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

BOSTON UNIVERSITY

RUTGERS THE STATE UNIVERSITY OF NEW JERSEY

Overview

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

- Use the redundancy from cycles of to detect and correct inconsistencies
- Vision-based control for coordination, homing, exploration
- Make intelligent use of resources through approximate computing
- Multiple heterogeneous robots share data and computational resources

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, loop closures, and coordinated collection from multiple agents.
Goal: Detect *spurious outliers* that can affect entire reconstruction.

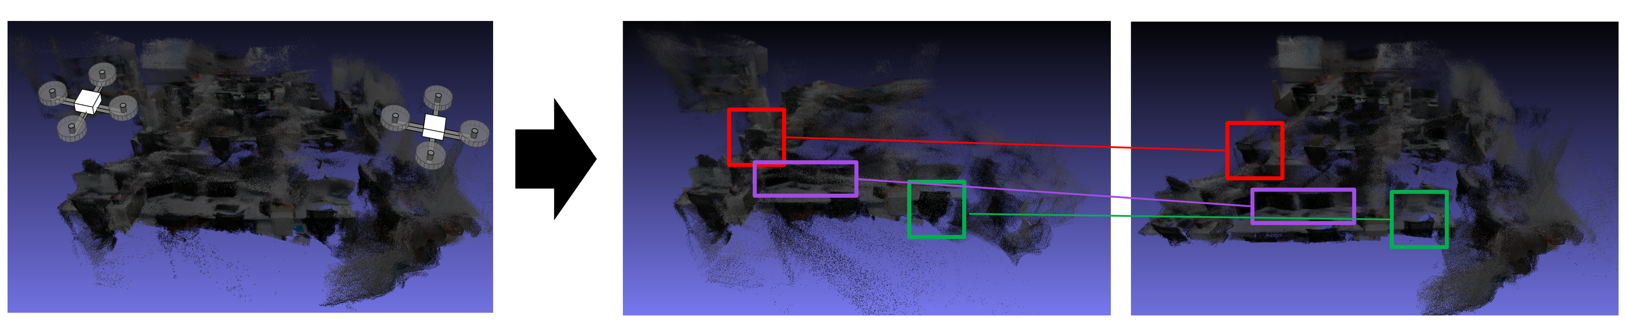

Principle and Results: On the theoretical side, we present an algorithm to compute the probability of recovering a good solution via robust optimization, for a specific graph instance and for known edge-wise outlier probabilities. Our solution is based on the dual simplex method from convex optimization

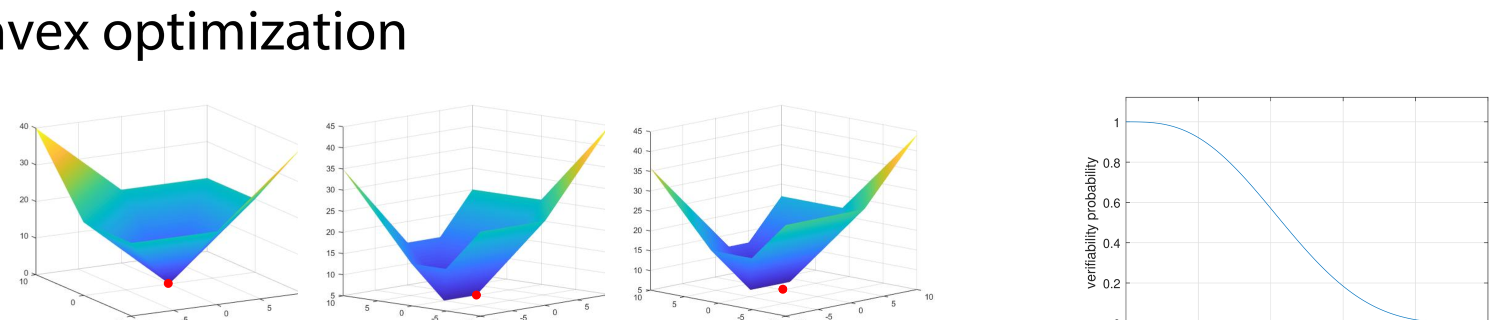

(a) U-verifiable (b) Verifiable (b) Non-verifiable

On the practical side, we can propose a new loopy graph statistical inference algorithm that uses ADMM to solve (exactly) the inference problem on each cycle separately, while iterating on Lagrange multipliers to enforce consistent inference results on the shared edges. We can perform inference with accuracy higher than Belief Propagation, especially for the case of a majority of outliers

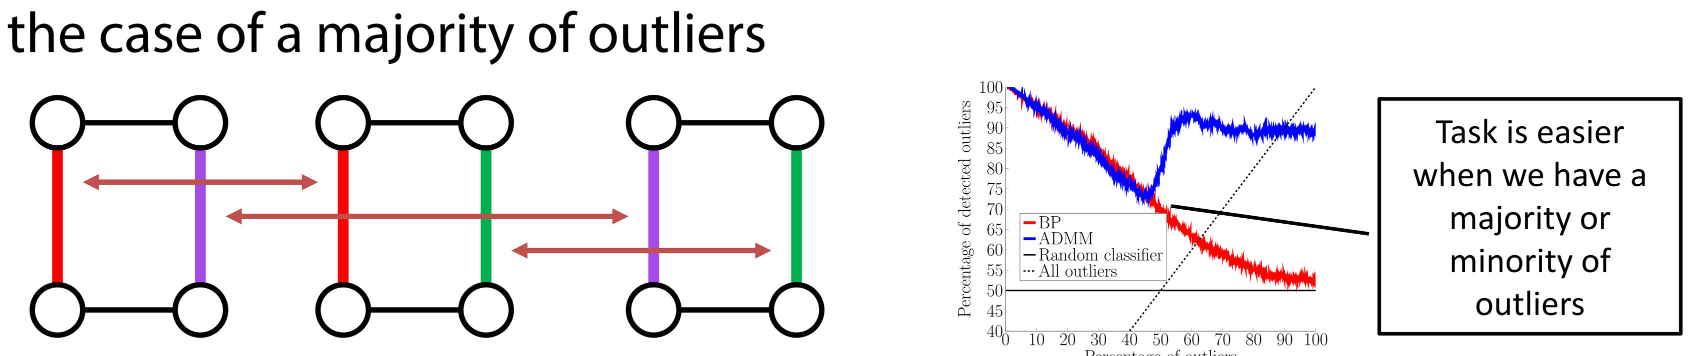

(a) Decomposition with cycles plus overlap constraints (b) Precision of inferred outliers, ADMM vs BP

Multi-Agent Vision-based Control

Motivation and Goal: Multi-agent mapping can benefit from vision-based low-level control of single agents and coordination between agents
Principle: Define control laws that are based on bearing measurements (no distances) and distributed (no central coordination). Use non-smooth Lyapunov theory and geometric arguments for global convergence.
Results: Algorithms for rigidity recovery, rendez-vous and formation control with directed graphs, visual homing with small FoV

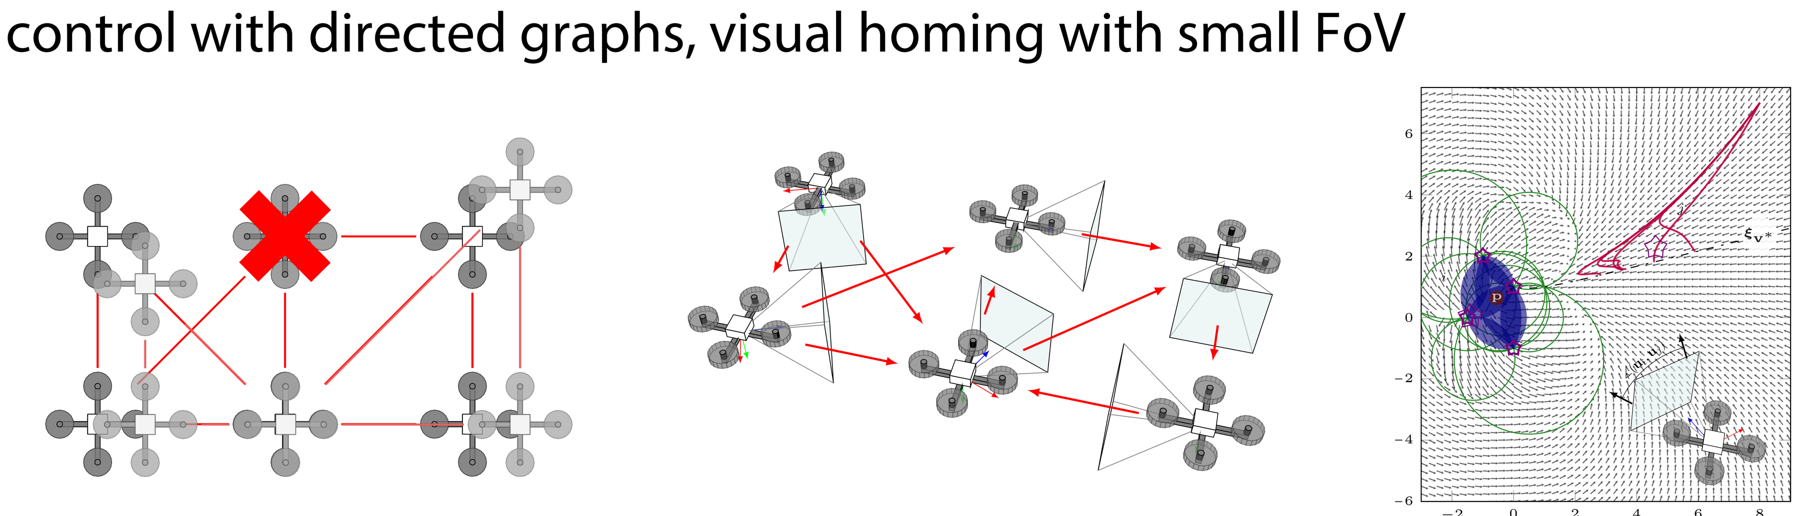

Exploiting Correlations in Streaming Video

Motivation: Mismatch between computations required for SotA machine learning algorithms (e.g., DNNs) and resources on mobile robots
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.

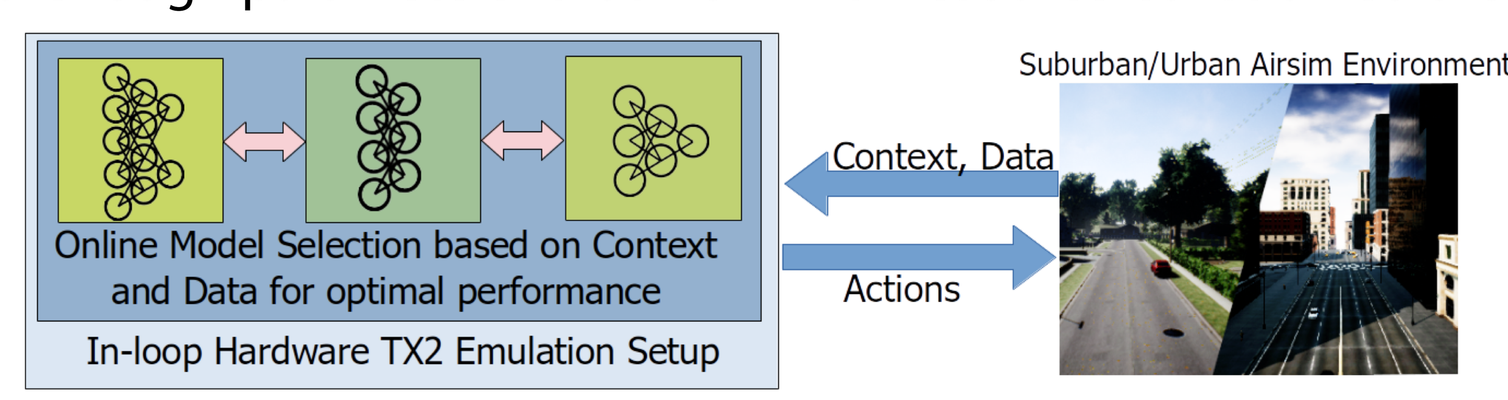

Principle and Results: First, synthesize multiple machine learning models with different speed/accuracy tradeoffs (e.g., DNN quantization and pruning). Online, use a Markov Decision Process (MDP) for an online optimization of the ratio between classifier score and time. We obtained a significant computational speedup with negligible loss of performance.

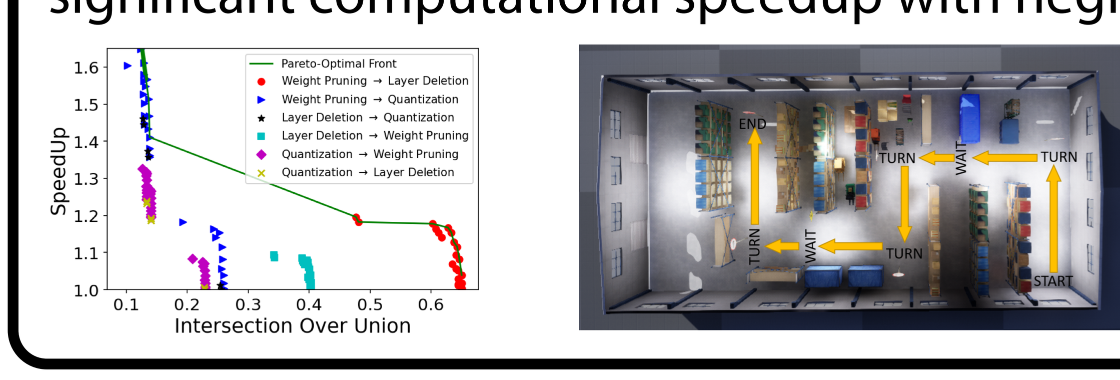

	Fixed	Ours
Average FPS	1.95 ±4.33%	35.72 ±27.28%
Average velocity [cm/s]	4.9 ±2.51%	32.21 ±2.85%

Multi-agent RL coordination for search

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 and Results: Distributed Advantage Actor Critic (DAAC) architecture based on the Deep Deterministic Policy Gradient and Recurrent Neural Networks. Learning is centralized, while rollout is completely distributed. Outperforms other learning-based approaches.

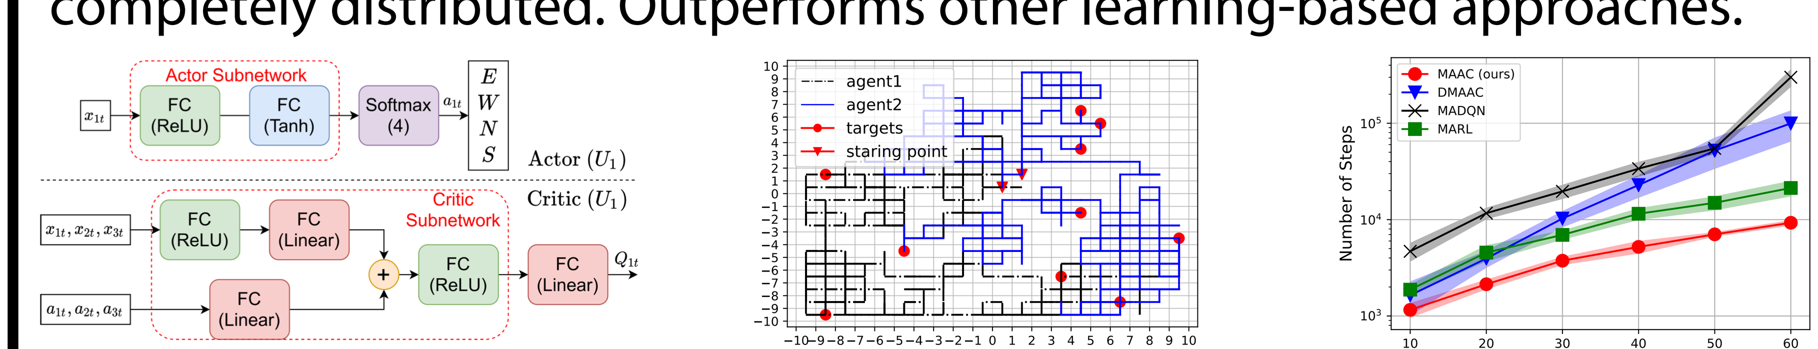

Single-agent RL exploration for SLAM

Motivation: Effectively tradeoff exploration vs accuracy
Goal: A Multi-Agent Deep RL framework to learn an effective exploration policy
Principle and Results: State (input to the controller) given by local view of the map, reward given by SLAM uncertainty. Work in progress, but some non-trivial results (e.g., moving backward)
Open source SLAM pipeline in Python, integrated with: Habitat-Sim, GTSAM, FBOW, TEASER++

https://github.com/armandok/pySLAM-D

2021 National Robotics Initiative (NRI) Principal Investigators' Meeting

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