Probably Unknown: Deep Inverse Sensor Modelling in Radar

Applied AI Lab, University of Oxford, UK

Rob Weston, Sarah Cen, Paul Newman, Ingmar Posner

For more information please email robw@robots.ox.ac.uk

Introduction

- Suppress noise artefacts in raw FMCW radar by **learning** sensor characteristics from data.
- Account for aleatoric sensor noise by modelling **heteroscedastic uncertainty** - uncertainty that varies with scene context. This allows regions of space which are likely to be occluded to be identified.
- **Self-supervised** using partial lidar labels allowing a robot to learn by simply traversing an environment.
- Outperforms classical CFAR filtering approaches in **detection performance.**
- Can be used as an **inverse sensor model** for probabilistic occupancy grid mapping.

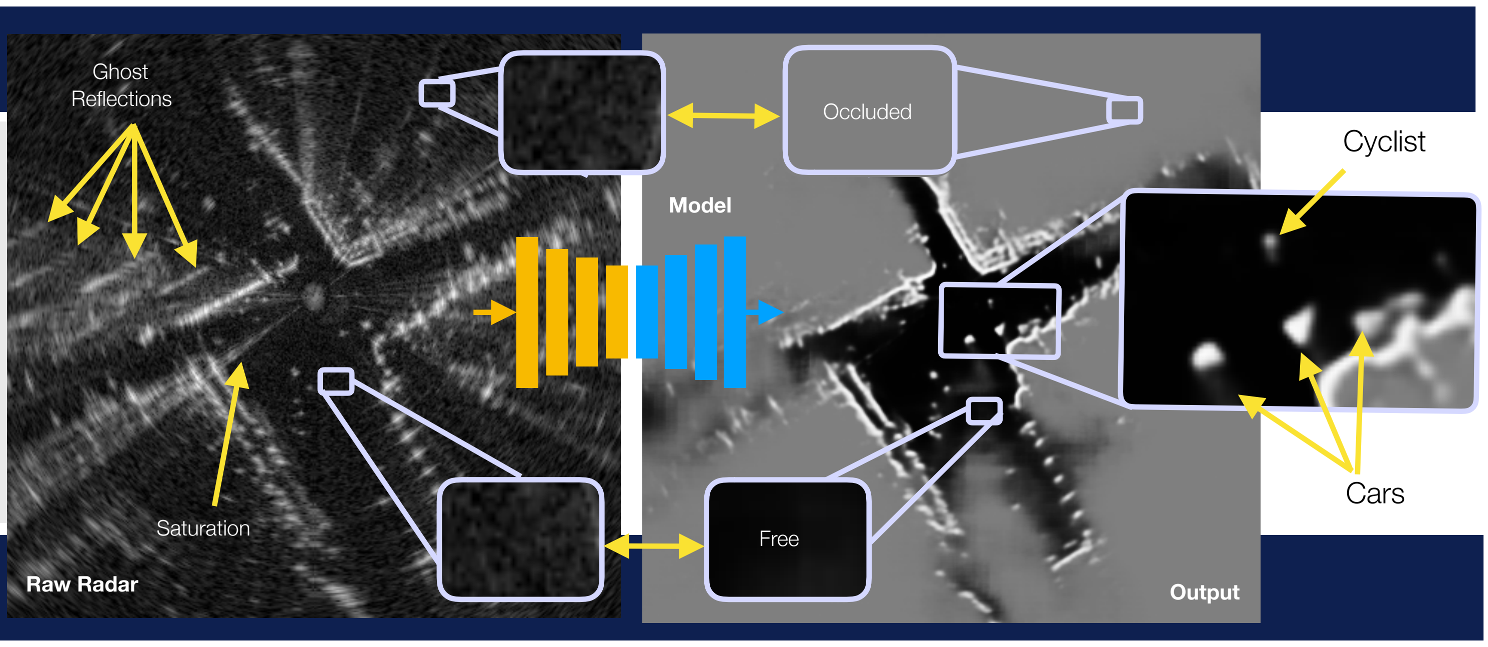

Results

Detection Performance

Outperforms classical CFAR filtering methods at detecting free and occupied space.

Method	Intersection over Union		
	Occupied	Free	Mean
CFAR (1D polar)	0.24	0.92	0.50
CFAR (2D Cartesian)	0.20	0.90	0.55
Static thresholding	0.19	0.77	0.48
Deep ISM (ours)	0.35	0.91	**0.63**

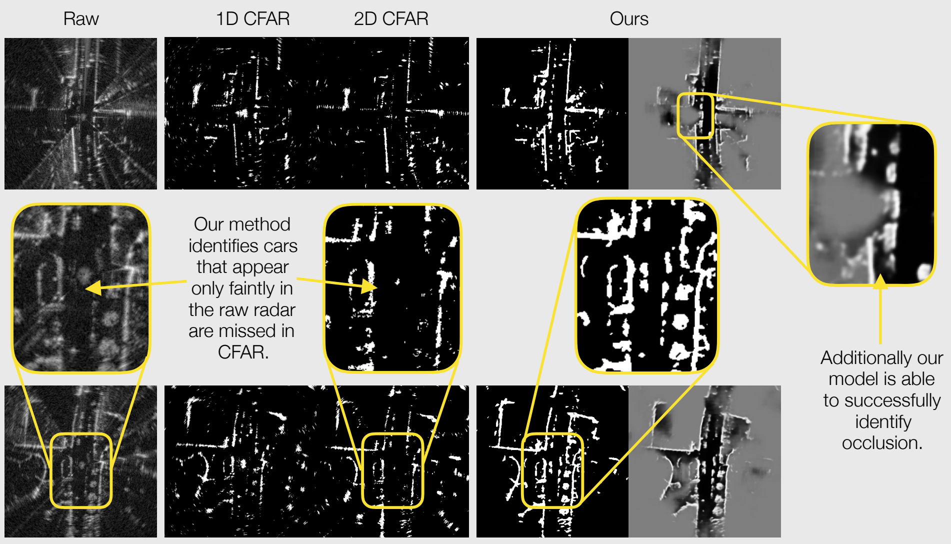

Heteroscedastic Uncertainty Prediction

The predicted uncertainty allows regions of space that are likely to be occluded to be identified.

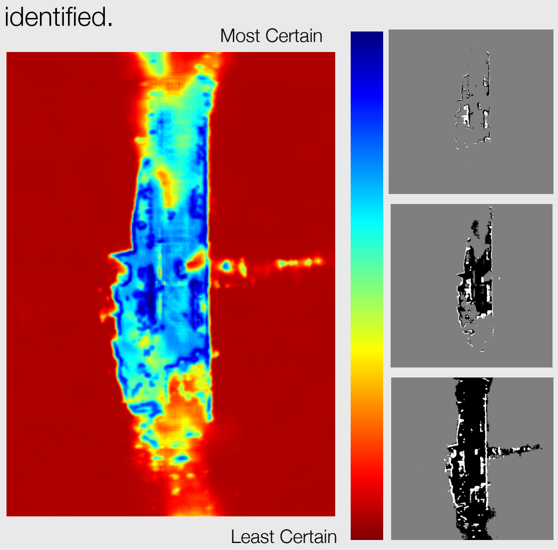

Inverse Sensor Model

Used as an Inverse Sensor Model an occupancy grid map can be constructed using a Binary Bayes filter.

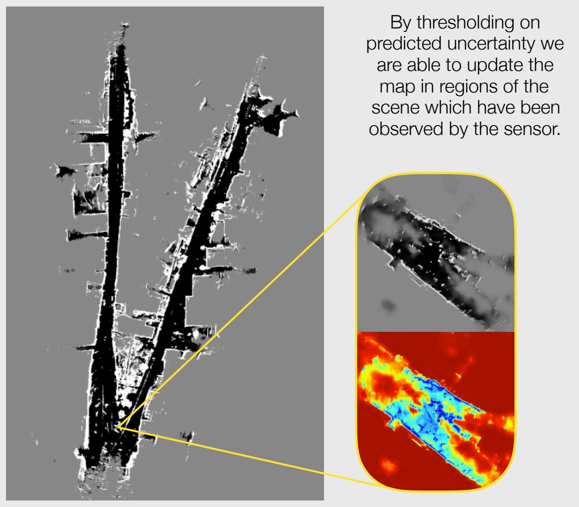

Approach

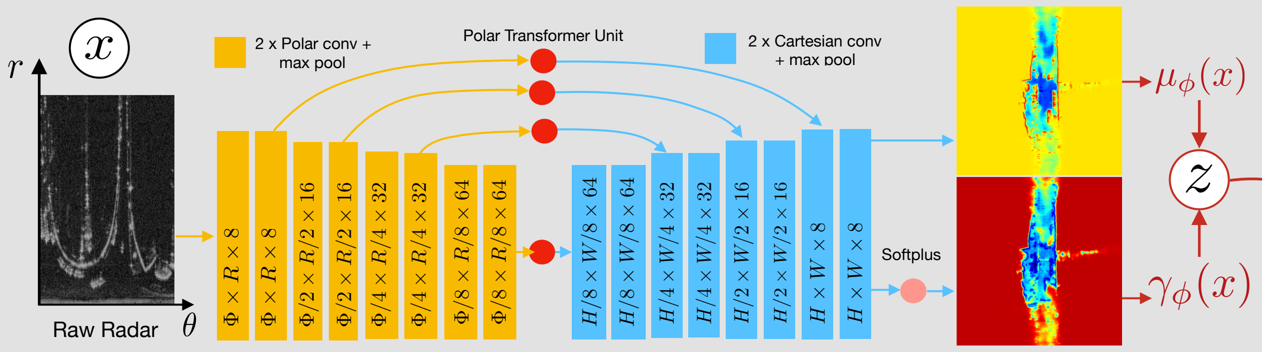

1. Model Aleatoric Uncertainty

Parameterise the mean **and** standard deviation of normally distributed logit z with a neural network.

$$p_\phi(z|x) = \mathcal{N}(z|\overbrace{\mu_\phi(x)}^{\text{Mean}}, \overbrace{\gamma_\phi(x)}^{\text{Standard Deviation}})$$

Predicted Logit

2.Training

Utilise partial labels generated from 3D lidar to **learn** about occupied and free space.

Force model to be **uncertain** where labels do not exist through **prior**.

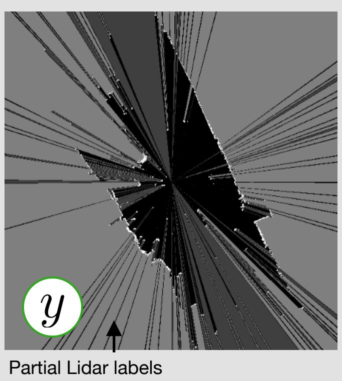

$$\phi^* = \arg\max_\phi \underbrace{\mathbb{E}_{p_\phi(z|x)}[\log \overbrace{p(y|z)}^{\text{Bern}(y|\sigma(z))}]}_{\text{Likelihood}} - \underbrace{d_{kl}[p_\phi(z|x) \| \overbrace{p(z)}^{\text{Norm}(z|0,I)}]}_{\text{Prior}}$$

No Loss

Free

Occupied

Occluded

3. Inference

Utilise **fast** and **accurate** analytic approximation to marginalise out the uncertainty associated with the predicted logic.

$$p(y|x) = \int p(y|z)\, p_\phi(z|x)\, dz$$

$$\approx \sigma\left(\frac{\mu_\phi(x)}{s_\phi(x)}\right)$$

$$s_\phi(x) = \left(1 + \left(\gamma_\phi(x)\sqrt{\pi/8}\right)^2\right)^{1/2}$$

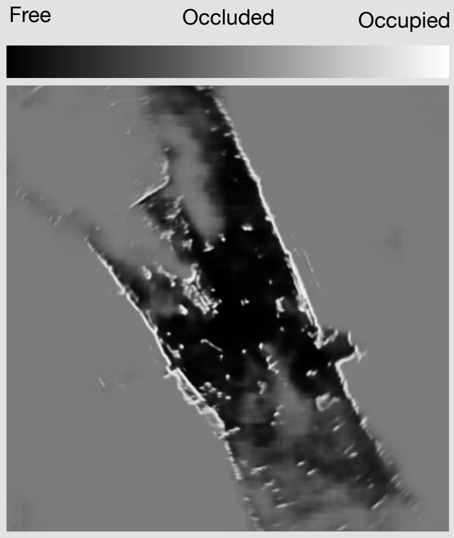

ORI Oxford Robotics Institute

University of Oxford