Towards Safe Autonomous Driving: Capture Uncertainty in Deep Object Detectors

Di Feng

Perception and Sensors for Autonomous Driving
Corporate Research, Robert Bosch GmbH, Stuttgart

Nov. 8, 2019
IROS 2019, Macau

Outline

1. Motivation
2. Uncertainties in object detection networks
3. Probabilistic LiDAR object detectors
4. Challenges

1. Motivation

Object detection

- Bounding box (2D or 3D) + Classification score
- Deep learning has advanced object detection
- Most object detectors are deterministic – we need **probabilistic** detectors!

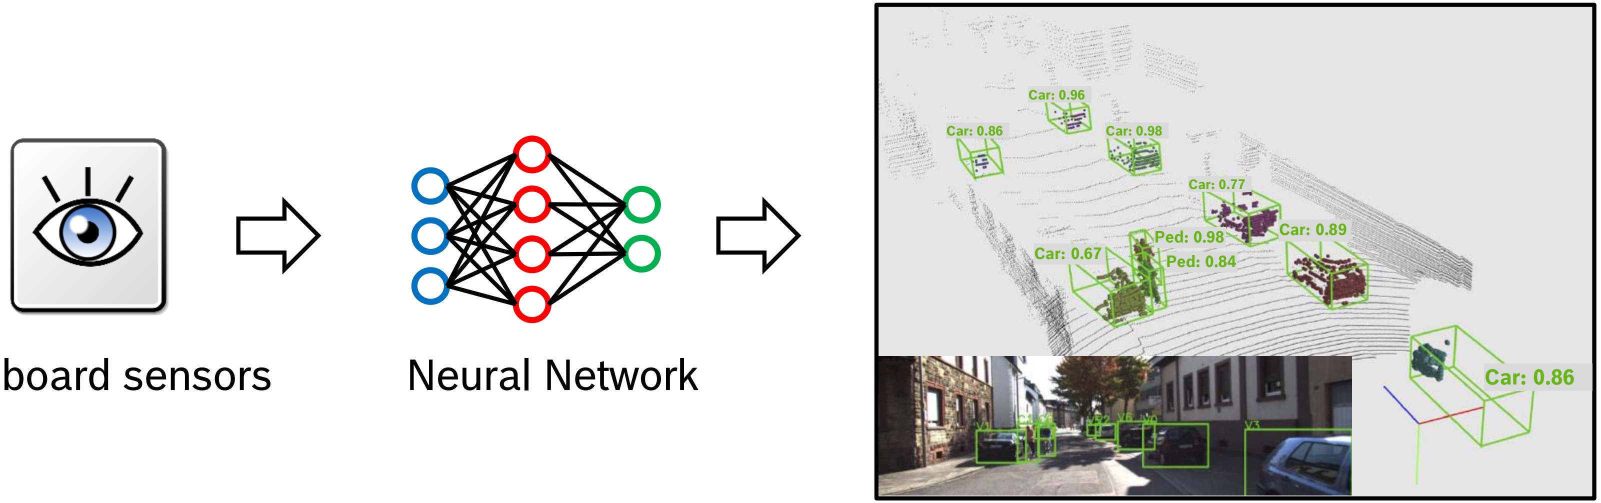

[Qi, et al., CVPR’18]

1. Motivation

Autonomous car in the wild

Adverse weather

Night drive

Unseen objects

https://www.flickr.com/photos/davidmoisan/3120533363/
https://www.flickr.com/photos/wackelijmrooster/4095146153
https://commons.wikimedia.org/wiki/Category:Embilipitiya

1. Motivation

Reliable uncertainty builds trust

“From an ecological and evolutionary perspective, humans may turn out to be good intuitive statisticians ...”

[Cosmides, et al., Cognition’96].

https://spectrum.ieee.org/automaton/robotics/artificial-intelligence/robot-companions-to-befriend-sick-kids-at-european-hospital

1. Motivation

Increasing robustness of the general system

On-board sensors

Object located	Yes (Prob. = 90%)
Object classified	“Pedestrian” (Prob. = 60%)
Camera signal	Uncertain
LiDAR signal	Certain

Probabilistic perception

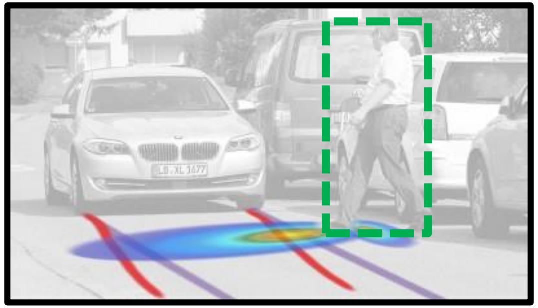

Prediction

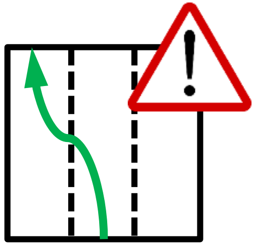

Decision making

1. Motivation

What can an ideal probabilistic object detector look like?

Front car recognized	Front car localized	Left car recognized	Left car localized
Likely	Uncertain	Very uncertain	Unlikely

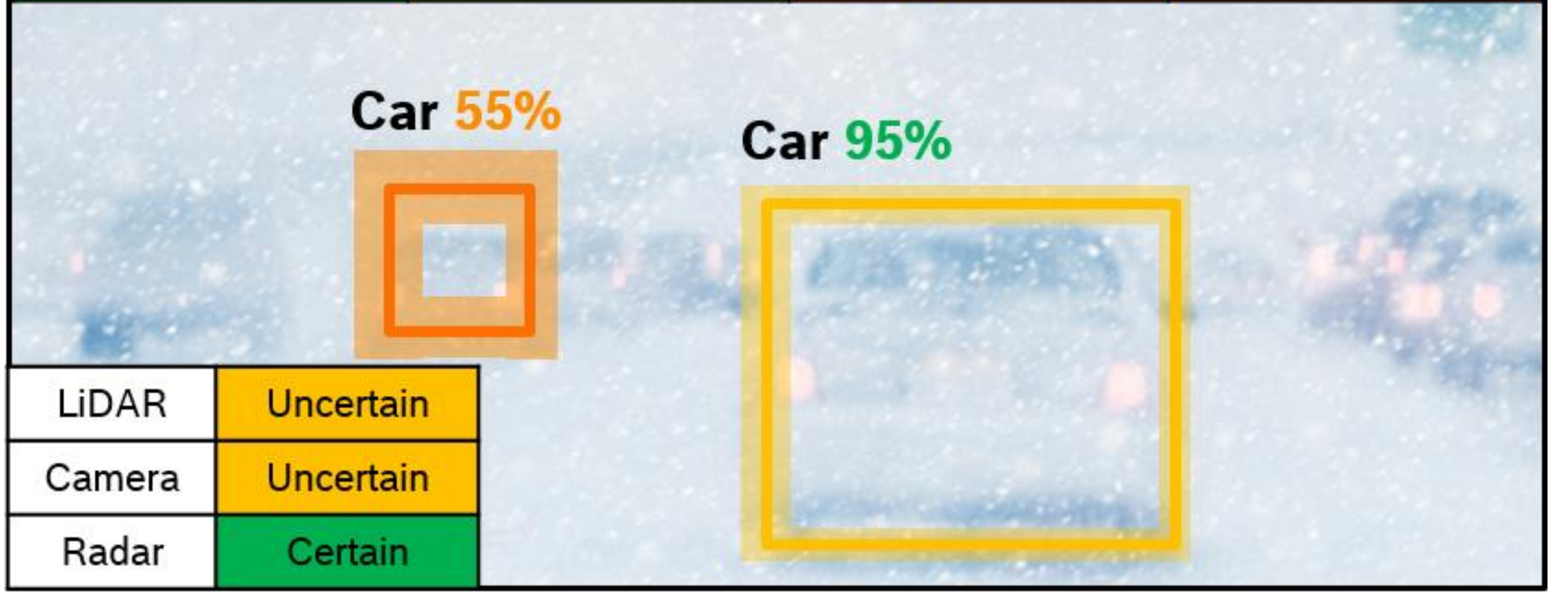

LiDAR	Uncertain
Camera	Uncertain
Radar	Certain

1. Motivation

An ideal object detector should model uncertainties ...

- **Holistic:** uncertainties in cls + reg
- **Well-calibrated:** represent empirical frequency
- **Explainable:**
  - reflect environmental noises
  - Comparable among sensors
  - reflect model deficiency
- **Useful**

Front car recognized	Front car localized	Left car recognized	Left car localized
Likely	Uncertain	Very uncertain	Unlikely

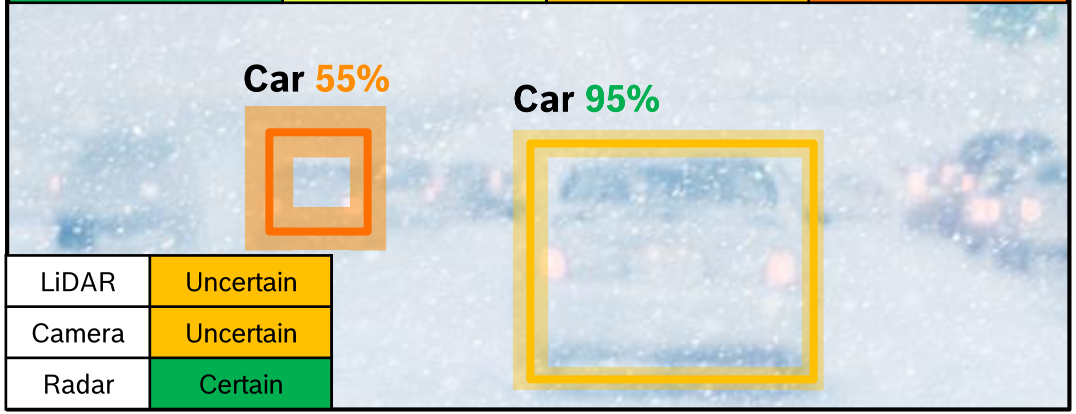

LiDAR	Uncertain
Camera	Uncertain
Radar	Certain

1. Motivation

Our attempts towards probabilistic object detectors

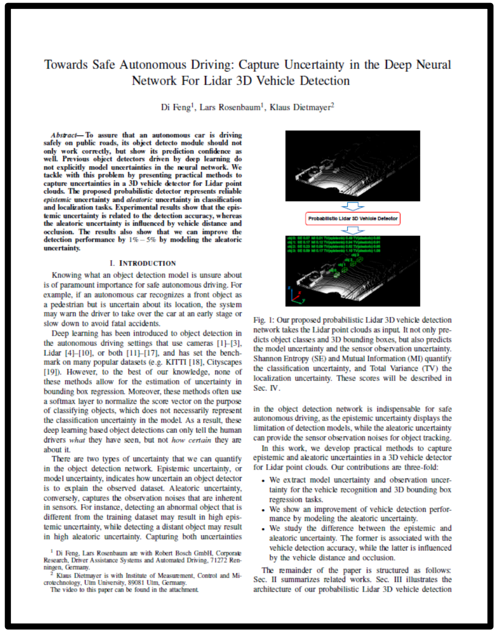
Towards Safe Autonomous Driving: Capture Uncertainty in the Deep Neural Network For Lidar 3D Vehicle Detection

Di Feng[1], Lars Rosenbaum[1], Klaus Dietmayer[2]

[Feng et al., ITSC’18]

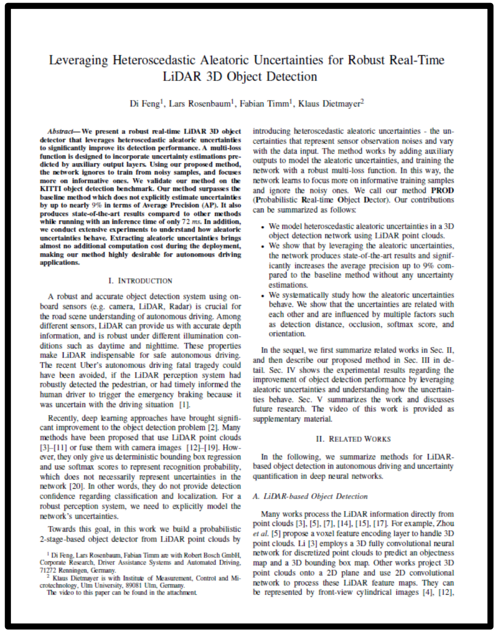
Leveraging Heteroscedastic Aleatoric Uncertainties for Robust Real-Time LiDAR 3D Object Detection

Di Feng[1], Lars Rosenbaum[1], Fabian Timm[1], Klaus Dietmayer[2]

[Feng et al., IV’19a]


Deep Active Learning for Efficient Training of a LiDAR 3D Object Detector

Di Feng[1,4], Xiao Wei[1,2], Lars Rosenbaum[1], Atsuto Maki[2], Klaus Dietmayer[4]

[Feng et al., IV’19b]

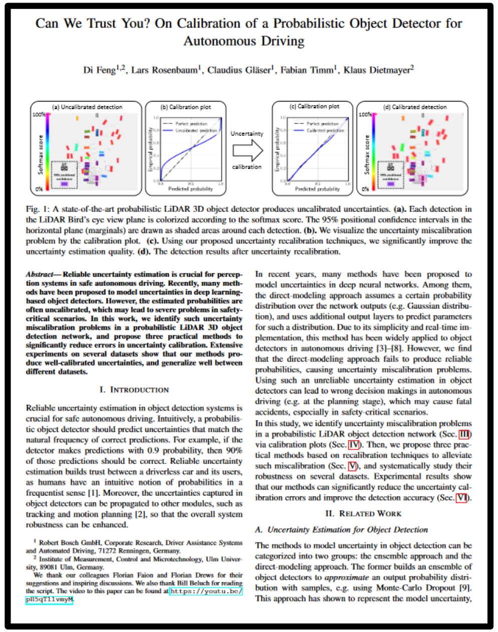
Can We Trust You? On Calibration of a Probabilistic Object Detector for Autonomous Driving

Di Feng[1,2], Lars Rosenbaum[1], Claudius Gläser[1], Fabian Timm[1], Klaus Dietmayer[2]

[Feng et al., IROS’19]

universität uulm

BOSCH

Outline

1. Motivation
2. **Uncertainties in object detection networks**
3. Probabilistic LiDAR object detectors
4. Challenges

2. Uncertainties in object detection networks

What kind of uncertainties can we model in object detection networks?

- Epistemic uncertainty: model’s capability to describe data

- Aleatoric uncertainty: observation noises (e.g. environment, sensors)

[Kendall et al., NeurIPS’17]

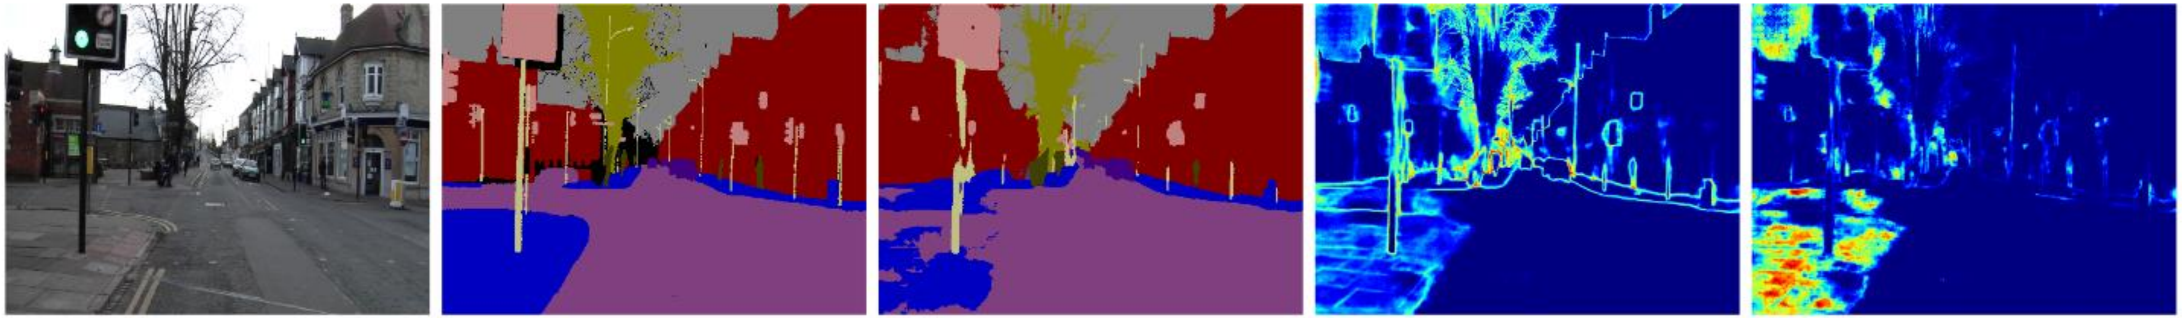

(a) Input Image (b) Ground Truth (c) Semantic Segmentation (d) Aleatoric Uncertainty (e) Epistemic Uncertainty

2. Uncertainties in object detection networks

Modeling uncertainties via Bayesian neural networks [MacKey, Neural'92]

$$\underbrace{\boxed{p(\mathbf{y}|\mathbf{x})}}_{\text{Predictive uncertainty}} = \int \underbrace{\boxed{p(\mathbf{y}|\mathbf{x}, \mathbf{W})}}_{\text{Aleatoric uncertainty}} \underbrace{\boxed{p(\mathbf{W}|\mathcal{D})}}_{\text{Epistemic uncertainty}} d\mathbf{W}$$

$\mathbf{x}$: Input vector

$\mathbf{y}$: Prediction output vector

$\mathcal{D}$: Training dataset

$\mathbf{W}$: Network weight variables

2. Uncertainties in object detection networks

Modeling uncertainties in object detection networks: a big work

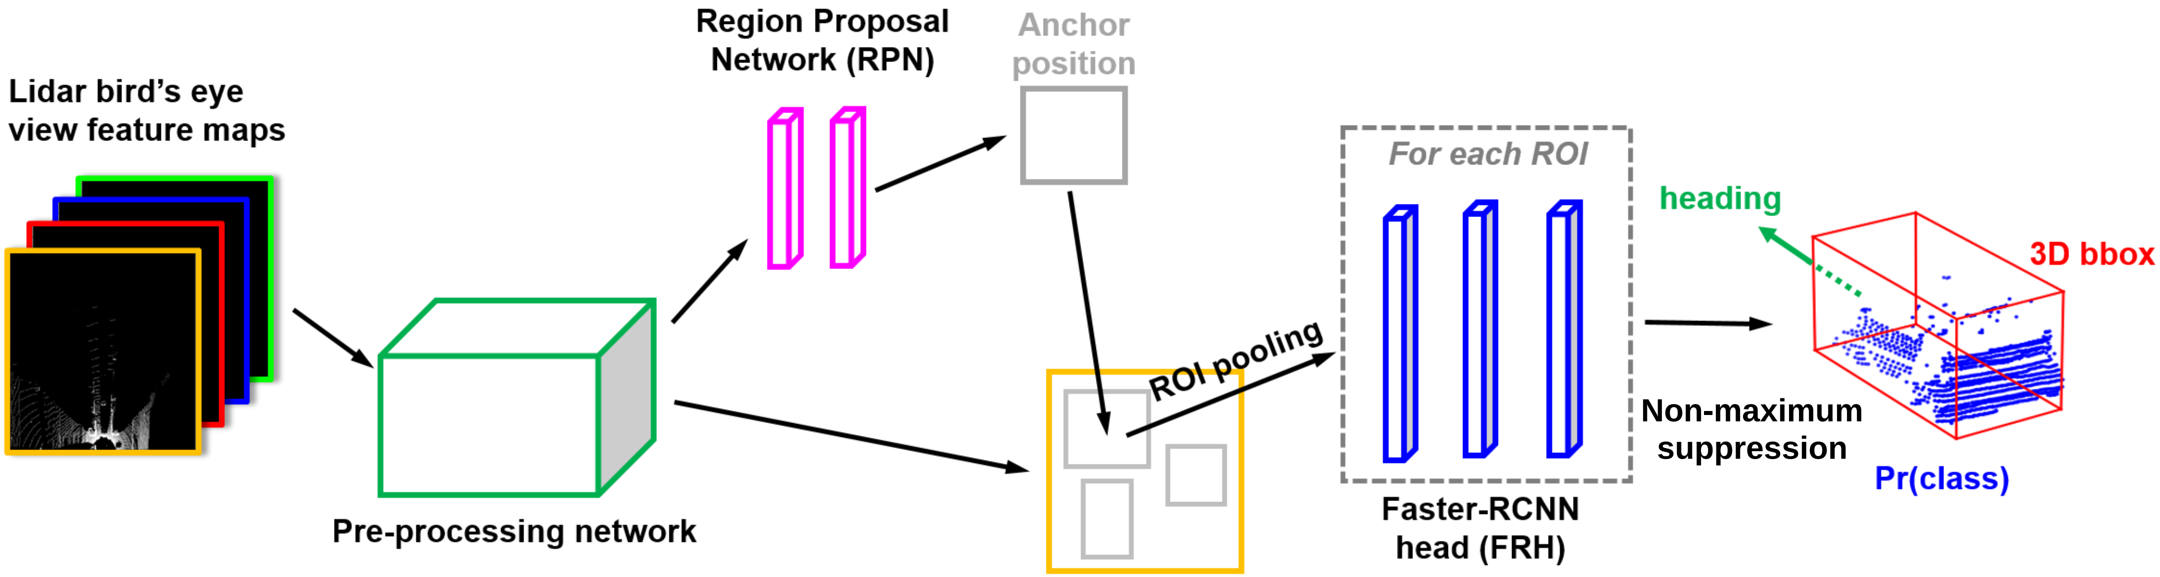

[Ren et al., NeurIPS'15]

2. Uncertainties in object detection networks

State of the art

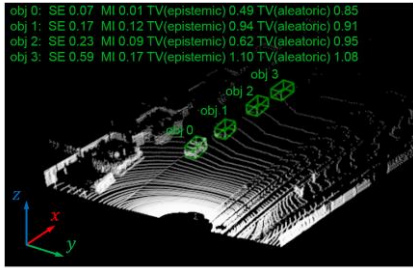

[Feng et al., ITSC'18]
Uncertainties in a LiDAR detector

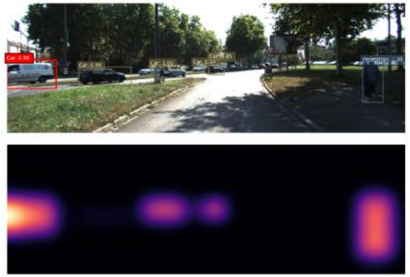

[Truong et al., ITSC'18]
Uncertainties in an image detector

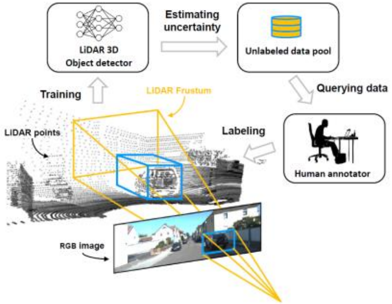

[Feng et al., IV'19b]
Active learning for a probabilistic detector

[Miller et al., ICRA'18]
Detection in open-set conditions.

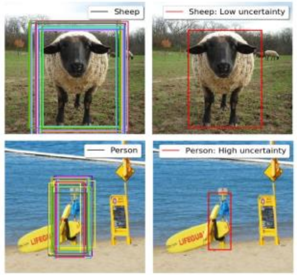

[Miller et al., ICRA'19]
Uncertainty and merging strategy

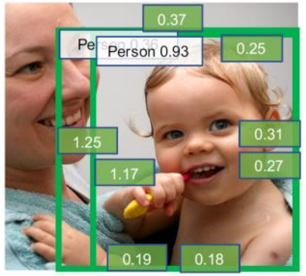

[He et al., CVPR'19]
Localization uncertainty and nms

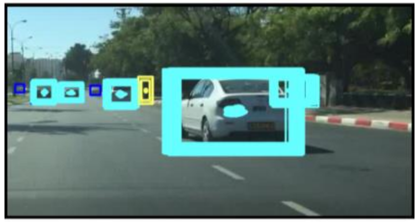

[Harakeh et al., 19]
Localization uncertainty and nms

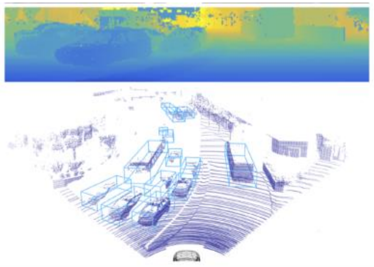

[Meyer et al., CVPR'19]
Localization uncertainty and nms

Outline

1. Motivation
2. Uncertainties in object detection networks
3. **Probabilistic LiDAR object detectors**
4. Challenges

3. Probabilistic LiDAR object detectors

Summary

- Bayesian neural network framework
  - Model-related uncertainties (epistemic)
  - Environmental noises (aleatoric)
- Two & One stage object detector
- Systematic analysis

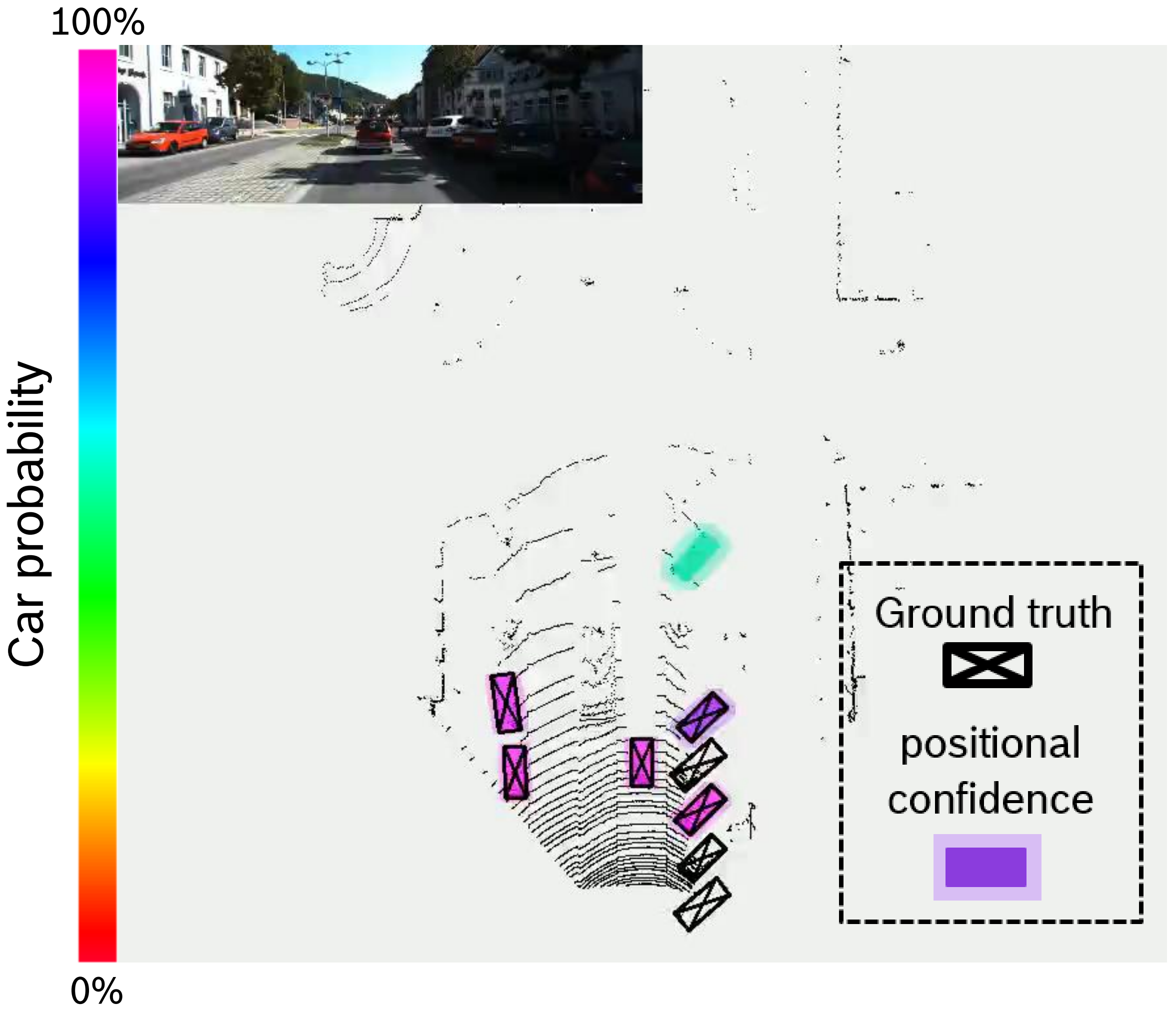

3. Probabilistic LiDAR object detectors

Epistemic and aleatoric uncertainties behave very differently [Feng et al., ITSC'18]

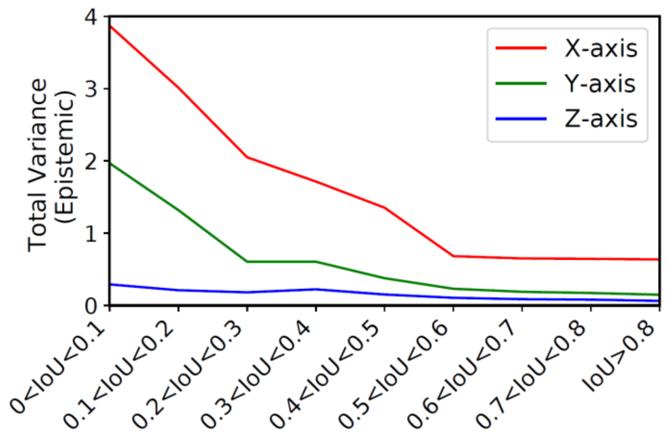

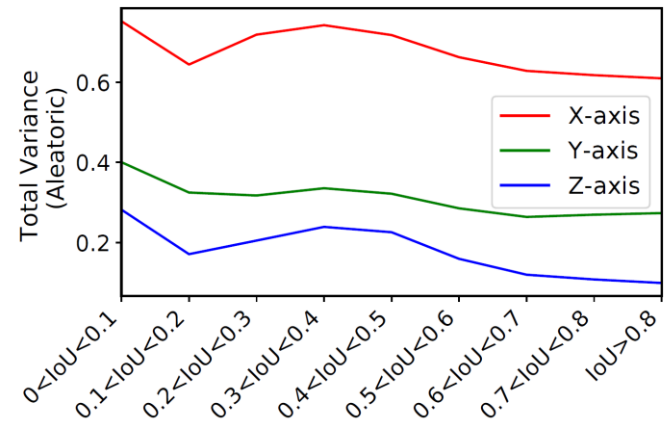

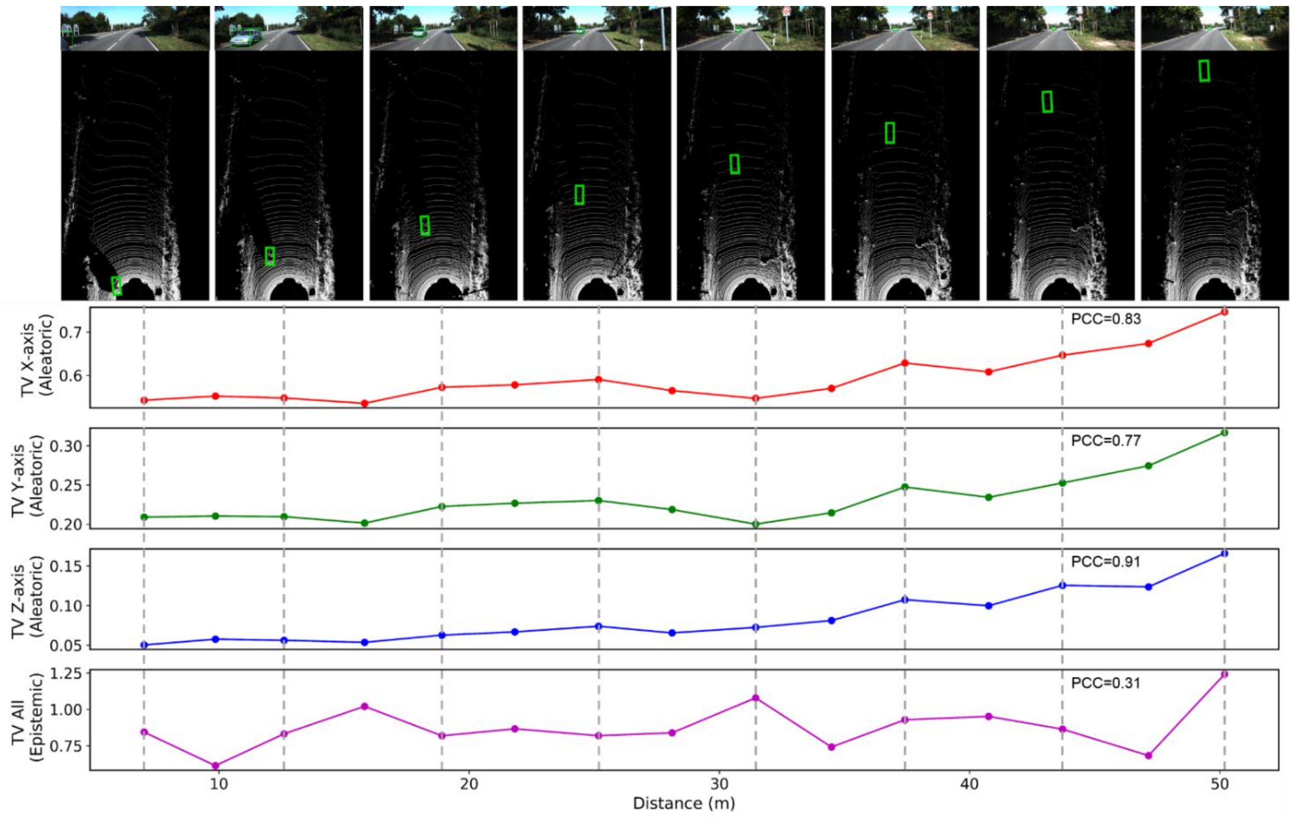

PCC: Pearson Correlation Coefficient

universität ulm BOSCH

3. Probabilistic LiDAR object detectors

Aleatoric uncertainty represents environmental noises [Feng et al., IV'19a]

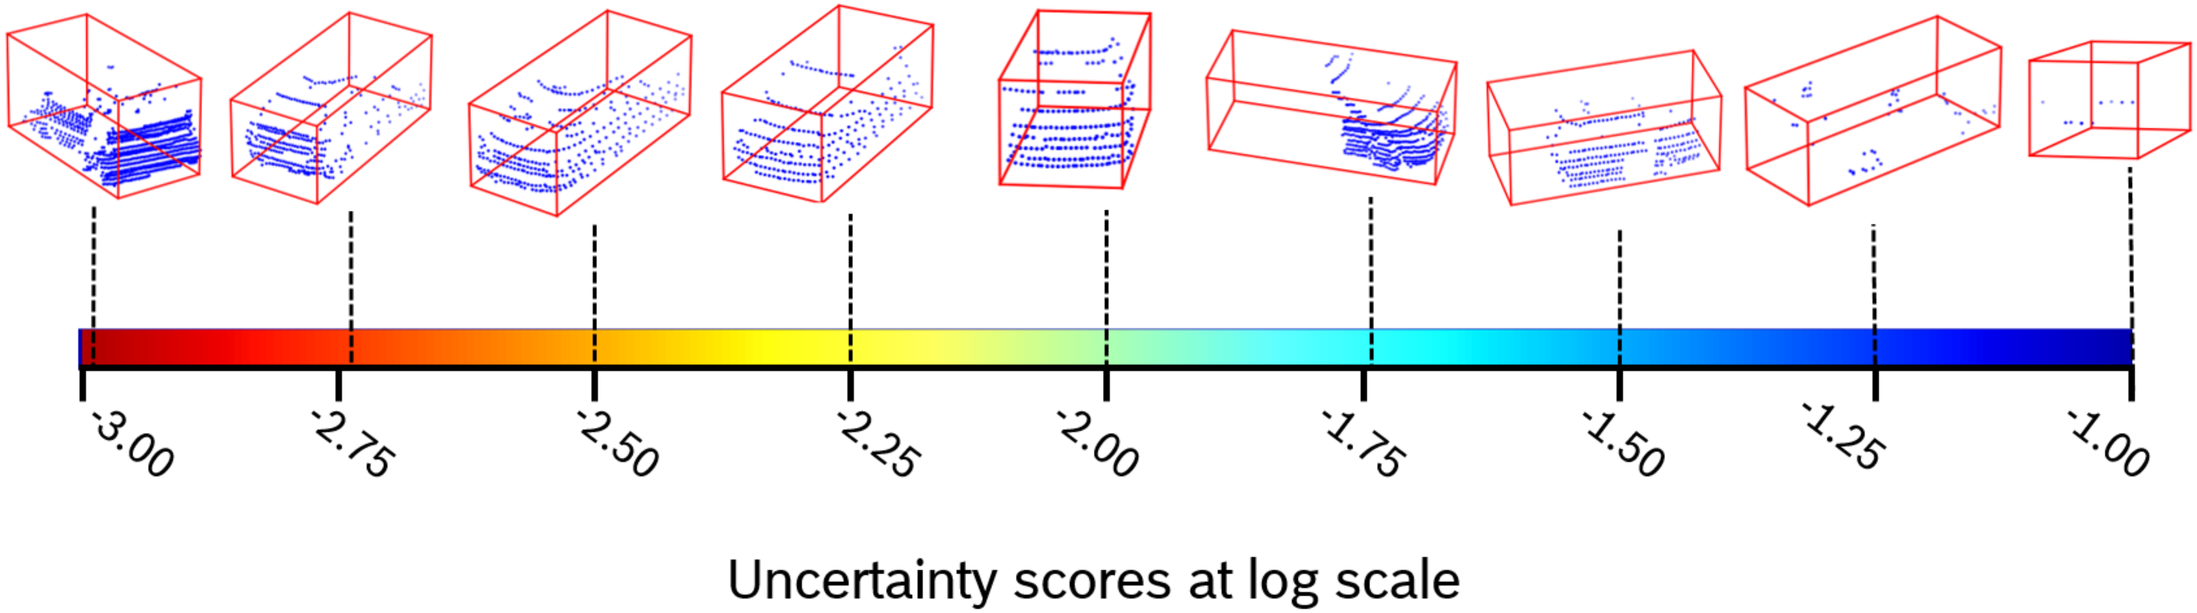

[Feng et al., IV’19a]

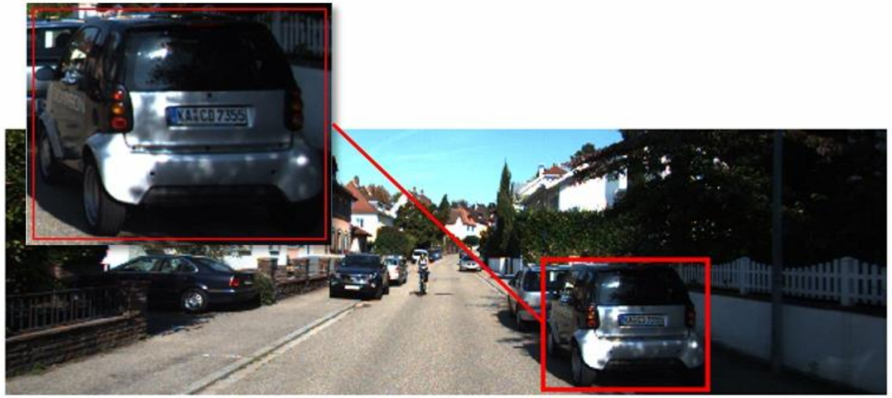

Detection in RGB image

Detection Information

Uncertainty score: **-2.92**
Softmax score: 0.87
Distance: 7.5 m
Occlusion: fully visible
Orientation: 92.1°

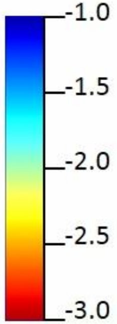

Uncertainty scores at log scale

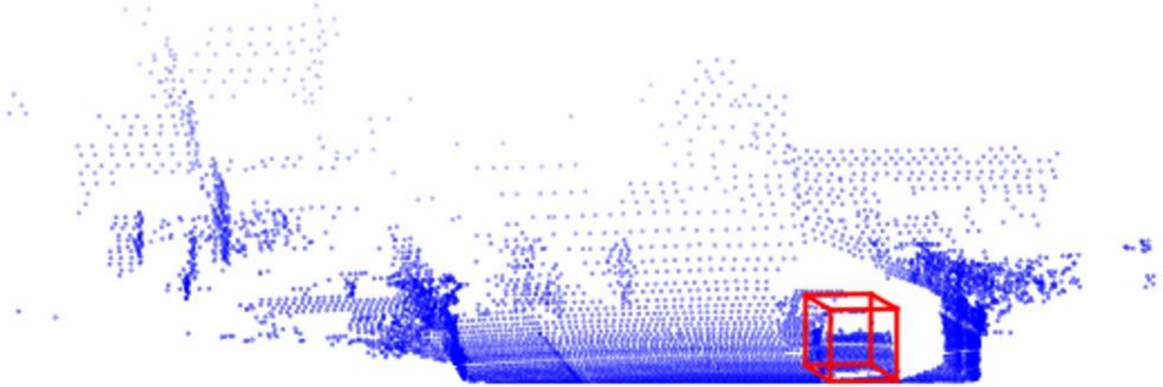

Detection in LiDAR point clouds

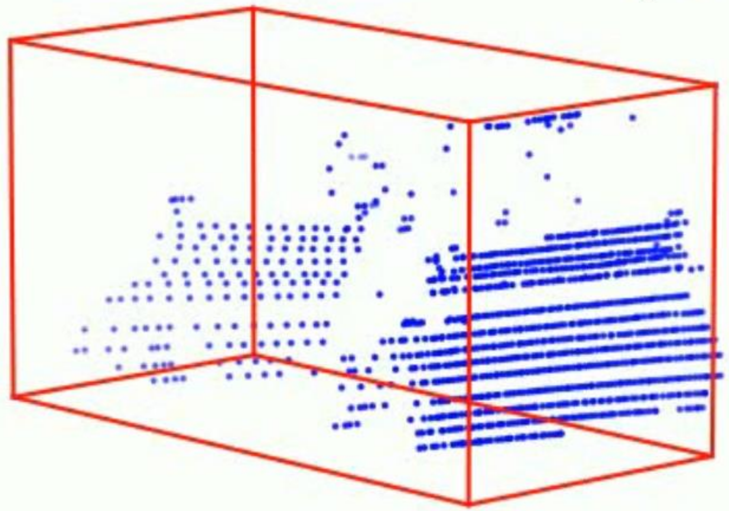

LiDAR points within the bounding box

3. Probabilistic LiDAR object detectors

Using aleatoric uncertainty to improve detection accuracy [Feng et al., IV'19a]

Comparison of 3D Car detection performance on KITTI *val* set [Geiger et al., CVPR'12]

Method	AP_{3D} (%)			AP_{BEV} (%)		
	Easy	Moderate	Hard	Easy	Moderate	Hard
F-PointNet (LiDAR)	69.50	62.30	59.73	-	-	-
PIXOR	-	-	-	86.79	80.75	76.60
VoxelNet	81.97	65.46	62.85	89.60	84.81	78.57
Baseline	71.50	63.71	57.31	86.33	76.44	69.72
Ours	**+7.31**	**+2.18**	**+7.88**	**+0.7**	**+0.71**	**+7.23**

* Baseline is the object detector without any uncertainty estimation.

[Qi et al., CVPR'18] [Zhou et al., CVPR'18] [Yang et al., CVPR'18]

universität uulm BOSCH

3. Probabilistic LiDAR object detectors

Using epistemic uncertainty to improve training efficiency [Feng et al., IV'19b]

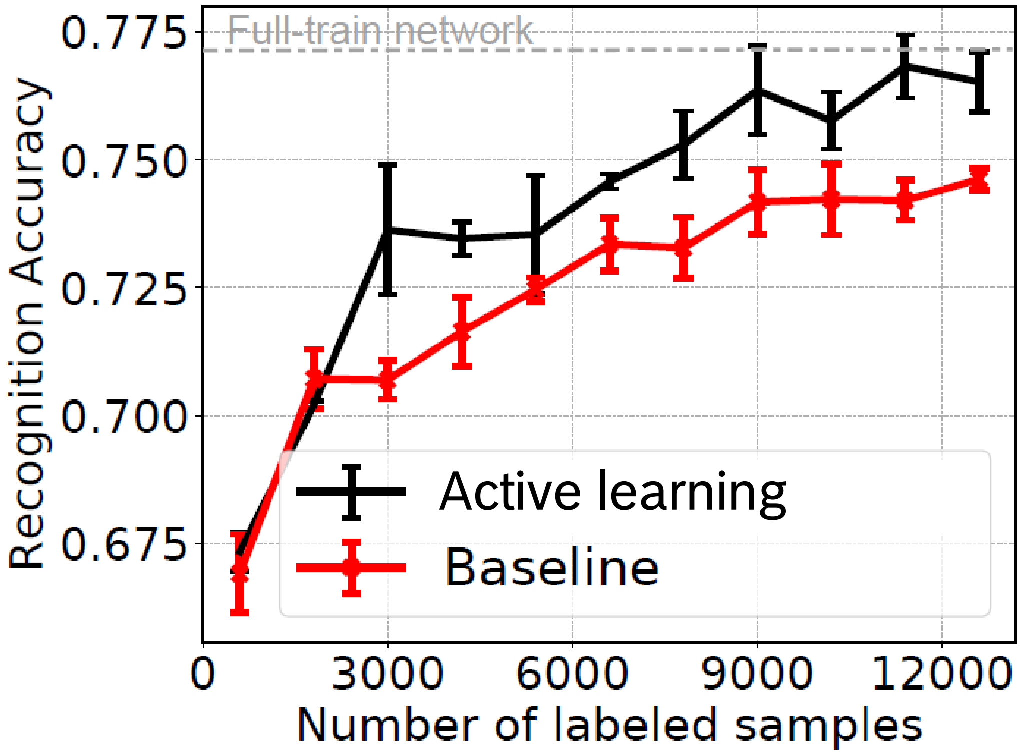

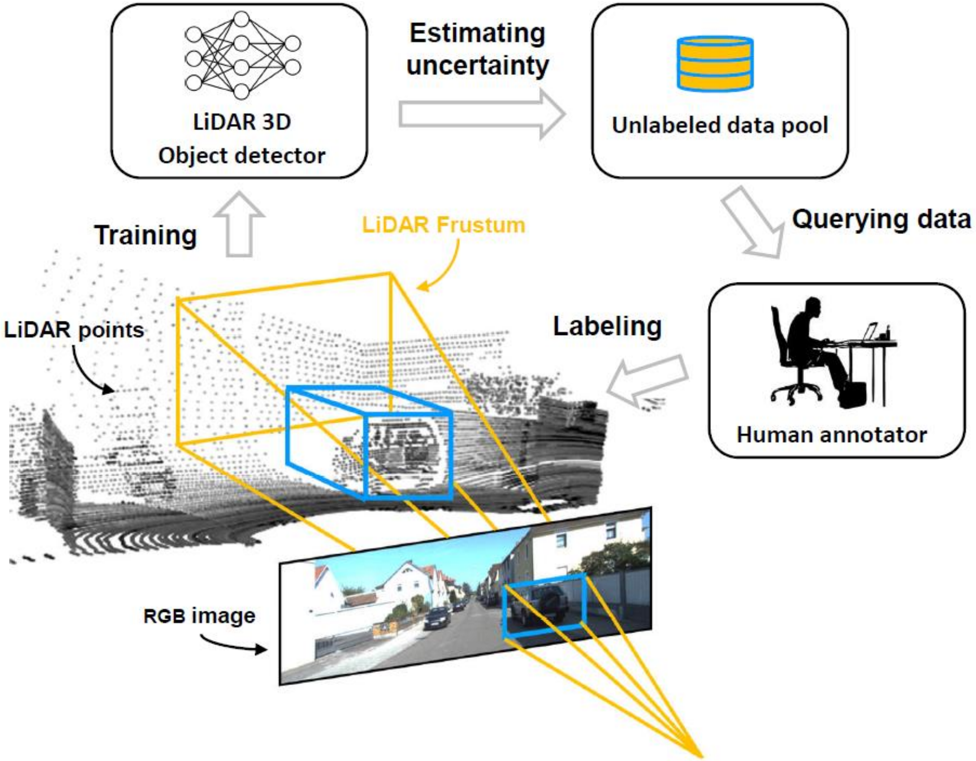

3. Probabilistic LiDAR object detectors

Can we trust uncertainty estimation? [Feng et al., IROS'19]

If a model makes predictions with 0.8 probability score, 80% of those predictions should be correct.

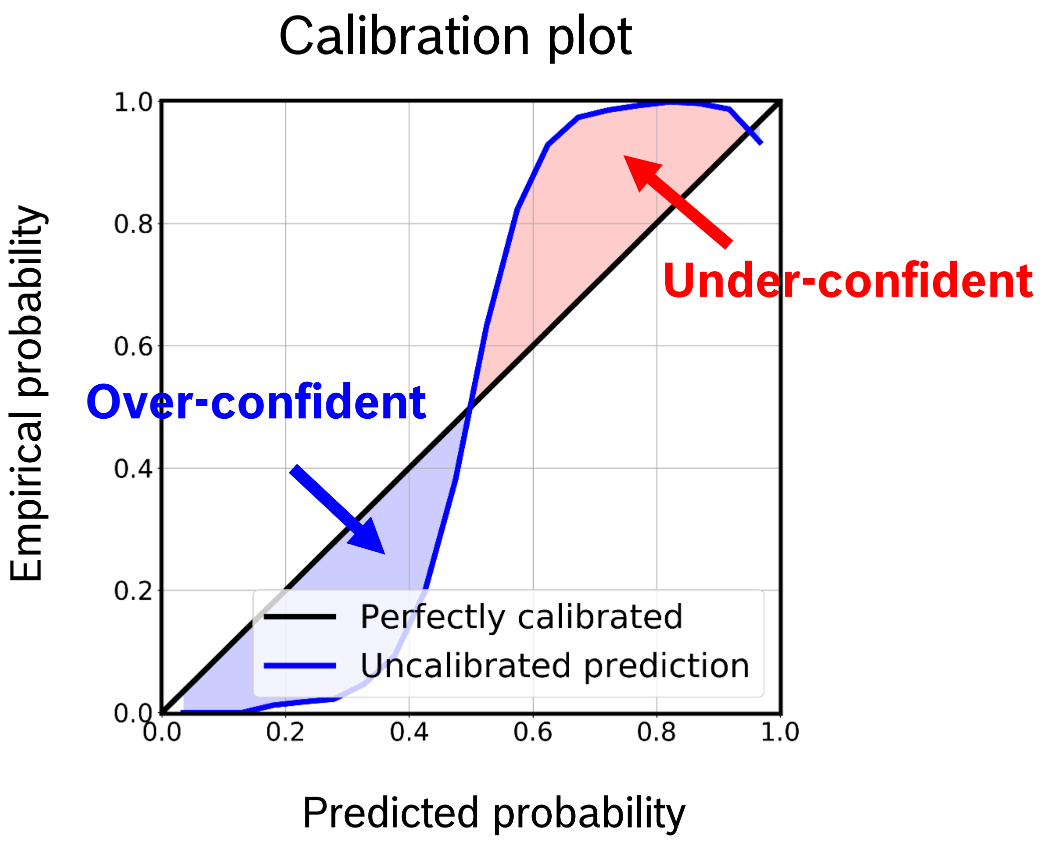

3. Probabilistic LiDAR object detectors

Importance of well-calibrated uncertainty

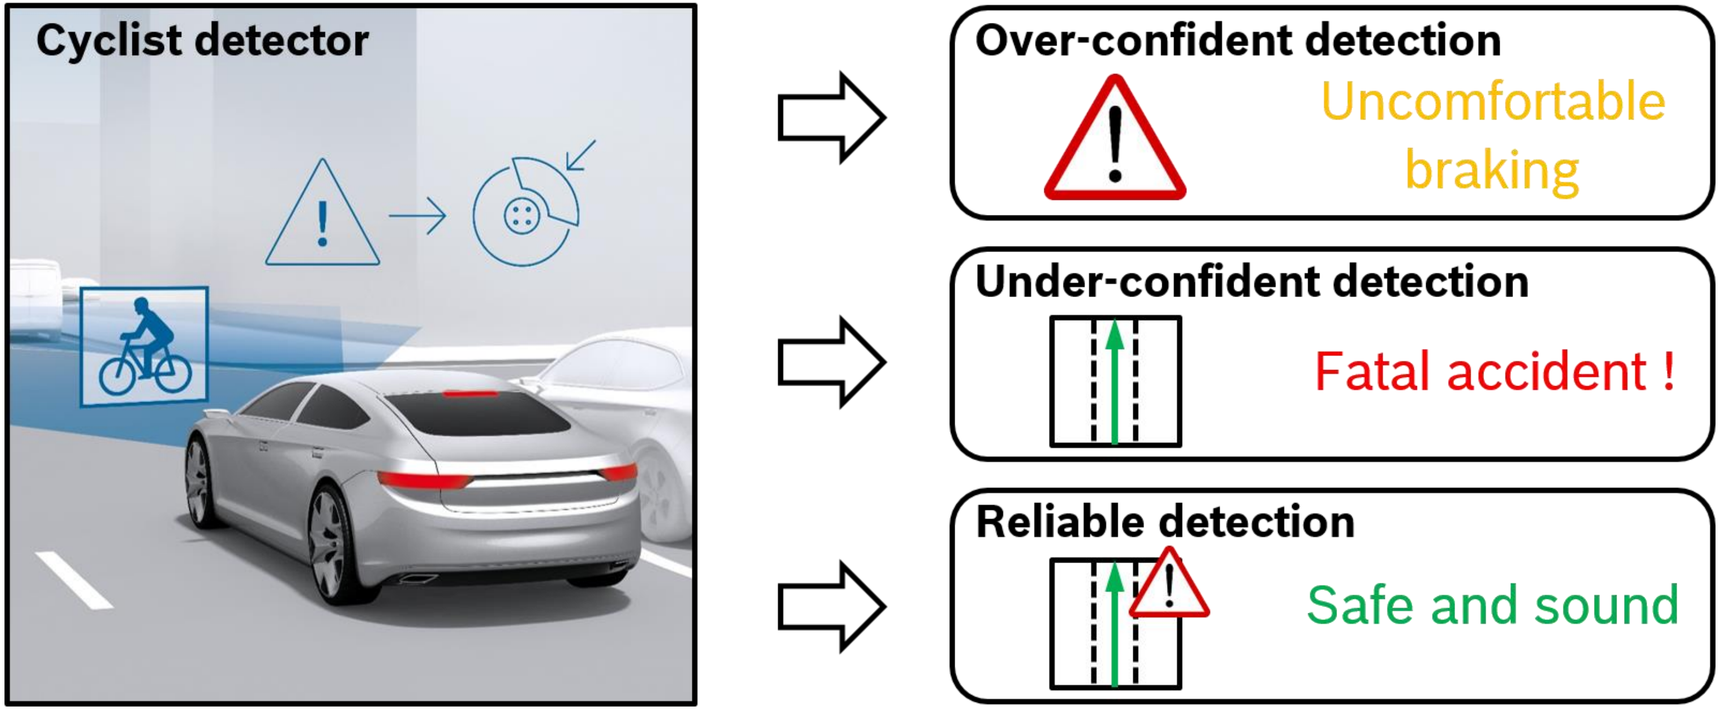

3. Probabilistic LiDAR object detectors

Identifying miscalibrated uncertainties in an one-stage detector [Feng et al., IROS'19]

(a) Classification

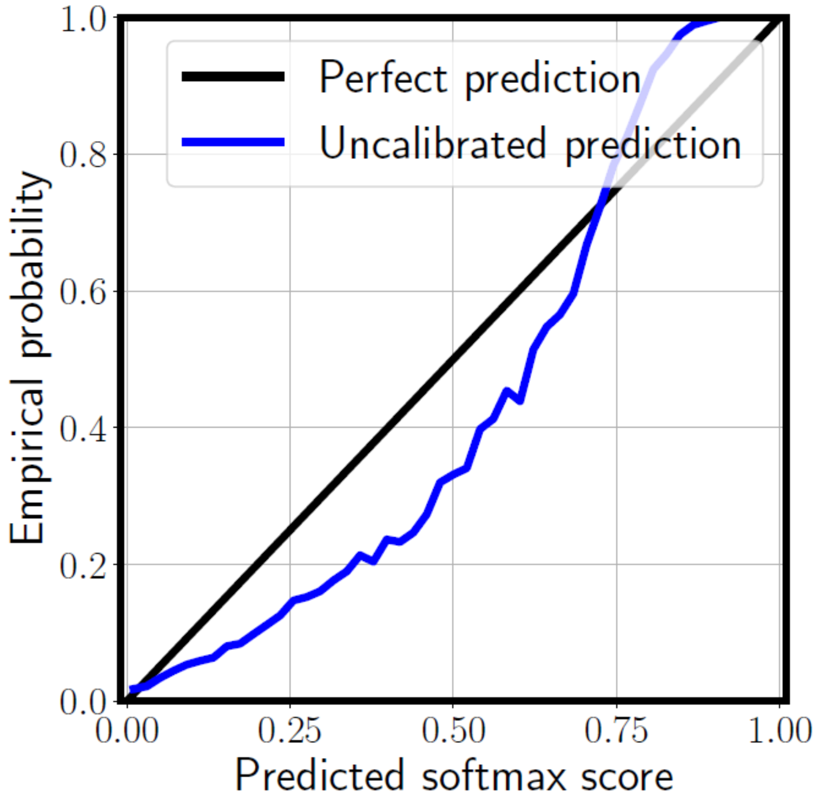

(b) Regression

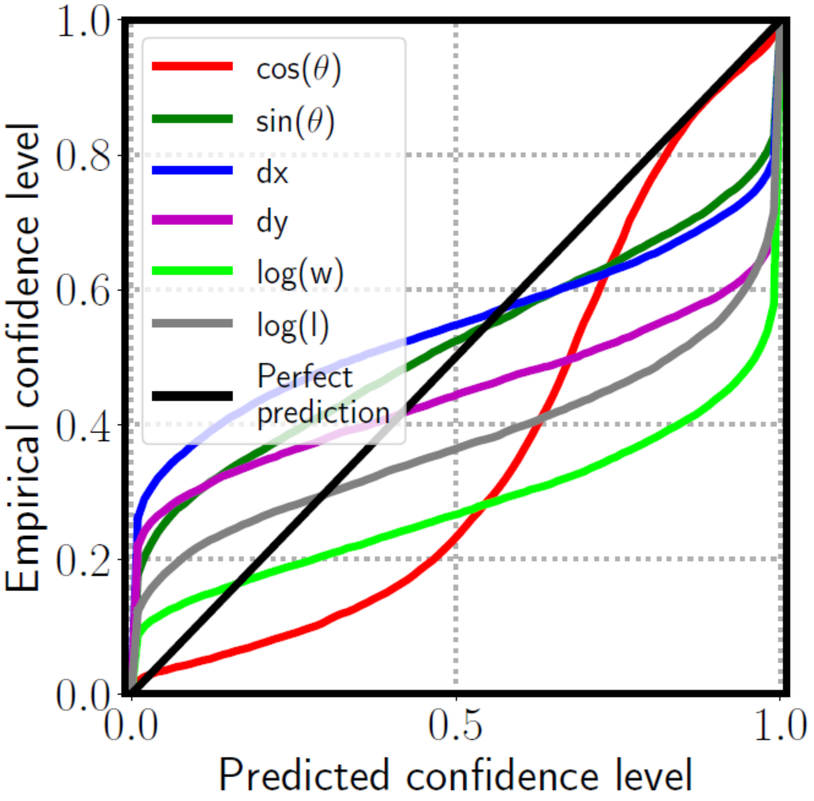

3. Probabilistic LiDAR object detectors

Proposing three uncertainty recalibration methods to largely reduce uncertainty calibration error

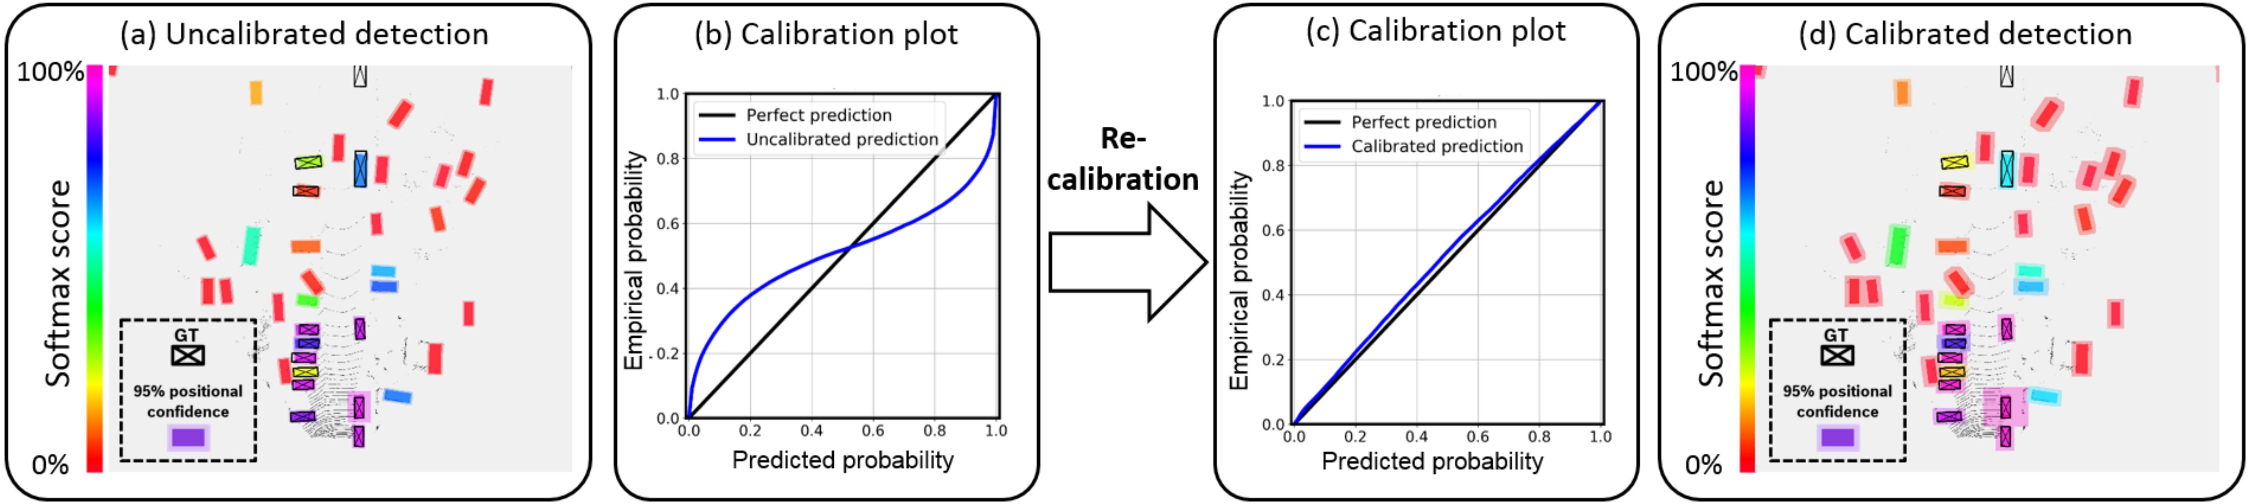

[Feng et al., IROS'19]

3. Probabilistic LiDAR object detectors

Recalibrating uncertainties – classification [Feng et al., IROS'19]

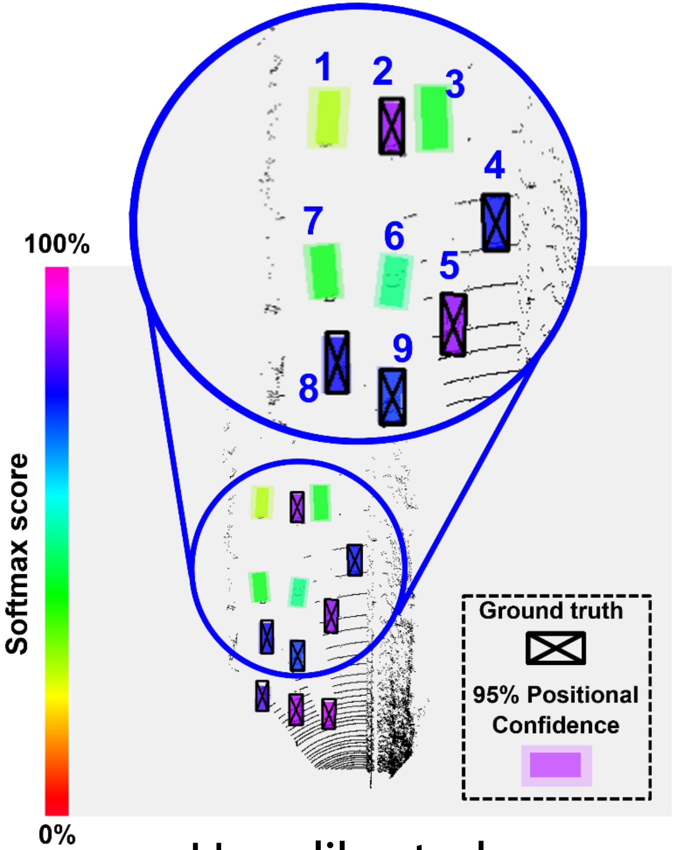

Uncalibrated

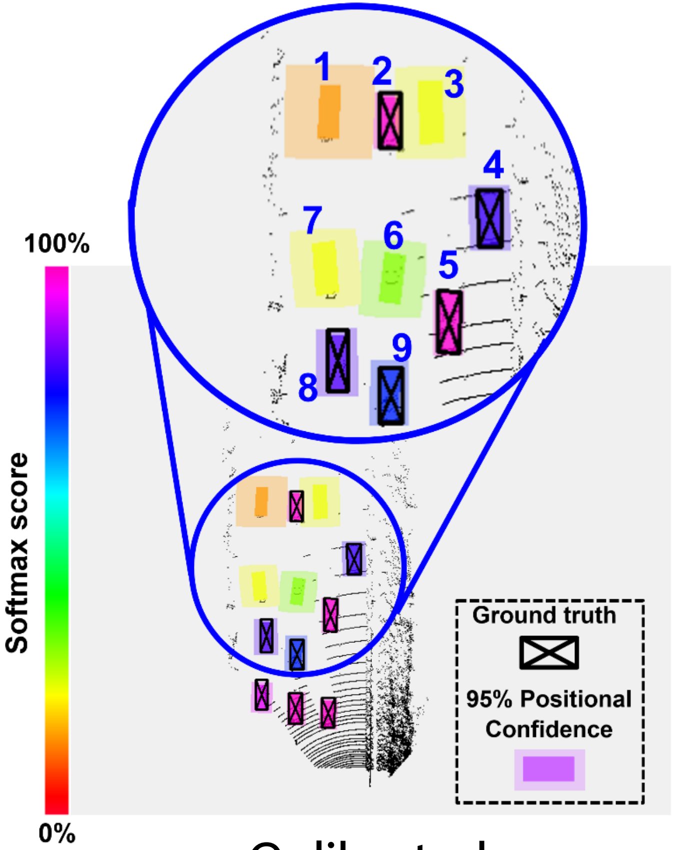

Calibrated

3. Probabilistic LiDAR object detectors

Recalibrating uncertainties – regression (marginal probability) [Feng et al., IROS'19]

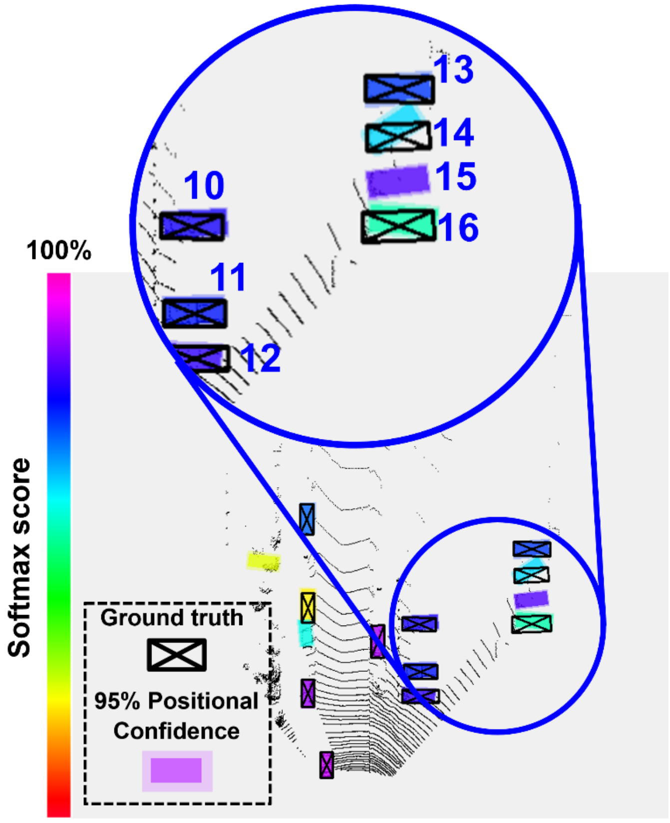

Uncalibrated

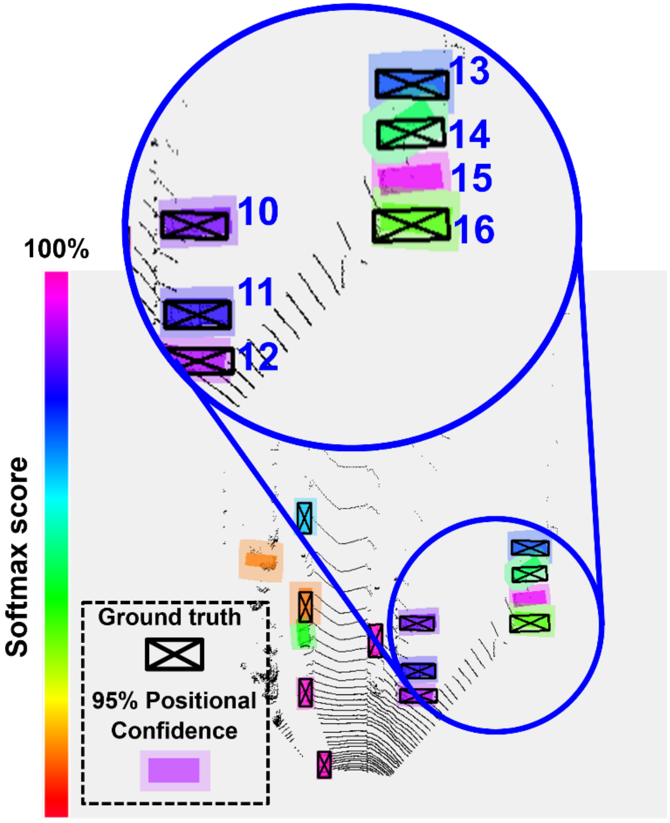

Calibrated

3. Probabilistic LiDAR object detectors

Our LiDAR object detectors model uncertainties:

- **Holistic:** cls+reg; in two/one-stage detectors; epistemic/aleatoric uncertainties
- **Well-calibrated:** after uncertainty recalibration [Feng et al., IROS’19]
- **Explainable:**
  - Reflect environmental noises such as distance & occlusion [Feng et al., ITSC’18 & IV’19a]
  - Reflect model’s accuracy [Feng et al., ITSC’18]
- **Useful**
  - Improve detection performance [Feng et al., IV’19a]
  - Improve training efficiency via active learning [Feng et al., IV’19b]

Problem solved? No!

Front car recognized	Front car localized	Left car recognized	Left car localized
Likely	Uncertain	Very uncertain	Unlikely

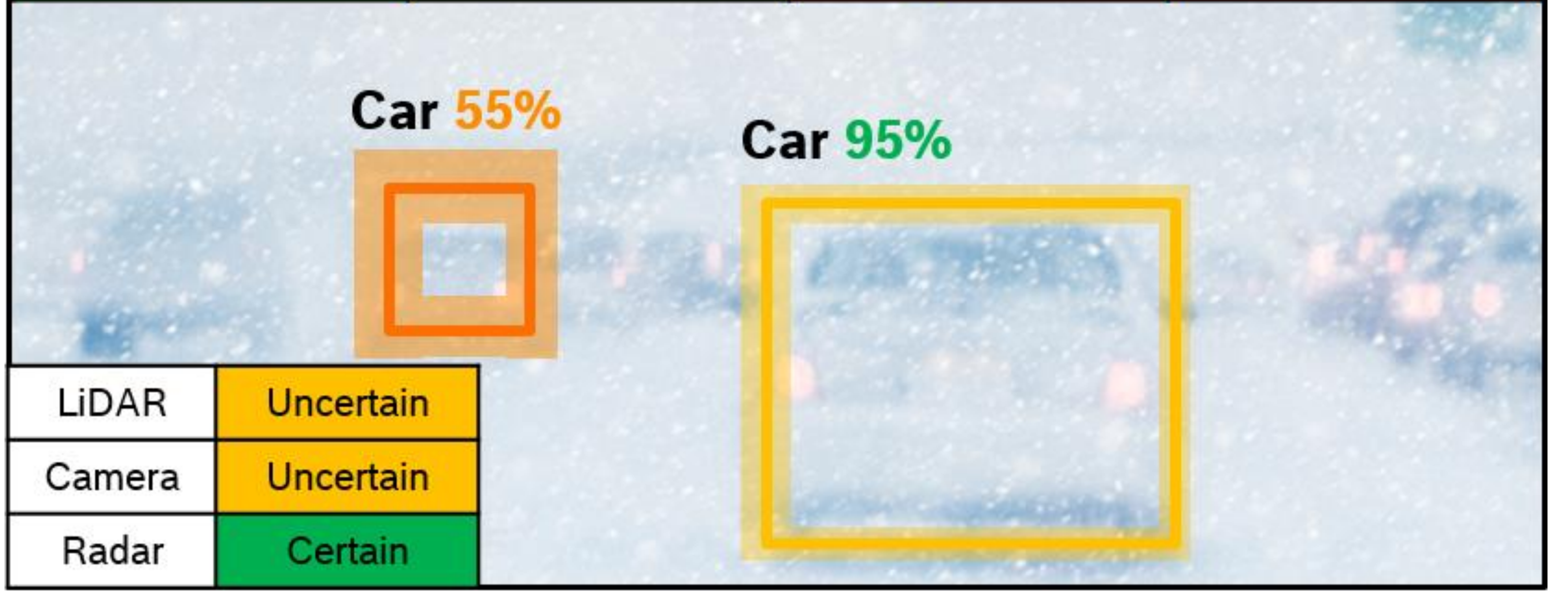

4. Challenges

Can we compare uncertainties in multi-modal perception systems?

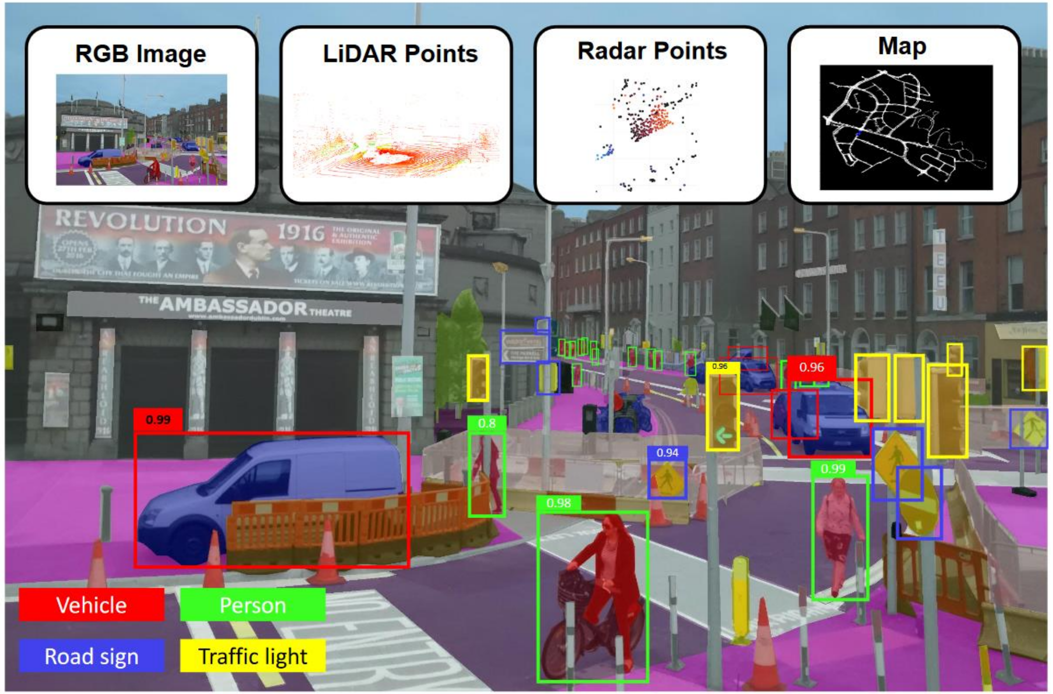

Feng et al., "Deep Multi-modal Object Detection and Semantic Segmentation for Autonomous Driving: Datasets, Methods, and Challenges." *IEEE Transactions on Intelligent Transportation Systems* (2019). Minor revision.

4. Challenges

Are those captured uncertainties useful?

On-board sensors

Object located	Yes (Prob. = 90%)
Object classified	"Pedestrian" (Prob. = 60%)
Camera signal	Uncertain
LiDAR signal	Certain

Probabilistic perception

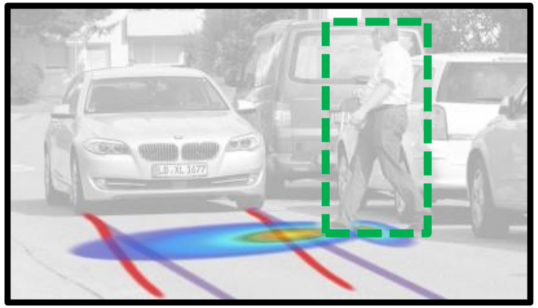

Prediction

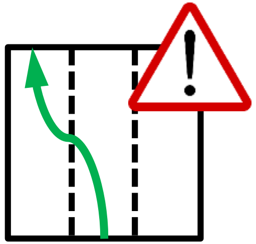

Decision making

- Can uncertainty improve the tracking performance?
- Where can we really see the benefit of uncertainty? (e.g. safety-critical scenarios)

THANK YOU

Perception and Sensors for Autonomous Driving
Bosch research
Di Feng | Lars Rosenbaum | Fabian Timm | Glaudius Gläser

Contact: Di.Feng@de.bosch.com

Collaborators
Christian Haase-Schütz | Heinz Hertlein | Klaus Dietmayer
Werner Wiesbeck | Maki Atsuto | Xiao Wei | Yifan Cao
Fabian Duffhauß

universität ulm BOSCH

References

[Qi et al., CVPR'18] Qi, Charles R., Wei Liu, Chenxia Wu, Hao Su, and Leonidas J. Guibas. "Frustum pointnets for 3d object detection from rgb-d data." In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 918-927. 2018.

[Cosmides et al., Cognition'96] Cosmides and J. Tooby, "Are humans good intuitive statisticians after all? Rethinking some conclusions from the literature on judgment under uncertainty," *Cognition*, vol. 58, no. 1, pp. 1–73, 1996.

[Feng et al., ITSC'18] Feng, Di, Lars Rosenbaum, and Klaus Dietmayer. "Towards safe autonomous driving: Capture uncertainty in the deep neural network for lidar 3d vehicle detection." In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pp. 3266-3273. IEEE, 2018.

[Feng et al., IV'19a] Feng, Di, Lars Rosenbaum, Fabian Timm, and Klaus Dietmayer. "Leveraging heteroscedastic aleatoric uncertainties for robust real-time lidar 3d object detection". In *30th IEEE Intelligent Vehicles Symposium*, 2019.

[Feng et al., IV'19b] Feng, Di, Xiao Wei, Lars Rosenbaum, Atsuto Maki, and Klaus Dietmayer. "Deep active learning for efficient training of a lidar 3d object detector." In *30th IEEE Intelligent Vehicles Symposium*, 2019.

[Feng et al., IROS'19] Feng, Di, Lars Rosenbaum, Claudius Gläser, Fabian Timm, and Klaus Dietmayer. "Can we trust you? On calibration of a probabilistic object detector for autonomous driving." In *IEEE/RSJ International Conference on Intelligent Robots and Systems Workshop,* 2019.

[Kendall et al., NeuralPS'17] Kendall, Alex, and Yarin Gal. "What uncertainties do we need in bayesian deep learning for computer vision?" In *Advances in neural information processing systems*, pp. 5574-5584. 2017.

[MacKey., Neural'92] MacKay, David JC. "Bayesian interpolation." *Neural computation* 4.3 (1992): 415-447.

References

[Ren et al., NeurIPS'15] Ren, Shaoqing, Kaiming He, Ross Girshick, and Jian Sun. "Faster R-CNN: Towards real-time object detection with region proposal networks." In *Advances in neural information processing systems*, pp. 91-99. 2015.

[Truong et al., ITSC'18] Le, Michael Truong, Frederik Diehl, Thomas Brunner, and Alois Knol. "Uncertainty Estimation for Deep Neural Object Detectors in Safety-Critical Applications." In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pp. 3873-3878. IEEE, 2018.

[Dimity et al., ICRA'18] Miller, Dimity, Lachlan Nicholson, Feras Dayoub, and Niko Sünderhauf. "Dropout sampling for robust object detection in open-set conditions." In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1-7. IEEE, 2018.

[Dimity et al., ICRA'19] Miller, Dimity, Feras Dayoub, Michael Milford, and Niko Sünderhauf. "Evaluating merging strategies for sampling-based uncertainty techniques in object detection." In *2019 International Conference on Robotics and Automation (ICRA)*, pp. 2348-2354. IEEE, 2019.

[He et al., CVPR'19] He, Yihui, Chenchen Zhu, Jianren Wang, Marios Savvides, Xiangyu Zhang. "Bounding box regression with uncertainty for accurate object detection". In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019.

[Harakeh et al., 19] Harakeh, Ali, Michael Smart, and Steven L. Waslander. "BayesOD: A Bayesian Approach for Uncertainty Estimation in Deep Object Detectors." *arXiv preprint arXiv:1903.03838*, 2019.

[Meyer, et al., CVPR'19] Meyer, Gregory P., Ankit Laddha, Eric Kee, Carlos Vallespi-Gonzalez, and Carl K. Wellington. "LaserNet: An Efficient Probabilistic 3D Object Detector for Autonomous Driving." in *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.

[Yang et al., CVPR'18] Yang, Bin, Wenjie Luo, and Raquel Urtasun. "Pixor: real-time 3D object detection from point clouds." In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 7652-7660. 2018.

References

[Zhou et al., CVPR'18] Zhou, Yin, and Oncel Tuzel. "Voxelnet: End-to-end learning for point cloud based 3d object detection." In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4490-4499. 2018.

[Geiger et al., CVPR'12] Geiger, Andreas, Philip Lenz, and Raquel Urtasun. "Are we ready for autonomous driving? the kitti vision benchmark suite." In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3354-3361. IEEE, 2012.