Gradually migrating Gaussians: a more stable product approximation CPHD filter for multisensor multitarget tracking

Santosh Nannuru[1], Mark Coates[1], Michael Rabbat[1], and Stephane Blouin[2]

[1]Department of Electrical and Computer Engineering, McGill University, Email: santosh.nannuru@mail.mcgill.ca, mark.coates@mcgill.ca, michael.rabbat@mcgill.ca

[2]DRDC Atlantic Research Centre, Email: stephane.blouin@drdc-rddc.gc.ca

June 2, 2014

Abstract

For the single sensor scenario, provided certain modeling assumptions are reasonable, the cardinalized probability hypothesis density (CPHD) filter is an effective and computationally tractable solution for multi-target tracking. The key assumptions are: (i) targets move and generate measurements independently; (ii) clutter is generalized Poisson and independent of the measurements; (iii) each measurement is generated by a single target or by clutter; (iv) each target generates only one measurement. The single-sensor CPHD filter has been tested extensively through numerical simulations and practical deployments.

There has been much less work addressing the extension of the filter (and other filters based on random finite set theory) to the multi-sensor setting. Two major approaches have been proposed: the heuristic iterated-corrector approximation [1] and the multisensor product approximation [2]. In the iterated-corrector CPHD filter, the measurements from each sensor are processed sequentially. The output PHD obtained after processing the measurements from the first sensor is used as the input PHD for processing the measurements of the second sensor and so on. Disturbingly, the final result depends on the order in which sensor data is processed, and in some settings the choice of order can significantly change the result [3].

The product approximation CPHD filter is more principled, does not impose significantly stronger assumptions and has a similar computational overhead. In several important practical settings, however, the joint multi-sensor likelihood can be very peaky compared to the prior distribution. In such cases, direct implementation of the multisensor product approximation CPHD filter using a Gaussian mixture model can lead to very high variance in the weights of the Gaussian components. The practical filter performance can deteriorate substantially.

In this work we address this issue of high variance in the Gaussian component weights for the product approximation CPHD filter. We perform the CPHD update in multiple small steps. This gradual assimilation of the multisensor information avoids the mismatch between the predicted PHD and the peaky joint likelihood. Unlike the iterated-corrector filter however, our method is unaffected by the order in which measurements are processed and the weights are correctly matched to the product approximation calculations. The multiple small steps move more Gaussian components towards the important regions of the posterior distribution, leading to reduced variance of the component weights and a more stable update.

References

[1] R. Mahler, "The multisensor PHD filter: II. Erroneous solution via Poisson magic," in *Proc. SPIE Int. Conf. Sig. Proc., Sensor Fusion, Target Recog.*, Orlando, FL, U.S.A., Apr. 2009.

[2] ——, "Approximate multisensor CPHD and PHD filters," in *Proc. Int. Conf. Inf. Fusion*, Edinburgh, U.K., Jul. 2010.

[3] S. Nagappa and D. E. Clark, "On the ordering of the sensors in the iterated-corrector probability hypothesis density (PHD) filter," in *Proc. SPIE Int. Conf. Sig. Proc., Sensor Fusion, Target Recog.*, Orlando, FL, U.S.A., Apr. 2011.

DRDC-RDDC-2015-N036