

TartanDrive 2.0: More Modalities and Better Infrastructure to Further Self-Supervised Learning Research in Off-Road Driving Tasks

Matthew Sivaprakasam¹, Parv Maheshwari², Mateo Guaman Castro¹, Samuel Triest¹, Micah Nye³,
 Steve Willits¹, Andrew Saba¹, Wenshan Wang¹, and Sebastian Scherer¹

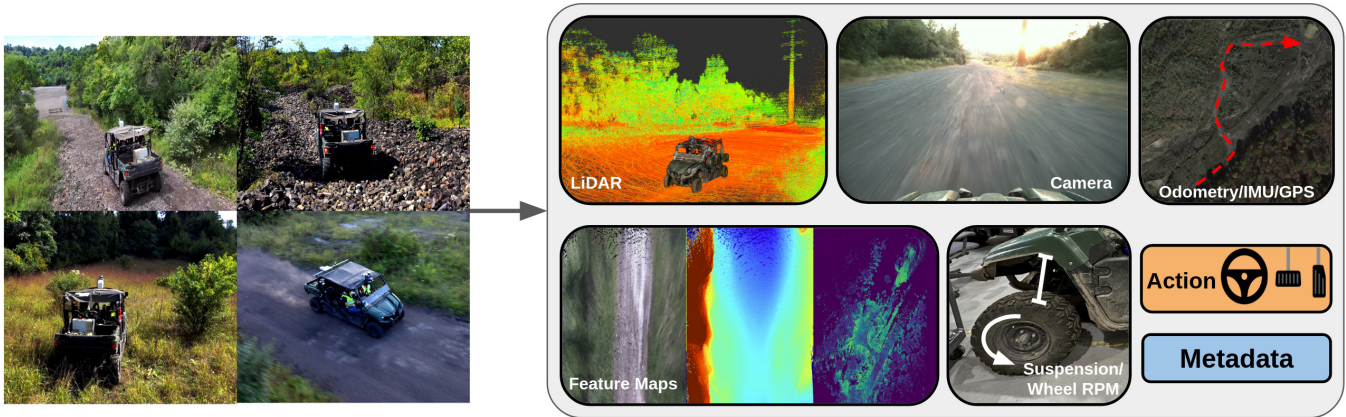


Fig. 1: We provide a new off-road driving dataset for self-supervised learning tasks. Multi-modal data is collected as the robot is driven through several types of terrain that present challenging scenarios to perception, planning, and control algorithms.

Abstract— We present TartanDrive 2.0, a large-scale off-road driving dataset for self-supervised learning tasks. In 2021 we released TartanDrive 1.0, which is one of the largest datasets for off-road terrain. As a follow-up to our original dataset, we collected seven hours of data at speeds of up to 15m/s with the addition of three new LiDAR sensors alongside the original camera, inertial, GPS, and proprioceptive sensors. We also release the tools we use for collecting, processing, and querying the data, including our metadata system designed to further the utility of our data. Custom infrastructure allows end users to reconfigure the data to cater to their own platforms. These tools and infrastructure alongside the dataset are useful for a variety of tasks in the field of off-road autonomy and, by releasing them, we encourage collaborative data aggregation. These resources lower the barrier to entry to utilizing large-scale datasets, thereby helping facilitate the advancement of robotics in areas such as self-supervised learning, multi-modal perception, inverse reinforcement learning, and representation learning. The dataset is available at <https://theairlab.org/TartanDrive2>.

I. INTRODUCTION

As the state of the art in autonomous driving improves, robots are expected to handle increasingly complex tasks. In off-road driving, this translates to situations such as a robot knowing whether the dark brown path in front of it is dry dirt or thick mud, whether it is worth it to take a shortcut through tall grass, and when to preemptively

increase velocity in order to make it up a hill. Navigating these decisions requires behaviors that are extremely difficult to robustly design and tune by hand. Many of the currently-researched solutions to this challenge involve the adoption of large neural networks and data-driven models. For the task of on-road driving in urban scenarios, there are several large datasets available [1–4]. These datasets were feasible not only because of the resources present at the institutions that collected them, but also because of the inherent everyday nature of on-road driving. There exist fewer datasets for off-road driving, primarily because the nature of off-road terrain makes gathering enough data uniquely difficult. Collecting samples in simulation is often insufficient due to the complexity of calculating the dynamics of terrain in complex environments. In fact, most real-world scenarios that are hard to accurately simulate are the same ones which produce complex situations for autonomous off-road agents, such as dense foliage and slippery surfaces. However, collecting data in real-life presents logistics challenges such as preserving driver safety and dealing with frequent vehicle damage and wear. These difficulties have resulted in a scarcity of available data for off-road driving.

While the number of off-road datasets has been growing [5–12], many of them are still limited in sample size, modality, or difficulty [13]. This is often due to their focus on specific tasks such as semantic segmentation, which requires manual labeling effort. Scaling up these datasets would come either at the cost of time or label quality. It is possible to use multiple datasets together to train a model, but the variability of off-road terrain often causes inconsistency in labels across datasets [14]. The difficulty of generating large

* This work was supported by ARL awards #W911NF1820218 and #W911NF20S0005.

¹ Robotics Institute, Carnegie Mellon University, [msivapra](mailto:msivapra@cmu.edu), [mguamanc](mailto:mguamanc@cmu.edu), [striest](mailto:striest@cmu.edu), [wenshanw](mailto:wenshanw@cmu.edu), [basti](mailto:basti@cmu.edu)

² Department of Mathematics, Indian Institute of Technology Kharagpur, parvmaheshwari2002@iitkgp.ac.in

³ Department of Mechanical Engineering, University of Pittsburgh, man172@pitt.edu

amounts of manually labeled data in off-road terrain suggests that self-supervised learning is essential to outperforming prior methods. Many works have shown that self-supervised methods can be used to learn strong representations [15–17], and some works have already shown how it can be used for various tasks in off-road driving environments [18–21]. We argue that in order to scale up the amount of data for off-road tasks, datasets should be designed to be task-agnostic and with self-supervised learning in mind.

In 2021, we released TartanDrive, a multi-modal off-road driving dataset designed for dynamics modeling [5]. In this work we present TartanDrive 2.0, a larger dataset geared towards self-supervised learning methods. The contributions of this dataset over the previous generation are as follows:

- 1) **More sensors and data:** In our original work, we argue that multiple modalities are essential for learning models for off-road terrain. To that end, we add LiDAR as an additional modality. We also collect seven hours over the five hours of data present in the original dataset, at higher average speeds and including some new areas, as well as areas present in the previous dataset that have dramatically changed over the past two years.
- 2) **Better infrastructure for end-users:** We have improved the infrastructure in order to improve ease-of-use and utility of the dataset, such as the ability to re-process the data based on user-specified configurations. We also implement a metadata system that filters and groups data by a variety of properties.
- 3) **Open-source data collection tooling:** Our data collection process has been constructed in a way that allows us to continually release more data over time as the seasons change and as we add more modalities. By releasing our tools and framework, others can collect their data and easily merge it with our own to create larger datasets.

II. RELATED WORK

Collecting general driving datasets is a common pursuit, the largest ones coming from companies gathering data in urban scenarios. For example, nuScenes by Motional is a dataset containing 1.4 million images and 390,000 LiDAR sweeps, with corresponding 3D object annotations [3,4]. The Waymo Open Dataset is another example with over a million images and LiDAR pointclouds [1,2]. Both datasets include infrastructure and tools to query and process these massive amounts of data in ways that support various downstream learning tasks and benchmarks that are not often seen in smaller datasets. There are also other efforts in collecting data for on-road environments, such as the RACECAR dataset that includes raw and processed sensor data from autonomous racecars driving at high speeds [22].

A number of earlier works have tried to circumvent the issues due to the lack of real-world off-road driving data by utilizing simulation environments. Tremblay et al. show how training a neural network on multiple modalities in simulation can allow a robot to predict its dynamics in




Fig. 2: The ATV used for data collection (left); The primary sensor payload on the vehicle (right)




Fig. 3: The 3D scan of the ATV. The rear half of the scan is cropped in the left picture to increase clarity of the front payload in the image.

unseen real-world data [23]. Sivaprakasam et al. collect data in simulation of a robot driving over obstacles in order to train a model that predicts the difference between a desired path and its resulting path [24].

Now, more off-road datasets have been collected. RUGD contains over 7,000 images in a variety of off-road terrain with manually-labeled semantic segmentation masks [7]. Rellis-3D includes annotations for 6,235 images and 13,556 scans from two different LiDARs, as well as the bags that they originally recorded which also contain IMU, GPS, and stereo image data. Sharma et al. created an off-road image dataset where, rather than creating semantic masks, they label regions in an image based on their traversability by different types of vehicles [8]. The first version of TartanDrive has also been available for two years and has been used for a number of tasks even outside the domain of off-road driving. Shah et al. have used this data as part of a bigger dataset that was used to train a large foundation model designed with vision-based robotic navigation in mind [25].

III. THE DATASET

The original TartanDrive dataset was designed for the task of dynamics prediction. When designing TartanDrive 2.0 we aimed to support a broader set of tasks in perception, planning, and control. Thanks to more sensor modalities and better infrastructure, the new data that we have collected

driving through diverse terrain and unique scenarios is more broadly applicable than before.

A. The Platform

We use the same Yamaha Viking All-Terrain Vehicle (ATV) as before (Fig. 2), first modified by Mai et al. [26]. We have modified it again in order to equip it with three LiDAR sensors. There are two Velodyne VLP-32 sensors mounted to the front of the roof of the vehicle, with one tilted downwards to increase coverage of the ground closer to the vehicle. There is also a Livox Mid-70, mounted under the MultiSense camera, which provides more information on objects directly in front of the vehicle. Using a Faro Focus Scanner, we generated a 3D pointcloud model of the ATV as shown in Fig. 3 and provide it with the dataset so that users can take their own measurements (e.g. distance between the wheel axle and some arbitrary point on the vehicle, or distance between the GPS antenna and the ground) for their own specific experiments.

B. Data Collection

The site for collecting data is the same location in western Pennsylvania as TartanDrive 1.0, consisting of terrain such as narrow paths, dense foliage, rocky terrain, dirt paths, and steep hills. It is worth noting that some areas of the site have changed significantly due to natural causes such as erosion or overgrowth as shown in Fig. 4. All the data included was acquired by a human tele-operating the vehicle. Each sequence of data is annotated with the following metadata:

- Driver ID and robot
- Number of people in the vehicle
- Date/time
- Context (e.g. data collection)
- Weather conditions (e.g. dry, damp, snow)
- Lighting conditions (e.g. sunny, overcast, sunset)
- Course ID or general location (from a pre-defined list)

as well as any other information relevant to that specific run. We emphasize that recording this metadata makes the dataset significantly easier to use, especially as the dataset becomes larger. At the end of the collection, a post-processor records information such as top speed, average speed, duration, and sensors present. During data collection, a co-pilot takes time-stamped annotations of relevant or unique events such as sensor failure or something uncommon such as a deer spotting in front of the vehicle. Additionally, in some runs the co-pilot periodically annotates a weak driving score ranging from 1-5 where 5 signifies ideal driving. Recording this metadata provides more structure in the data which significantly improves the utility of the dataset when used in combination with our post-processing pipelines.

C. Raw Data

Most of the sensors on our platform have already been detailed [5,27], but all raw data is briefly summarized below:




Fig. 4: The data collection site has changed significantly since our last dataset released in 2021. For example, some older dirt paths (top, shown in red) are now covered in tall grass (bottom).




Fig. 5: Example coverage provided by our LiDAR sensors. Purple and yellow points come from the two Velodynes, and cyan points from the Livox.

1) *Pointclouds*: We record incoming pointclouds from two Velodyne VLP-32 LiDAR sensors and a Livox Mid-70. We also provide extrinsics so that they can be merged into a single pointcloud. An example of the coverage they provide is shown in Fig. 5. All LiDARs are configured to run at 10Hz.

2) *Images*: A Carnegie Robotics MultiSense S21 is used to provide stereo images, specifically greyscale images from both cameras as well as RGB images from the left camera at 10Hz each.

3) *IMU and Pose*: A NovAtel PROPAK-V3-RT2i GNSS provides IMU data at 100Hz. This is fused with incoming GPS data to provide a pose estimate at 50Hz. The MultiSense also provides IMU data at 400Hz.

4) *Teleoperation*: The driver uses a joystick controller to send steering commands, which are then recorded alongside the values of the current steering angle (actuation on steering commands is not instantaneous). A Racepak G2X Pro Data Logger is used to provide the positions of the acceleration and brake pedals.

5) *Proprioceptive Information*: The Racepak is also used to record RPMs for each wheel and suspension shock travel for the rear two wheels. Shock travel data for the front two

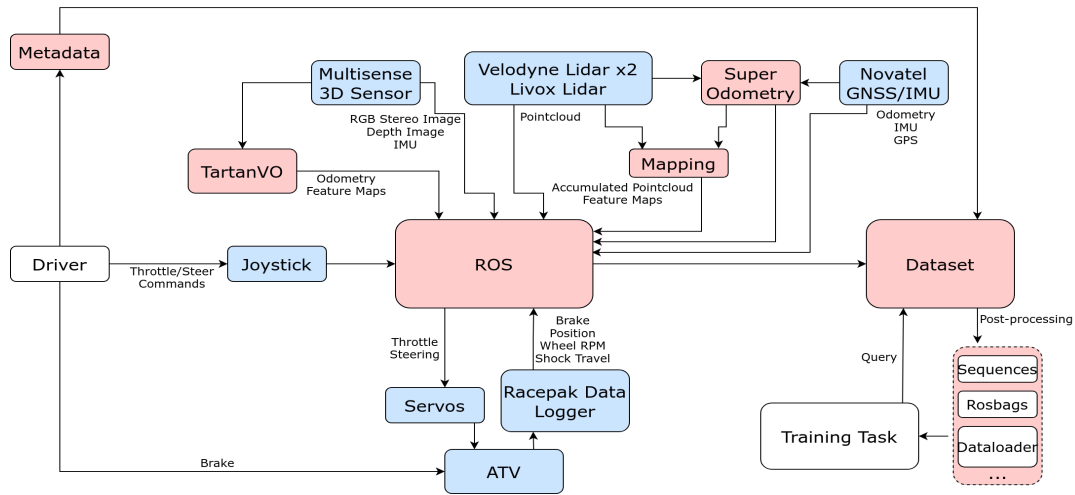


Fig. 6: The high-level flow of information in our data collection process, with software in red and hardware in blue.

wheels is omitted due to frequent damage to the sensors during aggressive driving maneuvers, but will be included in future data releases.

D. Post-Processed Data

In order to increase utility and ease-of-use of the dataset, we provide the raw data and the outputs from some existing modules that have been integrated into our software stack (Fig. 6) and commonly use as inputs to our own algorithms:

1) *TartanVO*: We run TartanVO [28] on the platform, which takes in the stereo images from the MultiSense as input. It outputs an odometry estimate, a predicted pointcloud, and top-down height and RGB maps, all at 10Hz.

2) *Roughness Cost*: We produce a roughness cost that describes the bumpiness of terrain as we drive over it, derived from a sliding window of Z-axis linear acceleration values from the IMU as described in Guaman Castro et al. [20].

3) *Odometry and Registered Pointcloud*: In addition to the odometry provided by the Novatel system, we also provide an estimated output by Super Odometry [29]. The high accuracy and frequency of this output allows for cleaner results in tasks such as pointcloud registration which is in turn important for other downstream tasks. It also serves as a baseline for others to benchmark their own odometry methods.

4) *Local LiDAR Maps*: Using the registered pointcloud provided by Super Odometry, we generate a birds-eye-view feature map 200x200m wide at .5m resolution. This map provides geometric information about the environment (Fig. 7), with features consisting of the following:

- Min/Max/Mean Height of Points
- Roughness
- SVD Features
- Estimated Ground Height
- Estimated Ground Slope (X, Y, Magnitude)

E. Data Pipelines

1) *Formatting*: We provide the data in two main formats, the first one being the rosbags that they were originally




Fig. 7: Using the registered pointcloud, we provide local maps that contain various geometric features.

recorded as. This allows users to test how their algorithms might perform in different environments in real-time. The second format is as a set of sequences similar to the KITTI format [30]. For each bag, a folder is created and a subfolder for each modality is initialized. The samples across all modalities are then timesynced and then placed in their respective subfolders (e.g. 'pointcloud_1/0000.bin' and 'image/0000.jpg' are associated with each other).

2) *Reconfiguration*: We have post-processed our data in a way that makes sense for our own algorithms, but our ATV platform is somewhat unique. Across different robots, the optimal parameters for elements such as sample frequency, map size, and map resolution vary. For example, a smaller robot with better agility might require a map with finer than .5m resolution. In order to make our data more directly applicable to other platforms, we provide scripts that allow users to regenerate the dataset with different parameters as they see fit.

3) *Utilities*: As previously mentioned, we collect metadata and some annotations during data collection. We provide scripts to take advantage of this metadata in order to filter the data by various elements and therefore increase the utility of the dataset for different learning tasks. For example, given the whole dataset, a user can easily create subsets grouped by components such as speed, GPS bounding box (we provide a GUI for easily generating bounding boxes),

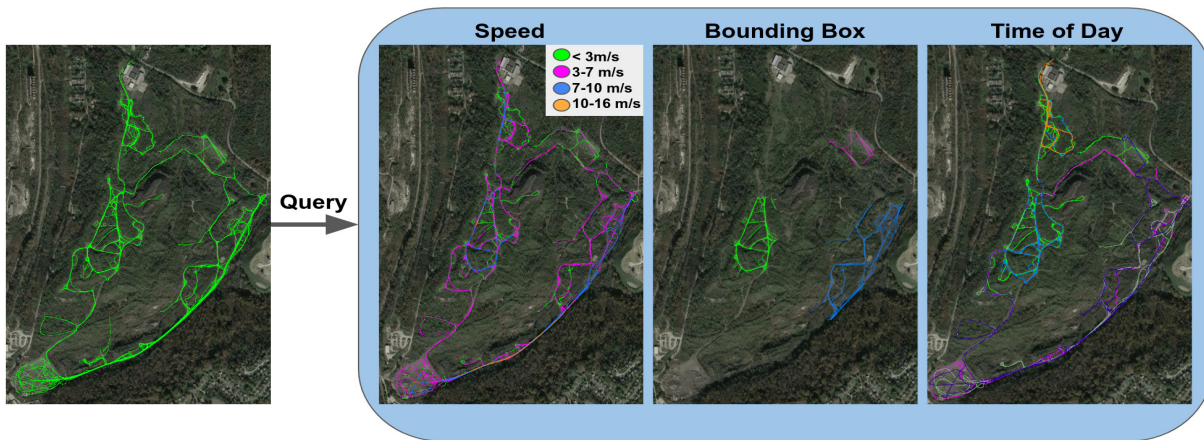


Fig. 8: Our dataset covers over 255 acres collected over seven hours. Using the accompanying metadata, the it can easily be split into groups by aspects such as speed, GPS bounding box, and time of day.




Fig. 9: In cases where conditions like poor lighting affect camera performance, LiDARs can allow a robot to still understand its surroundings.

lighting conditions, or driver (Fig. 8). This is especially useful for testing model generalizability by training on one location/condition, and testing on another.

4) *Common Framework*: The infrastructure we have developed has streamlined our data collection procedures to make it easy to train models on processed datasets as well as test our algorithms in real-time. By continuing to follow the same procedures, the data we collect in the future can easily be merged into our existing datasets. Likewise, other researchers can use our tools when collecting their own off-road driving data which could eventually lead to a larger multi-site, and multi-vehicle dataset.

IV. IMPACT ON CURRENT AND FUTURE RESEARCH

While we do not provide any explicit supervision labels (e.g. segmentation masks, classification labels) outside of annotations, the several modalities and actions that we provide facilitate a number of self-supervised learning tasks. Our original TartanDrive has already benefited off-road research immensely, and we believe the categories listed in Table I



Fig. 10: An example of the overlap between a single scan from all 3 LiDARS and our camera. Points are colored by forward distance from the vehicle.

are only a subset of the research topics TartanDrive 2.0 can bolster.

A. Perception - Cross-Modal Supervision:

Our data was used by Guaman et al. to visually predict a bumpiness cost supervised by an IMU-derived metric [20]. With the addition of LiDAR in our dataset, we can advance algorithms by providing information in scenarios where cameras fail (Fig. 9) and learning from a much more accurate source of odometry and depth (Fig. 10). For example, Chen et al. accumulates near-range LiDAR measurements to learn long-range visual traversability model [19]. Meng et al. train a model that uses images to predict feature maps supervised by lidar inputs [18]. However, these works rely on off-road driving datasets that are not publicly available. By releasing data in a similar domain with the same modalities, we lower the barrier to entry for expanding on this field of research, and allow a common point of comparison.

Dataset	State	Action	Image	Pointcloud	Heightmap	RGBmap	IMU	Wheel RPM	Shocks	Metadata	Labels
RUGD [7]	No	No	Yes	No	No	No	No	No	No	No	Yes
Rellis 3D [6]	Yes	Yes	Yes	Yes	No	No	Yes	No	No	No	Yes
Wild-Places [11]	Yes	No	No	Yes	No	No	No	No	No	No	Yes
Montmorency [31]	Yes	Yes	Yes	Yes	Yes	No	Yes	No	No	No	Yes
Verti-Wheelers [10]	Yes	Yes	Yes	No	No	No	Yes	Yes	No	No	No
TartanDrive 1.0 [5]	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes	No	No
TartanDrive 2.0 (Ours)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No

TABLE I: Overview and comparison of various off-road driving datasets

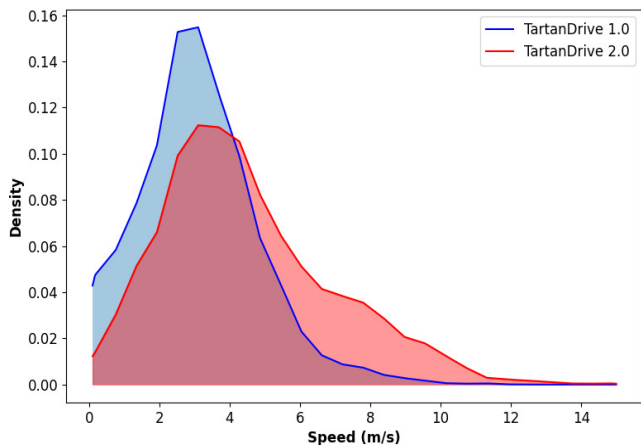


Fig. 11: In TartanDrive 2.0 we collect a higher concentration of data at higher speeds than in TartanDrive 1.0, reaching up to 15m/s.

B. Perception - Map Completion:

Prediction of occluded and sparsely sensed areas enables faster safe navigation in occluded area [18,32]. Our generated accumulated maps are being used as a supervisory signal in our ongoing research in map completion and can also serve as a common benchmark for other approaches.

C. Perception - Ground Height Estimation:

An understanding of the support ground surface in off-road terrain is essential for successful navigation. We enable methods of learning this feature by providing multiple modalities as inputs and a 3D model of the car that can be used alongside odometry estimates to provide supervision of the true ground surface.

D. Planning - Learning from Demonstration:

The actions derived from human teleoperation coupled with sensor inputs can be used to learn models that can reason about what areas are more traversable than others. Triest et al. treats teleoperated driving as expert demonstrations and use inverse reinforcement learning to birds-eye-view costmaps from the LiDAR feature maps [21].

E. Controls - Aggressive Maneuver Driving:

The unstructured nature of complex terrain makes it difficult to drive aggressively at high speeds without losing control. Our new dataset contains a higher proportion of speeds beyond 7m/s than before. Some of our ongoing research on learning vehicle dynamics models is supported by this data.

V. CONCLUSIONS AND FUTURE WORK

We have presented TartanDrive 2.0, which provides significant improvements over the original TartanDrive. We emphasize how the inclusion of more modalities can benefit research in self-supervised learning, and how more data and better infrastructure allow it to scale up to the needs of large models. The tools that we release with the dataset have streamlined our data collection, facilitating our plans of collecting more data on our existing site as the seasons change and then releasing that data over time.

Despite the major improvements in this generation of our dataset, there remains a number of directions to pursue in order to further increase its quality and utility. For example, while LiDAR sensors are useful in much of modern robotics research, there has also been an increasing amount of research in using passive sensors for perception [33,34]. To that end it would be helpful to add sensors like thermal cameras in the near future. Additionally, while self-supervision is the focus of this work, it may be worth hand-labelling (e.g. semantic masks) a subsets of the data for validation.

We also plan to include audio and language in our data. There is an increasing amount of research in foundation models and vision-language models [35,36], some of which have already demonstrated improvement in some robotics applications [37–40]. We hypothesize that learning the relationship between driver dialogue and their actions and/or the resulting environmental interactions might be supplemental in training large neural networks. Apart from recording dialogue, we also plan to record audio of the interactions between tires and the ground as we drive through different terrains, as Zurn et al. have shown that this information can be used to learn an understanding of complex terrain [41].

Additionally, our test site is very large and complex, but we have found that our own algorithms are close to enabling the ATV to successfully traverse much of the terrain. This brings up the issue of overfitting, and whether or not our algorithms would be able to reach the same level of performance in unseen, challenging terrain. By releasing our tools alongside the dataset we hope that others will contribute data from other locations in order to collaboratively form an even larger dataset. However, regardless of external contributions, we plan to continue collect data throughout various seasons and in other off-road sites with different types of challenging terrain that will allow ourselves and others to test generalizability and online adaptation of our algorithms.

REFERENCES

- [1] P. Sun, H. Kretschmar, X. Dotiwalla *et al.*, “Scalability in perception for autonomous driving: Waymo open dataset,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [2] S. Ettinger, S. Cheng, B. Caine *et al.*, “Large scale interactive motion forecasting for autonomous driving: The waymo open motion dataset,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021, pp. 9710–9719.
- [3] H. Caesar, V. Bankiti, A. H. Lang *et al.*, “nuscenes: A multimodal dataset for autonomous driving,” in *CVPR*, 2020.
- [4] W. K. Fong, R. Mohan, J. V. Hurtado *et al.*, “Panoptic nuscenes: A large-scale benchmark for lidar panoptic segmentation and tracking,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 3795–3802, 2022.
- [5] S. Triest, M. Sivaprakasam, S. J. Wang *et al.*, “Tartandrive: A large-scale dataset for learning off-road dynamics models,” in *2022 International Conference on Robotics and Automation (ICRA)*, 2022, pp. 2546–2552.
- [6] P. Jiang, P. Osteen, M. Wigness, and S. Saripalli, “Relis-3d dataset: Data, benchmarks and analysis,” 2020.
- [7] M. Wigness, S. Eum, J. G. Rogers *et al.*, “A rugd dataset for autonomous navigation and visual perception in unstructured outdoor environments,” in *International Conference on Intelligent Robots and Systems (IROS)*, 2019.
- [8] S. Sharma, L. Dabir, T. Hannis *et al.*, “Cat: Cava traversability dataset for off-road autonomous driving,” *IEEE Access*, vol. 10, pp. 24 759–24 768, 2022.
- [9] D. Maturana, P.-W. Chou, M. Uenoyama, and S. Scherer, “Real-time semantic mapping for autonomous off-road navigation,” in *Field and Service Robotics*, M. Hutter and R. Siegwart, Eds. Cham: Springer International Publishing, 2018, pp. 335–350.
- [10] A. Datar, C. Pan, M. Nazeri, and X. Xiao, “Toward wheeled mobility on vertically challenging terrain: Platforms, datasets, and algorithms,” 2023.
- [11] J. Knights, K. Vidanapathirana, M. Ramezani *et al.*, “Wild-places: A large-scale dataset for lidar place recognition in unstructured natural environments,” 2023.
- [12] A. Reinke, M. Palieri, B. Morrell *et al.*, “Locus 2.0: Robust and computationally efficient lidar odometry for real-time 3d mapping,” *IEEE Robotics and Automation Letters*, pp. 1–8, 2022.
- [13] A. Shaban, X. Meng, J. Lee *et al.*, “Semantic terrain classification for off-road autonomous driving,” in *Proceedings of the 5th Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, A. Faust, D. Hsu, and G. Neumann, Eds., vol. 164. PMLR, 08–11 Nov 2022, pp. 619–629. [Online]. Available: <https://proceedings.mlr.press/v164/shaban22a.html>
- [14] Z. Yang, Y. Tan, S. Sen *et al.*, “Uncertainty-aware perception models for off-road autonomous unmanned ground vehicles,” *CoRR*, vol. abs/2209.11115, 2022. [Online]. Available: <https://doi.org/10.48550/arXiv.2209.11115>
- [15] J.-B. Grill, F. Strub, F. Altché *et al.*, “Bootstrap your own latent - a new approach to self-supervised learning,” in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell *et al.*, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 21 271–21 284. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2020/file/f3ada80d5c4ee70142b17b8192b2958e-Paper.pdf
- [16] M. Oquab, T. Darcet, T. Moutakanni *et al.*, “Dinov2: Learning robust visual features without supervision,” 2023.
- [17] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, “A simple framework for contrastive learning of visual representations,” in *Proceedings of the 37th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, H. D. III and A. Singh, Eds., vol. 119. PMLR, 13–18 Jul 2020, pp. 1597–1607. [Online]. Available: <https://proceedings.mlr.press/v119/chen20j.html>
- [18] X. Meng, N. Hatch, A. Lambert *et al.*, “Terrainnet: Visual modeling of complex terrain for high-speed, off-road navigation,” 2023.
- [19] E. Chen, C. Ho, M. Maulimov *et al.*, “Learning-on-the-drive: Self-supervised adaptation of visual offroad traversability models,” 2023.
- [20] M. G. Castro, S. Triest, W. Wang *et al.*, “How does it feel? self-supervised costmap learning for off-road vehicle traversability,” 2023.
- [21] S. Triest, M. G. Castro, P. Maheshwari *et al.*, “Learning risk-aware costmaps via inverse reinforcement learning for off-road navigation,” 2023.
- [22] A. Kulkarni, J. Chrosniak, E. Ducote *et al.*, “Racecar – the dataset for high-speed autonomous racing,” 2023.
- [23] J.-F. Tremblay, T. Manderson, A. Noca *et al.*, “Multimodal dynamics modeling for off-road autonomous vehicles,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 1796–1802.
- [24] M. Sivaprakasam, S. Triest, W. Wang *et al.*, “Improving off-road planning techniques with learned costs from physical interactions,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 4844–4850.
- [25] D. Shah, A. Sridhar, N. Dashora *et al.*, “ViNT: A Foundation Model for Visual Navigation,” *arXiv pre-print*, 2023. [Online]. Available: <https://arxiv.org/abs/2306.14846>
- [26] J. Mai, “System design, modelling, and control for an off-road autonomous ground vehicle,” Master’s thesis, Carnegie Mellon University, Pittsburgh, PA, July 2020.
- [27] M. Sivaprakasam, S. Triest, M. G. Castro *et al.*, “Tartandrive 1.5: Improving large multimodal robotics dataset collection and distribution,” in *ICRA2023 Workshop on Pretraining for Robotics (PT4R)*, 2023. [Online]. Available: <https://openreview.net/forum?id=7Y1pnhFJ23>
- [28] W. Wang, Y. Hu, and S. Scherer, “Tartanvo: A generalizable learning-based vo,” *Conference on Robot Learning (CoRL)*, 2020.
- [29] S. Zhao, H. Zhang, P. Wang *et al.*, “Super odometry: Imu-centric lidar-visual-inertial estimator for challenging environments,” in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 8729–8736.
- [30] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, “Vision meets robotics: The kitti dataset,” *International Journal of Robotics Research (IJRR)*, 2013.
- [31] J.-F. Tremblay, M. Béland, F. Pomerleau *et al.*, “Automatic 3d mapping for tree diameter measurements in inventory operations,” *arXiv preprint arXiv:1904.05281*, 2019.
- [32] R. Schmid, D. Atha, F. Schöller *et al.*, “Self-supervised traversability prediction by learning to reconstruct safe terrain,” 2022.
- [33] A. Srivastav and S. Mandal, “Radars for autonomous driving: A review of deep learning methods and challenges,” 2023.
- [34] M. Brenner, N. H. Reyes, T. Susnjak, and A. L. C. Barczak, “RGB-d and thermal sensor fusion: A systematic literature review,” *IEEE Access*, vol. 11, pp. 82 410–82 442, 2023. [Online]. Available: <https://doi.org/10.1109/Access.2023.3301119>
- [35] OpenAI, “Gpt-4 technical report,” 2023.
- [36] A. Radford, J. W. Kim, C. Hallacy *et al.*, “Learning transferable visual models from natural language supervision,” in *Proceedings of the 38th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, M. Meila and T. Zhang, Eds., vol. 139. PMLR, 18–24 Jul 2021, pp. 8748–8763. [Online]. Available: <https://proceedings.mlr.press/v139/radford21a.html>
- [37] A. Brohan, N. Brown, J. Carbajal *et al.*, “Rt-1: Robotics transformer for real-world control at scale,” in *arXiv preprint arXiv:2212.06817*, 2022.
- [38] S. Vemprala, R. Bonatti, A. Bucker, and A. Kapoor, “Chatgpt for robotics: Design principles and model abilities,” Microsoft, Tech. Rep. MSR-TR-2023-8, February 2023. [Online]. Available: <https://www.microsoft.com/en-us/research/publication/chatgpt-for-robotics-design-principles-and-model-abilities/>
- [39] M. Shridhar, L. Manuelli, and D. Fox, “Cliport: What and where pathways for robotic manipulation,” in *Conference on Robot Learning*. PMLR, 2022, pp. 894–906.
- [40] D. Shah, B. Osiński, S. Levine *et al.*, “Lm-nav: Robotic navigation with large pre-trained models of language, vision, and action,” in *Conference on Robot Learning*. PMLR, 2023, pp. 492–504.
- [41] J. Zürn, W. Burgard, and A. Valada, “Self-supervised visual terrain classification from unsupervised acoustic feature learning,” *IEEE Transactions on Robotics*, vol. 37, no. 2, pp. 466–481, 2021.