# Uncertainty interpretations for the robustness of object detection in self-driving vehicles


**Filipa M. Ramos Ferreira**

filiparamos@fe.up.pt

–

Supervisor(s): Rosaldo J. F. Rossetti[1]

1. Artificial Intelligence and Computer Science Laboratory / Faculty of Engineering of the University of Porto


**MIT Portugal**
2022 Annual Conference

## Introduction

**Autonomous driving** has been established as one of the major desirable technological breakthroughs with both industry and academia racing towards its realisation. In this sense, **deep learning** has been a key factor as neural networks demonstrate high accuracy levels in a variety of autonomy-critical tasks. **Perception** tasks, central to any autonomous-vision based system, have directly benefited largely from such innovations. Nevertheless, the accuracy rises being reported in literature are not necessarily followed by a rise in the confidence in these systems. For this reason, it becomes paramount to study augmented **interpretability** of neural networks, especially through the quantification of uncertainty in perception models that could become a decisive factor in autonomous pipelines. Even so, research efforts on object detection do not generally provide uncertainty estimation for predictions, focusing only on improving dataset-dependent metrics. What is more, when available, this uncertainty can not be used in downstream tasks as the incorporation of these estimates has not been too explored in literature (Bhatt et al. 2020). On this note, this research project aims to conceptualise uncertainty interpretations and quantification methodologies for deep learning based object detection for autonomous driving applications. Furthermore, it is intended to provide a case study of the use of perception uncertainty estimates in **downstream layers** of the AV pipeline, such as decision-making, in order to showcase whether these estimates enable error recovery (e.g. misclassifications, missing objects) and thus enabling further fault-tolerant, confident behaviour.

## Literature Review

Uncertainty quantification techniques may be sub-divided into Bayesian and non-Bayesian methods.

Bayesian Methods

- Monte Carlo Dropout (Gal et al. 2015)
- Stochastic Batch Normalisation (Teye et al. 2018)
- Bayes by Backprop (Blundell et al. 2015)
- Test-time Data Augmentation (Ayhan et al. 2018)

Non-Bayesian Methods

- Deep Ensembles (Lakshminarayanan et al. 2017)
- Split-head Networks (Tran et al. 2020)
- Calibration (Guo et al. 2017)
- Direct modelling (Gast et al. 2018)

### Object detection

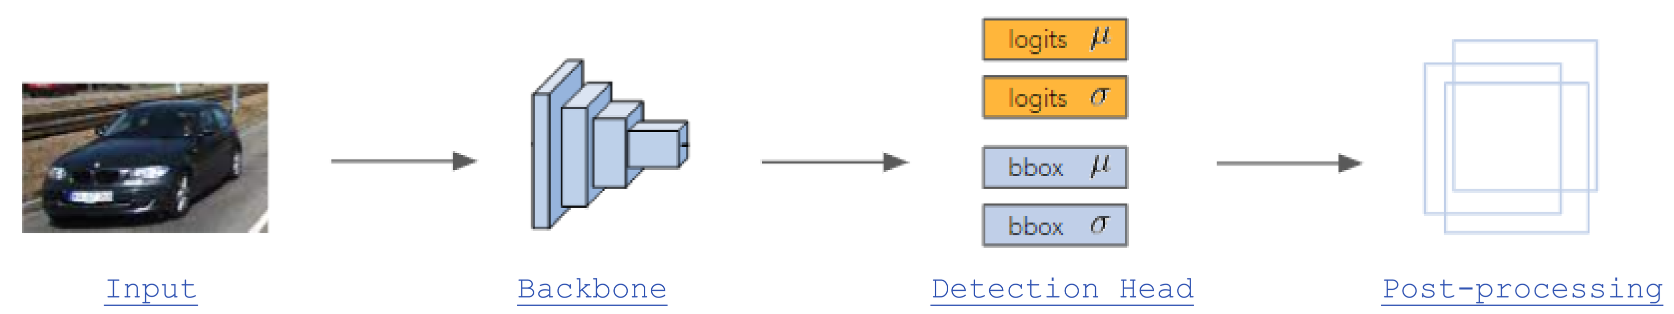


Backbone:
✓ MC-Dropout (Gal et al. 2015)
✓ Ensembles (Lakshminarayanan et al. 2017)

Detection Head:
✓ MC-Dropout (Gal et al. 2015)
✓ Ensembles (Lakshminarayanan et al. 2017
✓ Direct Modelling (Gast et al. 2018)

Post-processing:
✓ Adapted-NMS (Meyer et al. 2019)
✓ Sample Statistics (Miller et al. 2018)
✓ Bayesian Inference (Harakeh et al. 2019)

## Goals

**1) Relationship between uncertainty interpretations and model expressiveness**

✓ Analysis of a model through uncertainty
✓ Analysis of the information carried by different uncertainty interpretations

**2) Study of new interpretations of uncertainty for object detection tasks**

✓ Conceptualisation of an existence uncertainty that describes the confidence of the network on whether or not it missed the detection of some object(s)
✓ Conceptualisation of other uncertainty interpretations and contexts

**3) Uncertainty applications in downstream layers**

✓ Conceptualisation of the incorporation of uncertainty estimates in a POMDP-based decision layer
✓ Analysis of recovery from perception errors

## Work Plan

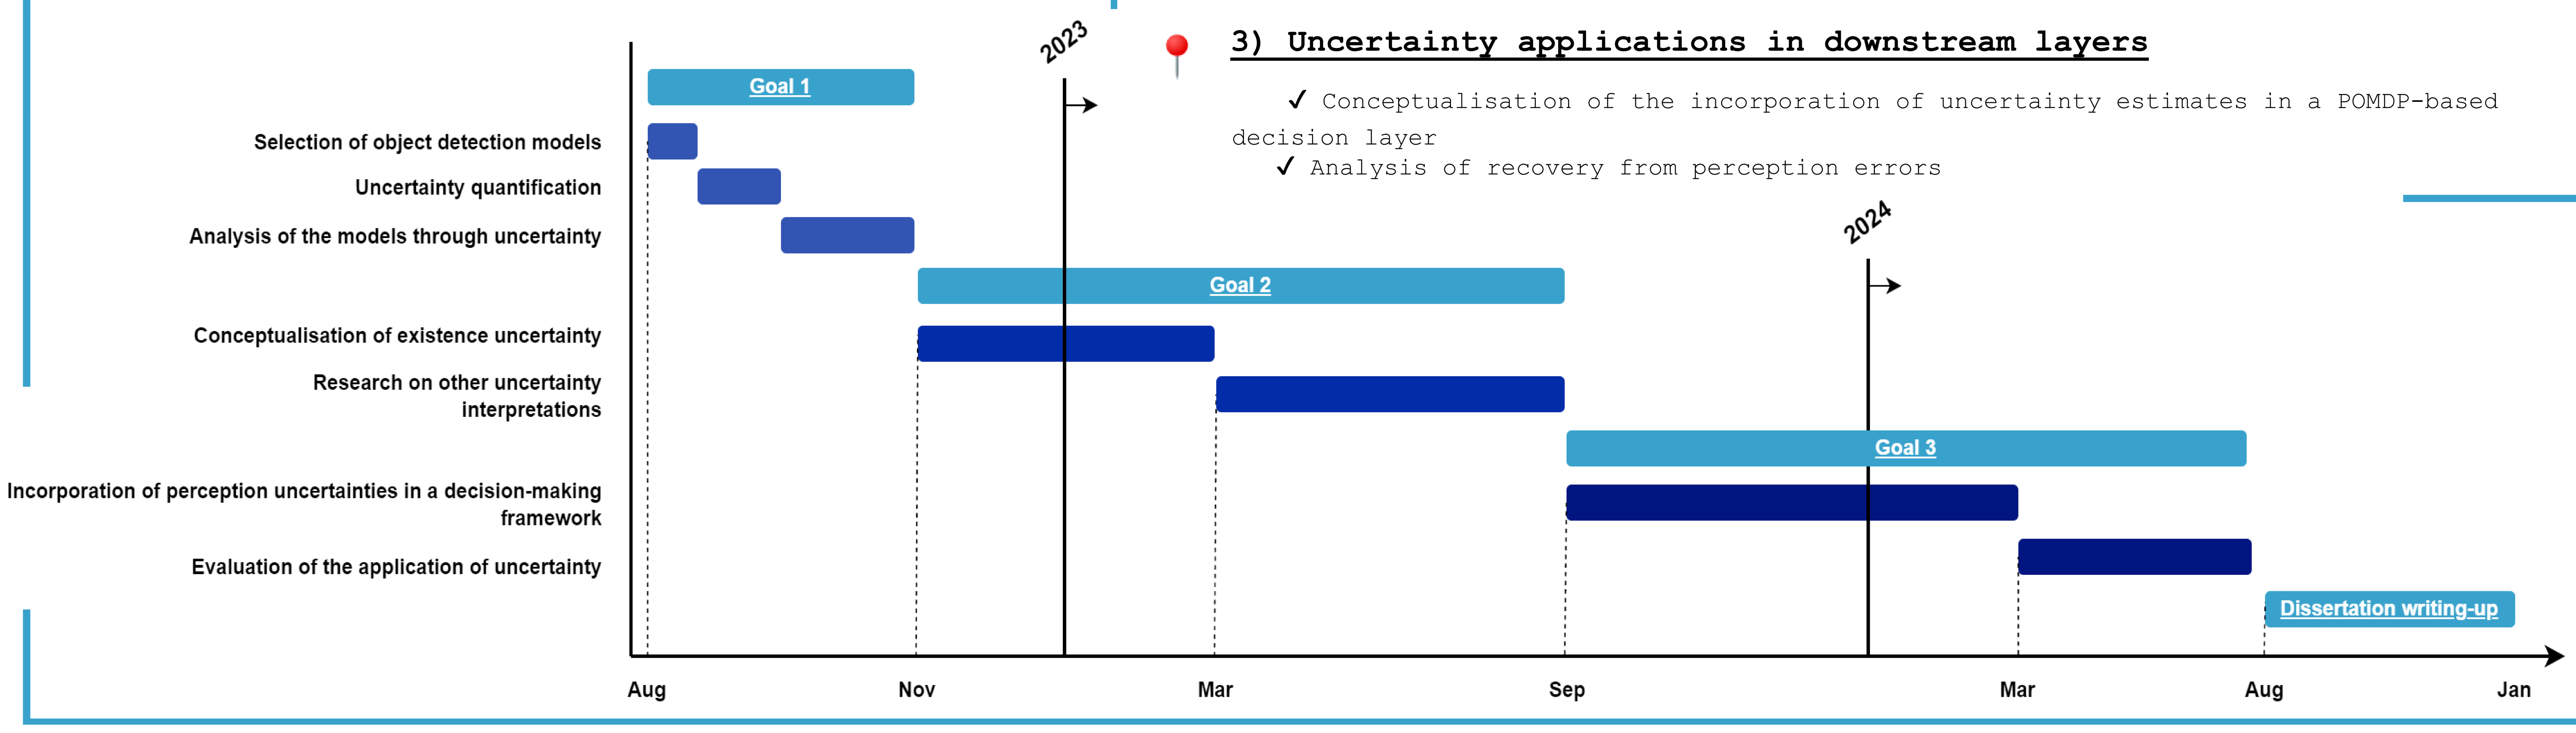


## Conclusions

This project challenges the currently limited interpretability of self-driving vehicles through the development of **uncertainty interpretations** and the demonstration of its application to downstream layers. We believe this project has the potential to enable the increased safety and sustainability of such vehicles, by demonstrating confidence on **fault-tolerant behaviour**, reducing accident events and preventing deaths.


(Tran et al. 2020) L. Tran et al., "Hydra: Preserving Ensemble Diversity for Model Distillation," arXiv, Jan. 2020.
(Gast et al. 2018) J. Gast and S. Roth, "Lightweight probabilistic deep networks," Jun. 2018, pp. 3369-3378, doi: 10.1109/CVPR.2018.00355.
(Blundell et al. 2015) C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra, "Weight Uncertainty in Neural Network," PMLR, Jun. 2015.

(Gal et al. 2015) Y. Gal and Z. Ghahramani, "Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning," *PMLR*, Jun. 2016.
(Guo et al. 2017) C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, "On Calibration of Modern Neural Networks," arXiv, 2017, doi: 10.48550/arxiv.1706.04599.
(Harakeh et al. 2019) A. Harakeh, M. Smart, and S. L. Waslander, "BayesOD: A Bayesian Approach for Uncertainty Estimation in Deep Object Detectors," *arXiv*, Mar. 2019.

(Ayhan et al. 2018) M. Ayhan and P. Berens, "Test-time Data Augmentation for Estimation of Heteroscedastic Aleatoric Uncertainty in Deep Neural Networks," MIDL, 2018.
(Teye et al. 2018) M. Teye, H. Azizpour, and K. Smith, "Bayesian Uncertainty Estimation for Batch Normalized Deep Networks," Proceedings of the 35th International Conference on Machine Learning, vol. 80, pp. 4907--4916, Jul. 2018.

(Miller et al. 2018) D. Miller, L. Nicholson, F. Dayoub, and N. Sunderhauf, "Dropout Sampling for Robust Object Detection in Open-Set Conditions," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, May 2018, pp. 3243-3249, doi: 10.1109/ICRA.2018.8460700.

(Lakshminarayanan et al. 2017) B. Lakshminarayanan, A. Pritzel, and C. Blundell, "Simple and scalable predictive uncertainty estimation using deep ensembles," *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pp. 6405-6416, Dec. 2017.

(Bhatt et al. 2020) Dhaivat Bhatt, Dishank Bansal, Gunshi Gupta, Hanju Lee, Krishna Murthy Jatavallabhula, and Liam Paull. Probabilistic object detection: Strengths, weaknesses, and opportunities. Workshop on AI for Autonomous Driving at the International Conference on Machine Learning, 2020.

(Meyer et al. 2019) G. P. Meyer, A. Laddha, E. Kee, C. Vallespi-Gonzalez, and C. K. Wellington, "Lasernet: an efficient probabilistic 3D object detector for autonomous driving," in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2019, pp. 12669-12678, doi: 10.1109/CVPR.2019.01296.



Funded by:

FCT Fundação para a Ciência e a Tecnologia
MIT Portugal

under the Doctoral Grant CRM:0017360 | MIT Portugal Program Sustainable Cities