ANPL Autonomous Navigation and Perception Lab

Bayesian Information Recovery from CNN for Probabilistic Inference

Dmitry Kopitkov and Vadim Indelman

TECHNION Israel Institute of Technology

Motivation

- State information recovery from stochastic sensor measurements is a fundamental task in robotics, challenged by a complex relation between state and measurements.
- Examples:
    - Simultaneous localization and mapping (SLAM)
    - Autonomous navigation/cars
    - Informative planning, active sensing
    - State transition in reinforcement learning

Related Work

- Measurement model is typically treated as given or hand-engineered
- Gaussian density assumption prevails, limiting state inference accuracy
- Images are handled through landmark detection/matching using hand-engineered feature detectors (e.g. SIFT). Unreliable due to mistakes in data association and uses only part of image information
- Recently more methods learn measurement likelihood in a supervised way via DL

Objective

- Infer robot pose ***x*** from odometry and image measurements by first learning image (representation) ***f*** likelihood
    - ***f*** is the output of a pre-trained CNN classifier
- Challenges:
    - ***x*** is unknown and only partially observed through measurements ***f***
    - Likelihood P(***f***|***x***) is very intricate with both first and second moments spatially changing
    - Opposite conditional P(***x***|***f***) is multimodal
- Contributions:
    - We learn likelihood from collected data via another NN
    - We combine it within Bayesian inference and infer robot's trajectory
    - No data association is required in our approach

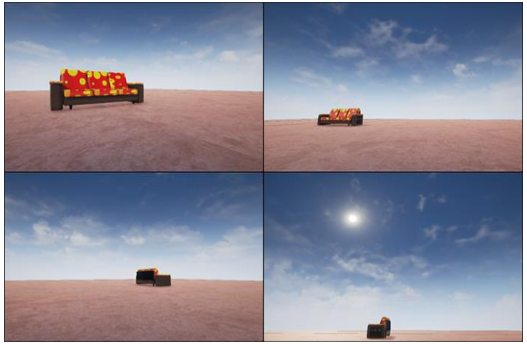

Key Idea

- Pre-deployment stage:
    - Image (feature representation) measurement likelihood P(***f***|***x***) is learned from training dataset via DL:

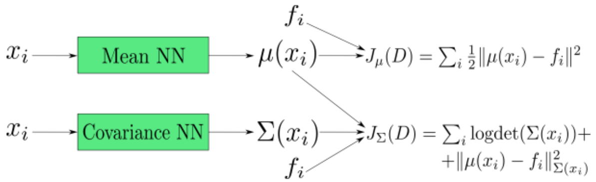

- Deployment stage:
    - Robot trajectory is estimated via Bayesian inference using the learned likelihood:

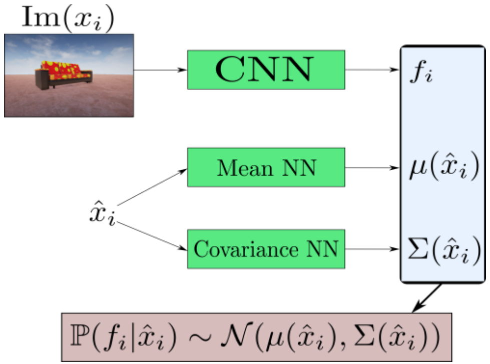

Bayesian Information Recovery from CNN for Probabilistic Inference

Dmitry Kopitkov and Vadim Indelman

TECHNION Israel Institute of Technology

Learned Likelihood

- “park bench” feature:

Ground truth:

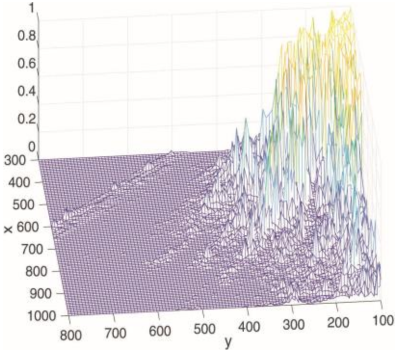

Learned mean:

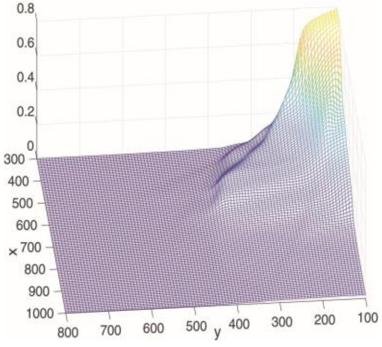

Learned std:

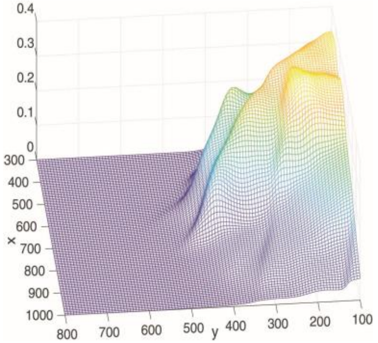

- “rocking chair” feature:

Ground truth:

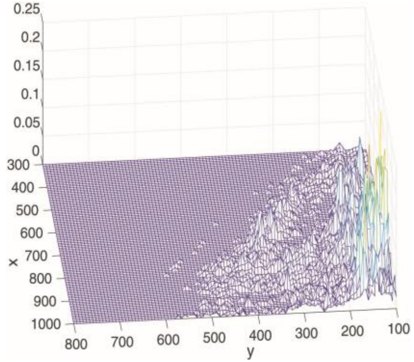

Learned mean:

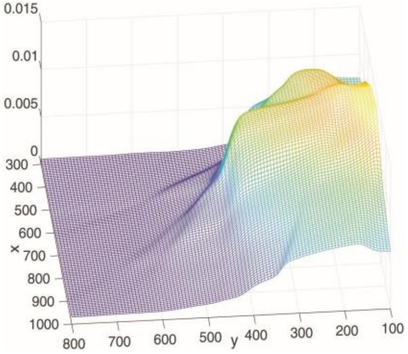

Learned std:

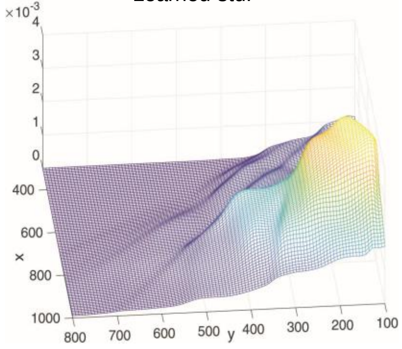

Validation area:

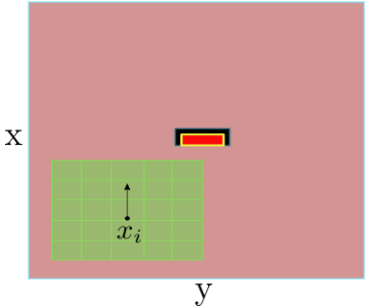

Bayesian Information Recovery from CNN for Probabilistic Inference

Dmitry Kopitkov and Vadim Indelman

Bayesian Formulation for Pose Estimation

- Belief over trajectory:

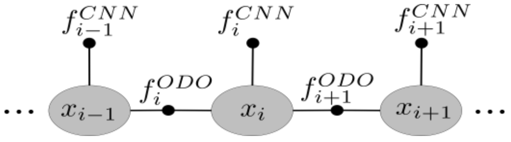

$$\mathbb{P}(X_k|history) \propto \prod_{j=1}^{n_f} f^j(X^j)$$

$$X_k^* = \arg\min_{X_k} \sum_{j=1}^{n_f} \operatorname{logdet}(\Sigma^j) + \sum_{j=1}^{n_f} \|h^j(X^j, z^j)\|_{\Sigma^j}^2$$

Results

- Estimated trajectory

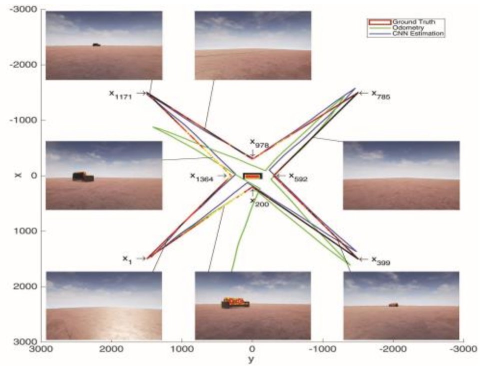

- Estimation error

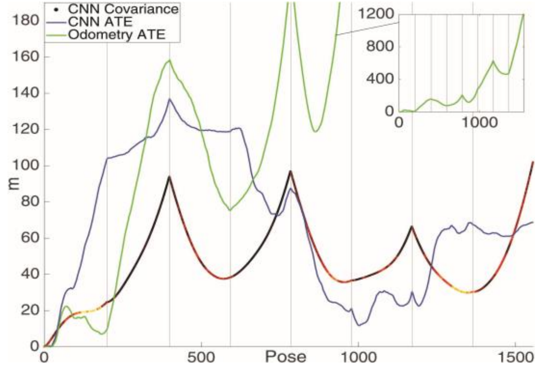

- Total error for 25 trained models under various inference configurations:

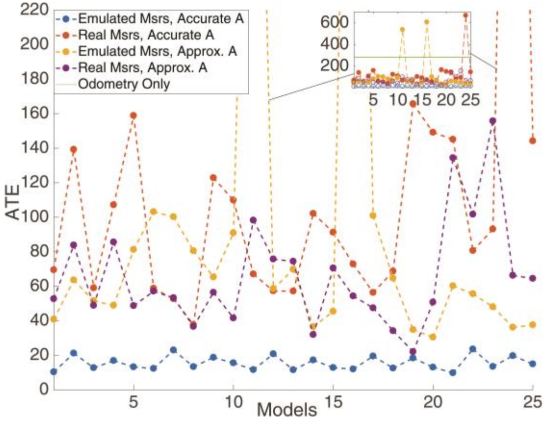

- Robust kernel inference

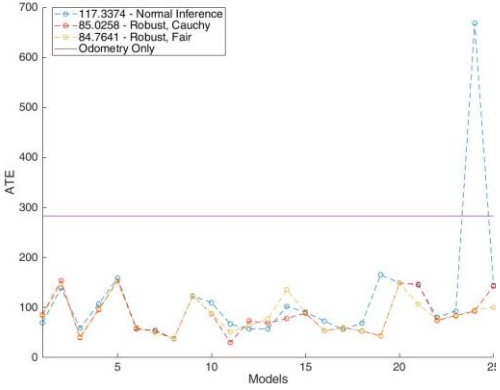

Conclusions

- State information can be recovered from DL-learned measurement likelihood
- The likelihood estimation was very rough
- Gaussian assumption is unrealistic
- Can we do better? (DeepPDF**)

** D. Kopitkov and V. Indelman, “Deep PDF: Probabilistic Surface Optimization and Density Estimation,” arXiv preprint arXiv:1807.10728, 2018.