

Application of adaptive sample size particle filtering to a chaotic dynamic system



David Livingston
Willamette University

Advisor: Haiyan Cheng




Introduction

A particle filter is a Sequential Monte Carlo numerical method for estimating the posterior density of a state space. Particle filters estimate the value of hidden state variables over time given limited observation data. Theoretical applications are wide, including weather forecasting and robot localization. In practice, particle filters often prove too slow and inefficient for real-life applications. Adaptive sampling has proved to be an effective method of increasing efficiency in some specific applications. We attempt to generalize adaptive sampling to more general, and more difficult, models.

Particle Filter

An individual particle represents a guess at the value of the unknown state. A particle filter uses an ensemble of many of these particles to create an estimate probability density for the state variable. This guess can be moved forward in time with a model representing how we think the variable progresses in reality. The accuracy of the guess can be improved based on information we gain about the variable (observations) through a process called resampling. During resampling, particles that are likely far from the truth are replaced with copies of more likely particles.




Adaptive Sampling

Particle filters typically use a constant number of particles at each step. Our confidence in the accuracy of the particle filter's estimation, however, is not necessarily constant. Variables can be easier to predict at certain times, requiring less particles for comparable accuracy. If we can determine how many particles are needed at any given time, the average number of particles used can be reduced. Adaptive sample size particle filtering determines the number of particles on-the-fly. One method is to stop sampling when the error between a current estimate and the actual posterior based on previous observations appears to be within acceptable bounds. Determining what error is acceptable is up to implementation.

KLD-Sampling

KLD sampling uses Kullback-Leibler Divergence as an estimator of error between estimate given current particles and the true posterior density. KLD is a common statistical measure of the information difference between two probability distributions



Before the particle filter begins, the state space must be discretized into some number of 'bins' for particles to fall into. The total number of particles is determined during resampling rather than beforehand. N_x , the number of particles needed for an acceptable estimate, is updated at each step.

KLD-sampling relies on being able to find a good discretization of the state space. Ideally, particles falling into a higher number of distinct bins corresponds to a higher uncertainty about the system state. This is difficult in applications where the state is many-dimensional and highly chaotic.


KLD sampling is the most widely known method of adaptive sampling, and many other techniques are based on this framework.. For these reasons, we chose to focus on KLD-sampling when generalizing to more complex models.

Models

We focused on two models, both designed by Edward Lorenz. Lorenz '63 and Lorenz '96 have both been commonly used as test suites for new particle filtering techniques. Both are highly chaotic. For Lorenz '96, we choose perhaps the most common parameters, with $N_{var}=40$ and $Forcing=8$. The resulting 40-dimensional model provides an excellent testbed for the efficacy of adaptive sampling in spaces with high dimensionality. The chaotic nature and high dimensionality of Lorenz '96 both provide tremendous difficulties for adaptive sampling. The three-dimensional Lorenz '63 system removes the latter of these difficulties, allowing us to investigate the effects of both.

Results

Tests on both Lorenz '63 and '96 showed that adaptive sampling resulted in negligible improvements in efficiency. When an adaptive particle filter and a control particle filter were run for roughly the same amount of time on the same parameters, the difference in accuracy between the two was not significant. The below figures show the average error (in RMSE) over 10 runs on a large time scale.



Our results indicate that adaptive sampling is not a promising method for improving particle filters on standard test models. In addition to providing no improvement in accuracy or efficiency, adaptive filters had several negative qualities as compared to the control filter. Runtime was less predictable, and naïve parameter settings can lead to wildly erratic or uncontrollably large particle counts.

References

- Fox, D. "Adapting the Sample Size in Particle Filters Through KLD-Sampling." International Journal of Robotics Research, 2003
 - Gordon et al. "Novel approach to nonlinear/non-Gaussian Bayesian state estimation", 1993
 - Lorenz, E. "Predictability- a problem partly solved." 1996
 - Lorenz, E. "Deterministic Nonperiodic Flow." Journal of the Atmospheric Sciences, 1963
- This material is based upon work supported by the National Science Foundation under grant No. DMS-1217073 and the Willamette SCRP program.