Solving the data association problem in multi-object tracking by Fourier analysis on the symmetric group

Risi Kondor
Gatsby Computational Neuroscience Unit
University College London
17 Queen Square, London WC1N 3AR
risi@gatsby.ucl.ac.uk

Andrew Howard and Tony Jebara
Computer Science Department, Columbia University
1214 Amsterdam Avenue, NY 10025, U.S.A.
{ahoward,jebara}@cs.columbia.edu

Abstract

In addition to modeling the position of individual targets, multi-object tracking must also address the combinatorial problem of matching objects to corresponding tracks. In general, maintaining a probability distribution over all $n!$ possibilities is clearly infeasible, while just maintaining an $n \times n$ matrix of "first order marginals" is a very impoverished representation. In this work we explain how to harness the theory of harmonic analysis on the symmetric group to get a hierarchy of approximations of increasing fidelity to this problem. Importatantly, not only are such band-limited approximations theoretically well justifiable, but they also admit efficient observations updates based on some ideas from Clausen's FFT for the symmetric group. Experiments show that our algebraic approach can outperform more conventional solutions to the data association problem while still remaining computationalliy feasible for n in the $30 \sim 40$ range.

The theory of group representations leads to an elegant generalization of Fourier transformation to arbitrary compact groups, even non-commutative one, such as the symmetric group $\mathbb{S}_n$ of permutations of n objects. Diaconis' well known book [1] has done a lot to popularize this theory in the statistics community, but its applicability to machine learning has until lately been largely unexplored.

The recent paper [3] by the present authors was the first to show the relevance of this theory to the data association problem, shortly followed by [2] which developed the framework even further. In general, the idea is that given a (continuous time) noise model on matchings, the space of distributions $p(\sigma)$ over permutations splits into different modes, and we find that the "low-frequency" modes correspond to Fourier components

$$\widehat{p}(\rho) = \sum_{\in \mathbb{S}_n} p(\sigma)\rho(\sigma)$$

indexed by certain specific irreducible representations $\{\rho\}$ of $\mathbb{S}_n$. We explain what these representations are, what their interpretation is in terms of k-ary interactions between targets/tracks, and how this all relates to random walks on the Cayley graph. Note that because $\mathbb{S}_n$ is a non-commutative group, the $\widehat{p}(\rho)$ Fourier components will actually be matrices.

Beyond just proposing a band limited model for data association, our framework must also include efficient filtering/inference algorithms that can handle everything we need to do with n in $O(n^K)$ time, preferably with a small K. We explain how the power of fast Fourier transforms can be harnessed to accomplish just that, implementing Bayesian online observation updates of the form

$$p\,(\,\sigma\,|\,O_{i\to j}\,) = \frac{p\,(\,O_{i\to j}\,|\,\sigma\,)\;\,p(\sigma)}{\sum_{\sigma'\in\mathbb{S}_n} p\,(\,O_{i\to j}\,|\,\sigma'\,)\;\,p(\sigma')}\,, \tag{1}$$

in just $O(n^{2k+1})$ steps, where k is the order of the "largest frequency" Fourier component that we maintain in the system. In practice this means that in a real time system it is feasible to maintain second order information (and thus outperform most conventional methods) for n in the $30 \sim 40$ range. The first and second order marginals can be computed from $\widehat{p}$ at similar cost, and the noise model comes almost for free, since it is diagonalized in Fourier space.

We performed experiments on real world air traffic control data and found that our new flexible method does indeed offer improvements over other approaches, such as [4]. In the experiments we tried to isolate the data association problem from the location model. An integrated model or an application of our methodlogy to more complicated dynamical systems of similar combined combinatorial/continuous character are interesting further directions to explore.

Recent developments have also put the issue of projecting back to the space of band-limited representations and of more general update rules, or in general, propagation of band-limited functions on more complicated graphical models, at the forefront of research (see [2]). Clearly, the approximation of dynamical systems with combinatorial elements, such as multi-object tracking, is a rich field for future research.

References

[1] Persi Diaconis. *Group Representation in Probability and Statistics*, volume 11 of *IMS Lecture Series*. Institute of Mathematical Statistics, 1988.

[2] Jonathan Huang, Carlos Guestrin, and Leonidas Guibas. Efficient inference for distributions on permutations. In *Advances in Neural Information Processing Systems 20*, 2008.

[3] Risi Kondor, Andrew Howard, and Tony Jebara. Multi-object tracking with representations of the symmetric group. In *Proceedings of the Eleventh International Conference on Artificial Intelligence and Statistics*, 2007.

[4] Brad Schumitsch, Sebastian Thrun, Gary R. Bradski, and Kunle Olukotun. The information-form data association filter. In *NIPS*, 2005.