Problems associated with the Probability Hypothesis Density Function approach for multi−target tracking

The problem of tracking and identifying multiple manuevering objects from a diverse collection of sensors has been a well studied problem for the past 20 years.

In general, we say that there is a finite set X of targets. At time step k, sensors generate a collection of (noisy) observations Z_k. Most standard approaches to this problem estimate X using a finite set of "tracks" X_k. (Each track typically contains additional information such as covariance estimates for the parameters of X_k). These standard approaches then must associate observations from Z_k with elements of their track set X_k, and then treat the problem as a single target tracking problem. Such approaches generally perform poorly in situations where there are a large number of objects close to one another.

Recently a new approach to this problem has been proposed, it is called the Probability Hypothesis Density Function (PHD). Essentially, the PHD approximates the true multi−target filtering densities with a Poisson point process (PPP) whose intensity measure μ has the property that for any set A, $\mu(A)$ is equal to the expected number of targets in A. One significant advantage to this approach is that it avoids the association step.

To implement the PHD, we have been using a particle system based approach. Recently particle systems have gathered quite an avid following as a computationally tractable and mathematically defensible means for solving nonlinear−filtering problems. Particle systems approximate the probability densities propagated by the nonlinear filtering equations using empirical distributions consisting of some finite number of particles. As such the proofs for such systems are conceptually similar to proofs for Monte Carlo integration and as such are straight forward, moreover the models are conceptually simple and straight forward to program. The formulation of a PHD is very similar to standard nonlinear−filtering problems and we have proven that minor extensions of these algorithms solve the PHD problem as well.

[aside. In our work, we always assume time is discrete, and state is a Cartesian product of a discrete space, with a d dimensional Euclidean space.]

In this workshop, we will focus on two problems of interest to us with regards to PHD. (And will bring a collection of additional problems which we are interested in.) We will provide a brief overview of Finite Set

Statistics (FISST) a unified mathematical framework for multi−target tracking developed by Dr. Ron Mahler of Lockheed Martin. We will then describe the PHD in more detail, and show how it approximates the true multi−target density functions. Finally we will cover the proofs of correctness for the PHD implementation.

The two specific problems that we will attack in the workshop are:

− While we can show our particle system approximations converge as the number of particles tends to infinity, what can we say about their error as a function of the number of particles N? In particular can we say anything about the error for small N. (Most error results we know of for such systems are asymptotic, and say nothing about smaller values of N.)

− Given a Poisson Point Process that estimates our target state, what is the best means by which to compute an estimate of that state to communicate to an end user? (An end user is not interested in seeing a Poisson point process density function, they want to know how many targets there are, and where they are.) What kinds of confidence can we place on our results which we can also convey to the end user?

Other problems of specific interest to us include:

− Our framework currently assumes observations are discrete. (i.e. Z is a finite set.) What if Z is a continuous function? Can our framework be extended to encompass that kind of observation?

− A standard problem in filtering is prediction and smoothing. Prediction estimates the state of the system at some point in the future. Smoothing takes multiple observations, and creates a better estimate of some time t in the past. At present, we have no way to perform smoothing using the PHD. Can this be accomplished?

− A related question: can we extend the PHD to filter on path space, instead of just state space?

− Based on our error estimates, are there particle propagation & resampling strategies that provably minimize errors? Most results in this area for standard nonlinear filtering minimize "single step" errors, is there a way in which long term errors can be minimized?

− Do general Large Deviation principles exist for our the PHD particle systems? We’re not interested in sharp rate functions, just reasonable lower bounds.