

Active Target Detection with Navigation Costs: A Randomized Benchmark

Sunav Choudhary*, Dhruva Kartik†, Naveen Kumar*, Srikanth Narayanan* and Urbashi Mitra*

*Dept. of Electrical Engineering, University of Southern California

3740 McClintock Ave, Los Angeles CA 90089

†Dept. of Electrical Engineering

Indian Institute of Technology, Guwahati, Assam, India

Email: {sunavcho,komathnk,ubli}@usc.edu, shri@siipi.usc.edu, dhruva.kartik@iitg.ac.in

Abstract

Detecting the presence of a target within a field, is a long-standing problem of interest in a variety of applications from environmental to military. We consider a framework in which a single autonomous vehicle (AV), physically samples a decaying separable field generated by the target, to determine the target location. The key metric of interest is the trade-off between the sample and computational complexity of target detection, the navigational cost for the AV and the detection error probability. What makes this problem non-trivial is the aspect of navigational cost and the emphasis on detection error probability rather than field reconstruction error. In the absence of these two features, standard low-rank matrix completion theory is order-optimal in terms of sample complexity for field reconstruction. We derive an active (adaptive) sampling and reconstruction strategy for target detection which takes the form of a *partial* low-rank matrix completion algorithm and analyze its detection error probability and sample complexity for exponentially decaying target signatures. We are able to simultaneously improve the sampling and computational complexity, respectively by a factor of $\frac{n \log n}{m \log m}$ and $\frac{R(n)}{R(m)}$, as compared to a brute force application of matrix completion, where $m \times m$ is the size of the sub-sampled $n \times n$ grid and for $p \in \{m, n\}$, $R(p)$ is the running time for a $p \times p$ rank-1 matrix completion problem from $\Theta(p \log^2 p)$ random measurements. The results are consistent with the intuition that detection should be ‘easier’ than estimation. We also examine the effect of the sampling pattern on the navigational cost for the AV subject to meeting a detection error probability. As a consequence of the strong concentration properties of the Euclidean Traveling Salesperson Problem, it turns out (somewhat surprisingly) that separation of navigational path planning and sampling pattern design is order optimal for uniform prior on the target location and a given detection error probability. It is at present unclear, whether such separation holds for a non-uniform prior on the target location or for other cost structures on the sampling process.