Vehicle Autonomy in Adverse Conditions: Radar, Optical & Fused Sensing

UDRC Theme Meeting, July 2020

Presented by Andy Wallace, Bernie Mulgrew

Thanks to Mehyrar Emambakhsh, Sap Mukerjee, Marcel Sheeny, Alireza Ahrabian, Yun Wu, Andreas Assmann, Marina Gashinova, Ed Hoare, Liam Daniel, Dom Phippen, Mikhail Cherniakov, David Wright, Shazad Gishikori, Gerald Buller, Abde Halimi, Sen Wang, Zhiyang Hong…

[dstl]

EPSRC
Engineering and Physical Sciences Research Council

Context: sensing for autonomy and assistance

UDRC OBJECTIVE (KTM): Situational Awareness for Land Vehicles: "Navigate and identify threats in complex (e.g. urban) environments with recommendations for action."

Scottish Parliament: Bringing Transformative Change to Scotland

"The development of driverless vehicles has the potential to transform journeys and transport management in Scotland and beyond. Autonomous cars are close to becoming a reality. *But are we prepared to have autonomous cars in all kinds of scenarios?"*

[dstl]

EPSRC
Engineering and Physical Sciences Research Council

Seeing though obscurants: which sensor?

	3D Landscaping	Object Classification	Range	Operation in Adverse Weather	Operation at night
Camera					 ->LWIR 
LiDAR				 (2)	
Radar	 (3)	 (1)			

TASKS

- 3D landscaping, i.e. perceiving surface (and material?) structure
- Actor and object recognition
- Predict behaviour of said actors, have situational awareness and make/recommend action

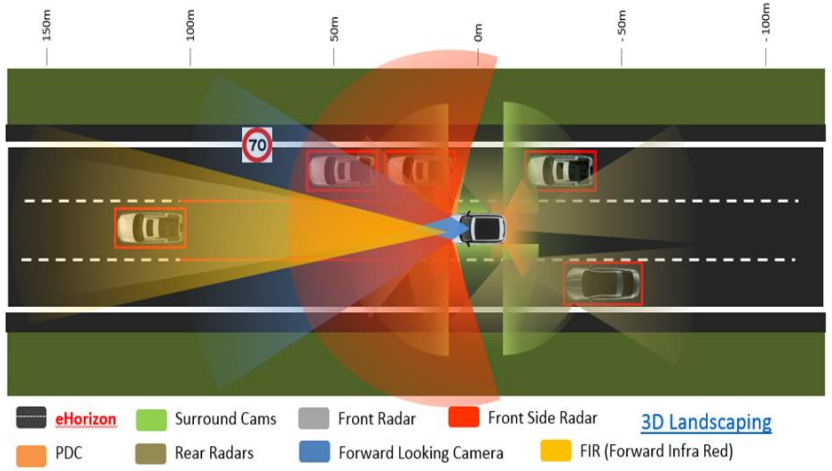

This talk has 3 sections, aimed at the red dots, (1,2,3)

1. Using and interpreting radar imagery

Test Facility for Land Vehicles

- The first problem we had is there appeared to be no suitable annotated dataset for radar imagery from a moving vehicle; we had to collect our own
- We collected data for experiments through fine and adverse weather using a LiDAR, Radar (79Ghz) and Stereo Camera.
- We labelled over three hours of data with 200,000 actors (cars, vans, pedestrians, bikes etc.)
- This can be used for recognition, fusion and behaviour experiments, and includes data collected in rain, snow, mist, and at night.

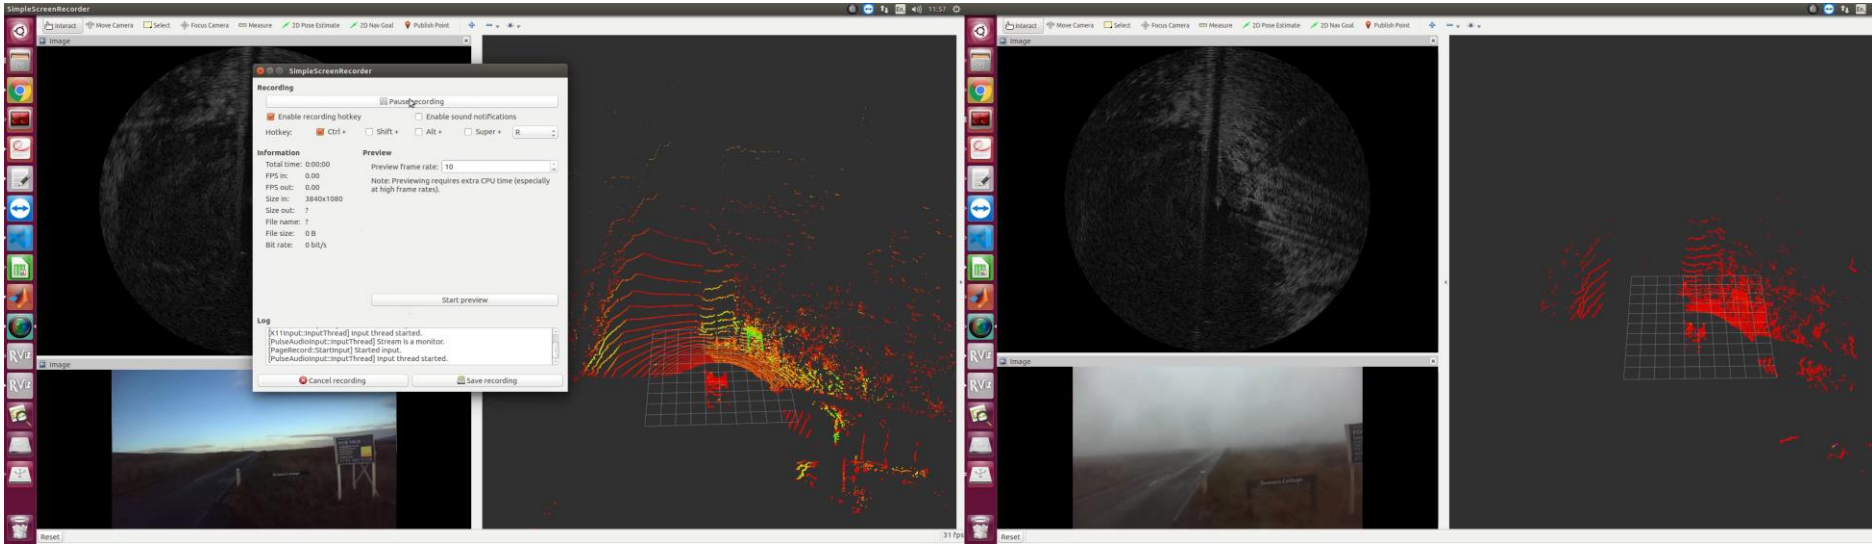

Vehicle detection in the wild (79GHz)

A lab study on object recognition (A-ConvNet, Transfer Learning, Data Augmention) in radar data showed high rates of accuracy for object recognition (80-90%) in 300GHz data (IET-RSN)

However, the wild data is challenging: A network was trained on a vehicle detection task, including cars, vans, trucks, buses, motorbikes and bicycles

It is important that motion is NOT used. Tracking and/or Doppler analysis can give better results BUT cars and pedestrians stop!

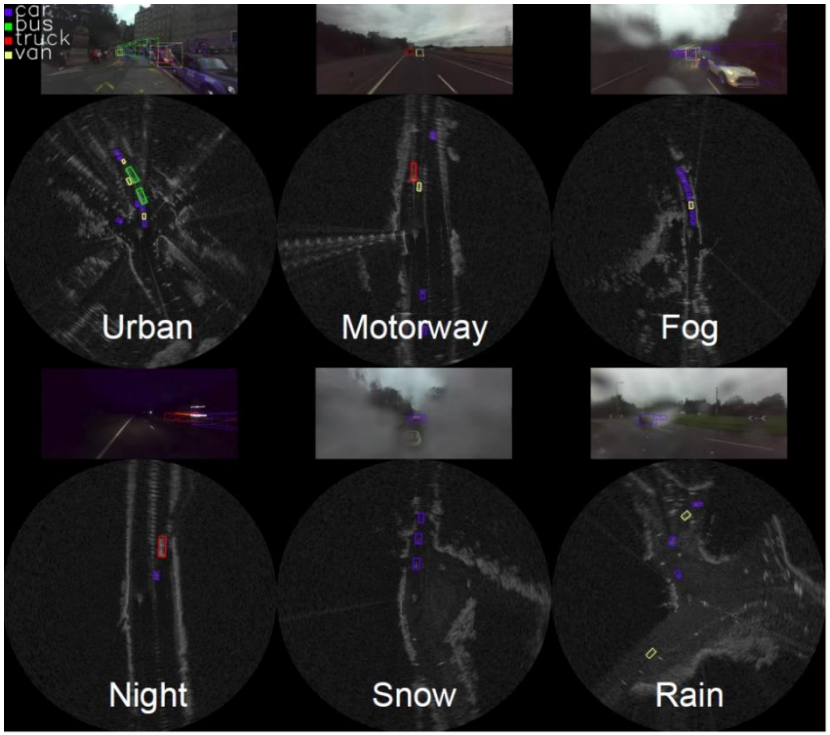

		Overall	Static	Motorway	Urban	Night	Rain	Fog	Snow
Radar	Faster RCNN ResNet-50 Trained on Good and Bad Weather	53.57	88.19	44.47	42.58	73.02	48.08	70.69	22.45
	Faster RCNN ResNet-50 Trained on Good Weather	52.77	88.08	49.03	35.05	64.03	43.50	62.02	27.63
	Faster RCNN ResNet-101 Trained on Good and Bad Weather	54.43	87.86	47.54	42.09	74.22	51.79	63.04	26.70
	Faster RCNN ResNet-101 Trained on Good Weather	52.90	87.98	46.44	36.26	64.40	42.51	56.99	17.77
Lidar	Faster RCNN ResNet-50 Trained on Good and Bad Weather	19.91	40.59	11.18	20.17	19.38	15.53	19.48	1.23
	Faster RCNN ResNet-50 Trained on Good Weather	17.49	40.87	11.03	17.35	14.13	9.31	15.02	1.82
	Faster RCNN ResNet-101 Trained on Good and Bad Weather	20.66	40.25	11.74	20.17	19.18	16.85	36.41	1.82
	Faster RCNN ResNet-101 Trained on Good Weather	17.81	41.80	10.85	20.84	13.39	16.17	13.07	1.30

Table 7: AP results on each scenario using rectangular bounding boxes.

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}}$$

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}}$$

$$\text{AP} = \int_0^1 p(r)\mathrm{d}r$$

Radar SLAM: This procedure (Zhiyang Hong and Sen Wang) using Graph Optimization and has been tested at night, in dense fog and heavy snowfall, in a GPS denied environment.

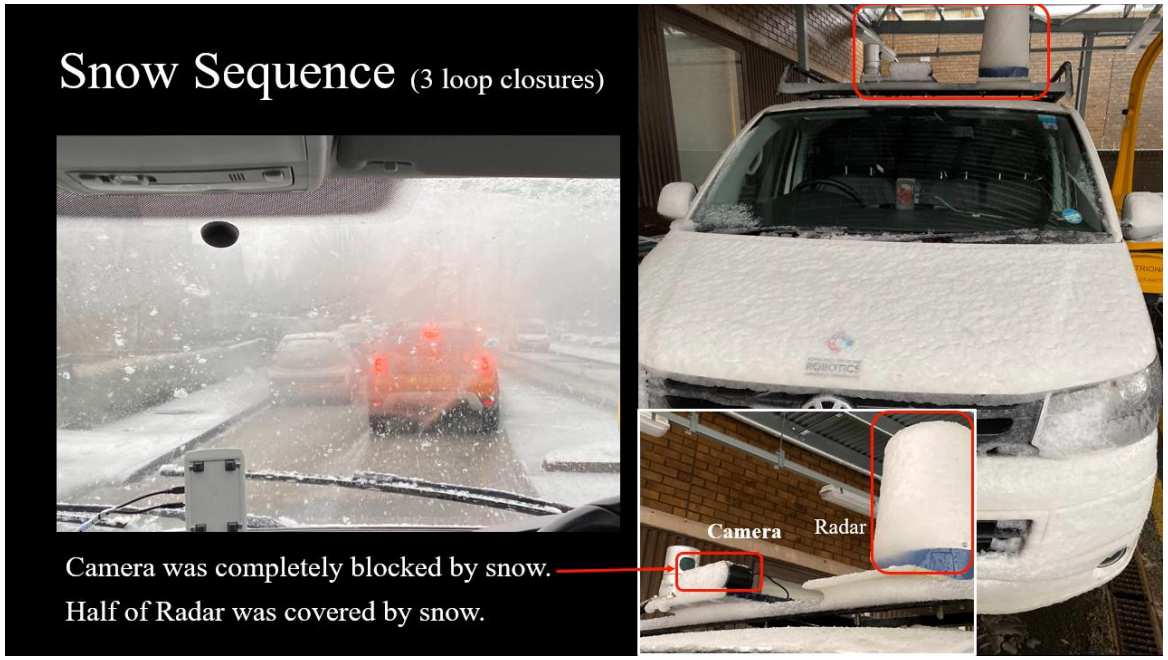

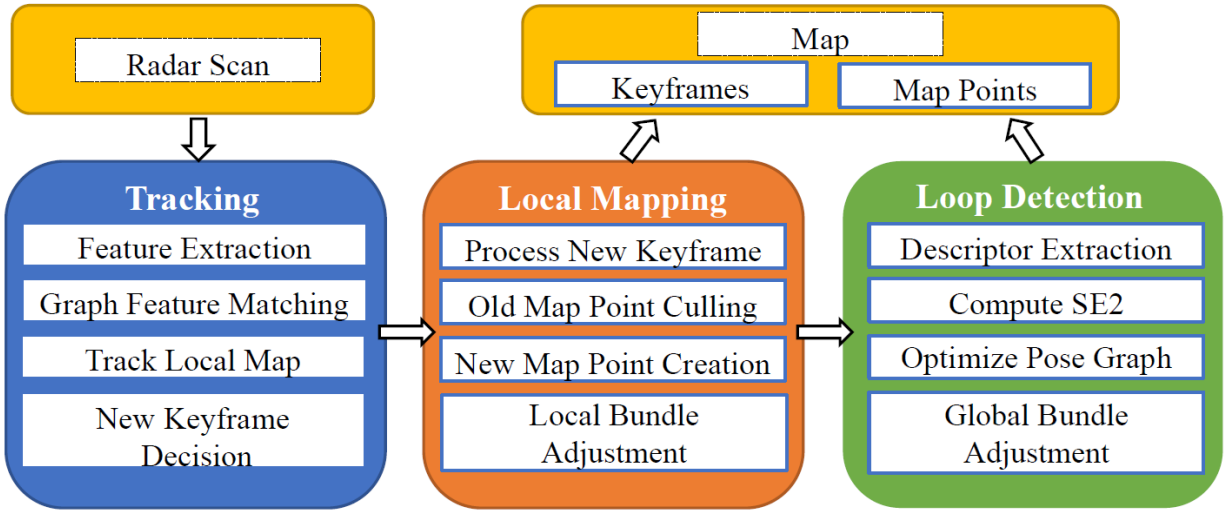

Predicting Future Behaviour

The intention is to predict the future motion (manoeuvre and trajectory) of detected vehicles

We used a benchmark (NGSIM) and our labelled wild radar data

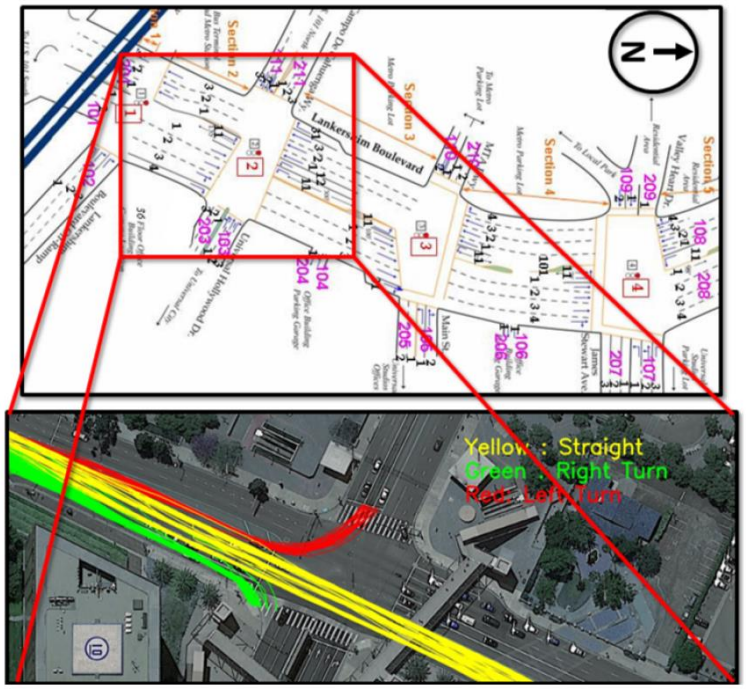

Fig. 3: Schematic diagram with junction and section locations of the NGSIM evaluation area, satellite view along with few sample trajectories on one selected junction, where yellow, green and red indicates vehicles performing straight, right and left maneuver at that junction, respectively.

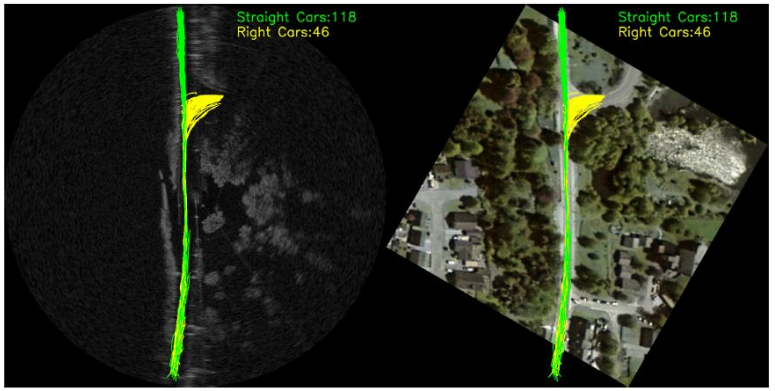

(a) Kingussie junction near Alvie Estate (T-junction)

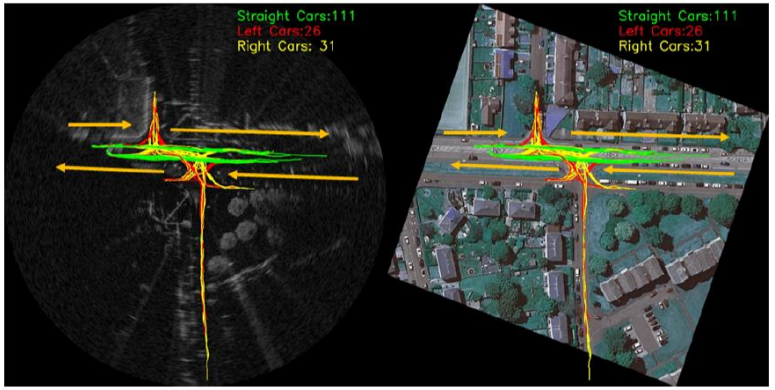

(b) Sighthill junction at Edinburgh city (four-way junction).

Video showing behaviour prediction (Sighthill): The network uses 30 previous frames to predict the next 30 frames on a rolling window

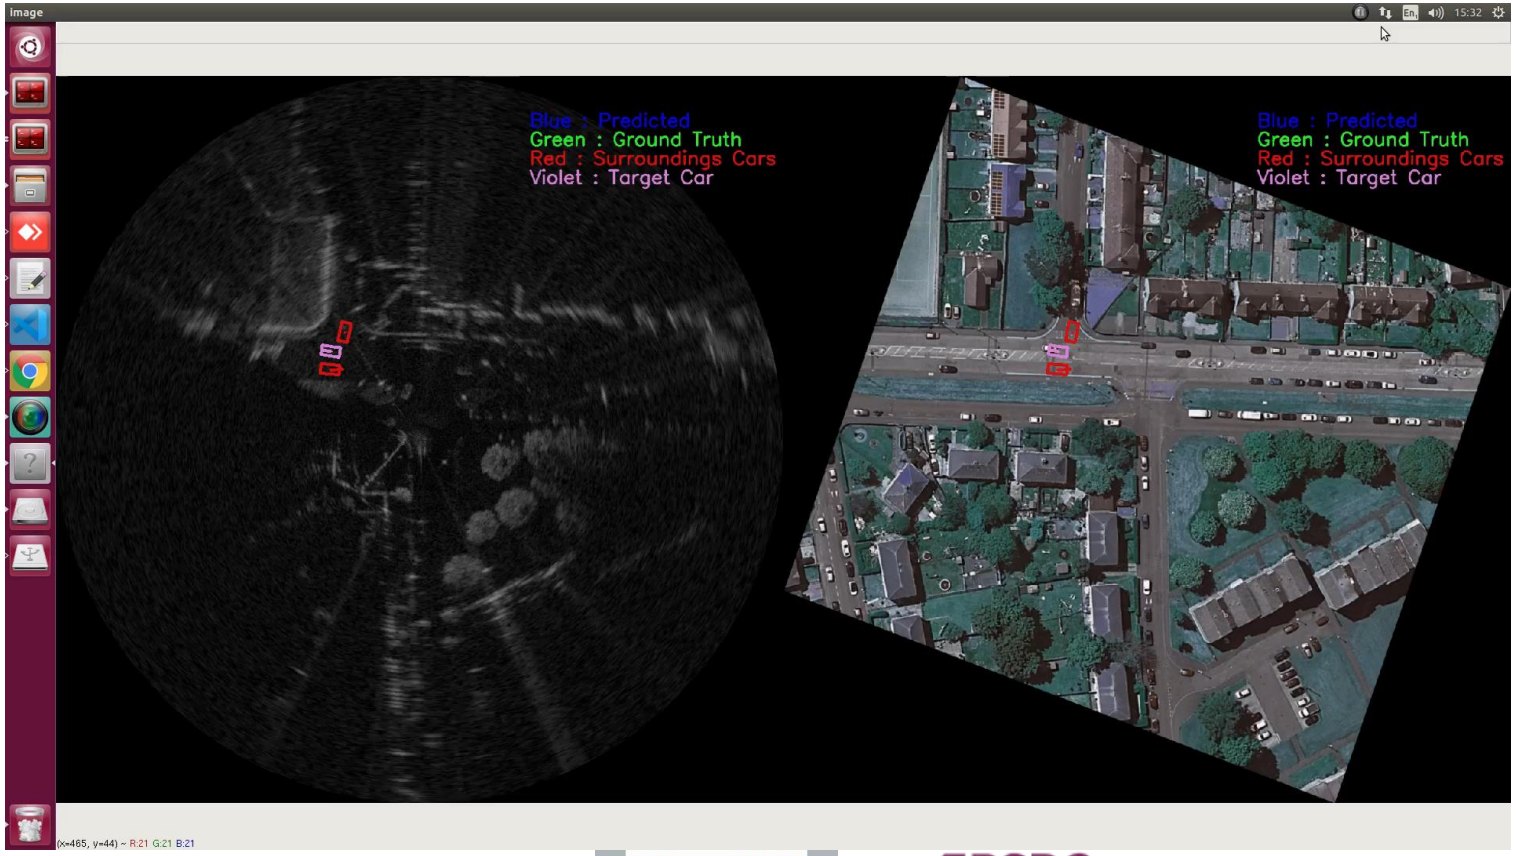

Predicting Future Behaviour

- We used various flavours of LSTM to encode the positions and velocities, both absolute and relative, of the surrounding vehicles, and some junction/lane data
- This was then used to predict future trajectories and manoeuvres within a fairly simple set of possibilities

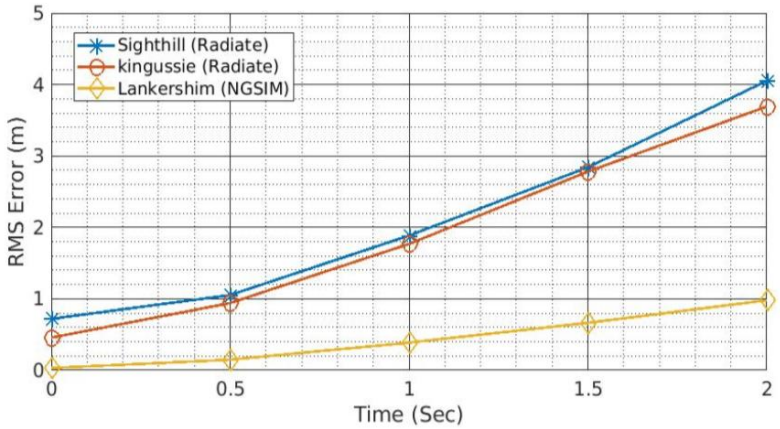

Fig. 7: The RMSE comparison of Lankershim (NGSIM), Kingussie (Radiate) and Sighthill (Radiate) dataset. This figure shows their average mean squared errors for the prediction time horizon from 1s to 2s.

- These results are comparable to SoA, but work is continuing

TABLE III: Confusion matrix for maneuver classification at Lankershim junction (NGSIM dataset).

		Actual Straight	Left	Right
Predicted	Straight	2346050	39300	35750
	Left	58050	962150	22950
	Right	30750	20700	350650

TABLE IV: Confusion matrix for maneuver classification at Kingussie junction (Radiate dataset).

		Actual Straight	Right	Left
Predicted	Straight	4755	6	–
	Right	4	3335	–
	Left	–	–	–

TABLE V: Confusion matrix for maneuver classification at Sighthill junction (Radiate dataset).

		Actual Straight	Left	Right
Predicted	Straight	2800	14	20
	Left	22	1400	8
	Right	46	10	1320

EPSRC
Engineering and Physical Sciences Research Council

Summary: interpreting radar data

- Recognising actors in radar data is possible, but far less accurate than video data (in good weather!), and there is considerable variation between different radar systems
- In bad weather, radar out-performs optical sensors, which fail rapidly
- The use of motion and/or background subtraction would improve results, but actors have 'difficult' motion patterns and backgrounds change
- Behaviour uses LSTM networks, which means that non-monotonic (e.g. stop-start) patterns may be used
- SLAM has been demonstrated in poor weather and GPS denied conditions, but it isn't clear how this would scale to larger networks.
- These components haven't been integrated: LSTM may give motion models for tracking, movement helps identify actors,

2. Lidar (and fusion) in challenging environments

The LiDAR problem in Obscuring Media

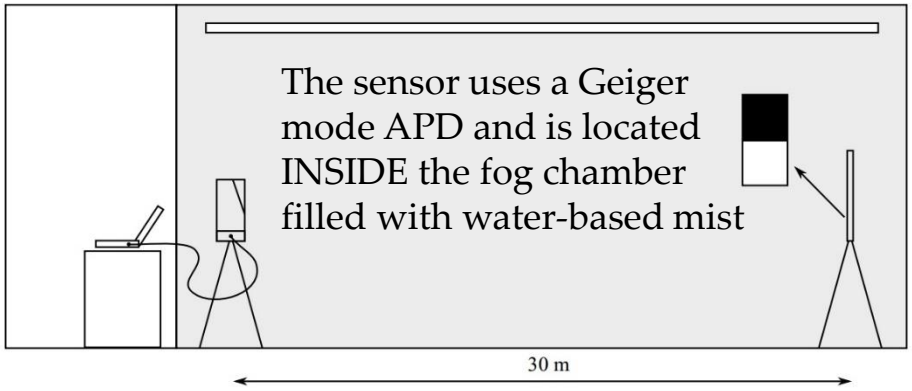

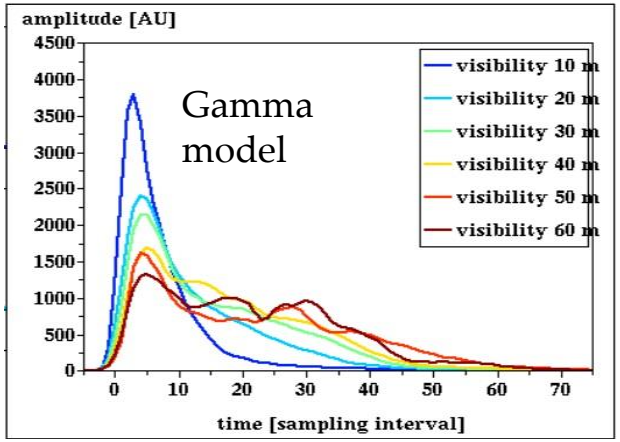

Mean waveforms, fog only

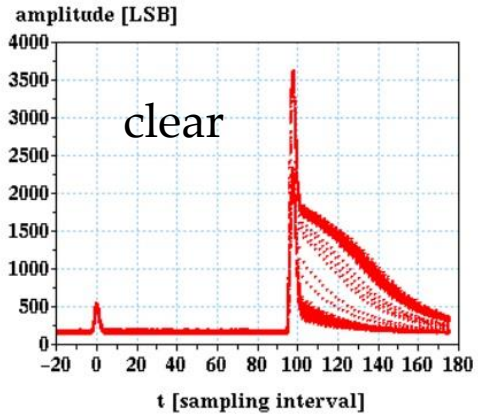

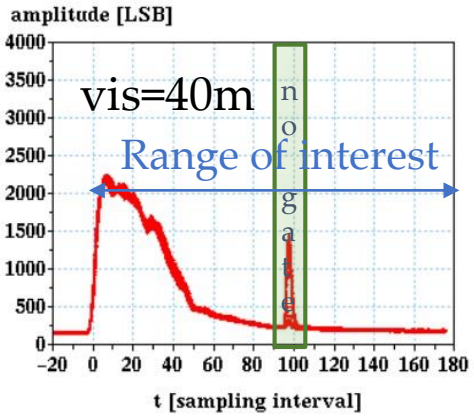

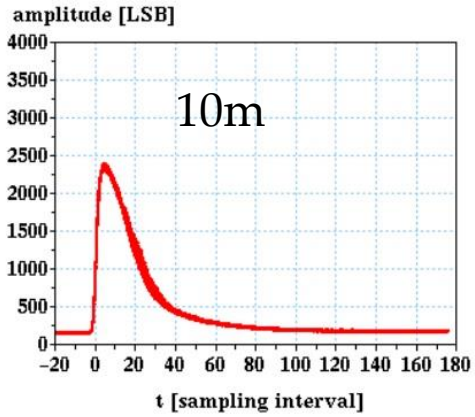

LiDAR returns from a Geiger mode APD FW sensor in a fog chamber from "Online waveform processing for demanding target situations", Pfennigbauer et al. SPIE, Vol. 9080 (see Carballo et al. for LiDAR dataset in 200m weather chamber that can house cars!)

dstl

EPSRC
Engineering and Physical Sciences Research Council

What do we know about LiDAR in and through obscurants?

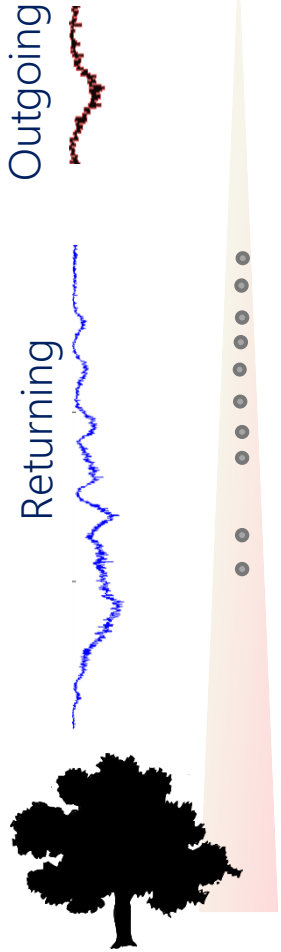

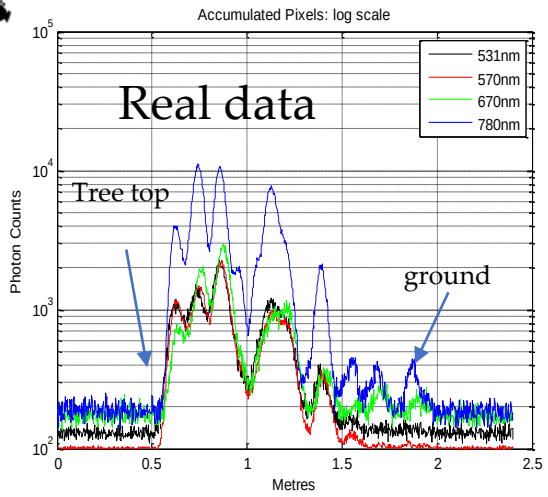

We can analyse the obscuring medium to measure bark and leaf area and NDVI

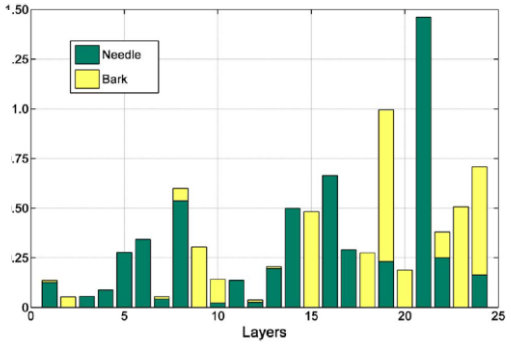

Fig. 15. Measured area profiles for leaf and bark as a function of canopy depth. These are measured as m^2/m^2 at the irregularly spaced layer positions shown in Fig. 6. This assumes a random distribution of areas in each layer so that subsequent layers have occluded material.

Synthetic data

We can detect vehicles or rogue vegetation below the tree canopy.

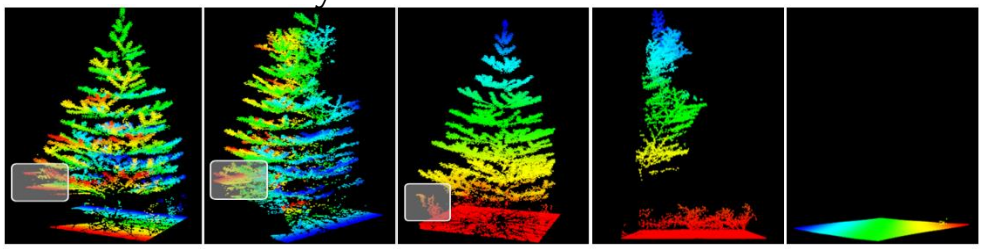

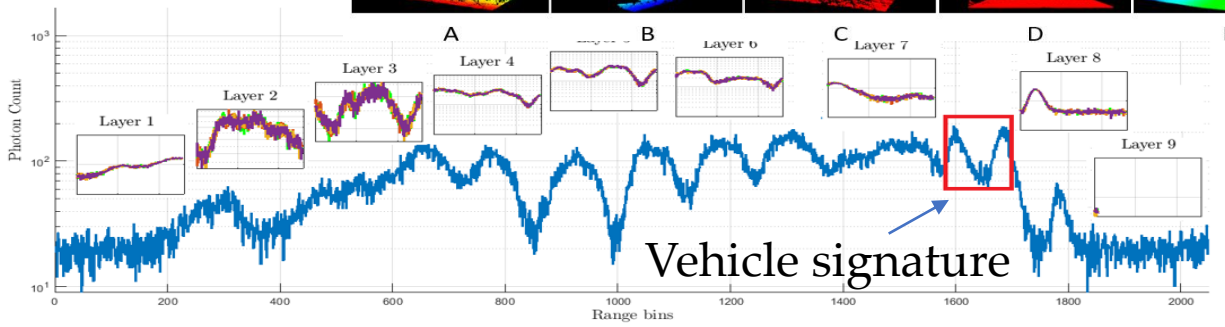

Multispectral!

F

Fusion: radar and LiDAR waveforms

- Most automotive radars have built in processing hardware/software and use relatively simple processing chains to detect one or two returns
- Through obscurants, we need full waveform LiDAR, and to recover weak signals "buried" in or behind the obscurant, and throughout the operational range (e.g. 1-200m)
- We use the radar waveform as a prior to improve the surface mapping
- An MCMC (or RJMCMC) chain explores the log-likelihood function

$$LL_r(y|k,\phi) = \sum_{i=1}^{i_{max}} -F(i;k,\phi) + y_i ln(F(i;k,\phi)) + ln(F(r)) \quad (5)$$

$$F(i;k,\phi) = \sum_{j=1}^{k} \beta_j f_{signal}(t_j) + \beta_{fog} f_{fog} + B_{bg}$$

where F(r) is either an exponential of the radar signal, or a normalized binary detection (CFAR)

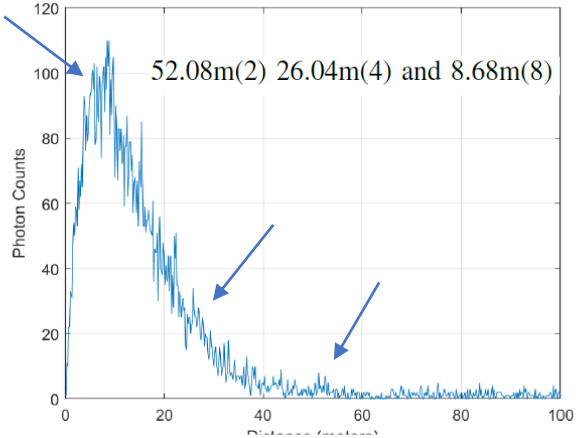

Simulations

Simulation: Results

	With radar	Without Radar
True Detections	129	99
% No Detections	0	3
% Single Detections	2	33
% Dual Detections	60	45
% Triple Detections	38	18

TABLE I: A true detection is classed as one within $0.25m$ of the true position. The number of times a single chain finds no, one, two or all returns is also tabulated as a percentage.

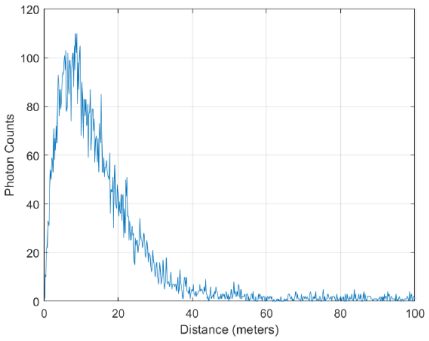

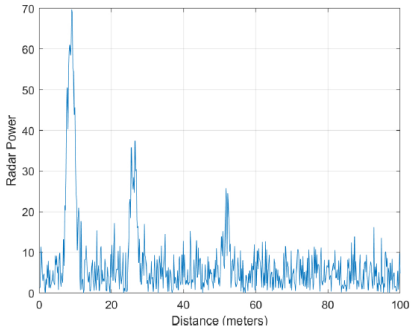

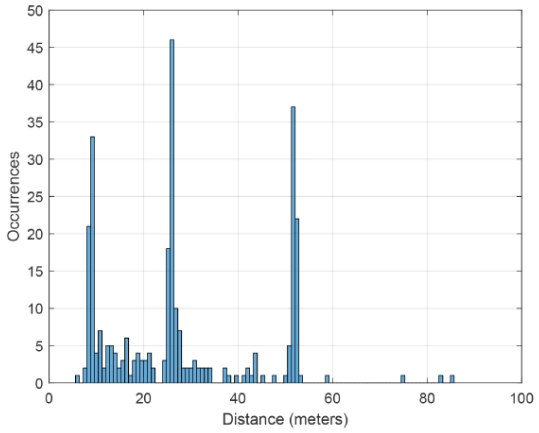

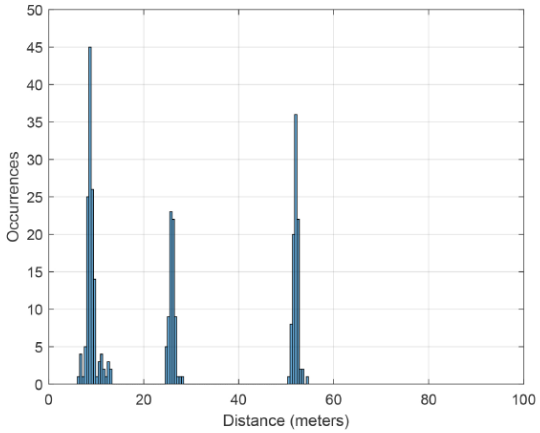

Fig. 5: Histogram of LiDAR signal detections without radar : Histogram of LiDAR signal detections with radar prior prior

Some results with real radar (CFAR) and LiDAR data, but simulated fog

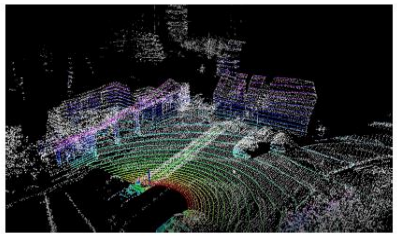

Fig. 7: Accumulated LiDAR 3D points using vehicle motion. The wall of the building in front and parked cars are clearly visible

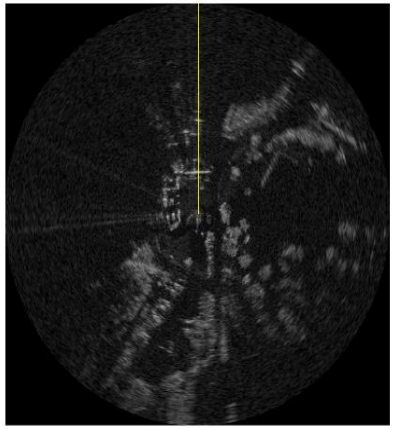

Fig. 8: The radar image; again the wall and cars are visible. The yellow line denotes the line of sight for the real radar power waveform in Fig. 9

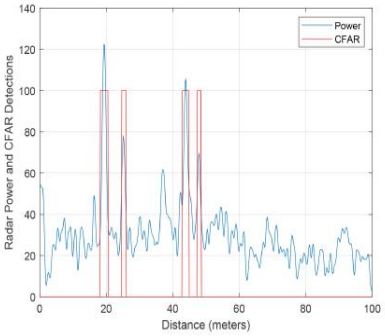

Fig. 9: Radar power spectrum for Navtech image captured at the location shown in Fig. 6. A result from CFAR detection (guard cells=10, training cells=300 false alarm rate =0.1) is also shown.

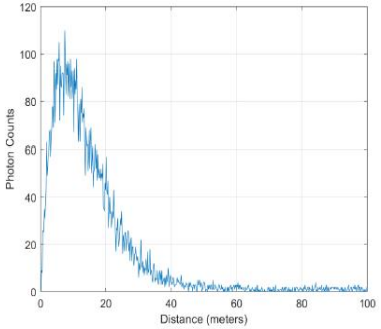

Fig. 10: LiDAR magnitude Spectrum corresponding to Navtech image. This uses real 3D data but generates waveform synthetically. The LiDAR return is at 19.44 meters.

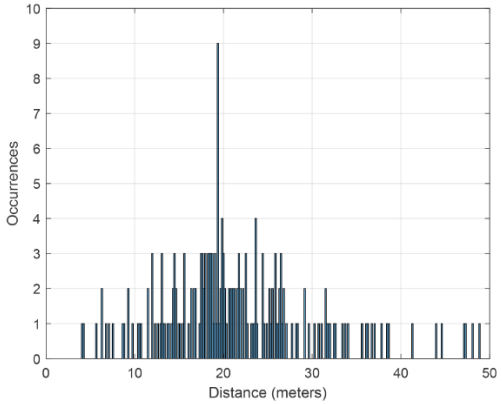

Fig. 12: Histogram of signal detections without radar waveform

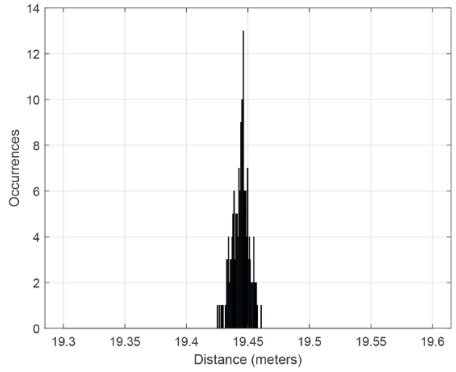

Fig. 11: Histogram of signal detections with radar waveform

dstl

EPSRC
Engineering and Physical Sciences Research Council

What happens when the medium is inhomogeneous?

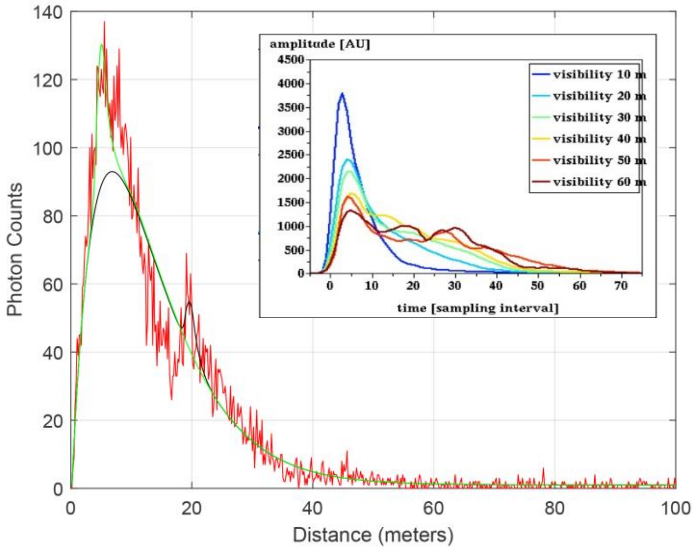

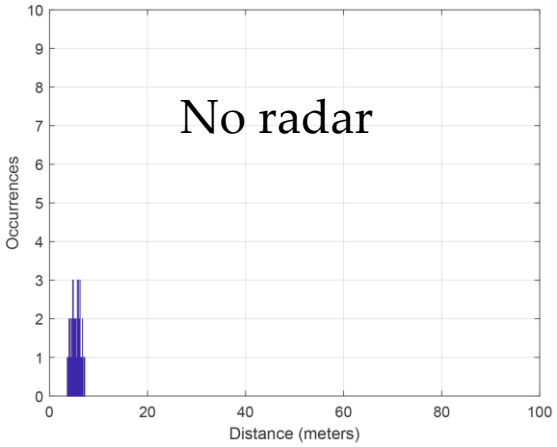

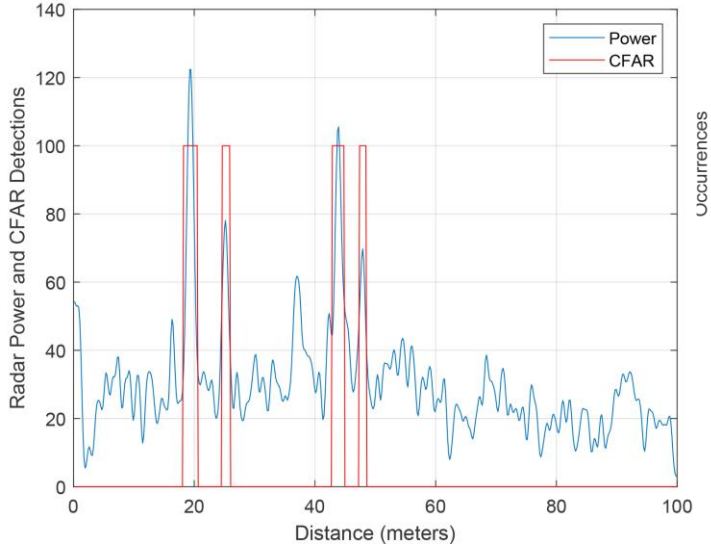

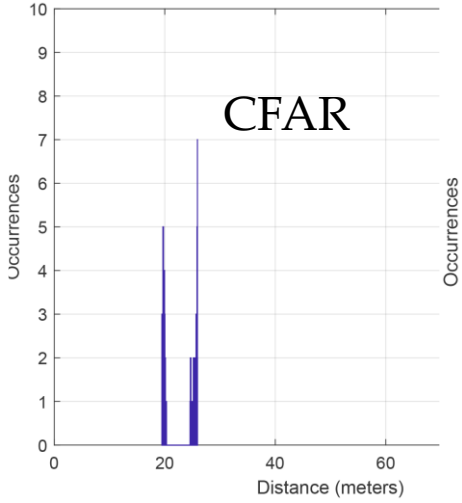

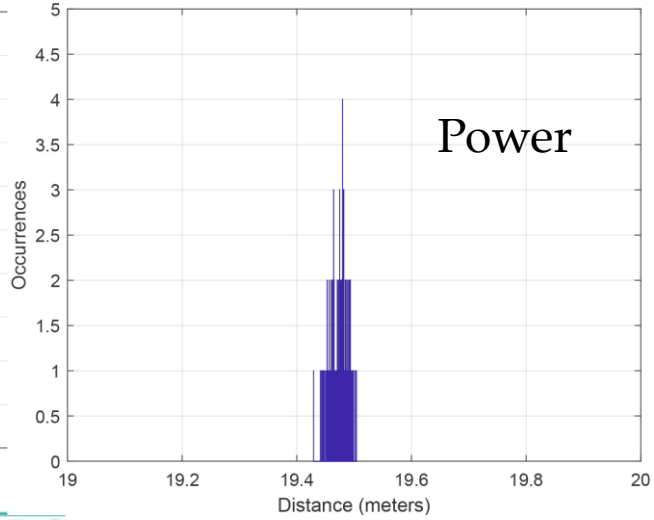

EPSRC
Engineering and Physical Sciences Research Council

Fusion of sparse LiDAR and dense camera data by convex optimization

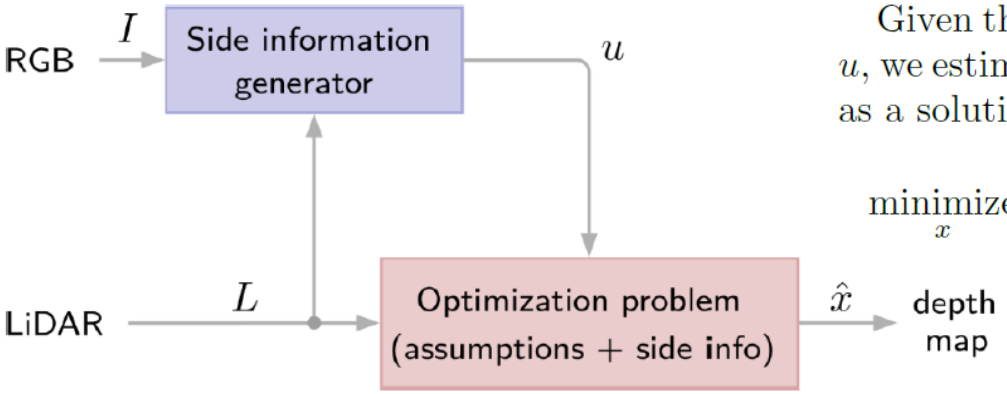

Given the sparse LiDAR map L and the side information u, we estimate the dense map of the scene, denoted $\hat{x} \in \mathbb{R}^n$, as a solution of

$$\underset{x}{\text{minimize}} \ \frac{1}{2}\|Sx - b\|_2^2 + \beta\|WHx\|_1 + \gamma\|Dx - u\|_1 \quad (1)$$

Fig. 1. Diagram of our method. Given an RGB image I and LiDAR data L, it generates a vector u that is used as side information in the optimization problem. The output is a dense depth map $\hat{x}$.

In adverse conditions, even with penetrative LiDAR, the camera image disappears, so could one use the full radar image instead to regularise the LiDAR?

[dstl]

EPSRC
Engineering and Physical Sciences Research Council

Summary: Key issues for automotive LiDAR in obscuring media

(see Wallace,Halimi, Buller, "Full Waveform LiDAR for Adverse Weather Conditions", IEEE Trans. on Vehicular Technology, , 69(7), 7064-77, 2020.)

1. Develop solid state arrays, but we need larger arrays (or mosaics) with better fill factors and better efficiency
2. Move to 1550nm. This better penetrates several common obscuring media, and allows higher, eye safe powers but is costly in terms of sensor development
3. Use full waveform analysis. This is essential to analyse the medium as well as the target, but increases data storage and/or algorithmic complexity.
4. Use sensor fusion as appropriate, at the full waveform data level but this has to be adaptive to conditions and true to physical models, supplemented by learning
5. For a full image, sub-sample, use spatial constraints, and include frame to frame processing. This also aids eye safety and reduces storage costs. However, the algorithms must be fast and have realistic constraints.
6. Parallel processing (SIMD or MIMD): this can lead to faster frame rates. If possible use dedicated hardware (ASIC, FPGA) but this can freeze or lead to non-optimal algorithmic development.
7. Avoid interference and malicious attack. Many LiDARS operating in the same space do interfere. More complex coding strategies may conflict with full wave analysis and lead to slower frame rates.

3. Forward looking SAR to improve mobile radar resolution

Radar imaging the road ahead: Forward looking SAR

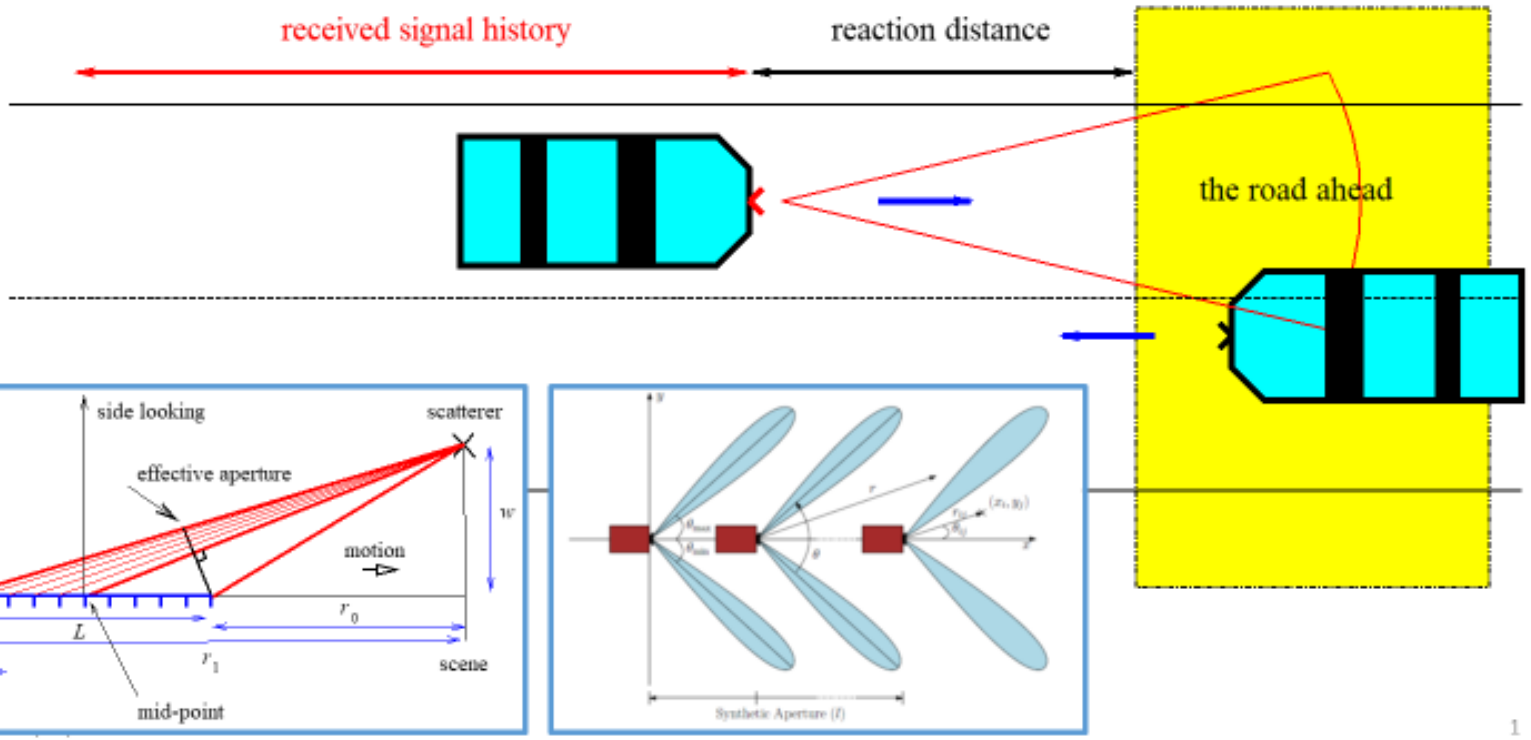

Practical Experiments

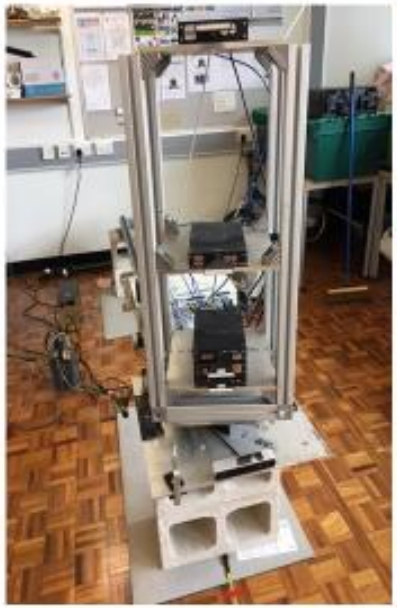


Forward looking SAR

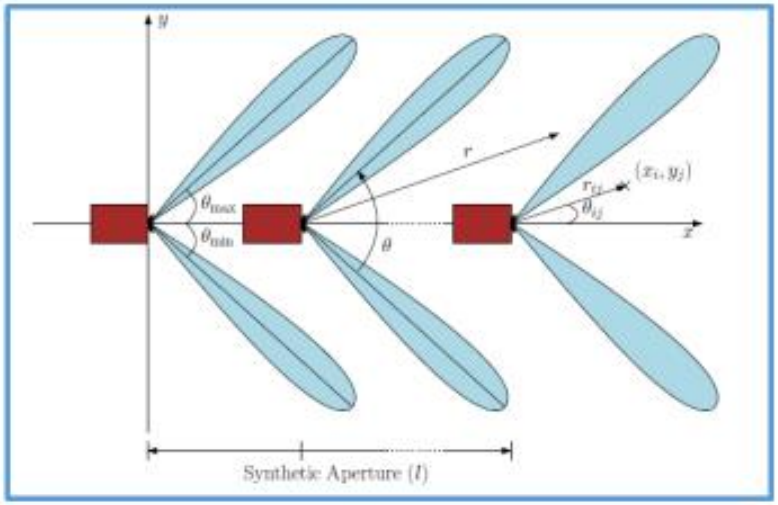

Raw radar data, fixed objects to left and right of scene, blurred moving objects in middle.

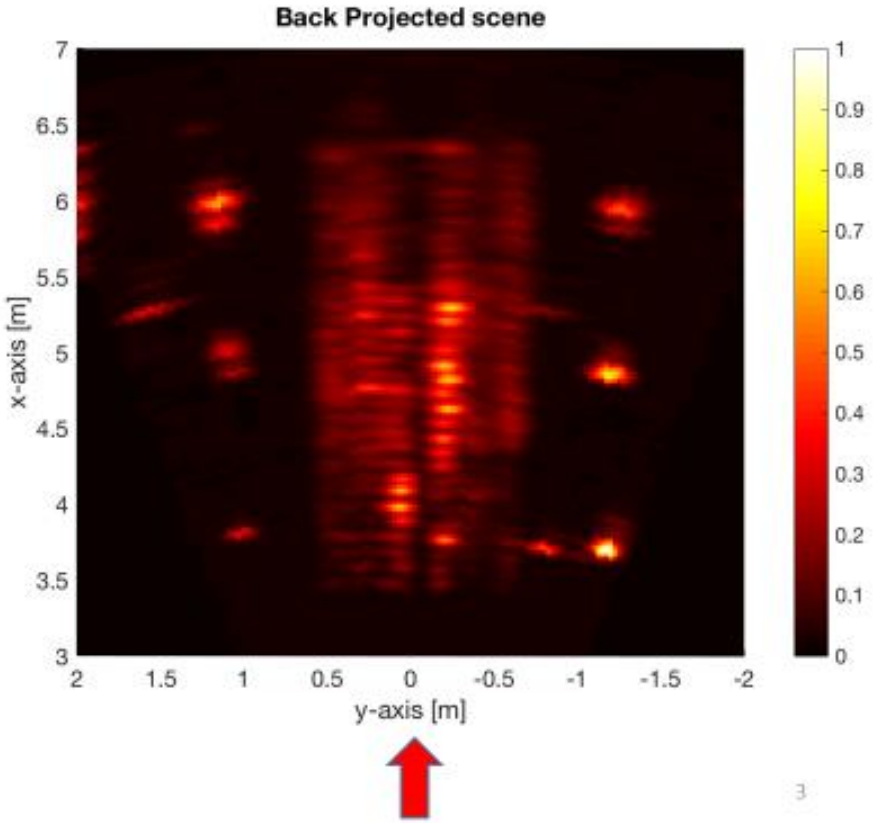

[dstl]

EPSRC
Engineering and Physical Sciences Research Council

Forward looking SAR

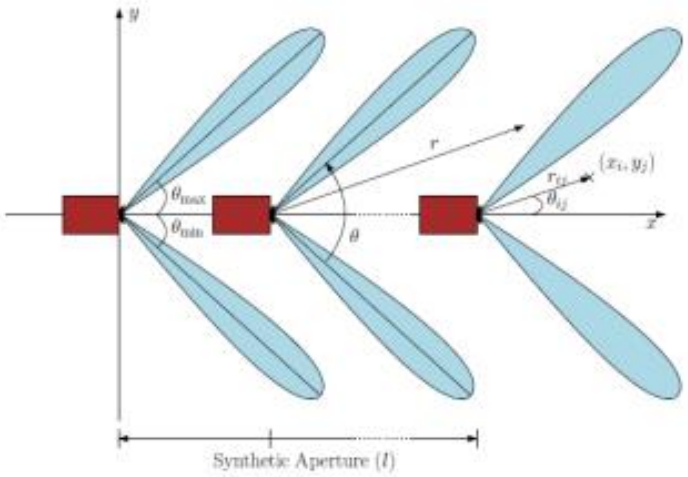

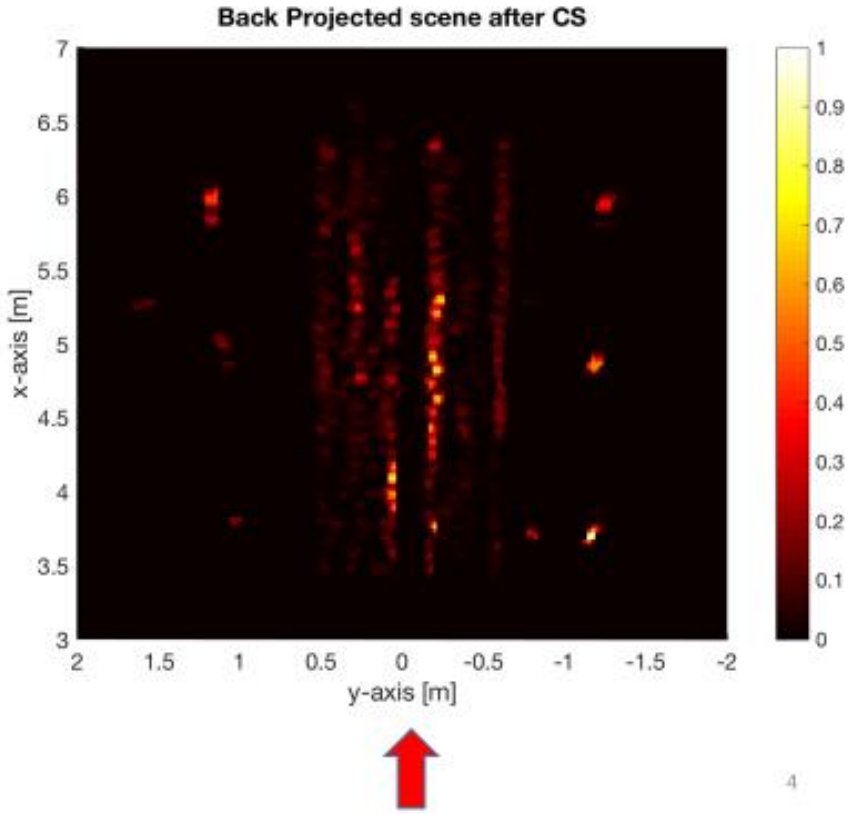

High resolution image of fixed objects, blurred moving objects ahead

[dstl]

EPSRC
Engineering and Physical Sciences Research Council

Forward looking SAR

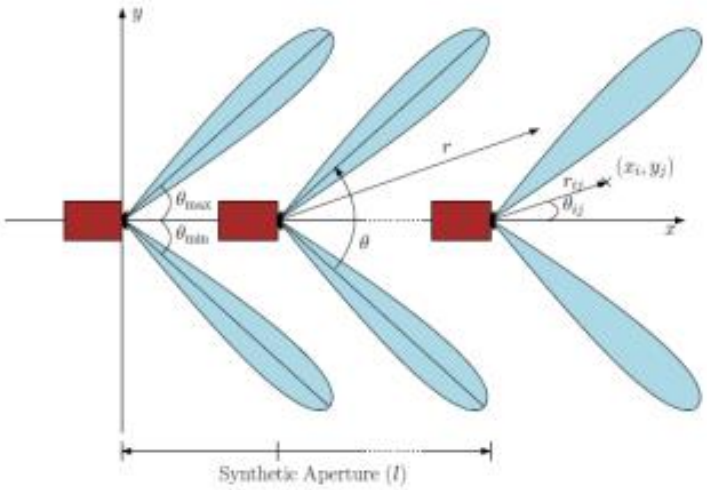

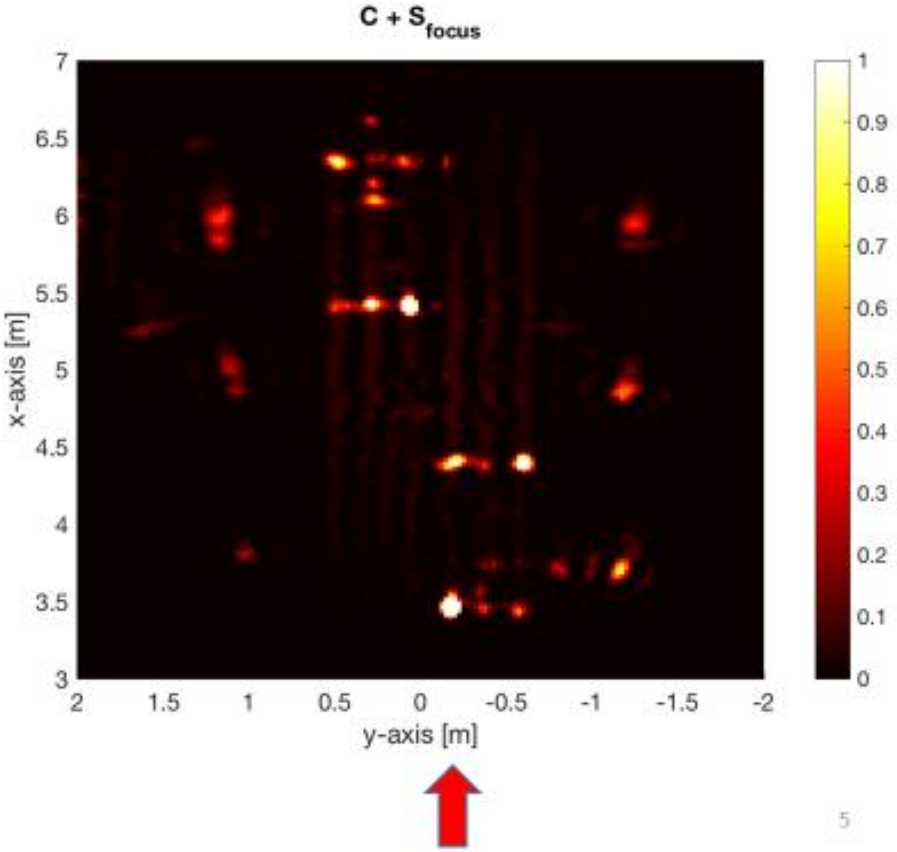

High resolution image of fixed objects, **high resolution image of moving objects** ahead

[dstl]

EPSRC
Engineering and Physical Sciences Research Council

Summary

- A forward scanning (FS) SAR methodology that combines scene scanning with synthetic aperture.
- Two imaging algorithms, modified back projection and compressed sensing based back projection.
- A non-parametric matrix decomposition approach to separate the dynamic part of the image from the clutter.
- An ADMM based iterative method to solve the optimisation problem.
- To focus the image, i.e., to build the synthetic aperture, we have proposed an extended DBSCAN method to achieve spatial segmentation along with the application of cross-correlation maximisation.

References

References: Radar image data analyis

M Sheeny, AM Wallace and S Wang, "300 GHz Radar Object Recognition based on Deep Neural Networks and Transfer Learning", IET Radar, Sonar and Navigation, Online, https://digital-library.theiet.org/content/journals/10.1049/iet-rsn.2019.0601

M Sheeny, AM Wallace and S Wang "RADIO: Parameterized Generative Radar Data Augmentation for Small Datasets", Applied Sciences, Volume 10; May 2020

S Mukherjee, S Wang and AM Wallace, "Interacting vehicle trajectory prediction with convolutional recurrent neural networks", ICRA, May 2020

Ziyang Hong, Yvan Petillot and Sen Wang. RadarSLAM: Radar based Large-Scale SLAM in All Weathers. arXiv:2005.02198

M. Sheeny, E de Pellegrin, S Mukherjee, A. Ahrabian, AM Wallace, S. Wang, "RADIATE: A Radar Object Recognition Dataset for Autonomous Cars in Adverse Weather", under review, 2020

References: Lidar and challenging environments

A Aßmann, B Stewart, J. Mota and AM Wallace, "Compressive super-pixel LiDAR for high frame rate 3D Depth Imaging" IEEE Global Conference on Signal and Image Processing, November, 2019

AM Wallace, A Halimi and GS Buller, "Full Waveform LiDAR for Adverse Weather Conditions", IEEE Transactions on Vehicular Technology, , 69(7), 7064-77, 2020.

AM Wallace, A McCarthy, C Nichol, X Ren, S. Morak[2], D Martinez-Ramirez, I. H. Woodhouse and GS Buller, "Design and of Evaluation of Multi-spectral LiDAR for the Recovery of Arboreal Parameters" IEEE Transactions on Geoscience and Remote Sensing, 52(8), 4942-4954, 2014.

A Ahrabian, JFC Mota and AM Wallace "Image-Guided Depth Estimation From Sparse Measurements using First- and Second-Order TV Minimization", under review

AM Wallace, A. Ahrabian, S. Mukherjee and T. Bemsibom "Detecting LiDAR Signals in Fog Using a Radar Prior", under review

References - Forward looking SAR and DBS

- Daniel, L.; Stove, A.; Hoare, E.; Phippen, D.; Cherniakov, M.; Mulgrew, B. & Gashinova, M. (2018), 'Application of Doppler beam sharpening for azimuth refinement in prospective low-THz automotive radars', *IET Radar, Sonar Navigation* **12**(10), 1121-1130.

- Gishkori, S.; Daniel, L.; Gashinova, M. & Mulgrew, B. (2019), 'Imaging for a Forward Scanning Automotive Synthetic Aperture Radar', *IEEE Transactions on Aerospace and Electronic Systems* **55**(3), 1420-1434.

- Gishkori, S.; Wright, D.; Daniel, L.; Gashinova, M. & Mulgrew, B. (2020), 'Imaging Moving Targets for a Forward-Scanning Automotive SAR', *IEEE Transactions on Aerospace and Electronic Systems* **56**(2), 1106-1119.