 UNSCENTED KALMAN FILTER

Doucet et al. [2] show that $p(x_t, d_t \mid x_{t-1}^{[i]}, d_{t-1}^{[i]}, z_t)$ is the "optimal" proposal distribution, that is, the distribution that minimizes the variance of the importance weights conditioned on $x_{t-1}^{[i]}, d_{t-1}^{[i]}$ and z_t. Unfortunately, it is difficult to sample from exactly this proposal

distribution, so we instead use UKFs [7] to approximate it. This approximation is similar to an unscented particle filter [15], but it takes into account the fact that some of our state variables are discrete while others are continuous.

Bayes' rule and our conditional independence assumptions imply

$$\begin{aligned} p(x_t, d_t \mid x_{t-1}^{[i]}, d_{t-1}^{[i]}, z_t) &= \eta^{[i]}\, p(z_t \mid x_t, d_t)\, p(x_t, d_t \mid x_{t-1}^{[i]}, d_{t-1}^{[i]}) \\ &= \eta^{[i]}\, p(z_t \mid x_t, d_t)\, p(x_t \mid x_{t-1}^{[i]}, d_t)\, p(d_t \mid d_{t-1}^{[i]}) \\ &= \eta^{[i]}\, \eta_{d_t}^{[i]}\, p(x_t \mid x_{t-1}^{[i]}, d_t, z_t)\, p(d_t \mid d_{t-1}^{[i]}) \end{aligned}$$

The normalizing constants in the above equations are

$$\begin{aligned} \eta^{[i]} &= 1/p(z \mid x_{t-1}^{[i]}, d_{t-1}^{[i]}) \\ \eta_{d_t}^{[i]} &= p(z_t \mid x_{t-1}^{[i]}, d_t) \end{aligned}$$

We can ignore $\eta^{[i]}$ since it doesn't depend on x_t or d_t. The discrete transition probability $p(d_t \mid d_{t-1}^{[i]})$ is known. So, we will use UKFs to approximate $\eta_{d_t}^{[i]}$ and $p(x_t \mid x_{t-1}^{[i]}, d_t, z_t)$ for each particle i and possible discrete transition d_t. To compute these approximations we need to examine each possible pair of i and d_t separately; this process can be computationally expensive, a complaint which we will return to in section 4.2.

The UKF is a recursive minimum mean square error estimator that often provides an improvement over the Extended Kalman Filter (EKF) for nonlinear models. The EKF linearizes the nonlinear process and measurement models using the first order terms of a Taylor series expansion. The UKF, on the other hand, does not approximate the nonlinear process and measurement models. It uses the actual models and instead approximates the distribution of the state variable as a Gaussian. The Gaussian approximation is specified using a minimal set of deterministically chosen samples called sigma points. Each sigma point is independently propagated through the process and measurement models, and the set of propagated sigma points is analyzed to provide a posterior Gaussian approximation. The process of calculating, propagating, and analyzing the sigma points is called an Unscented Transform or UT; see [7] for details.

In our case, we use an UT to approximate $p(x_t \mid x_{t-1}^{[i]}, d_t, z_t)$ as a Gaussian. This approximation will usually be excellent because we are conditioning on a single previous state and a single possible fault. Given this approximation, we can compute analytically the values for $\mu_{d_t}^{[i]}$ (the mean of x_t), $P_{d_t}^{[i]}$ (the covariance of x_t), and $\eta_{d_t}^{[i]}$ (the observation likelihood). Only the last of these quantities is not given by the standard Kalman filter equations. It is:

$$\eta_{d_t}^{[i]} \propto \frac{1}{|S_{d_t}^{[i]}|^{\frac{1}{2}}} \exp(-\tfrac{1}{2}\,(z_t - \hat{z}_{d_t}^{[i]})^T (S_{d_t}^{[i]})^{-1} (z_t - \hat{z}_{d_t}^{[i]}))$$

where $S_{d_t}^{[i]}$ is the innovation covariance and $\hat{z}_{d_t}^{[i]}$ is the predicted observation (both computed from the UT).

Once we have finished our UTs, we can sample from our proposal distribution by first drawing d_t and then x_t according to (5–6), then computing importance weights via (4).

$$\hat{p}(d_t \mid d_{t-1}^{[i]}) \sim p(d_t \mid d_{t-1}^{[i]})\, \eta_{d_t}^{[i]} \tag{5}$$

$$\hat{p}(x_t \mid x_{t-1}^{[i]}, d_t) \sim \mathcal{N}(\mu_{d_t}^{[i]}, P_{d_t}^{[i]}) \tag{6}$$

In order to find $\eta_{d_t}^{[i]}$ for every d_t and i, we must compute an UKF for each particle and every possible next discrete state transition. Given that there are potentially a large number of faults to transition to at any step, this may be inefficient. We propose to use the Variable Resolution Particle Filter to address this.

4.2 VARIABLE RESOLUTION PARTICLE FILTER

The Variable Resolution Particle Filter (VRPF) utilizes abstract particles that may represent individual states or sets of states. With this method, a single abstract particle can simultaneously track multiple states. A limited number of samples are therefore sufficient for representing large state spaces. A bias-variance tradeoff is made to abstract and refine states dynamically to change the resolution. As a result reasonable posterior estimates can be obtained with a relatively small number of samples.

We use a multi-layered hierarchy to represent the variable resolution state space model. Each physical (non-abstract) state is a leaf of the hierarchy. Sets of states with similar state transition and observation models are aggregated together at each level in the hierarchy. So, in addition to the physical states S_k, the variable resolution model uses abstract states A_j that represent sets of one or more physical states. We use domain knowledge to create a multi-layer hierarchy of abstract states:

$$A_j = \begin{cases} \{S_k\} & \text{if } j \text{ is a leaf} \\ \cup_{i \in \mathrm{children}(j)} A_i & \text{otherwise} \end{cases} \tag{7}$$

In the example above, the VRPF would aggregate the wheel faults on the same side of the rover together into an abstract fault state. Given a fault, the abstract state representing the side on which the fault occurs would have high likelihood and the samples in this state would be assigned a high importance weight. This would result in multiple copies of these samples on resampling proportional to weight. Once there are sufficient particles to populate all the refined states represented by the abstract state, the resolution of the state would be changed to the states representing the individual wheel faults. At this stage, the correct hypothesis is likely to be included in this particle based approximation at the level of the individual states and hence the correct fault is likely to be detected.

For any fixed resolution of the state space, a posterior distribution is computed by projecting the abstract samples onto the physical states, computing the posterior sample set given the measurement, and determining the abstract states that the posterior samples belong to. To vary the resolution of the state space, the VRPF uses a bias-variance tradeoff. Suppose that the VRPF is currently tracking states at one level of the hierarchy. A decision to abstract to the next coarser resolution is made if the combination of bias and variance in the abstract state A_j is less than the combination of bias and variance of all its children A_i, as shown below:

$$b(A_j)^2 + v(A_j) \leq \sum_{i \in \mathrm{children}(j)} [b(A_i)^2 + v(A_i)] \tag{8}$$

On the other hand if the state space is currently at the coarser resolution and the reverse of equation (8) holds, then the VRPF refines to the next finer resolution. To avoid hysteresis, all abstraction decisions are considered before any refinement decisions. Details on the VRPF may be found in [16].

We create a variable resolution state space model for the discrete fault states. This reduces the number of discrete transitions that need to be considered when computing $\eta_{d_t}^{[i]}$. As our experimental results show, the resulting computational savings are significant.

5 EXPERIMENTAL RESULTS

The problem domain for our experiments involves diagnosing locomotion faults in a physics based simulation of a six wheel rover. Figure 1(a) shows a snapshot of the rover in the Darwin2K [10] simulator. The particle set representing the state consists of N particles, where each particle $[d_t^{[i]}, x_t^{[i]}]$ is a hypothesis about the current state of the system.

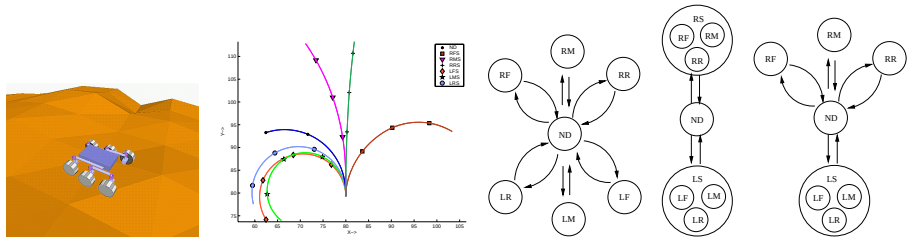

(a) (b) (c) (d) (e)

Figure 1: *(a) Snapshot from the dynamic simulation of the six wheel rocker bogie rover in the simulator, (b) An example showing the normal trajectory (ND) and the change in the same trajectory with a fault at each wheel. (c) Original discrete state transition model. The discrete states are: Normal driving (ND), right and left, front, middle and rear wheel faulty (RF, RM, RR, LF, LM, LR) (d) Abstract discrete state transition model. The states, RF, RM and RR have been aggregated into the Right Side wheel faulty states and similarly LF, LM and LR into Left Side wheel faulty states (RS and LS). (e) State space model where RS has been refined. All states have self transitions that have been excluded for clarity.*

$d_t^{[i]}$ is the discrete fault or normal state and $x_t^{[i]}$ is the multi-dimensional continuous state representing the change in position and orientation of the rover. Each discrete fault state has a different observation and predictive model for the continuous state. The probability of a discrete state is determined by the density of samples in that state.

The Markov model representing the discrete state transitions consists of 7 states. As shown in figure 1(c) the normal driving (ND) state may transition back to the normal driving state or to any one of six fault states: right front (RF), right middle (RM), right rear (RR), left front (LF), left middle (LM) and left rear (LR) wheel stuck. Each of these faults cause a change in the rover dynamics, but the faults on each side (right and left) have similar dynamics.

Figure 2 shows a comparison of the error from monitoring the state using a classical particle filter and a particle filter that uses an UKF proposal. The X axis shows the number of particles used, the Y axis shows the KL divergence from an approximation of the true posterior computed using a large number of samples. For the experiment in figure 2, the continuous measurements were the absolute rover position. 1000000 samples were used to compute an approximation to the true distribution. The KL divergence is computed over the entire length of the data sequence and is averaged over multiple runs over the same data set.[1] The data set included normal operation and each of the six faults. Figure 2(a) demonstrates that using an UKF proposal dramatically improves the performance. Figure 2(b) shows the KL-divergence along the Y axis and wall clock time[2] along the X axis. Both filters were coded in matlab and share as many functions as possible.

Given that the three wheels on each side of the rover have similar dynamics, we constructed a hierarchy for the VRPF that clusters the fault states on each side together. Figure 1(d) shows this hierarchical model, where the abstract states right side fault (RS), and left side fault (LS) represent sets of states {RF, RM, RR} and {LF, LM, LR} respectively. The

[1]The results are an average over 50 to 5 runs with repetitions decreasing as the sample size was increased.

[2]Given that this is wall clock time, it should be taken with a grain of salt. The tests were run on lab machines used by multiple users. We hope to have minimized the influence of other processes by running the tests at night.

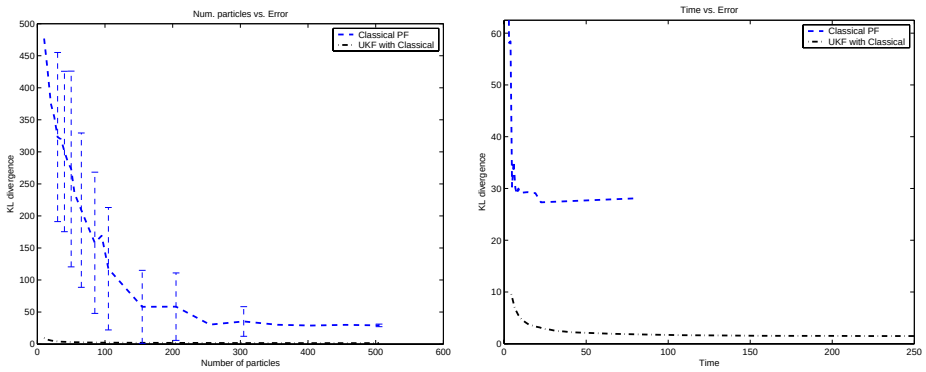

(a) (b)

Figure 2: *Comparison of the KL divergence from the true distribution for the classical particle filter and a particle filter that uses an UKF proposal, against (a) number of particles used, (b) wall clock time.*

highest level of abstraction therefore consists of nodes $\{$ND, RS, LS$\}$. Figure 1(e) shows how the state space in figure 1(d) would be refined if the bias in the abstract state RS given the number of particles outweighs the reduction in variance over the specialized states RF, RM and RR at a finer resolution.

When particle filtering is performed with the VRPF, the particles are initialized at the highest level in the abstraction hierarchy, i.e., in the abstract states ND, RS and LS. If a RF fault occurs, this is likely to result in a high likelihood of samples in RS. These samples will multiply, which may then result in the bias in RS exceeding the reduction in variance in RS over RF, RM and RR, thus favoring tracking at the finer resolution. Additional observations should then assign a high likelihood to RF.

Figure 3 shows a comparison of the performance of the classical particle filter, a particle filter with an UKF proposal (UPF) and a VRPF with an UKF proposal distribution (VR-UF). The improvement in performance of the VR-UF over UPF is expected to be even greater when the variable resolution state space model is larger and results in a larger reduction in the size of the state space than the simple experiment we present. This is because a UKF needs to be computed for every possible discrete state transition. The discrete state transitions between the abstract state in the VR-UF are smaller than the discrete state transitions between the physical states used in UPF.

6 CONCLUSIONS

We present an efficient method for monitoring hybrid state spaces and demonstrate the applicability of the approach for fault diagnosis on rovers. Unlike a number of existing methods, our approach is valid even when the process is non-linear. It is based on particle filters and has all the advantages of particle filters which include the ability to represent non-parametric posteriors, non-linear processes, and easy extension to an anytime approach. The main drawback of particle filters is that a large number of samples may be needed for reasonable approximations. Our approach uses an Unscented Kalman Filter to focus particles in regions of the state space with high predictive likelihood which requires a comparatively smaller number of samples for performance comparable to a classical par-

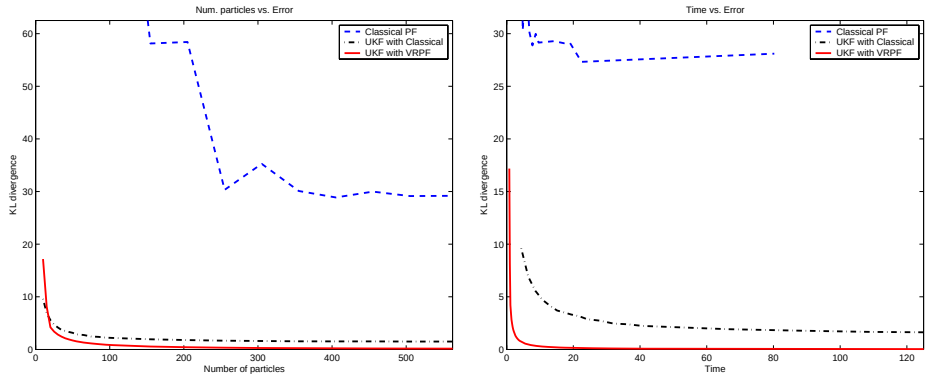

(a) (b)

Figure 3: *Comparison of the KL divergence from the true distribution for the classical particle filter, a particle filter with UKF proposal and the VRPF with UKF proposal, for a hybrid state space against (a) the number of particles used, (b) wall clock time.*

ticle filter. In addition it uses the Variable Resolution Particle Filter that maintains samples in different regions of the state space at dynamically varying resolution to minimize the number of next state transitions that must be considered when computing the predictive likelihood using an UKF approximation. Our experimental results show a significant improvement in performance with this approach. Although we used the UKF approximation at each time step, one may choose to use it only when a transition is made from a normal to fault state since this transition introduces a high uncertainty in the continuous state estimate.

References

[1] D. Avitzour. A stochastic simulation bayesian approach to multitarget tracking. IEEE Transaction on Aerospace and Electronic systems, 1995.

[2] A. Doucet. On sequential monte carlo sampling methods for bayesian filtering, 1998.

[3] A. Doucet, N. de Freitas, and N.J. Gordon, editors. *Sequential Monte Carlo Methods in Practice*. Springer-Verlag, 2001.

[4] N.J. Gordon, D.J Salmond, and A.F.M Smith. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. In *IEE Proceedings-F*, volume 140(2), pages 107–113, 1993.

[5] N.J. Gordon, D.J Salmond, and A.F.M Smith. Novel approach to nonliner/non gaussian bayesian state estimation. volume 140 of *IEE Proceedings-F*, pages 107–113, 1993.

[6] M. Isard and A. Blake. Condensation: conditional density propagation for visual tracking. *International Journal of Computer Vision*, 29(1):5–28, 1998.

[7] S. J. Julier and J. K. Uhlmann. A general method of approximating nonlinear transformations of probability distributions. Technical report, Robotics Research Group, Department of Engineering Science, University of Oxford, 1996.

[8] K. Kanazawa, D. Koller, and S. Russell. Stochastic simulation algorithms for dynamic probabilistic networks. pages 346–351, 1995.

[9] G. Kitagawa. A monte carlo filtering and smoothing method for non-gaussian nonlinear state space models. Proceedings of the 2nd U.S.-Japan Joint Seminar on Statistical Time Series Analysis, 1993.

[10] P.C. Leger. *Darwin2K: An Evolutionary Approach to Automated Design for Robotics*. Kluwer Academic Publishers, 2000.

[11] S. McIlraith. Diagnosing hybrid systems: A bayesian model selection approach. Proceedings of the Eleventh International Workshop on Principles of Diagnosis (DX'00), pages 140–146, June 2000.

[12] D. B. Rubin. *Using the SIR algorithm to simulate posterior distributions*. Oxford University Press, 1988.

[13] R. Y. Rubinstein. Simulation and monte carlo method. In *John Wiley and Sons*, 1981.

[14] M. A. Tanner. *Tools for Statistical Inference*. Springer Velag, 1996.

[15] R. van der Merwe, N. de Freitas, A. Doucet, and E. Wan. The unscented particle filter. Technical report, Cambridge University Engineering Department, Cambridge, England, August 2000.

[16] V. Verma, R. Simmons, and S. Thrun. Variable resolution particle filter. International Joint Conference of Artificial Intelligence, 2003.