Heritage Method for Track Extraction Using the Hybrid Intensity and Likelihood Ratio Tracking (iLRT) Filter

Roy L. Streit and Bryan Osborne

Metron Inc.
streit@metsci.com
osborne@metsci.com

Abstract: The PHD / Intensity filter requires some kind of post-processing procedure to extract target state estimates and their associated areas of uncertainty (AOUs). This is true for both Gaussian sum and sequential Monte Carlo, or particle, implementations of the filter. The method used here for the track extraction process is most directly applicable to particle implementations. Extension to the Gaussian sum filter implementation is not considered.

The particle filters use resampling methods that yield particles that retain knowledge of the associated sampling history, that is, the particles are candidate target trajectories. The basic idea behind the hybrid intensity and likelihood ratio tracking (iLRT) filter is to use the PHD / Intensity filter as a kind of "proposal density" to identify candidate particle trajectories and subsequently, in a separate auxiliary step, to evaluate how well each trajectory "fits" the data using a cumulative log-likelihood ratio (CLLR) score. The PHD / iFilter is an approximation to the real multitarget tracking problem, but this issue is mitigated because particle trajectories are optimally scored by the CLLR.

Target tracks, complete with time histories, are estimated from those trajectories whose CLLR scores exceed a specified threshold. The above-threshold trajectories are the basis of a clustering scheme called the heritage method. The name is appropriate because clustering is achieved by associating to the same target those trajectories that have a common origin when traced backward in time. The heritage method thus partitions the set of all above-threshold particle trajectories into disjoint sets, and each set in the partition corresponds to a single target. Track state estimates and AOUs are obtained from the trajectories in each set of the partition. Multiple target tracks are therefore extracted without explicitly associating measurements to targets and also without explicitly using a metric clustering method.

The CLLR is an optimum target detection statistic, that is, it requires the fewest number of measurements on average to declare target presence or absence on a specified path. All other statistical tests with the same type I and type II error probabilities require more measurements, on average, than the CLLR.

The iLRT filter is presented for the usual kinematic data, as well as kinematic data that are accompanied by additional "feature" data. The use of feature data in an intensity filter is relatively new, but such data poses no algorithmic difficulty. The primary limitation of feature data is the practical one of estimating the feature pdfs (probability density functions) conditioned on target and on clutter. The feature pdfs used here are notional only.

The iLRT filter and its interpretation as an image processing system are discussed, together with the details of the resampling method and CLLR scoring for the iLRT both with and without features. The heritage method for track extraction is the focus

of the paper. The strengths and weaknesses of the method are discussed via simulated multistatic sonar examples.