Boreas: A Multi-Season Autonomous Driving Dataset

Keenan Burnett¹, David J. Yoon¹, Yuchen Wu¹, Andrew Zou Li¹, Haowei Zhang¹, Shichen Lu¹, Jingxing Qian¹, Wei-Kang Tseng¹, Andrew Lambert², Keith Y.K. Leung², Angela P. Schoellig¹, Timothy D. Barfoot¹

Abstract

The Boreas dataset was collected by driving a repeated route over the course of one year, resulting in stark seasonal variations and adverse weather conditions such as rain and falling snow. In total, the Boreas dataset contains over 350km of driving data featuring a 128-channel Velodyne Alpha-Prime lidar, a 360° Navtech CIR304-H scanning radar, a 5MP FLIR Blackfly S camera, and centimetre-accurate post-processed ground truth poses. At launch, our dataset will support live leaderboards for odometry, metric localization, and 3D object detection. The dataset and development kit are available at boreas.utias.utoronto.ca.

Keywords

Autonomous vehicle, camera, dataset, GPS, IMU, lidar, radar, snow, winter

1 Introduction

To date, autonomous vehicle research and development has focused on achieving sufficient reliability in ideal conditions such as the sunny climates observed in San Francisco, California or Phoenix, Arizona. Adverse weather conditions such as rain and snow remain outside the operational envelope for many of these systems. Additionally, a majority of self-driving vehicles are currently reliant on highlyaccurate maps for both localization and perception. These maps are costly to maintain and may degrade as a result of seasonal changes. In order for self-driving vehicles to be deployed safely, these short-comings must be addressed.

To encourage research in this area, we have created the Boreas dataset, a large multi-modal dataset collected by driving a repeated route over the course of one year. The dataset features over 350km of driving data with stark seasonal variations and multiple sequences with adverse weather such as rain and falling snow. Our data-taking platform, shown in Figure 1, includes a 128-beam lidar, a 5 MP camera, and a 360° scanning radar. Globally-consistent centimetre-accurate ground truth poses are obtained by postprocessing GNSS, IMU, and wheel encoder data along with a secondary correction subscription. At launch, our dataset will support benchmarks for odometry, metric localization, and 3D object detection.

This dataset may be used to study the effects of seasonal variation on long-term localization. Further, this dataset enables comparisons of vision, lidar, and radarbased mapping and localization pipelines. Comparisons may include the robustness of individual sensing modalities to adverse weather or the resistance to map degradation.

The main contributions of this dataset are as follows:

• Data collected on a repeated route over the course of one year including multiple weather conditions.



Figure 1. Our platform, *Boreas*, includes a Velodyne Alpha-Prime (128-beam) lidar, a FLIR Blackfly S camera, a Navtech CIR304-H radar, and an Applanix POS LV GNSS-INS.

- A unique, high-quality sensor configuration including a 128-beam lidar and 360° radar.
- Post-processed GNSS/IMU data to provide accurate ground truth pose information.
- A live and open leaderboard for odometry, metric localization, and 3D object detection.
- 3D object labels collected in sunny weather.

Corresponding author:

Keenan Burnett, University of Toronto, 4925 Dufferin Street, Toronto, Ontario, Canada.

Email: keenan.burnett@robotics.utias.utoronto.ca

¹University of Toronto

² Applanix, Trimble

2

Table 1. Related datasets. Lead: public leaderboard. Size: For perception datasets, size is given as the number of annotated
frames and the number of annotations (3d boxes). GT: ground truth pose source. (A): automotive radar. (N): 360° Navtech radar.
RTK (Real-Time Kinematic) uses a GPS base station and differential measurements to improve GPS accuracy. RTX uses data from
a global network of tracking stations to calculate corrections. This can be used to achieve cm-level accuracy without a base station
(Applanix 2022). [†] Waymo's Mid-Range, Short-Range proprietary 3D lidar. [‡] The Oxford Robotcar dataset contains one sequence
with snow on the ground but that sequence has no falling snow.

Name	Lead	Size	Camera	Lidar	Radar	GT	Night	Rain	Snow	Seasons
Perception										
ApolloScape Huang et al. (2018)	1	144k 70k boxes	2x9.2MP	1x64C	×	GPS/IMU	1	1	X	×
Argoverse Chang et al. (2019)	1	22k 993k boxes	7x2.3MP +2x5MP	2x32C	×	GPS/IMU	1	×	X	×
CADC Pitropov et al. (2021)	×	7.5k 372k boxes	8x1.3MP	1x32C	×	GPS/IMU + RTK	x	x	1	×
KITTI (Object) Geiger et al. (2013)	1	15k 200k boxes	4x1.4MP	1x64C	×	GPS/IMU + RTK	×	×	x	×
nuScenes Caesar et al. (2020)	1	40k 1.4M boxes	6x1.4MP	1x32C	✓(A)	GPS/IMU + Lidar Loc	1	1	x	×
RADIATE Sheeny et al. (2021)	×	44k 200k boxes	2x0.25MP	1x32C	✔(N)	GPS/IMU	1	1	1	×
Waymo OD Sun et al. (2020)	1	230k 12M boxes	5x2.5MP	1(MR [†]) 4(SR [†])	×	GPS/IMU	1	1	x	×
Boreas-Objects-V1	1	7.1k 320k boxes	1x5MP	1x128C	✔(N)	GPS/IMU	×	×	×	×
Localization										
KITTI (Odometry) Geiger et al. (2013)	1	39km 22 seqs	4x1.4MP	1x64C	×	GPS/IMU + RTK	×	x	×	×
Complex Urban Jeong et al. (2019)	×	451km 40 seqs	2x1.9MP	2x16C + 2x1C	×	SLAM	X	×	X	×
Oxford RobotCar Maddern et al. (2017)	×	1000km 100 segs	3x1.2MP +3x1MP	1x4C + 2x1C	×	GPS/IMU + RTK	1	1	X‡	1
Oxford Radar Barnes et al. (2020)	×	280km 32 seqs	3x1.2MP +3x1MP	2x32C + 2x1C	✔(N)	GPS/IMU + VO	×	1	x	×
MulRan Kim et al. (2020)	×	124km 12 seqs	×	1x64C	✔(N)	SLAM	x	x	x	×
Boreas	1	350km 44 seqs	1x5MP	1x128C	✓(N)	GPS/IMU + RTX	1	1	1	1

2 Related Work

Several datasets have been produced to meet the increasing demand driven by autonomous driving research. Many of the published datasets focus on perception, particularly 3D object detection and semantic segmentation of images and lidar pointclouds. However, these datasets tend to lack variation in weather and season. Further, many of these datasets do not provide radar data. Automotive radar sensors are robust to precipitation, dust, and fog thanks to their longer wavelength. For this reason, radar may play a key role in enabling autonomous vehicles to operate in adverse weather. The Boreas dataset addresses these shortcomings by including a 360° scanning radar, and data taken during various weather conditions (sun, cloud, rain, night, snow) and seasons.

Another significant fraction of datasets focus on the problem of localization, usually odometry. The Boreas dataset includes both a high-density lidar (128-beam) and a 360° scanning radar. The combination of these sensors and the significant weather variation contained in this dataset enables detailed comparisons between the localization capabilities of these two sensing modalities. This is something that previous datasets were not able to support due to either not having a radar sensor or insufficient

weather variation. Furthermore, our post-processed ground truth poses are sufficiently accurate to support a public leaderboard for odometry and metric localization. To the best of our knowledge, a public leaderboard for metric localization is not offered elsewhere. For a detailed comparison of related datasets, see Table 1.

3 Data Collection

The majority of the Boreas dataset was collected by driving a repeated route near the University of Toronto Institute for Aerospace Studies (UTIAS) over the course of one year. Figure 2 illustrates the seasonal variations that were observed over this time. Figure 3 compares camera, lidar, and radar measurements in three distinct weather conditions: falling snow, rain, and sun. The primary repeated route will be referred to as the Glen Shields route and is depicted in Figure 4. Additional routes were also collected as either a single standalone sequence or a small number of repeated traversals. The Glen Shields route can be used for research related to long-term localization while the other routes allow for experiments that test for generalization to previously unseen environments. The frequency of different metadata tags is displayed in Figure 5.



Figure 2. This figure depicts one year of seasonal changes in the Boreas dataset. Each image represents a camera image that was taken on a different day. The sequences are sorted in chronological order from left to right and top to bottom, starting in November, 2020 and finishing in November, 2021. Note that the sequences are not evenly spaced in time.

4 Sensors

Table 2 provides detailed specifications for the sensors used in this dataset. Figure 6 illustrates the placement of the different sensors on Boreas.

5 Dataset Format

5.1 Data Organization

The Boreas dataset is divided into *sequences*, which include all sensor data and ground truth poses from a single drive. Sequences are identified by the date and time at which they were collected with the format boreas-YYYY-MM-DD-HH-MM. The data for each sequence is organized as shown in Figure 7.

5.2 Timestamps

The name of each file corresponds to its timestamp. These timestamps are given as UNIX epoch times in microseconds. All sensors were synchronized to the UTC time reported by the Applanix POS LV. The Velodyne lidar was synchronized using a standard hardwired connection to the Applanix POS LV carrying a PPS signal and NMEA messages. The camera was configured to emit a square-wave pulse where the rising edge of each pulse corresponds with the start of a new camera exposure event. The Applanix POS LV was then configured to receive and timestamp these event signals. Camera timestamps were then corrected in post using the recorded event times and exposure values: $t_{camera} = t_{event} + (exposure(t))/2$. The data-recording computer was synchronized to UTC time in a fashion similar to the



Figure 3. Weather variation in the Boreas dataset. Note that the lidar pointcloud becomes littered with detections associated with snowflakes during falling snow and that the radar data remains relatively unperturbed across the weather conditions.

Velodyne, using an RS-232 serial cable carrying a PPS signal and NMEA messages. The Navtech radar synchronizes its local clock using NTP. Since the data-recording computer publishing the NTP time is synchronized to UTC time, the radar is thereby also synchronized to UTC time.

For lidar pointclouds, the timestamp corresponds to the temporal middle of the scan. Each lidar point also has a timestamp associated with it. These point times are given in seconds relative to the middle of the scan. For radar scans, the timestamp also corresponds to the the middle of the scan: $\lfloor \frac{M}{2} \rfloor - 1$ where M is the number of azimuths. Each scanned radar azimuth is also timestamped in the same format as the filename.

5.3 File Formats

Camera images are rectified and anonymized by default. We use Anonymizer to blur license plates and faces (Understand

2022). Images are stored in the commonly-used png format. Lidar pointclouds are stored in a binary format to minimize storage requirements. Our devkit provides methods for working with these binary formats in both C++ and Python. Each point has six fields: [x, y, z, i, r, t] where (x, y, z) is the position of the point with respect to the lidar, *i* is the intensity of the reflected infrared signal, r is the ID of the laser that made the measurement, and t the point timestamp explained in Section 5.2. Raw radar scans are stored as 2D polar images: M azimuths x R range bins. We follow Oxford's convention and embed timestamp and encoder information into the first eleven columns (bytes) of each polar radar scan. The first eight columns represent a 64-bit integer, the UNIX epoch timestamp of each azimuth in microseconds. The next two columns represent a 16-bit unsigned integer, the rotational encoder value. The next column is unused but preserved for compatibility with the Oxford format. See (Barnes et al. 2020) for further details on the Navtech sensor



(a) OpenStreetMap data



(b) Mapbox satellite data

Figure 4. The Glen Shields route in Toronto, Ontario, Canada.



Figure 5. Frequency of metadata tags in the Boreas dataset.

and this file format. The polar radar scans can be readily converted into a top-down Cartesian representation, as shown in Figure 3, using our devkit.

Note that measurements are not synchronous as in other datasets (KITTI, CADC), which means that measurements with the same index do not have the same timestamp. However, given the timestamps and relative pose information, different sensor measurements can still be fused together. Lidar pointclouds are not motion-corrected, but we do provide methods for removing motion distortion in our devkit. Navtech radar scans suffer from both motion distortion and **Table 2.** Sensor specifications. [†]Position accuracy changes over time as a function of the number of visible satellites. [†]These numbers represent expected accuracy in nominal conditions. [‡]Our Navtech radar's firmware was upgraded partway through the project, older sequences have a range resolution of 0.0596m, and a range of 200m.

Sensor	Specifications
Applanix POS LV 220	 2-4cm RTX accuracy (RMS)[†] 200 Hz
Navtech CIR304-H Radar	 0.0438m range solution[‡] 0.9° horizontal resolution 250m range[‡] 4 Hz
FLIR Blackfly S Camera (BFS-U3-51S5C)	 2448x2048 (5 MP) 81° HFOV x 71° VFOV 10 Hz
Velodyne Alpha-Prime Lidar	 128 beams 0.1° vertical resolution (variable) 0.2° horizontal resolution 360° HFOV x 40° VFOV 300m range (10% reflectivity) ~ 2.2M points/s 10 Hz



Figure 6. Boreas sensor placement. Distances are given in metres. Measurements shown are approximate. Refer to the calibrated extrinsics contained in the dataset for precise measurements.

Dopler distortion, Burnett et al. (2021a) and Burnett et al. (2021b) provide methods to compensate for these effects.

```
boreas-YYYY-MM-DD-HH-MM
     applanix
        camera_poses.csv
        imu.csv
        gps_post_process.csv
        lidar_poses.csv
        radar_poses.csv
     calib
        camera0_intrinsics.yaml
        P_camera.txt
        T sens1 sens2.txt
     camera
        <timestamp>.png
     lidar
        <timestamp>.bin
     radar
        <timestamp>.png
     route.html
     video.mp4
```

Figure 7. Data organization for a single Boreas sequence.

6 Ground Truth Poses

Ground truth poses are obtained by post-processing GNSS, IMU, and wheel encoder measurements along with corrections obtained from an RTX subscription using Applanix's POSPac software suite. Positions are reported in metric UTM coordinates. For each sequence, applanix/qps_post_process.csv contains the postprocessed ground truth in the Applanix frame at 200Hz. Note that the Applanix frame is y-forwards, x-right, zup. For all poses, we use ENU (x-East, y-North, z-Up) as the origin. We follow the convention used at UTIAS for describing rotations, translations, and $4 \times$ 4 homogeneous transformation matrices (Barfoot 2017). Each sensor frame's ground truth is stored as a row in applanix/<sensor>_poses.csv with the following format: $[t, x, y, z, v_x, v_y, v_z, r, p, y, \omega_z, \omega_y, \omega_x]$ where t is the epoch timestamp in microseconds that matches the filename, $\mathbf{r}_{e}^{se} = [x \ y \ z]^{T}$ is the position of the sensor s with respect to ENU as measured in ENU, $\mathbf{v}_e^{se} = [v_x \ v_y \ v_z]^T$ is the velocity of the sensor with respect to ENU, (r, p, y) are the roll, pitch, and yaw angles, which can be converted into a rotation matrix between the sensor frame and ENU. $\omega_s^{se} =$ $[\omega_x \ \omega_y \ \omega_z]^T$ are the angular velocities of the sensor with respect to the ENU frame as measured in the sensor frame.

The pose of the sensor frame is then: $\mathbf{T}_{es} = \begin{bmatrix} \mathbf{C}_{es} & \mathbf{r}_{e}^{se} \\ \mathbf{0}^{T} & 1 \end{bmatrix} \in SE(3)$ where $\mathbf{C}_{es} = \mathbf{C}_{1}(\text{roll})\mathbf{C}_{2}(\text{pitch})\mathbf{C}_{3}(\text{yaw})$ (Barfoot 2017). We also provide post-processed IMU measurements in applanix/imu.csv at 200Hz in the Applanix frame that include linear acceleration and angular velocity.

The RMS position error reported by Applanix is typically less than 5cm in nominal conditions but can be as high as 20-40cm in urban canyons. Figure 8 shows the RMS errors estimated by the Applanix POSPac software. Accuracy can change depending on the visibility of satellites. Note that these values represent global accuracy and that relative pose estimates are more accurate over short time horizons.



Figure 8. Post-processed RMS position, velocity, and orientation error vs. time reported by Applanix's POSPac software for a Glen Shields sequence collected on 2021-09-07.

7 Calibration

7.1 Camera Intrinsics

Camera intrinsics are calibrated using MATLAB's camera calibrator (Mathworks 2022a) and are recorded in camera0_intrinsics.yaml. Images located under camera/ have already been rectified. The rectified camera matrix **P** is stored in P_camera.txt. To project lidar points onto a camera image, we use the pose of the camera \mathbf{T}_{ec} at time t_c and the pose of the lidar \mathbf{T}_{el} at time t_l to compute a transform from the lidar frame to the camera frame given by $\mathbf{T}_{cl} = \mathbf{T}_{ec}^{-1}\mathbf{T}_{el}$. Each point in the lidar frame is then transformed into the camera frame with $\mathbf{x}_c = \mathbf{T}_{cl}\mathbf{x}_l$, where $\mathbf{x}_l = [x \ y \ z \ 1]^T$. The projected image coordinates are then obtained using:

$$\begin{bmatrix} u \\ v \end{bmatrix} = \mathbf{D} \mathbf{P} \frac{1}{z} \mathbf{x}_c \tag{1}$$

where
$$\mathbf{D} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix}$$
, $\mathbf{P} = \begin{bmatrix} f_u & 0 & c_u & 0 \\ 0 & f_v & c_v & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$. (2)

7.2 Sensor Extrinsics

The extrinsic calibration between the camera and lidar is obtained using MATLAB's camera to lidar calibrator (Mathworks 2022b). The results of this calibration are illustrated in Figure 9. To calibrate the rotation between the lidar and radar, we use correlative scan matching via the Fourier Mellin transform (Checchin et al. 2010). Several lidar-radar pairs were collected while the vehicle was stationary at different locations. The final rotation estimate is obtained by averaging the results from several measurement pairs (Burnett 2020). The translation between the lidar and radar is obtained from the CAD model of the roof rack.



(a) Perspective View



(b) Colored Pointcloud

Figure 9. Lidar points projected onto a camera image using the camera-lidar calibration. (a) Lidar points are colored based on their longitudinal distance from the vehicle. (b) Lidar points are given RGB color values based on their projected location on the camera image.

The results of the radar-to-lidar calibration are shown in Figure 10. The extrinsics between the lidar and the Applanix reference frame were obtained using Applanix's in-house calibration tools. Their tool outputs this relative transform as a by-product of a batch optimization aiming to estimate the most likely vehicle path given a sequence of lidar pointclouds and post-processed GNSS/IMU measurements. All extrinsic calibrations are provided as 4x4 homogeneous transformation matrices under the calib/ folder.

8 3D Annotations

We provide a set of 3D bounding box annotations for a subset of the Boreas dataset, obtained in sunny weather. We refer to this as the Boreas-Objects-V1 dataset. Annotations were obtained using the Scale.ai data annotation service (Scale 2022). In total, 7111 lidar frames were annotated at 5Hz, resulting in 326,180 unique 3D box annotations. Since the lidar data was collected at 10Hz, the annotations may be interpolated between frames to double the number of annotated frames at a slightly lower fidelity. The data is divided into 53 continuous scenes where each scene is 20-70 seconds in duration. The scenes are then divided into 37 training scenes and 16 test scenes where the ground truth labels have been withheld for the benchmark. Figure 11 displays two statistics for our annotations.

We use the same folder structure as in Figure 7 but with an additional folder, labels/. Similar to KITTI, annotations for a particular frame are stored in a text file with the same filename (timestamp) as the lidar frame. Each row of a label file corresponds to a different 3D box annotation with the format: [uuid, type, d_x , d_y , d_z , x, y, z, yaw]. The uuid is a unique ID for a particular object track that is consistent across frames within a particular scene. The type is the semantic class for an object that can be one of:



Figure 10. Lidar measurements are drawn in red using a bird's eye view projection with the ground plane removed. Radar targets are first extracted from the raw radar data and then are draw as blue pixels. The two sensors have been aligned using the radar-to-lidar calibration.

{Car, Cyclist, Pedestrian, Misc}. The Car class includes coupes, sedans, SUVs, vans, pick-up trucks, and ambulances. The Cyclist class includes people riding motorcycles, but excludes parked bicycles. The Misc class includes other vehicle types such as buses, industrial trucks, streetcars, and trains. Objects are labelled within a rectangular area centered on the lidar +/- 75m in both dimensions. Bounding box locations (x, y, z) and orientations (yaw) are given with respect to the lidar frame. (d_x, d_y, d_z) represent the bounding box dimensions (length, width, and height). Figure 12 shows an example of what our 3D object annotations look like for lidar, camera, and radar.



Figure 11. 3D annotation statistics for Boreas-Objects-V1.



Figure 12. Examples of 3D annotations in the Boreas-Objects-V1 dataset.

9 Benchmark Metrics

At launch, we plan to support online leaderboards for odometry, metric localization, and 3D object detection. For odometry, we use the same metrics as the KITTI dataset (Geiger et al. 2013). The KITTI odometry metrics average the relative position and orientation errors over every subsequence of length (100m, 200m, 300m, ..., 800m). This results in two metrics, a translational drift reported as a percentage of path length and a rotational drift reported as degrees per metre travelled. For 3D object detection, we also defer to the KITTI dataset by reporting the mean average precision (mAP) on a per-class basis. For cars, a 70% overlap counts as a true positive, and for pedestrians, 50%. We do not divide our dataset based on difficulty levels.

The purpose of our metric localization leaderboard is to benchmark mapping and localization pipelines. In this scenario, we envision a situation where one or more repeated traversals of the Glen Shields route are used to construct a map offline. Any and all data from the training sequences may be used to construct a map in any fashion.

Then, during a test sequence, the goal is to perform metric localization between the live sensor data and the pre-built map. Localization approaches may make use of temporal filtering and can leverage the IMU if desired but GNSS information will not be available. The goal of this benchmark is to simulate localizing a vehicle in real-time and as such methods may not use future sensor information in an acausal manner.

Our goal is to support both global and relative map structures. Only one of the training sequences will be specified as the map sequence used by the benchmark. For 3D localization, users must choose either the lidar or camera as the reference sensor. For 2D localization, only the radar frames are used as a reference. For each (camera—lidar—radar) frame s_2 in the test sequence, users will specify the ID (timestamp) of the (camera—lidar—radar) frame s_1 in the map sequence that they are providing a relative pose with respect to: $\hat{\mathbf{T}}_{s_1,s_2}$. We then compute root-mean squared error (RMSE) values for the translation and rotation as follows:

$$\mathbf{T}_{e} = \mathbf{T}_{a,s_{1}} \mathbf{T}_{s_{1},s_{2}} \hat{\mathbf{T}}_{s_{1},s_{2}}^{-1} \mathbf{T}_{a,s_{1}}^{-1} = \begin{bmatrix} \mathbf{C}_{e} & \mathbf{r}_{e} \\ \mathbf{0}^{T} & 1 \end{bmatrix}$$
(3)

$$\mathbf{r}_e = \begin{bmatrix} x_e & y_e & z_e \end{bmatrix}^T \tag{4}$$

$$\phi_e = \arccos\left(\frac{\operatorname{tr} \mathbf{C}_e - 1}{2}\right) \tag{5}$$

where \mathbf{T}_{s_1,s_2} is the known ground truth pose, and \mathbf{T}_{a,s_1} is the calibrated transform from the sensor frame to the applanix frame (x-right, y-forwards, z-up). x_e, y_e, z_e are then the lateral, longitudinal, and vertical errors respectively. We calculate RMSE values for x_e, y_e, z_e, ϕ_e .

Users will also have the option of providing 6×6 covariance matrices Σ_i for each localization estimate. A pose with uncertainty is described as $\mathbf{T} = \exp(\boldsymbol{\xi}^{\wedge})\overline{\mathbf{T}}$ where $\boldsymbol{\xi} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$ (Barfoot 2017). Given $\hat{\mathbf{T}}_i = \hat{\mathbf{T}}_{s_1, s_2}(t_i)$, we compute an average consistency score c for the localization and covariance estimates:

$$\boldsymbol{\xi}_{i} = \ln \left(\mathbf{T}_{i} \hat{\mathbf{T}}_{i}^{-1} \right)^{\vee} = \begin{bmatrix} \rho_{1} & \rho_{2} & \rho_{3} & \psi_{1} & \psi_{2} & \psi_{3} \end{bmatrix}^{T}$$
(6)

$$c = \left(\sum_{i=1}^{N} \frac{\boldsymbol{\xi}_i^T \boldsymbol{\Sigma}_i^{-1} \boldsymbol{\xi}_i}{N \operatorname{dim}(\boldsymbol{\xi}_i)}\right)^{1/2} \tag{7}$$

A consistency score close to 1 is ideal. c < 1 means that the method is over-confident, c > 1 means that the method is conservative. Note that the above metrics will be averaged across the test sequences.

10 Development Kit

As part of this dataset, we provide a development kit to make it easier for new users to get started. The primary purpose of the devkit is to act as a wrapper around the entire dataset to be used in Python. This allows users to query frames and the associated ground truth for either odometry, localization, or 3D object detection. We also provide convenience methods for removing motion distortion from pointclouds, working with polar radar scans, and converting to and from Lie algebra and Lie group representations. The devkit also provides many different ways to visualize the sensor data. Finally, we provide several introductory tutorials as Jupyter notebooks that include projecting lidar onto a camera frame and visualizing 3D boxes. Evaluation scripts used by our benchmark will be stored in the devkit, allowing users to validate their algorithms before submission to the benchmark. The development kit can be found at boreas.utias.utoronto.ca.

11 Conclusion

In this paper, we presented Boreas, a multi-season autonomous driving dataset that includes over 350km of driving data collected over the course of one year. The dataset provides a unique high-quality sensor suite including a Velodyne Alpha-prime (128-beam) lidar, a 5MP camera, a 360° Navtech radar, and accurate ground truth poses obtained from an Applanix POS LV with an RTX subscription. We also provide 3D object labels for a subset of the Boreas data obtained in sunny weather. The primary purpose of this dataset is to enable further research into long-term localization across seasons and adverse weather conditions. Our website will provide an online leaderboard for odometry, metric localization, and 3D object detection.

Acknowledgements

We would like to thank Goran Basic for his help in designing and assembling the roof rack for Boreas. We also thank General Motors for their donation of the Buick vehicle. The Amazon Open Data Sponsorship program supports this project by hosting the Boreas dataset. This work was also partially financially supported by a Natural Sciences and Engineering Research Council (NSERC) grant.

References

Applanix (2022) www.applanix.com.

- Barfoot TD (2017) *State estimation for robotics*. Cambridge University Press.
- Barnes D, Gadd M, Murcutt P, Newman P and Posner I (2020) The Oxford Radar RobotCar dataset: A radar extension to the Oxford RobotCar dataset. In: *IEEE International Conference* on Robotics and Automation (ICRA). pp. 6433–6438.
- Burnett K (2020) Radar to Lidar Calibrator.
- Burnett K, Schoellig AP and Barfoot TD (2021a) Do we need to compensate for motion distortion and doppler effects in spinning radar navigation? *IEEE Robotics and Automation Letters* 6(2): 771–778.
- Burnett K, Yoon DJ, Schoellig AP and Barfoot TD (2021b) Radar odometry combining probabilistic estimation and unsupervised feature learning. In: *Robotics: Science and Systems (RSS)*.
- Caesar H, Bankiti V, Lang AH, Vora S, Liong VE, Xu Q, Krishnan A, Pan Y, Baldan G and Beijbom O (2020) nuScenes: A multimodal dataset for autonomous driving. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (CVPR). pp. 11621–11631.
- Chang MF, Lambert J, Sangkloy P, Singh J, Bak S, Hartnett A, Wang D, Carr P, Lucey S, Ramanan D et al. (2019) Argoverse: 3d tracking and forecasting with rich maps. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (*CVPR*). pp. 8748–8757.
- Checchin P, Gérossier F, Blanc C, Chapuis R and Trassoudaine L (2010) Radar scan matching slam using the Fourier-Mellin

- Geiger A, Lenz P, Stiller C and Urtasun R (2013) Vision meets robotics: The KITTI dataset. *The International Journal of Robotics Research* 32(11): 1231–1237.
- Huang X, Cheng X, Geng Q, Cao B, Zhou D, Wang P, Lin Y and Yang R (2018) The ApolloScape dataset for autonomous driving. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*. pp. 954–960.
- Jeong J, Cho Y, Shin YS, Roh H and Kim A (2019) Complex urban dataset with multi-level sensors from highly diverse urban environments. *The International Journal of Robotics Research* 38(6): 642–657.
- Kim G, Park YS, Cho Y, Jeong J and Kim A (2020) MulRan: Multimodal range dataset for urban place recognition. In: *IEEE International Conference on Robotics and Automation (ICRA)*. pp. 6246–6253.
- Maddern W, Pascoe G, Linegar C and Newman P (2017) 1 Year, 1000 km: The Oxford RobotCar dataset. *The International Journal of Robotics Research* 36(1): 3–15.

Mathworks (2022a) Using the Single Camera Calibrator App.

Mathworks (2022b) Lidar and Camera Calibration.

Pitropov M, Garcia DE, Rebello J, Smart M, Wang C, Czarnecki K and Waslander S (2021) Canadian adverse driving conditions dataset. *The International Journal of Robotics Research* 40(4-5): 681–690.

Scale (2022) www.scale.com.

- Sheeny M, De Pellegrin E, Mukherjee S, Ahrabian A, Wang S and Wallace A (2021) RADIATE: A radar dataset for automotive perception in bad weather. In: *IEEE International Conference* on Robotics and Automation (ICRA). pp. 1–7.
- Sun P, Kretzschmar H, Dotiwalla X, Chouard A, Patnaik V, Tsui P, Guo J, Zhou Y, Chai Y, Caine B et al. (2020) Scalability in perception for autonomous driving: Waymo open dataset. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 2446–2454.
- Understand (2022) Understand.ai Anonymizer.