

OORD: The Oxford Offroad Radar Dataset

Matthew Gadd, Daniele De Martini, Oliver Bartlett, Paul Murcutt, Matt Towlson, Matthew Widodo,
Valentina Muşat, Luke Robinson, Efimia Panagiotaki, Georgi Pramatarov, Marc Alexander Kühn,
Letizia Marchegiani, Paul Newman, and Lars Kunze

Oxford Robotics Institute (ORI), University of Oxford, UK

✉ mattgadd@robots.ox.ac.uk  [oxford-robotics-institute.github.io/oord-dataset](https://github.com/oxford-robotics-institute/oord-dataset)
😊 huggingface.co/mttgdd/oord-models  github.com/mttgdd/oord-dataset

Abstract—There is a growing academic interest as well as commercial exploitation of millimetre-wave scanning radar for autonomous vehicle localisation and scene understanding. Although several datasets to support this research area have been released, they are primarily focused on urban or semi-urban environments. Nevertheless, rugged offroad deployments are important application areas which also present unique challenges and opportunities for this sensor technology. Therefore, the *Oxford Offroad Radar Dataset (OORD)* presents data collected in the rugged Scottish highlands in extreme weather. The radar data we offer to the community are accompanied by GPS/INS reference – to further stimulate research in radar place recognition. In total we release over 90 GiB of radar scans as well as GPS and IMU readings by driving a diverse set of four routes over 11 forays, totalling approximately 154 km of rugged driving. This is an area increasingly explored in literature, and we therefore present and release examples of recent open-sourced radar place recognition systems and their performance on our dataset, including a learned neural network.

I. INTRODUCTION

Compared to cameras and LiDARs, radar systems offer distinct advantages, including extended-range perception and resilience to lightning and various weather conditions. Scanning radar-based motion estimation [1]–[8], pose estimation [9]–[11], place recognition [9], [12]–[21], and full simultaneous localisation and mapping (SLAM) [22], [23] are increasingly popular in the literature. This dataset is designed to support that work, and to prompt an extension into naturalistic, rugged environments, which are underexplored. These natural environments are important because off-road autonomous vehicles have many essential applications across industries operating on rugged terrain. This technology is essential for tasks such as agricultural operations, search and rescue missions, and exploration in remote areas, where human intervention may be impractical or hazardous.

In deploying autonomous machines to these sites, place recognition is vital for navigation and localisation. For this, more robust sensors are crucial to enhancing place recognition accuracy and reliability. Radar in particular can handle challenging conditions like low-light environments. We therefore present an offroad radar dataset with a focus on relocalisation, with Fig. 1 showing an example location from our collection site. Our contributions are as follows:

- 1) **Offroad scanning radar data** We release the first radar dataset focusing on off-road, difficult terrain and naturalistic environments,



Fig. 1: Our data collection site is in the Scottish Highlands, with comprehensive coverage of *Ardverikie Estate*, close to the historic boundary between *Lochaber* and *Badenoch*. This image is taken on *Lochan na h-Earba* from our (Dataset 10) and (Dataset 11) datasets (Sec. III), showing unpaved terrain over uneven landscape next to the vehicle and the inclement weather. Images, GPS traces, and example radar scans for more specific areas of *Ardverikie Estate* are provided in Fig. 3.

- 2) **Poor weather and lighting conditions** including collection after thick snowfall and in total darkness (being in the wilderness),
- 3) **Place recognition ground truth** Our raw data is referenced against good GPS as a ground truth for the place recognition task,
- 4) **Open-sourced radar place recognition implementations** For the first time, we release open-source a set of trained weights for a deep neural network which solves the radar place recognition,
- 5) **Software development kit** We release software tools for making the use of this dataset in some pre-existing, non-learned open-source radar place recognition implementations easy, as well as pipelines to accelerate new learning solutions.

II. RELATED WORK

Radar datasets Datasets featuring a similar class of millimetre-wave scanning radar include *Marulan* [24] by Peynot *et al*, which is furthermore also collected in natural environments, as is ours. This data, however, only has range of up to 40 m. Several recent radar datasets use the same radar class – from the same manufacturer¹ – as used in our work, including the *Oxford Radar RobotCar Dataset* [25] by Barnes *et al*, *MulRan* [12] by Kim *et al*, *RADIATE* [26] by Sheeny *et*

¹Navtech Radar: navtechradar.com

al, *Boreas* [27] by Burnett *et al*, and *OSDaR23* [28] by Tagiew *et al*. The *Oxford Radar RobotCar Dataset* [25] is purely urban and features many repeat traversals, useful in investigating place recognition. The ground truth provided in [25], however, is prepared for radar odometry rather than place recognition. *MulRan* [12] presents more diversity in scenery, having been collected in several distinct parts of Daejeon, South Korea. We similarly attend to generalisation-sensitive training/testing data and collect at many sites, albeit offroad and more widespread across the countryside rather than a city. Our repeat traversals are also more numerous than for *MulRan*. *RADIATE* [26], presenting vehicle detection and tracking ground truth, is similarly lacking in overlapping traversals and not geared towards place recognition, although a diversity in weather makes this dataset interesting for exploring this sensor’s robustness. Closest to our work is *Boreas* by Burnett *et al* [27], who collect multi-season data including a *Navtech CIR304-H* scanning radar. Our work is in part inspired by this, with ours featuring highly non-planar offroad driving as opposed to purely urban. We present data over 154 km of driving. For comparison, *MulRan* – also focused on global localisation – presents approximately 124 km of driving data (Table III in [12]).

Radar place recognition Owing to its inherent resilience in sensing capabilities, radar has attracted considerable attention for its adeptness in the face of adverse weather, demonstrating a capacity to navigate and operate effectively in inclement conditions. This is because radar systems typically use electromagnetic waves in the radio or microwave frequency range. Unlike visible light, these waves can penetrate through various weather elements like rain, snow, fog, and clouds without significant distortion or attenuation. Radar place recognition has thus been explored in [9], [13], [14], [16], [17], [21], [29]–[33]. This dataset was designed to support radar place recognition in off-road environments, where place recognition with camera and LiDAR in off-road environments has been investigated in e.g. [34]–[38].

III. OFF-ROAD ROUTES & COLLECTION SITES

Fig. 3 shows detailed examples of one of the four routes we collected data over, at the *Ardverikie Estate*. Data has been gathered in the region encompassing *Loch Laggan* and *Lochan na h-Earba*. The position of these lochs can be particularly well understood by referring to the caption of Fig. 3.

Loch Laggan, a freshwater loch located about 10 km west of *Dalwhinnie* in the Scottish Highlands, extends in a nearly northeast to southwest direction for approximately 11 km. The name *Lochan na h-Earba* actually refers to two lochs positioned south of *Loch Laggan*. These lochs are situated in a slender glen that runs from southwest to northeast, running roughly parallel to *Loch Laggan*. The *Binnein Shuas* range of hills separates them from *Loch Laggan*.

Fig. 3 is arranged as follows: *Top Left*: Photograph of dataset collection location (handheld camera, not on-board imagery). *Top Middle*: GPS trace of route driven. *Top Right*: Ground truth match matrix for a pair of trajectories at that site. *Bottom*: Example radar scans from various points along the route (Cartesian).

We release data from four areas of the *Ardverikie Estate*. These feature distinct landscape (therefore typical radar re-



Fig. 2: On-board camera images for each released foray in Sec. III. These images will be released alongside the radar scans, and are included here to give a sense of the inclement collection conditions.

turns) as well as driving conditions. We refer to the sites as *Two Lochs*, *Bellmouth*, *Maree*, and *Hydro*. Each of these is discussed below.

Route 1: Bellmouth This is a 8.98 km route. The route is completely landlocked, heading away from and then back towards *Loch Laggan*. The ascent and descent are reflected by a somewhat symmetrical elevation change – climbing for the first half of the route, and descending for the second half. The vehicle is always on a gravel track. We performed four outings along this route. The first two these outings are prior to any snow falling in the *Ardverikie Estate* area. The next two of these outings were after significant snow dump (Fig. 2(a)).

Route 2: Hydro This is a 13.65 km route. This route reaches the highest elevation of all of our collection sites – ascending a gravel track to maximum elevation and then descending along the same route. For parts of the route, we travel along a mountain river. The snow coverage is significant (Figs. 2(b) and 2(c)), but not as thick as for *Maree* below.

Route 3: Maree This is a 16.87 km route. The route is mountainous, in the elevated area to the south of *Loch Laggan*. There are two ascents and descents – with a descent halfway through the journey, and another at the end of the journey. The vehicle is primarily on a gravel track. Both outings feature heavy snow, with a completely covered track. In one outing the vehicle is driving in total darkness (with no artificial illumination beyond the vehicle’s spotlights as this is a very remote and isolated natural collection site). The outings are in the same direction, but within each outing itself we double back on parts of the route and return to the starting point along the same route we took to begin the route.

Route 4: Two Lochs This is a 20.62 km route across the entire estate. The route is alongside two lochs (*Loch Laggan* and *Lochan na h-Earba*), see Fig. 3 (*Top Middle*, where the lochs are on top of and the bottom of the image, respectively). Referring to the route in an anti-clockwise sense, while going alongside *Loch Laggan* the vehicle is on tarmac – it is otherwise on unpaved road for the rest of the route. The two outings on this route are completed *after* snow has fallen (in previous days), so melted somewhat, but while it is still very visible on the ground etc. The outings are in the opposite direction as shown by the off-diagonal matches in the ground truth matrix in Fig. 3 (*Top Right*).

Summary: Datasets provided In summary, the datasets provided are listed below, including the date-string for the collection time and the site that they were collected at.

- (Dataset 1): 2021-11-25-12-01-20 (Bellmouth)
- (Dataset 2): 2021-11-25-12-31-19 (Bellmouth)

- (Dataset 3): 2021-11-26-15-35-34 (Bellmouth)
- (Dataset 4): 2021-11-26-16-12-01 (Bellmouth)
- (Dataset 5): 2021-11-27-14-37-20 (Hydro)
- (Dataset 6): 2021-11-27-15-24-02 (Hydro)
- (Dataset 7): 2021-11-27-16-03-26 (Hydro)
- (Dataset 8): 2021-11-28-15-54-55 (Maree)
- (Dataset 9): 2021-11-28-16-43-37 (Maree)
- (Dataset 10): 2021-11-29-11-40-37 (Two Lochs)
- (Dataset 11): 2021-11-29-12-19-16 (Two Lochs)

Some of these are shown in Fig. 2. We note that, beyond the weather and illumination effects discussed for each route above, some forays have lens adherents on the on-board camera (e.g. Fig. 2(a) in particular), motivating the use of radar in this sort of domain and for the localisation task.

IV. SENSORS & DATASET FORMAT

Radar This sensor is the same as that used for the *Oxford Radar RobotCar Dataset* [25], i.e. a Navtech CTS350-X Millimetre-Wave frequency-modulated continuous wave radar, 4 Hz, 400 measurements per rotation, 163 m range, 4.38 cm range resolution, 1.8° beamwidth. These are the same settings as for the *Oxford Radar RobotCar Dataset* [25], which should support future investigations of domain differences complementary with respect to radar settings. Also, the majority of driven tracks in our new dataset are in densely forested areas with many trees within the 163 m range. The radar was mounted at the centre of the vehicle aligned to the vehicle axes. Radar scans are stored as lossless-compressed PNG files in polar form with each row representing the sensor reading at each azimuth and each column representing the raw power return at a particular range.

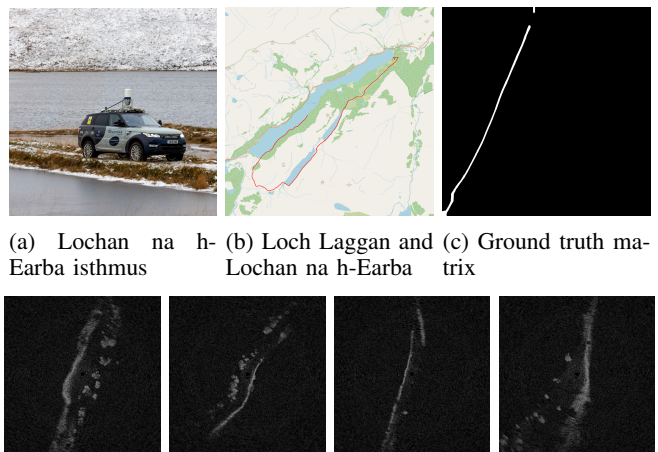
GPS/INS We use a Microstrain 3DM-RQ1-45 GPS/INS. This sensor has direct satellite and inertial measurements, and computes position, velocity, and attitude. To provide accurate readings, it uses a triaxial accelerometer, gyroscope, magnetometer, and temperature sensors, as well as a pressure altimeter. Furthermore, dual on-board processors run an Extended Kalman Filter (EKF).

Vehicle setup For evaluating place recognition (Sec. V below), the relative GPS/INS poses suffice. The radar mounted is mounted on the roof of the vehicle – but forwards, towards the front windshield. The GPS/INS is also mounted on the roof, but backwards, towards the rear windshield. The radar and GPS/INS data packets are timestamp-synchronised as per [39].

V. SOFTWARE & EXPERIMENTAL TOOLS

Downloads & data loaders Similarly to [25], we provide data loaders to work with the radar data format, which as per Sec. IV are released in raw polar format, but can be easily transformed into Cartesian format with our SDK. For working with driven sequences, we provide a basic iterable `Dataset` class which returns each radar scan paired with its *closest* (in a timestamp sense) GPS/INS reading (Sec. IV). We also provide a batch download script to rapidly access the data.

Radar place recognition evaluation We evaluate place recognition performance by comparing “distance matrices” in the ground truth space given by GPS, as shown by example



(a) Lochan na h-Earba isthmus (b) Loch Laggan and Lochan na h-Earba (c) Ground truth map (d) Ex. scan #1 (e) Ex. scan #2 (f) Ex. scan #3 (g) Ex. scan #4

Fig. 3: The *twolochs* dataset is the most extensive we provide, and is collected over a journey featuring (1) a tarmac road under tree cover alongside a *Loch Laggan* (top), (2) a gravel track over the shoulder of a mountain, (3) a partially submerged track fording a loch, (4) a gravel track alongside *Lochan na h-Earba*, and (5) a gravel track descending a mountainside to the start location.

in Fig. 3 (*Top Right*) to distance matrices composed of e.g. pairwise distances between vectors representing radar scans (see Sec. VI) or some other similarity measurement. These matrices have heights given by the number of radar scans in a “query” foray and widths similarly for a “reference” foray. Highlighted regions are such that those radar scan pairs have locations given by GPS which are within 25 m (this threshold being configurable in our SDK). Off-diagonal revisits (for e.g. *Maree*, not shown) represent traversals in the opposite direction. This is especially important in the radar modality as the sensor has a 360° field-of-view.

Then, $\text{Recall}@N$ refers to the percentage of query scans which result in *at least one* successful localisation (i.e. resulting in a returned location within 25 m of the query location) when accepting as candidates all of the N nearest neighbours in the embedding space. This is a commonly used metric and threshold, as in visual place recognition [40], [41]. Note that in Fig. 3 the true positives are “thickened” for visualisation purposes, to 100 m, but we experiment in Sec. VII with 25 m.

Radar place recognition with neural networks As will be discussed below, we evaluate a family of learned methods on our dataset. This includes a pretrained and frozen backbone feature extractor along with a pooling layer which is fine-tuned on radar data from this dataset, and we release these weights at huggingface.co/mttgdd/oord-models, with this being, to our knowledge, the first open-sourced learned radar place recognition system.

Open-source radar place recognition methods As below in Sec. VI, our SDK includes either (1) new implementations of or (2) imports, as submodules, of other open-source radar place recognition systems [12], [20], [33]. Experiments on novel systems can be quickly run over all pairs of trajectories discussed in Sec. III, and prepared in the format of Tab. I, by a common `.yaml` configuration format.

| | (Dataset 2)-vs-(Dataset 1) | (Dataset 5)-vs-(Dataset 6) | (Dataset 5)-vs-(Dataset 7) | (Dataset 8)-vs-(Dataset 9) | (Dataset 11)-vs-(Dataset 10) |
|--------------|------------------------------|------------------------------|------------------------------|------------------------------|------------------------------|
| (Baseline 1) | 94.29 / 97.53 / 98.24 | 91.15 / 93.57 / 93.78 | 94.06 / 95.80 / 96.00 | 91.76 / 97.22 / 98.42 | 72.96 / 82.28 / 84.71 |
| (Baseline 2) | 98.45 / 98.73 / 98.94 | 92.65 / 93.89 / 93.94 | 95.34 / 95.85 / 95.85 | 94.28 / 99.15 / 99.40 | 85.82 / 87.69 / 88.25 |
| (Baseline 3) | 98.87 / 99.79 / 99.93 | 93.51 / 94.10 / 94.16 | 96.00 / 96.16 / 96.16 | 96.88 / 99.06 / 99.32 | 84.27 / 87.36 / 88.25 |
| (Baseline 4) | 97.67 / 98.52 / 98.87 | 93.78 / 94.26 / 94.37 | 95.70 / 99.59 / 99.95 | 97.14 / 99.06 / 99.44 | 5.41 / 10.98 / 16.23 |
| (Baseline 5) | 90.49 / 97.25 / 98.10 | 92.76 / 93.83 / 93.83 | 94.72 / 95.85 / 95.85 | 89.11 / 97.91 / 98.72 | 4.53 / 10.10 / 15.51 |
| (Baseline 6) | 90.70 / 97.04 / 98.45 | 91.96 / 93.62 / 93.83 | 93.14 / 95.49 / 95.75 | 87.92 / 96.67 / 98.25 | 17.55 / 40.07 / 50.44 |
| (Baseline 7) | 94.50 / 98.38 / 98.80 | 92.28 / 93.83 / 93.83 | 93.70 / 95.54 / 95.80 | 90.61 / 97.27 / 98.42 | 23.18 / 47.19 / 58.22 |
| (Baseline 8) | 74.91 / 92.39 / 95.28 | 82.79 / 91.31 / 92.60 | 83.91 / 92.78 / 94.62 | 65.16 / 86.93 / 92.10 | 14.74 / 35.87 / 47.08 |
| (Baseline 9) | 72.73 / 90.49 / 94.71 | 81.29 / 90.40 / 91.96 | 84.07 / 93.14 / 94.31 | 69.43 / 87.15 / 91.97 | 9.71 / 21.80 / 31.29 |

TABLE I: Recall@1/5/10 localisation success rate (%) for all example methods defined in Sec. VI and discussed in Sec. VII over 5 pairs of trajectories from our dataset.

VI. EXAMPLES

- (Baseline 1): **RingKey** as a component of *ScanContext* [42] averages point cloud contents at a fixed distance around the vehicle (thus being rotation invariant). Note, for this we do not use the orientation refinement of *ScanContext*, as distances between row-reduced vectors and orientation scores are separate, i.e. these form a hierarchy of localisers. To be clear, this is different to (Baseline 2) below, where the maximum circular cross-correlation is the score used for similarity between scans (with no initial score used to reduce the candidate set).
- (Baseline 2): **RaPlace** [20] measures the similarity score between radar scans by using Radon-transformed sinogram images and cross-correlation in the frequency domain. This gives rigid transform invariance during place recognition, and suppresses the effects of radar multipath and ring noises.
- (Baseline 3): **Open-RadVLAD** [33] uses only the polar representation. Also, for partial translation invariance and robustness to noise, it uses only a 1D Fourier Transform along radial returns. It also achieves rotational invariance and a very discriminative descriptor space by building a vector of locally aggregated descriptors (VLAD).
- (Baseline 4): Here we train radar-specific neural network models – using **ResNet18** [43] as a backbone feature extractor and **NetVLAD** [40] as a pooling layer (with 64 clusters). Inputs are polar radar scans of 128×128 resolution and embeddings are 128-dimensional. The radar returns are replicated 3 times at the input in order to form a 3-channel image. We start with pretrained weights learned against **ImageNet** [44] and then, in an unsupervised learning step, initialise the **NetVLAD** cluster centres by k-means++ [45] over the deep features extracted from the reference trajectory radar frames.
- (Baseline 5), (Baseline 6), (Baseline 7), (Baseline 8), (Baseline 9): Here we show that the plug-and-play ability of our dataset with respect to widely available pretrained networks, to accelerate research in this area. We use specifically the `pytorch/vision:v0.10.0` models available at pytorch.org/hub as well as the `gmberton/cosplace` models available at github.com/gmberton/CosPlace. This covers **MobileNet** [46] (Baseline 6), **GoogleNet** [47] (Baseline 7), **AlexNet** [48] (Baseline 8), and **VGG19** [49] (Baseline 9), which are all trained for *ImageNet* classification and thus not specialised for place recognition but represent inputs in terms of wide variety of concepts and features.

It also covers **CosPlace (ResNet18-512)** [41] (Baseline 5) which is specifically trained for place recognition – but not specifically for radar, in contrast to (Baseline 4).

VII. RESULTS

As can be seen in Tab. I, considering that all of each of the routes is at minimum 9km in length, all radar-specific methods perform very well, with in excess of 90% for even Recall@1 i.e. when relying on only a single nearest neighbour in embedding space look up. Overall the best performing methods are (Baseline 3) **Open-RadVLAD** [33], (Baseline 4) **ResNet18-NetVLAD**, and (Baseline 2) **RaPlace** [20].

There is in fact a precipitous drop in localisation success rate for all methods in (Dataset 10)-vs-(Dataset 11) in `twolochs`. We attribute this to `twolochs` (Sec. III) being the longest route (there therefore being more candidate scans to match to – i.e. the search space being larger). In Fig. 3 this can be seen in the ground truth matrix (*Top Right*) by the large region to the right of the white strip, where the query trajectory ends at *Loch Laggan* and therefore only overlaps with the reference trajectory (*Middle Right* in Fig. 3) for approximately 40% of the route, with the reference trajectory continuing north-east all the way along both bodies of water in *Loch Laggan*. This may also be due to the long periods driven along either *Loch Laggan* or *Lochan na h-Earba* where there is a dearth of distinctive scenery in an entire hemisphere of the radar scan. An example scan from this type of scene can be seen in the third Cartesian frame in Fig. 3 (*Bottom*). Therefore, `twolochs` is a useful test setting for recognition with vast maps and with featureless scenery. For this challenging route, only radar-specific methods perform well and we have (Baseline 2) **RaPlace** performing best, with (Baseline 3) **Open-RadVLAD** matching it in Recall@10. (Baseline 4) **ResNet18-NetVLAD** performs poorly on this pair of trajectories, likely due to clustering of the reference trajectory features, many of which are irrelevant to the query trajectory.

VIII. CONCLUSION

We have presented a novel radar dataset, carefully gathered under unique and challenging conditions. This dataset has been crafted to catalyse advancements in the emerging field of radar place recognition. We explored the use of our dataset in this task over a series of comprehensive experiments and evaluations, carried out across various open-source radar place recognition systems. This was not only to showcase the dataset but also to establish a robust platform for future research in this evolving area.

REFERENCES

- [1] S. H. Cen and P. Newman, "Precise ego-motion estimation with millimeter-wave radar under diverse and challenging conditions," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 6045–6052.
- [2] D. Barnes, R. Weston, and I. Posner, "Masking by moving: Learning distraction-free radar odometry from pose information," *arXiv preprint arXiv:1909.03752*, 2019.
- [3] Y. S. Park, Y.-S. Shin, and A. Kim, "Pharao: Direct radar odometry using phase correlation," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 2617–2623.
- [4] P.-C. Kung, C.-C. Wang, and W.-C. Lin, "A normal distribution transform-based radar odometry designed for scanning and automotive radars," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 14 417–14 423.
- [5] K. Burnett, A. P. Schoellig, and T. D. Barfoot, "Do we need to compensate for motion distortion and doppler effects in spinning radar navigation?" *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 771–778, 2021.
- [6] D. Adolfsson, M. Magnusson, A. Alhashimi, A. J. Lilienthal, and H. Andreasson, "Cfear radarodometry-conservative filtering for efficient and accurate radar odometry," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 5462–5469.
- [7] K. Burnett, D. J. Yoon, A. P. Schoellig, and T. D. Barfoot, "Radar odometry combining probabilistic estimation and unsupervised feature learning," *arXiv preprint arXiv:2105.14152*, 2021.
- [8] R. Aldera, M. Gadd, D. De Martini, and P. Newman, "What goes around: Leveraging a constant-curvature motion constraint in radar odometry," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 7865–7872, 2022.
- [9] D. Barnes and I. Posner, "Under the radar: Learning to predict robust keypoints for odometry estimation and metric localisation in radar," in *International Conference on Robotics and Automation (ICRA)*, 2020.
- [10] H. Yin, Y. Wang, L. Tang, and R. Xiong, "Radar-on-lidar: metric radar localization on prior lidar maps," in *2020 IEEE International Conference on Real-time Computing and Robotics (RCAR)*. IEEE, 2020, pp. 1–7.
- [11] H. Yin, R. Chen, Y. Wang, and R. Xiong, "Rall: end-to-end radar localization on lidar map using differentiable measurement model," *IEEE Transactions on Intelligent Transportation Systems*, 2021.
- [12] G. Kim, Y. S. Park, Y. Cho, J. Jeong, and A. Kim, "Mulran: Multimodal range dataset for urban place recognition," in *International Conference on Robotics and Automation (ICRA)*, 2020.
- [13] Ş. 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)*. IEEE, 2020, pp. 4358–4364.
- [14] W. Wang, P. P. de Gusmo, B. Yang, A. Markham, and N. Trigoni, "RadarLoc: Learning to Relocalize in FMCW Radar," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2021.
- [15] H. Yin, X. Xu, Y. Wang, and R. Xiong, "Radar-to-lidar: Heterogeneous place recognition via joint learning," *Frontiers in Robotics and AI*, 2021.
- [16] J. Komorowski, M. Wysocka, and T. Trzcinski, "Large-scale topological radar localization using learned descriptors," in *Neural Information Processing: 28th International Conference, ICONIP 2021, Sanur, Bali, Indonesia, December 8–12, 2021, Proceedings, Part II 28*. Springer, 2021, pp. 451–462.
- [17] M. Gadd, D. De Martini, and P. Newman, "Contrastive learning for unsupervised radar place recognition," in *2021 20th International Conference on Advanced Robotics (ICAR)*. IEEE, 2021, pp. 344–349.
- [18] S. Lu, X. Xu, H. Yin, Z. Chen, R. Xiong, and Y. Wang, "One ring to rule them all: Radon Sinogram for place recognition, orientation and translation estimation," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 2778–2785.
- [19] K. Cait, B. Wang, and C. X. Lu, "Autoplace: Robust place recognition with single-chip automotive radar," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 2222–2228.
- [20] H. Jang, M. Jung, and A. Kim, "Raplac: Place recognition for imaging radar using radon transform and mutable threshold," *arXiv preprint arXiv:2307.04321*, 2023.
- [21] J. Yuan, P. Newman, and M. Gadd, "Off the Radar: Variational Uncertainty-Aware Unsupervised Radar Place Recognition for Introspective Querying and Map Maintenance," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023.
- [22] Z. Hong, Y. Petillot, and S. Wang, "Radarslam: Radar based large-scale slam in all weathers," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020.
- [23] D. Wang, Y. Duan, X. Fan, C. Meng, J. Ji, and Y. Zhang, "Maroam: Map-based radar slam through two-step feature selection," *arXiv preprint arXiv:2210.13797*, 2022.
- [24] T. Peynot, S. Scheduling, and S. Terho, "The marulan data sets: Multi-sensor perception in a natural environment with challenging conditions," *The International Journal of Robotics Research*, vol. 29, no. 13, pp. 1602–1607, 2010.
- [25] 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 *International Conference on Robotics and Automation (ICRA)*, 2020.
- [26] M. Sheeny, E. De Pellegrin, S. Mukherjee, A. Ahrabian, S. Wang, and A. Wallace, "Radiate: A radar dataset for automotive perception in bad weather," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 1–7.
- [27] K. Burnett, D. J. Yoon, Y. Wu, A. Z. Li, H. Zhang, S. Lu, J. Qian, W.-K. Tseng, A. Lambert, K. Y. Leung *et al.*, "Boreas: A multi-season autonomous driving dataset," *The International Journal of Robotics Research*, vol. 42, no. 1-2, pp. 33–42, 2023.
- [28] R. Tagiew, M. Köppel, K. Schwalbe, P. Denzler, P. Neumaier, T. Klockau, M. Boekhoff, P. Klasek, and R. Tilly, "Osdar23: Open sensor data for rail 2023," *arXiv preprint arXiv:2305.03001*, 2023.
- [29] M. Gadd, D. De Martini, and P. Newman, "Look around you: Sequence-based radar place recognition with learned rotational invariance," in *2020 IEEE/ION Position, Location and Navigation Symposium (PLANS)*. IEEE, 2020, pp. 270–276.
- [30] Z. Hong, Y. Petillot, A. Wallace, and S. Wang, "Radarslam: A robust simultaneous localization and mapping system for all weather conditions," *The International Journal of Robotics Research*, 2022.
- [31] D. Adolfsson, M. Karlsson, V. Kubelka, M. Magnusson, and H. Andreasson, "Tbv radar slam: trust but verify loop candidates," *arXiv preprint arXiv:2301.04397*, 2023.
- [32] H. Jang, M. Jung, and A. Kim, "RaPlace: Place Recognition for Imaging Radar using Radon Transform and Mutable Threshold," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023.
- [33] M. Gadd and P. Newman, "Open-RadVLAD: Fast and Robust Radar Place Recognition," *arXiv preprint arXiv:2401.15380*, 2024.
- [34] E. Pepperell, P. I. Corke, and M. J. Milford, "All-environment visual place recognition with smart," in *2014 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2014, pp. 1612–1618.
- [35] M. Bosse and R. Zlot, "Place recognition using regional point descriptors for 3d mapping," in *Field and Service Robotics: Results of the 7th International Conference*. Springer, pp. 195–204.
- [36] H. Oh, N. Chebrolu, M. Mattamala, L. Freißmuth, and M. Fallon, "Evaluation and deployment of lidar-based place recognition in dense forests," *arXiv preprint arXiv:2403.14326*, 2024.
- [37] J. Knights, K. Vidanapathirana, M. Ramezani, S. Sridharan, C. Fookes, and P. Moghadam, "Wild-places: A large-scale dataset for lidar place recognition in unstructured natural environments," in *2023 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2023, pp. 11 322–11 328.
- [38] T. Föhst, M. Arndt, K. Berns, C. C. Gava, and R. F. Vassallo, "Off-road place recognition using fused image features," in *ISR 2010 (41st International Symposium on Robotics) and ROBOTIK 2010 (6th German Conference on Robotics)*. VDE, 2010, pp. 1–6.
- [39] A. Harrison and P. Newman, "Ticsync: Knowing when things happened," in *2011 IEEE International Conference on Robotics and Automation*. IEEE, 2011, pp. 356–363.
- [40] R. Arandjelovic, P. Gronat, A. Torii, T. Pajdla, and J. Sivic, "Netvlad: Cnn architecture for weakly supervised place recognition," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [41] G. Berton, C. Masone, and B. Caputo, "Rethinking visual geolocalization for large-scale applications," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 4878–4888.
- [42] G. Kim and A. Kim, "Scan context: Egocentric spatial descriptor for place recognition within 3d point cloud map," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018.
- [43] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [44] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 2009, pp. 248–255.
- [45] D. Arthur and S. Vassilvitskii, "K-means++ the advantages of careful seeding," in *eighteenth annual ACM-SIAM symposium on Discrete algorithms*, 2007, pp. 1027–1035.

- [46] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "Mobilenetv2: Inverted residuals and linear bottlenecks," in *IEEE conference on computer vision and pattern recognition*, 2018.
- [47] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *IEEE conference on computer vision and pattern recognition*, 2015.
- [48] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," *Advances in neural information processing systems*, vol. 25, 2012.
- [49] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.