Image Classification On The Edge:
Faster Multi-Camera Object Tracking

Nick Nordlund (Yale University), **Heesung Kwon** (ARL), **Geeth de Mel** (IBM UK), **Leandros Tassiulas** (Yale University).

Problem Setting

An object of interest must be located and tracked in real time as it moves through an area covered by a dense video surveillance network.

- At any time, multiple cameras may observe the object of interest from different angles.
- Cameras have some local processing capabilities but they achieve more accurate image classification results by offloading video streams to a central cloud server.
- Central cloud becomes backlogged when too many cameras offload video data.

Method

Leverage the spatial and temporal relationships of the surveillance network to

- Estimate the general location of the object of interest using noisy image classification on the edge
- Provide an estimate for a prior distribution to input into a linear quadratic estimator

Model

- Video Cameras are generated randomly and observe some cluster of nearby grid cells
- Random Gaussian noise is added to the true observation to simulate cheap image classifiers

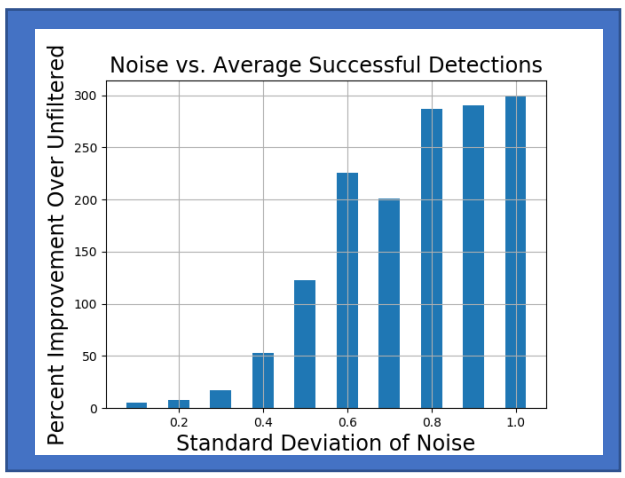

Figure 1: How much improvement refining the estimate gives over selecting the max-value noisy sensor reading

Algorithm

- At time t, there is an estimate of the state from the previous time step x_{t-1}
- Get prior estimate for x_t from random walk
- Gather measurement data from noisy image classifiers at top-k cameras based on x_{t-1} for posterior estimate of state
- Combine prior and posterior estimates for final estimate of x_t

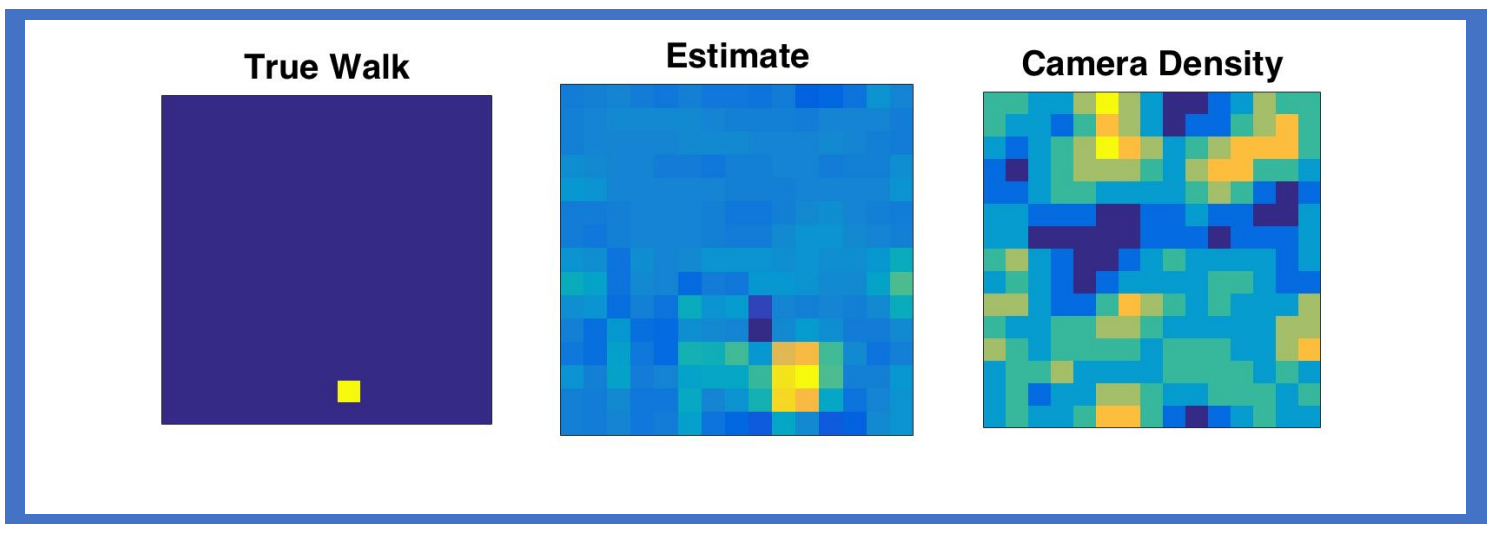

Figure 1: Where the target actually is, where the system thinks it is, and how the cameras are distributed in space

Distributed Analytics and Information Science International Technology Alliance
Annual Fall Meeting, September 2018