Robust, scalable, distributed semantic mapping for co-robots

Roberto Tron (Boston University), Dario Pompili (Rutgers University)

https://sites.bu.edu/tron/robust-scalable-distributed-semantic-mapping

Overview

Goal: enable multiple co-robots to robustly and efficiently map and understand the environment despite common problems such as fast motion

- Multiple heterogeneous robots share data and computational resources
- Incorporate semantic information (object detections) into mapping.
- Use the redundancy from cycles of to detect and correct inconsistencies
- Make intelligent use of resources through approximate computing

ADMM statistical outlier identification

Motivation: In modern mapping solutions, we often need to find links between different parts of a dataset, e.g., to handle tracking and loop closures. Object detections provide additional data association clues.

Goal: Obtain robust correspondences across parts of the map, despite gross noise (outliers), and provide localized estimates of the probability of errors.

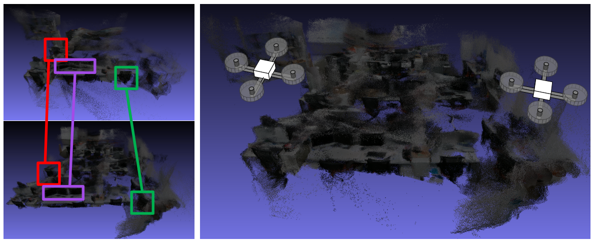

Principle: Exploit the redundancy contained in cycles of measurements, and use Expectation Maximization for estimating the probability that each measurement is an outlier

Example: estimate absolute rotations from relative ones.

- *Evidence*: closure error on cycles (assumes Gaussian noise).
- *Inference*: EM procedure which estimates the covariances for inlier/outliers (M-step), then use a novel method of inference based on Alternating Direction Method of Multipliers (ADMM, E-step)

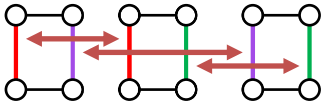

(a) Decomposition with cycles plus constraints on overlapping edges

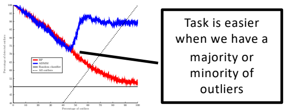

Task is easier when we have a majority or minority of outliers

(b) Precision of inferred outliers with ADMM vs BP

ADMM approach: Solve (exactly) the inference problem on each cycle separately, while iterating on Lagrange multipliers to enforce consistent inference results on the shared edges

Results: We can perform inference with accuracy higher than Belief Propagation, especially for the case of a majority of outliers

Exploiting Correlations in Streaming Video

Goal: Leverage spatial and temporal correlation in subsequent frames of videos from drones with overlapping Field of Views (FoVs) to identify "good-enough parameters" that work with resource-constrained drones.

Temporal Correlation: Offline learning with Markov Decision Process (MDP)

- Identify offline parameters (e.g., # of object proposals) that can reduce computations with a tolerable loss of accuracy
- Learn selection policy based on an MDP
- *State:* classifier score, current parameter selection; *Actions:* choice of Pareto-optimal parameters; *Reward:* ratio between classifier score and time

Spatial Correlation: Collect data only from a subset of the network cameras

- Use low-complexity metrics, e.g., intersection between the histograms of two images to quantify data correlation between different camera nodes

Results: Decrease in execution time by 20-70% for accuracy of 100-98% on video datasets

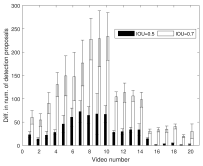

(a) Average difference in the no. of object-detection proposals between subsequent frames of the videos by varying the Intersection Over Union (IOU)

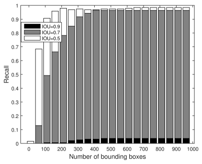

(b) Variation of recall with no. of object-detection proposals.

Multi-Agent Coordination

Motivation: Traditional Reinforcement Learning (RL) cannot scale to large state and action spaces

Goal: Develop a Multi-Agent Deep RL. framework that overcomes the above limitations for real-time coordination of drones.

Principle: We propose Distributed Advantage Actor Critic (DAAC) architecture based on the Deep Deterministic Policy Gradient where the actor-critic architecture is realized on a group of drones instead of threads on a computer.

Applications: Novel mapping and search strategies.

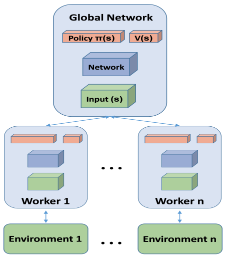

(c) Distributed Advantage Actor Critic (DAAC) architecture for parallel training, where each worker is a drone.

Theoretical characterization of a-priori robustness in the presence of outliers

Motivation: It is empirically known that by optimizing a robust cost function when solving a SLAM or other localization problem, we can effectively remove the effect of outliers; however, there have been no previous work on theoretical conditions predicting when this is possible.

(a) When we minimize a non-robust cost (L2 norm), a single outlier biases the entire estimation

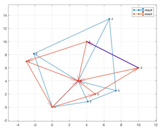

(b) When we minimize a robust cost (L1 norm), the effects of a single outlier are canceled

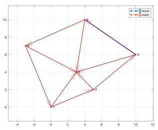

To provide our results, we assume here the simple translation-only case

Principle: By using convex optimization theory and the dual simplex method, we can characterize when, for a specific instance, we can reach the true solution (*uniquely verifiable*), a cost-equivalent solution (*verifiable*), or neither (*non-verifiable*). This depends only on the support of the outliers and the topology of the graph. We can then compute a *verifiability probability* curve.

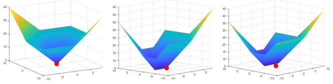

(a) U-verifiable (b) Verifiable (b) Non-verifiable

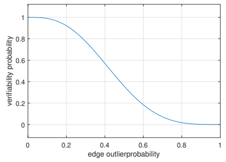

2020 National Robotics Initiative (NRI) Principal Investigators' Meeting
FEBRUARY 27 - 28, 2020 | ARLINGTON, VIRGINIA

Award ID#: IIS-1734454/IIS-1734362
(09/2017-08/2020)