

2.6 – Handling change during operation

Practical guidance –automotive

Authors: Dr Matthew Gadd, University of Oxford

The vagaries of autonomous driving

Due to the complexity of the underlying activity and the variation in its context, there are various sources of changes which manifest for an autonomous driving operation. We understand these to include:

1. Weather conditions
2. Lighting
3. Dynamic objects and agents
4. Internal settings and parameters

These have safety and efficiency implications, where autonomy-enabling tasks which rely on sensor observations and system calibration can have their performances perturbed or totally degraded.

In the face of such variability we propose several related approaches to change, detailed below.

Change-immune sensing

While extremely fast and high-resolution, LiDAR is sensitive to weather conditions (see Fig. 2), especially rain and fog, and cannot see past the first surface encountered. Vision systems are versatile and cheap and naturally understandable by human operators, but easily impaired by scene changes (see Fig. 1), like poor lighting or the sudden presence of snow.

Radar overcomes the challenges to these technologies because it is a long-range, onboard system that performs well under a variety of lighting (it is immune) and atmospheric conditions, and it is rapidly becoming more affordable than LiDAR. Due to its long wavelength (which allows it to penetrate certain materials) and beam spread, radar can return multiple readings from the same transmission and generate a rich representation of the world.



Figure 1: Vision-based techniques are susceptible to drastic appearance and illumination variation, caused by inclement weather conditions or seasonal changes -- heavy precipitation in this example.




Figure 2: The main limitation of LiDAR sensors for outdoor environments is that their measurements are affected by adverse environmental conditions, such as rain and fog. In this example, taken from [11], LiDAR range decreases when rain rate increases from 0 mmph (left), 9 mmph (middle), and 17 mmph (right).

Adapting to change

Despite their susceptibility to change, vision and LiDAR are, however, still affordable and information rich technologies and for this reason feature heavily in academic and commercial autonomous vehicle research and development. Therefore, 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.



Figure 3: Instead of mapping from a difficult condition (night) to a presumed ideal condition (e.g. sunny), we learn to map from any adverse appearance to a shared, common appearance that while not seeming natural actually works better for a segmentation task.

In our work we therefore develop methods which exploit the predictable effects of weather to “hallucinate” clear-condition images from a troublesome condition (e.g. blizzard, sandstorm). To avoid human bias in hand-selecting a clean seasonal appearance and to curtail the explosion in the required number of pairwise mappings, a composite canonical appearance which maximises effective robot behaviour can be learned, as we have showed in [1] and illustrated in Fig. 3.

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.

An example from our work taken from [2] and illustrated in Fig. 4 is the boost to localisation performance observed when using real radar examples sampled along a vehicle’s trajectory rather than synthesising pairs. Radar is highly “stochastic”, in that the complexity of the measurement process introduces fast modes of change in scans, even if the vehicle is

stationary. We leverage these in order to naturally regularise the adaption of neural network weights while learning radar scan representations.




Figure 4: Radar scans registered to one another by vehicle motion which in the absence of noises should exactly overlap and not differ in content. However, the radar scan formation process is subject to unpredictable changes due to many sensing artefacts. We exploit this for bonafide “true positive” examples which nevertheless have distinct scan content, for better neural network optimisation.

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.

In this area, our work on image segmentation in [3] and illustrated in Fig. 5 has resulted in a selective segmentation system which can identify which portions of images do not belong to the set of known classes that the system was trained to identify. To achieve this, we use a self-supervised procedure which incorporates “out-of-distribution” patches from the vast store of unlabelled imagery available online. The system is self-supervised in the sense that we do not know the class of the out-of-distribution patches, but we do know in a binary sense that those pixels are unknown. This allows the neural network which performs image segmentation to identify image regions when deployed live as not being sufficiently represented (“known”) in the training corpus - indicating a “distributional change”, which should be interpreted as a change in the operating conditions as compared to design and implementation of the system.

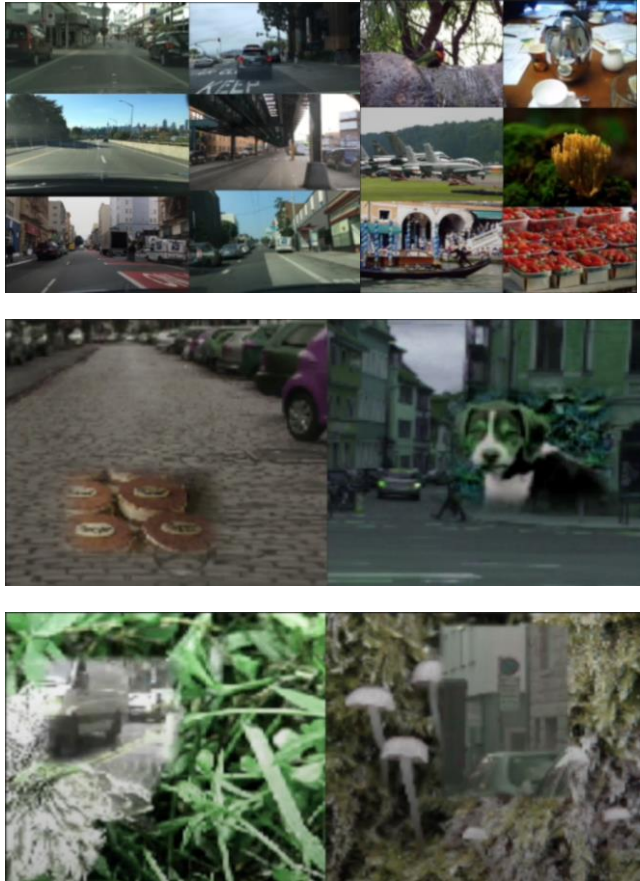


Figure 5: Top left: In-distribution images from typical autonomous driving scenes. Top right: Out-of-distribution images. Middle: Out-of-distribution patches inserted into in-distribution scenes. Bottom: In-distribution patches inserted into out-of-distribution scenes.

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.

In our work, we have found in [4] that it is best to collect datasets for localisation and motion estimation with radar over many repeat visits over the trajectory of interest, due to the stochastic nature of the radar measurement process (see above).

This has informed our trial planning for our upcoming dataset [5], where additionally as shown in Fig. 5 we have collected data across seasonal changes which are not exactly repeatable (in England) but which must be handled robustly if encountered.



Figure 5: Datasets featuring typical and unusual change are crucial for the future success of autonomous vehicles which can handle change during operation. Featured top is an example of typical scene change (day to night). Featured bottom is an example of scene change (clear to snow) with a different periodicity (seasonal as opposed to diurnal) which also may be less common (at least in certain parts of the world).

To embrace or to reject change?

When actively addressed, change in the environment as scanned by sensors may be embraced (e.g. as for training learned systems, see above) or rejected for more robust system performance (e.g. as for selective segmentation, see above).

More examples of change rejection in our work include [6] where dynamic obstacles occluding road boundaries are rejected from the semantic understanding of the road layout (changed from a pre-recorded semantic kerb map, for instance) for better lateral localisation in lane-keeping scenario, and [7] where perturbations in radar scans which are used for motion estimation are identified and smoothed, [8] where instead a learned neural network attention mask is used in a similar same task, and [10] where defects in 3D scans are corrected occurring to their change as measured against a higher quality sensor (used only for training).

In terms of embracing change, our work in [9] is an important bridge towards implementing in the radar domain the “experience-based navigation” methodology of [12]. Here, a simple but elegant algorithm requires new memories to be recorded only when already-stored memories are insufficient to explain the current sensor observation. While as we have laid out above radar is less susceptible to appearance change, our work in [9] will allow us to correct for the small amount of drift that may occur due to unavoidable changes such as structural variation as cities change over extended time-scales.

Summary

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. Dealing with

change can be tackled at the system conception phase (design, coding, training, etc) and in the field. Algorithms and hardware are equally important. This area is one of the core challenges to universal autonomy.

References

- [1] H. Porav, T. Bruls, and P. Newman, "Don't worry about the weather: Unsupervised condition-dependent domain adaptation," in IEEE Intelligent Transportation Systems Conference (ITSC), 2019
- [2] M. Gadd, D. De Martini, P. Newman, "Unsupervised Place Recognition with Deep Embedding Learning over Radar Videos", Workshop on Radar Perception for All-Weather Autonomy at the IEEE International Conference on Robotics and Automation (ICRA), 2021
- [3] D. Williams, M. Gadd, D. De Martini, and P. Newman. "Fool Me Once: Robust Selective Segmentation via Out-of-Distribution Detection with Contrastive Learning," in Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) 2021.
- [4] D. Barnes, M. Gadd, P. Murcutt, P. Newman, and I. Posner, "The Oxford Radar RobotCar Dataset: A Radar Extension to the Oxford RobotCar Dataset," in Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), Paris, 2020
- [5] M. Gadd, D. De Martini, L. Marchegiani, L. Kunze, and P. Newman, "Sense-Assess-eXplain (SAX): Building Trust in Autonomous Vehicles in Challenging Real-World Driving Scenarios," in Proceedings of the IEEE Intelligent Vehicles Symposium (IV), Workshop on Ensuring and Validating Safety for Automated Vehicles (EVS AV), 2020
- [6] 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, 2020.
- [7] R. Aldera, D. De Martini, M. Gadd, and P. Newman, "What Could Go Wrong? Introspective Radar Odometry in Challenging Environments," in IEEE Intelligent Transportation Systems (ITSC) Conference, Auckland, New Zealand, 2019.
- [8] R. Weston, M. Gadd, D. De Martini, P. Newman, and I. Posner, "Efficient Learned Radar Odometry," in IEEE International Conference on Robotics and Automation (ICRA), (Philadelphia, PA, USA), September 2022.
- [9] D. De Martini, M. Gadd, and P. Newman, "kRadar++: Coarse-to-fine FMCW Scanning Radar Localisation," Sensors, Special Issue on Sensing Applications in Robotics, vol. 20, no. 21, p. 6002, 2020
- [10] S. Saftescu, 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
- [11] Goodin, C., Carruth, D., Doude, M. and Hudson, C., 2019. Predicting the Influence of Rain on LIDAR in ADAS. Electronics, 8(1), p.89.

[12] W. Churchill and P. Newman, "Practice Makes Perfect? Managing and Leveraging Visual Experiences for Lifelong Navigation," in Proc. IEEE International Conference on Robotics and Automation (ICRA), Minnesota, USA, 2012