ASSURING AUTONOMY
INTERNATIONAL PROGRAMME

DEMONSTRATOR PROJECT

Final report

SAX
Sense-Assess-Explain

MARCH 2022

Sense-Assess-eXplain (SAX): Building Trust in Autonomous Vehicles in Challenging Real-World Driving Scenarios

Assuring Autonomy International Programme

Technical Report

Matthew Gadd[1], Daniele De Martini[1], Daniel Omeiza[2], Paul Newman[1], and Lars Kunze[1]

***Abstract*—The Oxford Robotics Institute (ORI) has completed the SAX demonstrator project with 141 hours and 3700 kilometres of driving data across urban, semi-urban, motorway, rural, and off-road driving scenarios, resulting in over 200 terabytes of sensor logs with different sensors and tens of thousands of labels for various autonomy-critical driving tasks. In the process, we have developed our research portfolio in robust sensing and scene understanding, including motion estimation, localisation, object detection, and explanation generation. This has resulted in about 30 academic publications. The dataset is in the process of finalisation and peer-review and should be released to the public later in 2022. In this technical report, we summarise our findings and describe our project's evolution in terms of our initial proposal and the broader programme.**

I. Introduction

This technical report describes how we have demonstrated world-leading research in mobile autonomy (including perception, localisation, and mapping) in challenging on-road and off-road driving scenarios addressing fundamental technical issues to overcome critical barriers to assurance and regulation for large-scale deployments of autonomous systems. This document is meant to accompany [1], where we gave an overview of the whole project aim and scope.

This report proceeds as follows. We describe our output and progress over the course of the project. We do this explicitly against the initial objectives and planned work packages and with reference to the Assuring Autonomy International Programme's body of knowledge, which we originally proposed contributing to in meeting the call criteria. Detail on a publication-level is front-heavy, and thereafter we provide rich links between sections, avoiding repetition.

II. Research Objectives

Our proposal was built around the following research objectives:

O1: To robustly and reliably sense and interpret the environment in severe and changing weather conditions, overcoming the limitations of classic sensing modalities.

O2: To continuously assess and optimise the performance of perception as well as navigation methods.

O3: To demonstrate a system capable of explaining in non-engineering and human terms what a robot/vehicle has seen and how it has influenced its decision making.

In the following sections we demonstrate how we have met these objectives.

[1] Oxford Robotics Institute, Dept. Engineering Science, University of Oxford, UK

[2] Dept. of Computer Science, University of Oxford, UK

III. Work Packages

Our work was arranged in five major work packages. WP1 meets **O1**, WP2 meets **O2**, and WP3 meets **O3**.

A. WP1: Alternative Sensing

Harsh weather and lighting conditions in particular pose non-trivial challenges to autonomous vehicle development, above all with the usage of traditional sensing systems, such as cameras and lidars. In this area, we investigated the usage of uncommon sensing modalities and configurations (such as scanning radars and audio) as well as external weather and map services – both for the vehicle's motion estimation and surface classification, mostly leveraging deep learning approaches.

1) WP1.A Radar-based Motion Estimation and Localisation: Here, we showed how Frequency-Modulated Continuous-Wave (FMCW) scanning radar can be a reliable sensor for motion estimation and localisation.

In the area of **motion estimation**, in [2] we demonstrated that readily-available odometry pipelines can be enhanced through introspection techniques to make them more robust in challenging, outdoor scenarios (see Sec. III-B1). In [3] we proposed a lightweight, learnt method for radar odometry based in a two-stage correlation operation. In [4], we used a

In the area of **localisation**, in [5], [6] we investigated the potential of radar for **place recognition** and in [7] we combined this with **metric pose estimation** in order to build a complete, radar-only localisation pipeline. Moreover, we followed these works in [8], [9] by proposing an unsupervised training of the place recognition network, requiring no place labels. In [10], [11], [12], [13], we explored how external services (linking to Sec. III-A3), e.g. readily-available **satellite imagery**, can be used as map proxy for radar sensing.

2) WP1.B Auditory Sensing (stretch): Here, we showed that audio can be a valuable sensor for **assessing road surfaces** – e.g. gravel, tarmac – and their status – e.g. dry, wet. In particular, in [14], we used audio to classify drivable surfaces in an outdoor environment, specifically grass and gravel. Here, the classification pipeline is used to cheaply create a vast amount of labels needed to train a supervised radar surface segmentation pipeline by fusing odometry information.

3) WP1.C Leveraging External Services: Here, we showed that ready-available satellite images can be very valuable as **cheap map proxies** for range sensors. In [10], we used a multi-stage approach to convert first a satellite image and its proximal radar scan into a synthetic radar image centred on the satellite data but resembling the radar scan; secondly, we register the two radar image, the live one and the synthetic, though a correlation approach. In [11], [13] we simplify the training procedure by relaxing the need of accurate ground-truth offsets between the satellite and radar images. In [12] we converted both LiDAR and satellite images into collections of 2D points for solving both place recognition and metric localisation between satellite imagery and LiDAR sensors.

4) WP1.D Data Collection: We carried out an extensive data collection; the aim was to cover as many **scenarios and weather conditions**, ranging from urban data in central London to off-road data in Scotland Highlands. Alongside the dataset, we produce **copious labels** for semantic scene understanding, object detection and odometry; please refer to Sec. IV-C for more details. Nevertheless, no dataset can comprehend all possible scenarios a deployed system could possibly face, so we investigated techniques for **synthesising new data** that can be used for bolstering robustness to unseen situations [15], [16] (see Sec. V-C for more details).

B. WP2: Performance Assessment

Here, we have shown that is possible to predict and estimate the performance of a robot system both in the case of motion estimation/localisation and perception.

1) WP2.A Predicting Localisation Performance: Embedding a form of estimation of performance and **confidence** in the odometry and localisation pipelines can bolster their robustness and prevent critical failures. We showed it in [2], where the eigenvectors used by the matching approach lead to a confidence score that can be used for **rejecting expected failure cases**. In [8], [9] we analysed how **distributional similarity** as a proxy for localisation confidence can increase the system performance. Moreover, in [4], we showed that **filtering** matches through a motion model before computing the motion estimation is greatly beneficial.

Finally, alongside the dataset (c.f. Sec. III-A4), we released odometry and localisation ground truth for future research in performance assessment of localisation.

2) WP2.B Estimating Model Confidence: The ability of a model to learn a mapping from input data to outputs, e.g. objects in the scene or odometry prediction, is not sufficient if it is not accompanied by the estimation of its confidence. We explicitly discussed this topic in [17] in the context of semantic segmentation, where the segmentation network has been augmented with a second task network that is trained in an unsupervised fashion to predict the pixel-wise model confidence on the output. The model confidence can then be exploited to **reject unconfident predictions**, improving model **introspection**.

Moreover, alongside the dataset (c.f. Sec. III-A4), we release plenty of semantic segmentation ground-truth masks for future research. For more details, refer to Sec. IV-C.

C. WP3: Causal Explanation

1) WP3.A Scenario-based Requirement Analysis: The safe deployment of autonomous vehicles (AVs) in real-world scenarios requires them to be **accountable and trustworthy**. Hence, it is necessary to explore the impact of explanations that provide information about an AV's behaviour—in essence explaining what a vehicle **has "seen"**, done and **might do** in a given scenario. In order to unpack the meanings and properties of explanations, we conducted a thorough comprehensive review of explanations generally in autonomous driving [18]. We formed a categorisation of the explanations studied in the driving literature, and then we explored how the explanations have been considered in the different AV operations (e.g., perception, localisation, planning, vehicle and control, and system management). We also provided **technical and regulatory recommendations** for more effective AV explainability.

In [19], we studied the different **explanations types** identified from our survey and applied them in different driving scenarios obtained in the real-world. A major objective of the work was to **identify scenarios** where explanations could be primarily useful. Driving scenarios were defined by a mix of road topology, road rules/traffic control signs, and vehicle actions with respect to other road participants. The intelligibility of explanations was evaluated based on the degree to which an explanation improves a person's understanding of an AV's operation after performing a set of tasks. Subsequently, in [20], we assessed the **human's perceptions of AVs** in the presence of causal and non-causal explanations. While results suggested that explanations could be useful for enhancing understanding of driving behaviours, they might not necessarily lead to a higher **calibration of trust in AVs**. We took a step further to obtain experts' thoughts on explainability, we conducted a CHI workshop [21] where we brought experts in the field to share their thoughts on the current state-of-the-art contributions on explainability in autonomous physical systems. Conclusions reached through this workshop aligned with the results from our previous studies.

2) WP3.B Semantic Scene Representation: Semantics can be a very powerful representation of the sensor's data, since it contains **high-level information** that is robust to change in appearance.

We showed results for semantic segmentation on visual data featuring **out-of-distribution** content in [17] (see Sec. III-B2) where the semantic result is accompanied by the relative model confidence.

In [22] we exploit object level, semantic information to map a city environment for LiDAR localisation achieving competitive accuracy results with very **compressed representations**.

Moreover, we showed **radar's potential** in semantic scene understanding: [23] uses vision and LiDAR as ground truth for urban segmentation, [24] uses a LiDAR map to learn a traversability map from single radar images and [14] uses audio to supervise segmentation maps of drivable surfaces in an outdoor environment (see Sec. III-A2).

Tangentially, we exploited semantic information in [16] as input for **image synthesis** to produce aligned realistic data for training more robust semantic segmenters.

Finally, we have annotated the delivered dataset with plenty of semantic masks in all sceneries. For more details, refer to

Sec. IV-C.

3) WP3.C Learning and Inference for Causal Explanation: Having constructed a clear understanding of the theories, applications and practices around explainable AI in autonomous driving (Sec. III-C1), we explored different ways to build **interpretable models for specific driving tasks**. In [25], we designed algorithms to generate **factual and counterfactual explanations** for the predictions of vehicle collision **risk models** learned on the planar time to collision (TTC) metrics. Results indicated that the AV developers can potentially apply this explainer tool for model debugging purposes.

We also conducted a comparative study on the different **explanation techniques** for deep convolutional neural networks (CNN). In [26], we compared the contextual importance and utility (CIU) algorithm with other commonly used techniques. CIU indicated better performances in effectively occluding the part of the test images that were not important for the deep CNN model. A user study conducted following this study also verified this result.

IV. Milestones and Deliverables

We use milestones and deliverables to give a broader view of the above work packages. MS1/D1 relates to **O1** and **O3**. MS2/D2 relates to **O3**. MS2/D3 relates to **O1**. MS3/D4 relates to **O3**.

A. MS1/D1: Taxonomy of challenging scenarios

The project investigated three types of driving scenarios: off-road; semi-urban and motorway; and dense urban scenarios. We collected data from different places using the same robotic platform and sensors. This will allow users of the resulting dataset to test and evaluate their algorithms and methods in a variety of challenging scenarios—from the city centre of London to the Highlands of Scotland.

- **Hounslow Hall Estate**, in Buckinghamshire, England. This site features three routes, mainly off-road, of increasing difficulty even under fair conditions.
- **Ardverikie Estate**, in the Scottish Highlands. This isolated estate features gravel and rock tracks and sandy beaches in highly variable lighting and weather conditions.
- **The New Forest**, a large area of unenclosed pasture land and forest in Southern England. Along the spectrum of driving difficulty, this site presents easy surfaces (public roads) and low traffic.
- **The Oxford Ring-Road**, in Oxfordshire. This site features driving at high-speed (up to 70 mph) in dense traffic around a network of A-, B-, and M-type motorways.

Furthermore, through a user study, we investigated a taxonomy of events including: normative, near-misses, collisions, and emergency events.

Our research themes contributing to success in these outcomes are Sec. III-A4, Sec. III-C1.

B. MS2/D2: Taxonomy of explanations for key stakeholders

We have considered explanation dimensions, types of explanations, and stakeholders.

For **explanation dimensions**, we consider causes filters, content type, model dependence, interactivity, system type, and scope.

For **types of explanations**, we consider causal explanations, contrastiveness and non-contrastiveness, as well as counterfactuals.

In terms of **stakeholders**, we consider end-users (passengers, pedestrians, pedestrians with reduced mobility, other road participants such as cyclists, and auxiliary drivers), technicians and engineers, as well as regulators and insurers.

Our research themes contributing to success in these outcomes are Sec. III-C3.

C. MS2/D3: Dataset of unconventional sensor

We selected a full sensor suite, capable of covering a wide range of application:

- A forward-facing stereo camera and a set of three two-lens stereo cameras facing forwards and backwards;
- A set of five single-beam, 2D lasers and a single, roof-mounted, multiple-beam 3D laser;
- A roof-mounted scanning radar;
- A bumper-mounted automotive radar;
- An Inertial Navigation System (INS);
- The internal Controller Area Network (CAN) bus signals;
- Four omnidirectional boom microphones on the two front and two – optional – back wheel arches;
- An in-cabin monocular camera for recording the driver's behaviour;
- An in-cabin, optional microphone to complement the in-cabin video under professional driving instruction.

Alongside, we provide a set of ground-truth annotations for a variety of tasks:

- Bounding box *object positions* for camera as well as for LiDAR (where we provide 3D boxes) and radar;
- Pixel-wise *segmentation*, with object classes dictated by the different scenarios, with corresponding CAN signalling;
- *Event-based audio-visual sequences*, where we uniquely provide synchronised audio and video, using in-cabin recording of a professional driving instructor (only for urban scenario);
- *Position tracking*, using a Leica Viva TS16 Total Station for precise millimetre-accurate position ground truth.

Our research themes contributing to success in these outcomes are Sec. III-A1, Sec. III-A4.

V. Body of Knowledge

In our contribution to the BoK, we addressed the Assurance Objective on Explanability. Here, we have considered explanations to be useful for:

- **Justification** of decisions taken by the vehicle, c.f. Sec. V-F;

- **Control**, detecting and mitigating performance drops, c.f. Sec. V-E;
- **Improvements** by predicting, circumventing, and explaining performance issues, c.f. Sec. V-G;
- **Discovery** by singling out factors which have a major impact on the performance of learnt models, c.f. Secs. V-C and V-D.

Feeding into all these streams are sensing and understanding requirements, c.f. Secs. V-A and V-B.

A. 2.2.1.1: Defining Sensing requirements

Our work in this area considers that different combinations of scene and possible threat will require different sensor payloads and even unusual sensors.

A diverse dataset Our dataset contains a broad combination of scenes – urban, rural and off-road – and hazards – mixed driving surfaces, adverse weather conditions, and other actors' presence Our publication output in this area includes [1], [27]. This work arises from **O1** and **O3** and was aimed at satisfying Secs. III-A4, IV-A and IV-C.

FMCW scanning radar We showed how AVs could utilise radar independently from other sensors for low-level autonomy tasks, ranging from odometry and localisation to scene understanding to path planning. Our publication output in this area includes [9], [13], [8], [12], [24], [14], [11], [5], [6], [10], [7], [23], [2]. This work arises from **O1** and was aimed at satisfying Secs. III-A1 to III-A3.

Audio Audio has the advantage to be inherently invariant to the scene illumination, although it contains only very punctual information. Our publication output in this area includes [14]. This work arises from **O1** and was aimed at satisfying Sec. III-A2.

CAN CAN contain critical information for several tasks, which can treat them either as sensory data – e.g. for driver identification – or as control signals – e.g. for training behavioural-cloning algorithms. Our publication output in this area includes [27]. This work arises from **O2** and was aimed at satisfying Sec. III-B1.

External services We challenge the definition of a sensor by including services provided by external operators, particularly satellite imagery. Our publication output in this area includes [12], [10], [11], [13]. This work arises from **O1** and was aimed at satisfying Sec. III-A3.

B. 2.2.1.2: Defining Understanding requirements

Our work in this area focuses on structured but data-driven algorithms for localisation, object detection, and planning/control.

Radar for robust understanding We find this sensor to be critical in inclement weather and at high-speed in motorway driving scenarios. Our publication output in this area includes [9], [13], [8], [12], [24], [14], [11], [5], [6], [10], [7], [23], [2]. This work arises from **O1** and was aimed at satisfying Secs. III-A1 to III-A3.

Implicit and explicit task understanding We consider lower-level autonomy-enabling tasks which require implicit understanding as well as higher-level autonomy-enabling tasks which provide explicit understanding of the world. Our publication output in this area includes [9], [4], [3], [2], [5], [8], [6], [23], [24], [14]. This work arises from **O1** and **O3** and was aimed at satisfying Secs. III-B1 and III-B2.

Cross-modal understanding There are certain scenarios in which one modality will perform best. There will be some scenarios in which a range of choices are available for which technology to deploy. Our publication output in this area includes [2], [23], [24]. This work arises from **O1** and **O2** and was aimed at satisfying Secs. III-A1 and III-B1.

C. 2.3: Implementing requirements using ML

Here, our work has focused on lack of enough data or labels and the cost of labelling.

Generative approaches We have found that generative approaches provide a way to synthesise new training data that the AV never experienced during the data collection. Our publication output in this area includes [15], [16]. This work arises from **O1** and was aimed at satisfying Secs. III-A4 and III-C2.

Weakly-supervised learning Our work has exploited augmented available labels by either projecting one sensor on top of another or retaining labels through time, significantly increasing the number of available annotations. Our publication output in this area includes [2], [23], [14], [24]. This work arises from **O1** and was aimed at satisfying Secs. III-A1 and III-A4.

Unsupervised Learning We have found that unsupervised learning has allowed us to train with much more raw sensor scans, mitigating entirely the need for ground truth or pseudo-labels, and exposing the network during training to much more realistic data. Our publication output in this area includes [11], [8], [13], [9]. This work arises from **O1** and was aimed at satisfying Secs. III-A1 and III-A4.

D. 2.3.1: Sufficiency of training

Our work in this area has focused on the sources of training data and types of annotations which are required for deployment in more diverse scenarios.

Dataset overview Our key contribution is the deployment of one platform in five places ranging from on-road to off-road environments. Our focus for data capture revolves around unusual sensing modalities, mixed driving surfaces, varied operational domains, and adverse weather conditions.

Annotation We have focused our annotation effort on bounding box object positions, pixel-wise segmentation, event-based audio-visual sequences, and position tracking.

Our publication output in this area includes [1], [27]. This work arises from **O1** and **O3** and was aimed at satisfying Secs. III-A4, IV-A and IV-C.

E. 2.6: Handling change during operation

Here, we focus on weather conditions, lighting, dynamic objects and agents, and internal settings and parameters.

Adapting to change Where it is still desirable to use camera and laser, our work has developed processing strategies which can understand, represent, and effect change on the sensor stream can be powerfully applied to "normalise" sensor data for better autonomous performance. Our publication output in this area includes [28], [15].

Learning from change Changes, being inevitable, should be maximally exploited as a rich source of information for learned systems. Therefore, these learned systems should be structured such that they are able to incorporate signals arising from change. Our publication output in this area includes [9], [8].

Identifying change We assert that even if change cannot be handled satisfactorily, it is still hugely beneficial for the system to understand that change has occurred. It is furthermore important for the interpretability of autonomous driving tasks that the change can be localised within the sensor observation. Our publication output in this area includes [17].

Representing change in training datasets It is crucial that publicly available datasets for autonomous driving feature change. This is true when developing learned systems as well as for more classical autonomy stacks. These changes should capture typical changes as well as more unusual deviations in scene appearance. Our publication output in this area includes [1], [27].

To embrace or to reject change? When actively addressed, change in the environment as scanned by sensors may be embraced or rejected for more robust system performance (e.g. as for selective segmentation, see above). Our publication output in the area of embracing change includes [7] and in the area of rejecting change includes [29], [4], [3], [30].

This work arises from primarily **O2** but also **O1** and was aimed at satisfying Secs. III-A1 and III-B1.

F. 2.8: Explainability

Here, we focus on *explanations* as key for building trust in autonomous systems such as autonomous vehicles (AVs) and identify different types of explanations in challenging driving scenarios and stakeholders for which explanations are relevant.

Need for Explainability, the needs for explanations differ across different stakeholders. While end-users, in-vehicle participants, among others would demand clear explanations in natural language, developers might prefer more technical explanations that would facilitate easy. debugging of their systems. Ultimately, the goal is to build to safe AVs that are transparent, accountable, and trustworthy. Our publication output in this area includes [18], [21].

Standards and Regulations we have identified several standards related to explainability in AVs. Our publication output in this area includes [18].

Stakeholders in AV Explainability we identified three groups of stakeholders, in particular: *End-Users* (i.e. passengers, auxiliary driver, pedestrians and other road users), *Developers and Technicians* (i.e. AV developers and automobile technicians) and *Regulators and Insurers* (i.e. system auditors, regulators, accident investigators and insurers). Our publication output in this area includes [18].

Explanations Categorisation We provide a categorisation of explanations based on the different methodologies identified in literature (*unvalidated guidelines*, *empirically derived* and *psychologically constructed*). We also identified the different explanation dimensions in AVs. Our publication output in this area includes [18].

Generating Explanations We consider explanations in visual and natural language formats and have developed natural language explanations for collision risks models. Our publication output in this area includes [25], [26].

Evaluating Explanations We proposed an interpretable, tree-based data representation approach to assess accountability in autonomous driving through explanations which reference to AV's actions, observations, and road rules along with a user study. Our publication output in this area includes [19], [20].

Activities Description We have identified multiple types and dimensions of explanations as well as the requirements for different stakeholders. Our publication output in this area includes [18].

This work arises from **O3** and was aimed at satisfying Secs. III-C1 and III-C3.

G. 3: Understanding and controlling deviations from required behaviour

We focus in this area on distributional shifts from training data, poorly modelled vehicle kinematics, sensor artefacts, environmental ambiguity, and scene distractors.

Vehicle modelling We have focused here on the fidelity of motion estimation being suspect to incorrect data association on a frame-by-frame basis. This is tied to environmental ambiguity and sensor artefacts. Our publication output in this area includes [4], [6].

Sensor artefacts Here, sensor failure modes and other noises perturb the vehicle's measurements which are processed by autonomy-enabling systems in order to make predictions and inform behaviour. Our publication output in this area includes [5], [9], [8].

Environmental ambiguity Specific to the way in which the sensor technology works, different parts of the environment may have very similar "appearance" which can be confusing for downstream tasks, leading to errors in scene understanding. Our publication output in this area includes [2], [17].

Scene distractors Perturbation of the sensor stream and therefore deviation from ideal performance in autonomy enabling tasks and autonomous behaviour itself is also possible due to the presence of unaccounted for and uncontrollable objects and actors in the scene. Our publication output in this area includes [29], [3].

This work arises from primarily **O2** but also **O1** and was aimed at satisfying Secs. III-A1 and III-B1.

VI. Conclusion

In this document, we discussed the results we achieved during the SAX project in world-leading research in mobile autonomy. We addressed fundamental questions to overcome critical barriers to assurance and regulation for large-scale deployments of autonomous systems in challenging on-road and off-road driving scenarios.

In particular, we present a universal view of AV sensing requirements and how uncommon sensing modalities can be suitable for overcoming challenging operational scenarios, showing that scanning radar, despite its age as a technology, is a tool that can rival vision and laser in scene understanding across all crucial autonomy-enabling tasks. We present an overview of AV training requirements and approaches to tackle the lack of specific sensing combinations or labels, and we released a vast dataset in a variety of scenarios and conditions, comprising a full sensor suite and manually-annotated labels for odometry, localisation, semantic segmentation and object detection. Change for autonomous vehicle operation is unavoidable, predictable to an extent but often unpredictable, a rich source of training information, crucial to be properly represented in training data, often correctable during operation or if not at least detectable. We identify different types of explanations in challenging driving scenarios. Moreover, we characterise several dimensions for explanations and identify different stakeholders for which explanations are relevant. Furthermore, we develop methods and provide guidance for generating explanations using vehicle perception and action data in dynamic driving scenarios.

References

[1] M. Gadd, D. de Martini, L. Marchegiani, P. Newman, and L. Kunze, "Sense–Assess–eXplain (SAX): Building Trust in Autonomous Vehicles in Challenging Real-World Driving Scenarios," in *2020 IEEE Intelligent Vehicles Symposium (IV)*, Oct. 2020, pp. 150–155, iSSN: 2642-7214.

[2] R. Aldera, D. D. Martini, M. Gadd, and P. Newman, "What Could Go Wrong? Introspective Radar Odometry in Challenging Environments," in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, Oct. 2019, pp. 2835–2842.

[3] R. Weston, M. Gadd, D. De Martini, P. Newman, and I. Posner, "Fast-MbyM: Leveraging Translational Invariance of the Fourier Transform for Efficient and Accurate Radar Odometry," in *IEEE International Conference on Robotics and Automation (ICRA)*, Philadelphia, PA, USA, September 2022.

[4] R. Aldera, M. Gadd, D. De Martini, and P. Newman, "Leveraging a Constant-curvature Motion Constraint in Radar Odometry," under Review.

[5] Ș. Săftescu, M. Gadd, D. De Martini, D. Barnes, and P. Newman, "Kidnapped Radar: Topological Radar Localisation using Rotationally-Invariant Metric Learning," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, May 2020, pp. 4358–4364, iSSN: 2577-087X.

[6] M. Gadd, D. De Martini, and P. Newman, "Look Around You: Sequence-based Radar Place Recognition with Learned Rotational Invariance," Apr. 2020, pp. 270–276.

[7] D. De Martini, M. Gadd, and P. Newman, "kRadar++: Coarse-to-Fine FMCW Scanning Radar Localisation," *Sensors*, vol. 20, no. 21, p. 6002, Jan. 2020. [Online]. Available: https://www.mdpi.com/1424-8220/20/21/6002

[8] M. Gadd, D. De Martini, and P. Newman, "Unsupervised Place Recognition with Deep Embedding Learning over Radar Videos," *arXiv:2106.06703 [cs]*, Jun. 2021, arXiv: 2106.06703. [Online]. Available: http://arxiv.org/abs/2106.06703

[9] M. Gadd, D. De Martini, and P. Newman, "Contrastive Learning for Unsupervised Radar Place Recognition," in *2021 20th International Conference on Advanced Robotics (ICAR)*, Dec. 2021, pp. 344–349.

[10] T. Y. Tang, D. De Martini, D. Barnes, and P. Newman, "RSL-Net: Localising in Satellite Images From a Radar on the Ground," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1087–1094, Apr. 2020.

[11] T. Y. Tang, D. Martini, S. Wu, and P. Newman, "Self-Supervised Localisation between Range Sensors and Overhead Imagery," *Robotics: Science and Systems*, 2020.

[12] T. Y. Tang, D. D. Martini, and P. Newman, "Get to the Point: Learning Lidar Place Recognition and Metric Localisation Using Overhead Imagery," vol. 17, Jul. 2021. [Online]. Available: http://roboticsproceedings.org/rss17/p003.html

[13] T. Y. Tang, D. De Martini, S. Wu, and P. Newman, "Self-supervised learning for using overhead imagery as maps in outdoor range sensor localization," *The International Journal of Robotics Research*, vol. 40, no. 12-14, pp. 1488–1509, Dec. 2021. [Online]. Available: http://journals.sagepub.com/doi/10.1177/02783649211045736

[14] D. Williams, D. De Martini, M. Gadd, L. Marchegiani, and P. Newman, "Keep off the Grass: Permissible Driving Routes from Radar with Weak Audio Supervision," in *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*, Sep. 2020, pp. 1–6.

[15] V. Mușat, I. Fursa, P. Newman, F. Cuzzolin, and A. Bradley, "Multi-weather city: Adverse weather stacking for autonomous driving," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 2906–2915.

[16] V. Mușat, D. De Martini, M. Gadd, and P. Newman, "Depth-SIMS: Semi-parametric Image and Depth Synthesis," in *IEEE International Conference on Robotics and Automation (ICRA)*, Philadelphia, PA, USA, September 2022.

[17] D. Williams, M. Gadd, D. De Martini, and P. Newman, *Fool me once: robust selective segmentation via out-of-distribution detection with contrastive learning*. IEEE, 2021. [Online]. Available: https://ora.ox.ac.uk/objects/uuid:eebd0d6a-e21b-4065-b09f-1f9afa3b6b07

[18] D. Omeiza, H. Webb, M. Jirotka, and L. Kunze, "Explanations in autonomous driving: A survey," *IEEE Transactions on Intelligent Transportation Systems*, 2021.

[19] D. Omeiza, H. Web, M. Jirotka, and L. Kunze, "Towards accountability: providing intelligible explanations in autonomous driving," in *2021 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2021, pp. 231–237.

[20] D. Omeiza, K. Kollnig, H. Web, M. Jirotka, and L. Kunze, "Why not explain? effects of explanations on human perceptions of autonomous driving," in *2021 IEEE International Conference on Advanced Robotics and Its Social Impacts (ARSO)*. IEEE, 2021, pp. 194–199.

[21] D. Omeiza, S. Anjomshoae, K. Kollnig, O.-M. Camburu, K. Främling, and L. Kunze, "Towards explainable and trustworthy autonomous physical systems," in *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*, 2021, pp. 1–3.

[22] G. Pramatarov, M. Gadd, D. De Martini, and P. Newman, "Box-Graph: Semantic Place Recognition and Pose Estimation from 3D LiDAR," under Review.

[23] P. Kaul, D. de Martini, M. Gadd, and P. Newman, "RSS-Net: Weakly-Supervised Multi-Class Semantic Segmentation with FMCW Radar," in *2020 IEEE Intelligent Vehicles Symposium (IV)*, Oct. 2020, pp. 431–436, iSSN: 2642-7214.

[24] M. Broome, M. Gadd, D. De Martini, and P. Newman, "On the Road: Route Proposal from Radar Self-Supervised by Fuzzy LiDAR Traversability," *AI*, vol. 1, no. 4, pp. 558–585, Dec. 2020. [Online]. Available: https://www.mdpi.com/2673-2688/1/4/33

[25] R. Nahata, D. Omeiza, R. Howard, and L. Kunze, "Assessing and explaining collision risk in dynamic environments for autonomous driving safety," in *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*. IEEE, 2021, pp. 223–230.

[26] S. Anjomshoae, D. Omeiza, and L. Jiang, "Context-based image explanations for deep neural networks," *Image and Vision Computing*, vol. 116, p. 104310, 2021.

[27] M. Gadd, D. De Martini, O. Bartlett, P. Murcutt, M. Towlson, L. Kunze, and P. Newman, "The Sense-Assess-eXplain Range-Rover RobotCar Dataset," in Preparation.

[28] H. Porav, T. Bruls, and P. Newman, "Don't worry about the weather: Unsupervised condition-dependent domain adaptation," in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*. IEEE, 2019, pp. 33–40.

[29] T. Suleymanov, M. Gadd, L. Kunze, and P. Newman, "LiDAR Lateral Localisation Despite Challenging Occlusion from Traffic," in *IEEE/ION Position, Location and Navigation Symposium (PLANS)*, Portland, OR, USA, April 2020, pp. 334–341.

[30] Ș. Săftescu, M. Gadd, and P. Newman, "Look Here: Learning Geometrically Consistent Refinement of Inverse-Depth Images for 3D Reconstruction," *International Journal of Pattern Recognition and Artificial Intelligence (IJPRAI)*, 2021.

ASSURING
AUTONOMY
INTERNATIONAL PROGRAMME