

TNS: Terrain Traversability Mapping and Navigation System for Autonomous Excavators

Tianrui Guan^{1*} Zhenpeng He¹ Ruitao Song¹ Dinesh Manocha² Liangjun Zhang¹

¹ Robotics and Auto-Driving Laboratory, Baidu Research

² University of Maryland, College Park

Abstract—We present a terrain traversability mapping and navigation system (TNS) for autonomous excavator applications in an unstructured environment. We use an efficient approach to extract terrain features from RGB images and 3D point clouds and incorporate them into a global map for planning and navigation. Our system can adapt to changing environments and update the terrain information in real-time. Moreover, we present a novel dataset, the Complex Worksite Terrain (CWT) dataset, which consists of RGB images from construction sites with seven categories based on navigability. Our novel algorithms improve the mapping accuracy over previous SOTA methods by 4.17 – 30.48% and reduce MSE on the traversability map by 13.8 – 71.4%. We have combined our mapping approach with planning and control modules in an autonomous excavator navigation system and observe 49.3% improvement in the overall success rate. Based on TNS, we demonstrate the first autonomous excavator that can navigate through unstructured environments consisting of deep pits, steep hills, rock piles, and other complex terrain features. Dataset, videos, and a full technical report are available at gamma.umd.edu/tns/.

I. INTRODUCTION

Excavators are one of the most common types of heavy-duty machinery used for earth-moving activities, including mining, construction, environmental restoration, etc. As the demand for excavators increases, many autonomous excavator systems [57, 24, 42] have been proposed for material loading tasks, which involve perception and motion planning techniques.

Some of the major issues in terms of using autonomous excavators are the development of robust perception and navigation sub-systems. In general, perception in unstructured environments such as excavation has many challenges. There have been many works related to unstructured environments, including perception and terrain classification [15, 49, 44] and navigation [28, 21, 35, 26]. Applications in unstructured, hazardous environments have even more difficulties in terms of robustness and limitations on the computational budget. For example, many accurate learning methods have been proposed to improve the perception capabilities, but we cannot assume access to large GPUs or clusters for excavators operating in hazardous environments. Instead, we need to develop robust methods with lower computational requirements.

Traversability is a term that encompasses both perception and navigation. It has been well-studied for decade, and there have been many works [8, 11, 58, 45, 30] on traversability estimation for planning and navigation. Terrain traversability is a binary value, or a probability score, measuring the difficulty




Fig. 1: Overview of our system TNS: **Top left**: Sensors on the excavator, including RGB cameras and Livox LiDAR. **Top right**: Detection region from a third-person perspective. **Middle left**: Frontal view captured by the camera. **Middle right**: Semantic segmentation output, where green, yellow, and maroon correspond to flat region, bumpy region, and rock, respectively. **Bottom left**: Colored point cloud with semantic labels. **Bottom right**: Terrain traversability output, where the traversability value decreases from green to grey. This output is used for automatic navigation in complex, outdoor environments.

of navigating a region through perception sensors like camera, LiDAR and IMU. Terrain traversability estimation is a critical step between perception and navigation. In many autonomous driving (AD) cases [35, 16], a method capable of detecting obstacles and distinguishing road and non-road regions is sufficient for navigation. On the other hand, in an unstructured, hazardous environment where off-road navigation is unavoidable, there are many factors that must be considered, including efficiency, adaptability, and safety. In such cases, not only a more detailed classification according to terrain features is needed, but also a continuous value for traversability is preferred to describe the complexity of the terrain and provide the best option for the navigation module. Therefore, we

* Work done during an internship at Baidu RAL.

need good techniques to detect traversable regions for reliable navigation in an unstructured scene.

Main Results: We present a terrain traversability mapping and navigation system (TNS) for traversability classification and autonomous navigation. We describe an efficient semantic-geometric fusion method to extract traversability maps. Our method leverages the physical and computational constraints of the robot, including maximum climbing degree, width of the body, run-time computational budget, etc. The novel aspects of our approach include:

- 1) We present a real-time terrain traversability estimation and navigation system (TNS) from 3D LiDAR and RGB camera inputs for mapping, planning, and navigation. We describe a novel learning-based geometric fusion solution that considers machine specifications and hardware limitations for terrain traversability prediction in unstructured environments. We show that our method is the state-of-the-art (SOTA) traversability mapping method on complex terrains. Our method outperforms previous SOTA methods by 4.17-30.48% in terms of mAcc and reduces the MSE by 13.8-71.4%.
- 2) We have integrated TNS with planning and control algorithms and evaluated the performance extensively in real-world settings on an autonomous excavator in various challenging construction scenes, as shown in Figure 1. We also elaborate on many non-trivial issues that came up during the implementation and evaluation and how we address them. We show that our TNS can safely navigate an excavator in unstructured environments and observe a 49% improvement in terms of planning success rate. We highlight the benefits of TNS as the first autonomous excavator that can navigate through complex, unstructured environments.
- 3) We present the Complex Worksite Terrain (CWT) dataset, which consists of 30 minutes of video and 669 RGB images in unstructured environments with seven different classes based on terrain types, traversable regions, and obstacles. We will release the CWT dataset in the public domain.

II. RELATED WORK

A. Field Robots and Systems

Field robots usually refer to machines that operate in off-road, hazardous environments. These include heavy-duty service robots for industrial usage in mining [43], excavation [57], agriculture [36], construction [31], etc. To satisfy industrial needs and save labor costs, many automated systems [57, 24, 42] have been developed for service robots in the field. These systems include modules for perception, planning and control. However, it remains a challenge to fully automate many tasks in unknown, unstructured environments.

B. Terrain Traversability Recognition

The concept of traversability, also referred to as “drivability,” “navigability,” etc. [32], has been studied for decades. There are many viewpoints on the problems and challenges

associated with traversability, and investigations into such topics have had different evaluation methods and goals. Many works focus on getting correct predictions of the terrain [29, 55, 47, 40, 25, 19, 7, 15, 17, 12] by some notion of ground truth based on human-labeled annotation, similar to the metrics of 2D and 3D semantic segmentation. Most of the methods mentioned above are based on visual features of the terrain, which sometimes lack the properties that enable real-world navigation due to recognition failure.

On the other hand, some works focus on obtaining traversability maps that result in the best navigation outcomes. There are plenty of works [33, 30, 2, 11, 5, 58] on classifying different terrains based on either material categories or navigability properties and demonstrate their mapping results through navigation outcomes.

However, those methods deal with structured roads or roads with clear path boundaries in unstructured environments. In more complex environments, point clouds obtained from LiDAR are used to extract geometric attributes of the surface, including slope, height variation, roughness, obstacles, etc., as proposed in [8, 6, 50, 60, 1, 16]. [41] uses both point cloud and RGB images to classify terrains with safe, risky, and obstacle labels in the 2D image plane for better performance. [39] presents a pipeline from perception to motion control and uses five different data sources for navigation, including range and intensity values from a 2D LiDAR and edge information from an RGB-D camera. [23, 22] analyze the terrain and create roadmaps for road safety. The works most similar to our proposed method are [11, 58, 45, 30], which focus on finding a better terrain representation for navigation in unstructured terrains.

C. Datasets for Unstructured Environments

Most recent developments in perception tasks like object detection and semantic segmentation focus on urban driving scene datasets like KITTI [14], Waymo [46], etc., which achieve high accuracy in terms of average precision. On the other hand, unstructured scenes like the natural environment, construction sites, and complicated traffic scenarios are less explored, for two primary reasons. First, there are fewer datasets with unstructured environments; second, perception and autonomous navigation in unstructured off-road environments are challenging due to unpredictability and diverse terrain types.

Recent efforts in off-road perception and navigation include RUGD [51] and RELLIS-3D [20], which are semantic segmentation datasets collected from a robot navigating in off-road and natural environments. These datasets contain scenes like trails, forests, creeks, etc. [38] is a construction dataset containing annotations of heavy-duty vehicles for detection, tracking, and activity classifications.

III. PERCEPTION FOR AUTONOMOUS EXCAVATORS

The road conditions in structured environments such as highways are usually navigation-friendly, so the core problem during navigation in structured environments is avoiding obstacles rather than determining which part of the surface

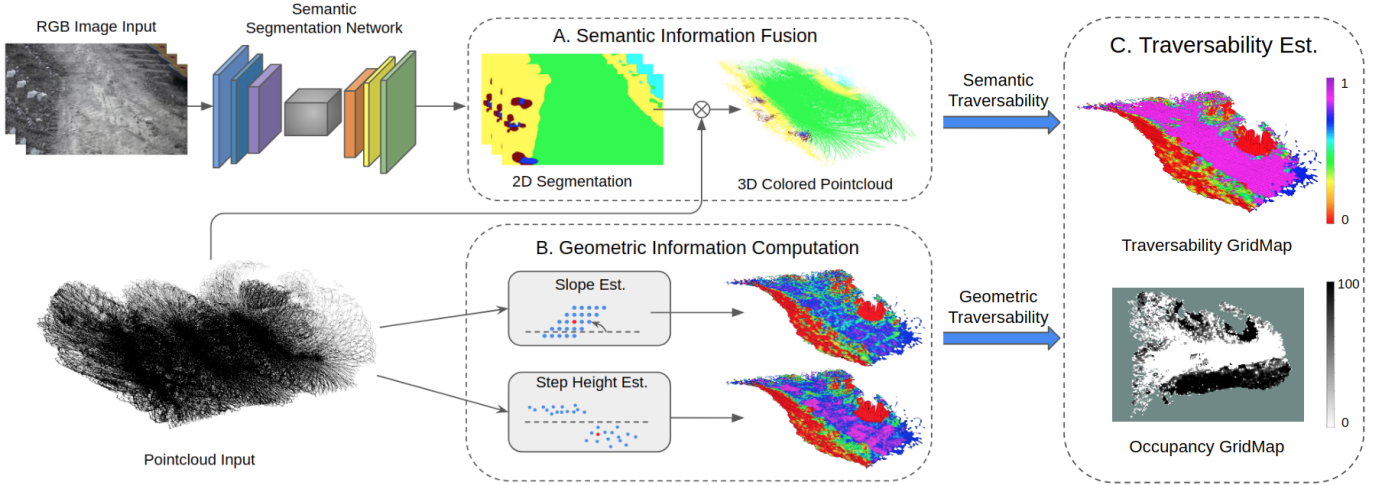


Fig. 2: **Overview of the perception module in TNS:** Our system takes RGB images and point clouds as inputs to infer traversability. We extract semantic information using segmentation and associate terrain labels with point clouds, as shown in A (top). We extract geometric information using slope and step height estimation, as shown in B (bottom). We produce a traversability grid map based on semantic and geometric information and convert it to a 2D occupancy map for path planning and navigation, as shown in C (right).

is easier and safer to navigate. In contrast, excavators are usually operated in unstructured and dangerous environments consisting of rock piles, cliffs, deep pits, steep hills, etc. Such an environment lacks any lane markings, and the arrangement of obstacles tends to be non-uniform. In addition, due to tasks like digging and dumping, the working conditions for excavators are constantly changing. Landfalls and cave-ins occur, potentially causing the excavator to tip over and injure the operator. Therefore, it is crucial to identify different terrains and predict safe regions for navigation. Furthermore, we need solutions with low computational requirements.

In our context, traversability [32] refers to the capability of a ground vehicle to reside over a region of terrain under an admissible state wherein it can enter given its current state. In order to solve navigation challenges for excavators as well as other working vehicles in unstructured terrain, we formulate the problem of obtaining an accurate traversability map representation as follows:

Problem Definition: Given sensor inputs S_1, S_2, \dots, S_h from h different sources over a time span T , the goal is to obtain a 2D grid map $T \in [0, 1]^{H \times W}$ with resolution r , where T corresponds to some region R of shape (Hr, Wr) . The maximum value corresponds to a non-traversable region and the minimum value corresponds to the most traversable region.

Metrics for Traversability Map: We need to consider the following measurements in excavator applications:

- **Accuracy:** Similar to [41, 11, 45], we use an ROC curve to measure the accuracy of the traversability prediction. In addition, the map output should fit the terrain closely, so we also use MSE (mean squared error) as a fitness measurement.
- **Performance:** [30, 58] use navigation outcome to measure their terrain traversability mapping algorithms, which include travel time, success rate, etc.
- **Energy constraints and run-time:** Due to the limitations of hardware and power supply on the excavator, energy




Fig. 3: **Overall pipeline of TNS for autonomous excavator navigation:** We show different components of TNS as blue blocks and use green and red blocks to represent the intermediate output and hardware, respectively.

efficiency and run-time computational budget should also be measured in a terrain traversability mapping method.

IV. TNS: SYSTEM ARCHITECTURE

In this section, we describe our system for terrain traversability mapping and navigation (TNS) in excavator applications, as shown in Figure 3. TNS takes a 3D point cloud stream from the LiDAR, an RGB camera stream from the RGB camera, and the corresponding poses of the excavator extracted from the GPS-RTK module. The goal of our proposed system is to identify safe, navigable regions for excavators and autonomously navigate the excavator based on the traversability map and the planned trajectory. The output of TNS includes a global map consisting of terrain information, including semantic information, geometric information, and a final traversability score, as well as the planned trajectory.

A. Traversability Mapping

The terrain is represented as an elevation grid map and is updated in real-time based on incoming point clouds and RGB images. Internally, each grid cell in the map stores the average height value of the latest p points within this cell, as well as overall information about those points like update time, slope, step height, and their semantic information. A traversability

score is calculated for each grid cell. In Figure 2, we present an overview of our perception approach. Our implementation is based on the open-source grid map library [13].

Segmentation and Mapping to Point Cloud: We use 2D semantic segmentation on unstructured terrains. Given an input RGB image $I \in \mathbb{R}^{3 \times H \times W}$, the goal is to generate a mask $P \in \{0, 1, \dots, N-1\}^{H \times W}$, where N is the number of classes. We use Fast-SCNN [34] after leveraging accuracy and efficiency, as shown in Table II.

After we get the segmentation prediction P , we use a timestamp to locate the corresponding point cloud C and use camera calibration matrices to find the correspondence of each point to the segmentation results and save the terrain label in the grid map cell.

Geometric Information Computation: In this section, we present details of slope and step height estimation and highlight how machine specifications are considered to calculate the geometric traversability score.

1) *Slope Estimation:* Each grid cell g is abstracted to a single point $p = \{x, y, z\}$, where x, y is the center of the cell in the global coordinate frame and z is the height value of the grid. The slope s in arbitrary grid cell g is computed by the angle between the surface normal and the z-axis¹ of the global coordinate frame:

$$s = \arccos(n^z), n^z \in [0, 1]$$

where n^z is the component of normal \vec{n} on the z-axis.

Similar to [8, 3], we use Principal Component Analysis (PCA) to calculate the normal direction of a grid cell. The covariance matrix C_{cov} of the nearest neighbors of the query grid cell is calculated as follows:

$$C_{cov} = \frac{1}{k} \sum_{i=1}^k (p_i - \bar{p}) \cdot (p_i - \bar{p})^T, C_{cov} \cdot \vec{v}_j = \lambda_j \cdot \vec{v}_j,$$

$$j \in \{0, 1, 2\}, \lambda_i < \lambda_j \text{ if } i < j,$$

where k is the number of neighbors considered in the neighborhood of g , $p_i = \{x, y, z\}$ is the position of the neighbor grid in the global coordinate frame, \bar{p} is the 3D centroid of the neighbors, λ_j is the j -th eigenvalue of the covariance matrix, and \vec{v}_j is the j -th eigenvector. The surface normal \vec{n} of grid g is the eigenvector \vec{v}_0 with the smallest absolute value of eigenvalue λ_0 .

The purpose of the slope estimation is to get the shape of the terrain and avoid navigating on a steep surface. For excavator applications, the width between the tracks or wheels is a good indicator of the navigation stability on rough terrain. Usually, when the area of a rough region is less than half the width between the excavator's tracks, the excavator can navigate through it without any trouble. Specifically in our excavator setup, the width of our excavator track is 0.6 m, so we chose the grid resolution $d_{res} = 0.2$ m and search the nearest eight neighbors, which covers the necessary area.

2) *Step Height Estimation:* The step height h is computed as the largest height difference between the center point p of the grid and its k' nearest neighbors:

$$h = \max(\text{abs}(p^z - p_i^z)), i \in [1, k']$$

Since slope is a description of variation in the terrain in a relatively small region, we choose to use a neighbor search parameter $k' = 7 * 7 > k$ that spans 1.4 m to measure height change in a larger scope. For excavator applications, the step height calculation guarantees that the track does not traverse areas with extreme height differences.

3) *Geometric Traversability Estimation:* Based on information about slope and step height of the terrain, we can calculate a geometric traversability score T_{geo} . According to the physical constraints of the robot, we create some critical values, $s_{cri}, s_{safe}, h_{cri}, h_{safe}$, as the thresholds for safety and danger detection. The purpose of those threshold values is to avoid danger when the surface condition exceeds the limits of the robot and to avoid more calculations when the surface is very flat. The formula for geometric traversability T_{geo} for each grid is:

$$T_{geo} = \begin{cases} 0 & s > s_{cri} \text{ or } h > h_{cri} \\ 1 & s < s_{safe} \text{ and } h < h_{safe} \\ \max(1 - (\alpha_1 \frac{s}{s_{cri}} + \alpha_2 \frac{h}{h_{cri}}), 0) & \text{otherwise} \end{cases}$$

where the weights α_1 and α_2 sum up to 1.

The step height estimation is complementary to slope estimation; it provides a global perspective, whereas slope is local terrain information. Combining these two specifications can help us remove noise in the map, such as bumps caused by dust, and ensure the robustness of the T_{geo} .

Traversability with Geometric and Semantic Fusion: In this section, we describe our algorithm for geometric-semantic fusion. From the semantic and geometric information, we use a continuous traversability score $T \in [0, 1]$ to measure how easily the surface can be navigated. This is especially relevant to off-road scenarios because we prefer flat regions over bumpy roads to save energy. Moreover, when an excavator is navigating on a construction site, being able to correctly identify different regions is critical to avoid hazardous situations like flipping over.

The overall traversability score T is calculated based on semantic terrain classes C_{sem} and geometric traversability T_{geo} on each grid:

$$T = \begin{cases} 0 & C_{sem} = \{\text{rock, excavator, obstacle, water}\} \\ 1 & C_{sem} = \{\text{flat}\} \text{ and } T_{geo} > 0 \\ T_{geo} & \text{otherwise} \end{cases},$$

This method is simple yet more effective than other comparably complicated fusion methods [45, 58], as demonstrated in Section VI and Section VII.

B. Traversability-based Planning

We modify Hybrid A* [27] to calculate a trajectory based on the traversability map output after the post-processing step. Hybrid A* is a global path planner based on a 2D occupancy grid map as an input for trajectory planning. The planner will

¹Up direction in the real world

generate a trajectory and send it to the motion controller, which guides the excavator to follow this trajectory.

The traditional Hybrid A* algorithm only considers the traveling distance and certain driving maneuvers (such as reversing, turning, etc.), not the ground condition and traversability. As a result, the autonomous excavators can be easily navigated to areas with low traversability in real-world applications with the traditional Hybrid A* planner. To solve the problem, we extend the Hybrid A* algorithm by introducing TNS and calculating the traversability cost. Specifically, we calculate the cost to the start of a vertex, which is the distance from the start state to the vertex with extra reversing or turning cost, and is weighted by the traversability value obtained from TNS. In the improved Hybrid A* algorithm, the cost to start $g(x)$ is increased by $k_{TNS} \cdot \delta l + g_{extra}$ when performing vertex expansion from the parent to the child vertices, where δl is the distance between the two nodes, g_{extra} is the extra penalty for reversing and turning, and k_{TNS} is the traversability weighting factor calculated by:

$$k_{TNS} = \frac{k_t T_t A_t + k_u T_u A_u}{A_t + A_u},$$

where A_t and A_u are the areas covered by the two tracks and between two tracks, respectively, of the excavator from the parent to the child vertices; T_t and T_u are the mean traversability value of areas A_t and A_u ; and k_t and k_u are two calibrating parameters.

C. Control and Navigation

The trajectory tracking controller is composed of a lateral trajectory tracking controller and a longitudinal speed controller:

- **Tracking Controller:** This module can adjust the steering of the robot for path following. It outputs the desired steering rate based on the heading error and the cross-track error. The cross-track error is defined as the distance between the point on the path closest to the reference point of the excavator. The control commands for the left and right tracks of the excavator are calculated using a lookup table according to the speed proportional and integral (PI) error metric and the desired steering rate. The tracking controller is developed based on [18].
- **Speed Controller:** This module can adjust the speed of the robot. The speed controller receives the actual speed from the sensor and calculates the PI error metric according to the desired speed.

D. Benefits over Prior Methods

Previous perception methods for traversability calculation only use geometric approaches [8, 50, 3, 4] in simple scenarios for mobile robot applications, or they can only navigate in an off-road environment with a clear visual path [11, 58, 45, 30]. Our system is the first one to focus on excavator navigation applications in very challenging environments consisting of pits, hills, rock piles, etc. without a clear pathways. In addition, our experiments and data are based on real-world scenarios in a construction site. Our method also adapts to the physical

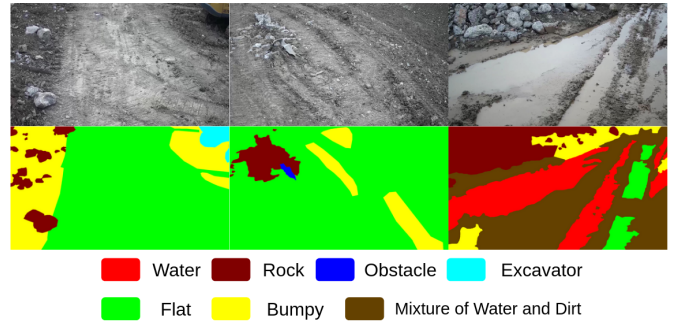


Fig. 4: **Complex Worksite Terrain (CWT) dataset:** We show a few samples from our CWT dataset (**top**) and corresponding annotations (**bottom**). All images are collected in unstructured environments with various terrain types.

constraints of excavators to determine threshold, resolution of the grid, and k neighbors.

We test our system TNS on an excavator based on the Autonomous Excavator System [57]. Note that previous AES systems mainly focus on digging tasks, while our system focuses on providing accurate mapping estimation and navigation in unstructured environments.

V. COMPLEX WORKSITE TERRAIN (CWT) DATASET

In this section, we present the Complex Worksite Terrain (CWT) dataset, which is collected at a construction site while an excavator is navigating through the work area. The hardware has the same setup as described in Section VII-A. We collect three videos (30 minutes in total) under different circumstances and annotate 669 images of size 1920×1080 according to terrain semantics. We only highlight the ontology and differences between CWT and other off-road datasets [20, 51], and provide details of the collection, class distribution, and analysis in the supplemental material.

The CWT dataset is annotated with seven labels based on terrain features and navigability, as shown in Table I. The annotation is decided based on the opinion of a team of excavator operators. In most cases, when flat surfaces are detected, they are preferable to other surfaces.

While the CWT dataset and other datasets like RUGD [51] and RELIS-3D [20] are collected in unstructured, outdoor environments, the CWT has several distinctions. As shown in Figure 4, the CWT dataset mostly consists of uneven terrain with unfavorable road conditions and covers many situations that might be encountered on a work site, including rock-piles, pits, stagnant water after rain, etc.

Types	Descriptions	Navigability	Distribution
Flat Region	Flat surfaces that most vehicles like cars can traverse.	Easy	41.76%
Bumpy Region	Bumpy surfaces that most vehicles can not traverse except working vehicles like excavators.	Medium	42.59%
Rock Pile	Very common on work-site; Need to be avoided most of the time.	Forbidden	6.51%
Water	Water might be trapped in deep trench after raining; Need to be avoided.	Forbidden	3.66%
Mixtures of Water and Dirt	Shallow water with mostly visible soil or dirt; Can be traversed.	Medium	4.90%
Excavator & Vehicles	Common vehicles that appear on work-site, like excavators.	Forbidden	0.35%
Obstacles	Uncommon objects that need to be avoided, like steel bar, sign block, etc.	Forbidden	0.23%

TABLE I: **CWT Ontology:** Classification of terrain features used in our approach

In addition, the CWT dataset focuses entirely on roads and terrains, and the annotation is based on terrain semantics instead of fine-grained semantics on every possible classes. Such annotation scheme is designed for the benefit of other downstream tasks, including planning and navigation for robots of any sizes, and excavation activities on hazardous terrains.

Overall, CWT presents many new challenges to the vision community to improve perception in hazardous environment, while providing support for autonomous robotics applications in dangerous environment. We demonstrate the difficulty of our dataset by showing the performances of several SOTA semantic segmentation methods on the CWT and existing off-road datasets like RELLIS-3D in Section VI-A. The CWT dataset can be accessed through this link.

VI. EXPERIMENTS AND EVALUATIONS

In Section VI-A, we show evaluation results for the semantic segmentation task on our CWT dataset and RELLIS-3D [20]. In Section VI-B, we evaluate our TNS on RELLIS-3D and show the benefits of our method compared to other SOTA mapping methods.

A. Perception Evaluation on the CWT Dataset

We show some evaluations using several SOTA segmentation methods on the CWT dataset and the RELLIS-3D dataset in Table II. The CWT dataset is a more challenging terrain dataset than RELLIS-3D. We also highlight the number of parameters and Giga-FLOPS (floating-point operations per second) as a measurement since energy efficiency is an important factor for robotic applications. The method and evaluation for segmentation is based on MMSeg [9].

B. Terrain Traversability Map Evaluation

In Table III, we evaluate the accuracy of our method and compare it with several SOTA traversability mapping methods on the RELLIS-3D dataset. We use the ground truth semantic labels from RELLIS-3D on a 3D point cloud and convert the labels to either 0 or 1 to indicate traversability on a grid map. During evaluation, we assume that the traversability map is based on the Clearpath Warthog, the same robot that collected the RELLIS-3D dataset: traversable regions like grass, dirt, concrete, and asphalt are set to 0, while puddles, bushes, and obstacles are set to 1. Even though our method outputs a continuous value between 0 and 1, we want to simplify the conversion between labels and traversability scores to avoid any biases.

1) *Comparisons*: Since many methods do not have publicly available codes, we implement their methods based on the papers, which can only run on an offline dataset and not in the real world. We compare our method with the following methods:

Dahlkamp et al. [11] use a Mixture of Gaussian Model to make a binary prediction on RGB images for traversable regions and make an inverse perspective transform to the world coordinates.

Methods	Params ↓	Dataset	mIoU ↑	mAcc ↑	Img Size	GFLOPs ↓
CGNet [53]	0.494 M	CWT	53.41	67.59	1920 x 1080	27.62
		RELLIS	65.9	79.25	1920 x 1200	30.67
Fast SCNN [34]	1.45 M	CWT	54.77	68.75	1920 x 1080	7.45
		RELLIS	69.27	80.99	1920 x 1200	8.03
Fast FCN [52]	68.7 M	CWT	41.68	51.85	1920 x 1080	1031.51
		RELLIS	68.24	79.21	1920 x 1200	1145.6
BiSeNetV2 [56]	14.77 M	CWT	54.37	67.05	1920 x 1080	97.51
		RELLIS	65.33	75.06	1920 x 1200	108.38
SETR* [59]	109.67 M	CWT	19.91	30.61	1920 x 1080	–
		RELLIS	65.53	76.57	1920 x 1200	–
		–	–	–	1024 x 512	337.46†
DPT* [37]	309.17 M	CWT	29.02	47.65	1920 x 1080	–
		RELLIS	55.38	66.23	1920 x 1200	–
		–	–	–	1024 x 512	424.87†
Segformer [54]	3.72 M	CWT	50.6	64.29	1920 x 1080	50.55†
		RELLIS	68.62	83.4	1920 x 1200	–

TABLE II: **Perception Accuracy on the CWT and RELLIS-3D [20] Dataset**: We list several SOTA semantic segmentation methods and train the model with 240K iterations. The CWT dataset has lower accuracy compared to the RELLIS dataset. * marks methods that do not converge well after 240K additional iterations. † marks the GFLOPs as an approximation and a lower bound.

Sock et al. [45] use a Linear Support Vector Machine for a 2-classes prediction and some mapping between terrain slope and a traversability score between 0 and 1. The final map is obtained through Bayes Fusion of terrain classification and slope information.

Zhao et al. [58] use a multi-class segmentation method based on RGB images and make projections onto a grid map for planning and navigation. Maturana et al. [30] use a distance transformation and update new observations with Bayes’s rule. **Geometric-based methods [8, 60]** only use geometric information from the point cloud for navigation tasks.

3D semantic segmentation [48, 10] methods are useful for classifying terrains. We obtain their inference results from the official repository of RELLIS-3D [20].

2) *Evaluation Metrics and Results*: We evaluate the traversability map based on offline data with four different metrics. In general, our method has better performance in terms of accuracy and MSE. Note that in the first three metrics, all traversability values are converted to either 0 or 1 for methods that have a continuous output. The metrics are described as follows:

Mean Accuracy: The average accuracy of traversable and non-traversable regions.

All Accuracy: Accuracy over all grids.

ROC (Receiver Operation Curve): Previous methods [11, 45] make binary predictions over each grid, so ROC is a common indicator of the performance through true positive and false positive rates, as shown in Figure 6.

MSE (Mean Squared Error): To describe how well the prediction fits the ground truth, we also calculate the average distance between the prediction and the ground truth over all grids.

VII. PERFORMANCE IN REAL-WORLD ENVIRONMENTS

In this section, we highlight the results on real-world environments and overall performance of our navigation system

Methods	Modality	Goal	Trav / Non-Trav Acc \uparrow	mAcc \uparrow	aAcc \uparrow	AUC \uparrow	MSE \downarrow
KPConv* [48]	LiDAR	3D segmentation	33.33 / 79.24	56.28	67.65	-	0.253
SalsaNet* [10]	LiDAR	3D segmentation	94.82 / 57.75	76.28	67.11	-	0.370
Chilian et al. [8]	LiDAR	Mapping & Navigation	66.19 / 88.29	77.24	82.17	0.790	0.155
Dahlkamp et al. [11]	RGB Camera	Navigation	4.41 / 99.91	52.16	54.42	0.751	0.123
Zhao et al. [58]	LiDAR + Stereo Camera	Mapping & Navigation	9.31 / 99.85	54.58	56.67	0.528	0.128
Sock et al. [45]	LiDAR + RGB Camera	Mapping & Navigation	1.93 / 99.93	50.93	53.21	0.590	0.156
TNS (ours)	LiDAR + RGB Camera	Mapping & Navigation	71.77 / 91.05	81.41	85.70	0.803	0.106

TABLE III: **SOTA comparisons:** We list several prior methods and highlight the benefits of our method on the RELIS [20] terrain dataset. Our method outperforms previous SOTA methods by 4.17-30.48% in terms of mAcc and reduces the MSE by 13.8-71.4%.

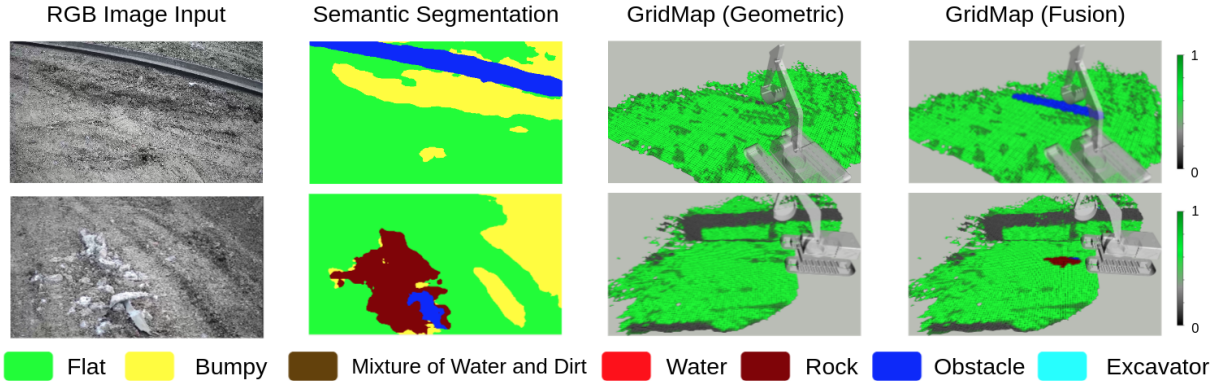


Fig. 5: **Visual results of TNS:** In the traversability map, the higher the traversability score, the easier it is for robots to navigate the corresponding terrain. More visual results are available in the supplementary material.

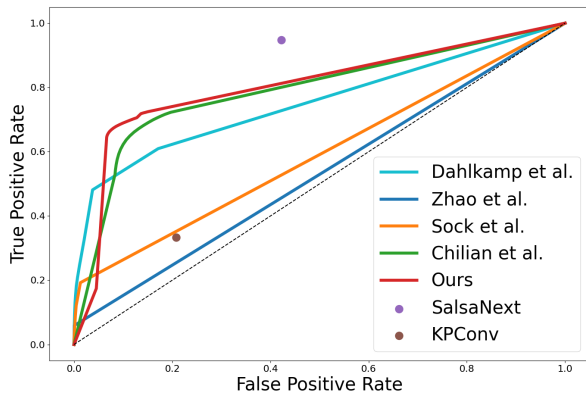


Fig. 6: **ROC plot:** We plot ROCs on several SOTA mapping/segmentation methods. LiDAR-based segmentation methods [10, 48] are trained on the point cloud labels, so they have the advantage of prior knowledge on the ground truth. In the real world, annotated 3D point cloud data would not be easily available for applications. We use a point in the ROC plot to represent those methods, as there is not a threshold to adjust.

based on TNS. We also compare its performance with a geometric-only method [8].

A. Hardware Setup

We use an XCMG XE490D excavator to perform our experiments. The excavator is equipped with a Livox-Mid100 LiDAR, an HIK web camera with FOV of 56.8 degrees with a pitch angle of 30.3 degrees to detect the environment, and a Huace real-time kinematic (RTK) positioning device to

provide the location. We run our code on a laptop with an Intel Core i7-10875H CPU, 16 GB RAM, and 6GB GeForce RTX 2060 on the excavator.

XCMG XE490D excavator has a maximum climbing angle of 35 degrees; the typical recommended climbing angles for any vehicle as a safe climbing angle is 10 degrees. Therefore, we set $s_{cri} = 35 \text{ deg}$ and $s_{safe} = 10 \text{ deg}$. In addition, we obtain an approximation of the maximum height allowed by s_{cri} and s_{safe} after expanding three times the resolution d_{res} along the surface to get:

$$h_{cri} = 3 \tan(s_{cri}) \times d_{res} = 0.35 \text{ m}$$

$$h_{safe} = 3 \tan(s_{safe}) \times d_{res} = 0.10 \text{ m}$$

B. Traversability Map Results and Analysis

In this section, we evaluate our system in the real world with visual results. In Figure 5, we show some typical scenarios excavators encounter to illustrate the advantages of geometric and semantic fusion. In those cases, the steel bar and stone were not captured by geometric calculation, while with semantic information, those obstacles can be detected.

C. Planning Based on Offline Traversability Map

Based on the resulting occupancy grid maps from the proposed TNS and geometric-only method [8], we randomly choose start and goal positions on an unoccupied grid with over 90 trials. The success rates of finding a valid path without collision for our TNS and the other method are 82.6% and 33.3%, respectively. We show some comparisons on planning




Fig. 7: **Planner output comparisons between geometric-only scheme [8] (top) and TNS (bottom):** We show planned trajectories with our modified Hybrid A* [27] planner. The planning is based on a global traversability map. We highlight some obstacles that are not observed by the geometric method (red) as well as some traversable regions that are falsely observed by the geometric method (blue).

Trajs	Type	Total Len (m)	Avg Err (m)	Min Err (m)
9	Straight	158.85	0.102	0.040
8	Small Turn	219.35	0.104	0.032
8	Sharp Turn	242.14	0.059	0.042

TABLE IV: **Real-world Experiments and Trials.** we have tested TNS on different types of trajectories, including straight paths, normal turns, and sharp turns. Our system can achieve 10cm tracking error accuracy for all these scenarios.

results in Figure 7. We use an occupied threshold t_{occ} of 0.6. The height of the cabin h_{cab} is 0.5 m, and the distance between two tracks d_{track} is 2.75 m for map post-processing and planner configuration.

D. Real-world Experiments and Trials

We test our system TNS on two construction sites with a total area of at least 200 m^2 . We summarize those trials in Table IV. We tested 3 types of trajectories, including going straight while avoiding lower traversability areas, making normal turns, and making sharp turns on the terrain. For all tests, the excavator was able to successfully reach the given target, which demonstrates the robustness of our system. Furthermore, the tracking error of all trajectories is within 10cm on average. For details of the testing site, please refer to the supplemental materials.

E. Run-time Analysis of Traversability Map

Our method consists of the following major parts, which contribute to the overall runtime of the system:

- **Segmentation** generates a pixel-wise semantic classification on each image in the RGB input stream.
- **Projection** casts the 2D segmentation result onto the 3D point cloud and assigns each point a semantic label through the calibration matrix.
- **Geometric traversability calculation** estimates and updates slope and step height based on point cloud data in a grid map representation.

In Table V, we give details of the run-time of each component in the system. The final fusion step is under 2 ms and does not contribute to the overall runtime of the method. Our

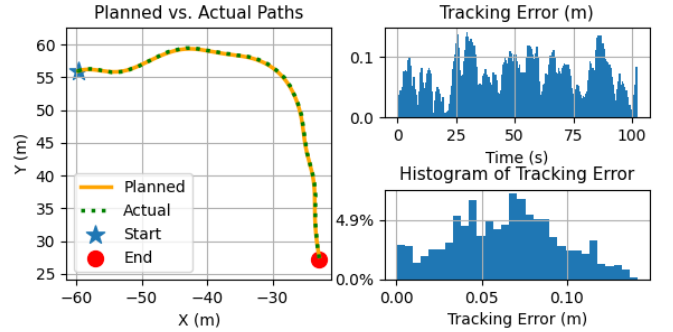


Fig. 8: **Tracking controller performance evaluation:** We plot one of the trajectory tracking results with planned and actual paths (left), tracking error vs. time (top-right), and the histogram of the tracking error (bottom-right).

Run-time (ms)	Max	Min	Mean
Segmentation	100.2	53.5	75.4
Projection	54.0	35.0	42.3
T_{geo} Calculation	38.0	9.0	22.1

TABLE V: **Runtimes of different modules.** Our method can be run in real-time and update the traversability map at a rate of 10 Hz.

method can update the traversability map at a rate of 10 Hz. Please refer to the video for more visual results of excavator navigation.

F. Controller Error Analysis

The tracking trajectory controller can maintain the excavator around the desired path with the maximum absolute lateral tracking error less than 15 cm in most of our test runs. In the case shown in Figure 8, the maximum tracking error is around 14 cm. This test run lasts for 102 seconds with an average speed of 0.5 meters per second. The total length of the trajectory is about 50 meters. The left plot shows the planned path from the improved hybrid A* planner and the actual path of the excavator. The excavator starts from the blue pentagram and ends at the red dot. The top-right plot represents the tracking error, which is the distance between the excavator and the closest point on the planned path. The bottom-right plot is the histogram of the tracking error, where the y-axis represents the percentage of each tracking error column. The tracking error is around 6 cm most of the time.

G. Analysis and Lessons Learned

In this section, we highlight some of the failures of and lessons learned from the design and evaluation of our system:

- **Perception errors** include segmentation error and Lidar measurement error. It is hard to find similar terrains or scenarios in existing datasets for annotations and supervised training, especially when the terrain becomes rougher and bumpier. To alleviate segmentation error, we collect and annotate some terrain data on construction sites with different terrain labels, including flat surface, bumpy surface, water puddle, obstacles, rocks, etc., aiming to improve perception accuracy in unstructured environments and enable such construction vehicle

applications. However, the Lidar measurement can be unreliable due to the dust in the air. To remove such noise, we use step height estimation and semantic fusion for more robust traversability predictions, as mentioned in Section IV-A.

- **Terrain roughness** is an issue in geometric-based traversability methods on a mobile robot [8]. However, it is less effective for large machines like excavators due to the scale difference. In our case, roughness can be partially modeled either through the slope and step height or captured by visual features from the RGB images. However, it could become an issue if the terrain is very uneven or has large rocks or obstacles.
- **Localization accuracy** directly impacts the quality of the system. In our experiment, the main reason for the localization inaccuracy is the drift of the RTK system on the altitude. In our open test field, the accuracy of the RTK system in latitude and longitude is around 5 cm, whereas the altitude accuracy is about 20 cm. To plan accurately and navigate over a period of time, we only use the most recent grid cells to calculate the traversability score because the drift is small. In addition, our attempt to use SLAM for localization failed because most features are quite uniform (similar hills, pits, rock piles, etc.), causing degraded performance and very low accuracy due to instability. In the future, we could build a more stable localization system to fuse RTK, LiDAR, and camera data.
- **Planner** needs to be adjusted to fully utilize the traversability map. We choose the Hybrid A* algorithm over the standard A* algorithm in our system to avoid sharp turns, which could cause damage or bumpiness to the ground surface. We adjust the Hybrid A* planner as described in IV-B to compute a smoother and safer path with continuous traversability map values. However, it is hard to guarantee that our planner will always generate a smooth path on arbitrary terrains.
- **Computational and power budget** is a major issue in the design of our perception and planning algorithm. Our traversability map computations and navigation module run on a laptop with an Intel Core i7-10875H CPU and a 6GB GeForce RTX 2060. Our implementation must be efficient and light-weight to run in real-time. Recently, many deep and reinforcement learning methods have been proposed for object detection and navigation, but they require a high-end GPU for efficient execution. We can't use such methods on our platform.
- **Safety:** In deploying the autonomous excavator system to the real world, safety is always the most critical consideration. We develop the terrain traversability mapping component to describe the complexity of the terrain and provide safe regions for the autonomous excavator to navigate. Our method can be combined with other safety strategies such as object detection, collision avoidance, etc., and maintain the stability of the excavator to ensure the safety of autonomous operation.
- **Excavator size** also governs the performance of our system. There are three broad classes of excavators:

compact excavator (less than 6 tons), standard excavator (7 – 45 tons) and large excavators (45 – 90 tons). The size of the excavator impacts the performance of the navigation system when computing a smooth trajectory and the resulting path. There is a relative trade-off between mobility and stability for different sizes. We have evaluated the performance of TNS on a large, 49-ton excavator. In general, developing autonomous excavation technology for larger excavators is more challenging.

VIII. CONCLUSIONS, LIMITATIONS, AND FUTURE WORK

In this paper, we present a terrain traversability mapping and navigation System (TNS) for autonomous excavation navigation. We highlight its application and benefits on difficult excavator navigation tasks in real-world scenarios. We use a novel learning-based geometric fusion solution and demonstrate its benefits over prior mapping algorithms. We also release the CWT dataset with challenging real-world scenes in unstructured construction sites for perception tasks.

Our work has some limitations. Due to safety issues, we are not able to extensively test our system in all types of scenarios, including cases with many human workers and other machines. We have only evaluated the performance on a large, 49-ton excavator. As part of our future work, we would like to improve the planner further and utilize the specifications of the excavator like a human operator. For example, the excavator should be able to run over small obstacles using the space between two tracks. In addition, we would like to evaluate the performance in different types of outdoor terrains. Our longer-term goal is to enable autonomy and collaborations among machines or with humans on construction sites. This requires several systems and modules working together, including autonomous excavation, autonomous navigation, and human machine interactions.

ACKNOWLEDGEMENT

This work was done as a summer intern at Baidu RAL. We appreciate the discussion and support from Baidu RAL team.

REFERENCES

- [1] Juhana Ahtiainen, Todor Stoyanov, and Jari Saari-nen. Normal distributions transform traversability maps: Lidar-only approach for traversability mapping in outdoor environments. *Journal of Field Robotics*, 34(3): 600–621, 2017. doi: <https://doi.org/10.1002/rob.21657>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/rob.21657>.
- [2] Mohammed Abdessamad Bekhti and Yuichi Kobayashi. Regressed terrain traversability cost for autonomous navigation based on image textures. *Applied Sciences*, 10(4), 2020. ISSN 2076-3417. doi: 10.3390/app10041195. URL <https://www.mdpi.com/2076-3417/10/4/1195>.
- [3] M. Bellone, A. Messina, and G. Reina. A new approach for terrain analysis in mobile robot applications. In *2013 IEEE International Conference on Mechatronics (ICM)*, pages 225–230, 2013. doi: 10.1109/ICMECH.2013.6518540.

- [4] Mauro Bellone, Giulio Reina, Nicola Giannoccaro, and Luigi Spedicato. 3d traversability awareness for rough terrain mobile robots. *Sensor Review*, 34, 03 2014. doi: 10.1108/SR-03-2013-644.
- [5] Mauro Bellone, Giulio Reina, Luca Caltagirone, and Mattias Wahde. Learning traversability from point clouds in challenging scenarios. *IEEE Transactions on Intelligent Transportation Systems*, 19(1):296–305, 2018. doi: 10.1109/TITS.2017.2769218.
- [6] Tim Braun, Henning Bitsch, and Karsten Berns. Visual terrain traversability estimation using a combined slope/elevation model. In Andreas R. Dengel, Karsten Berns, Thomas M. Breuel, Frank Bomarius, and Thomas R. Roth-Berghofer, editors, *KI 2008: Advances in Artificial Intelligence*, pages 177–184, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg. ISBN 978-3-540-85845-4.
- [7] R. Omar Chavez-Garcia, Jérôme Guzzi, Luca M. Gambardella, and Alessandro Giusti. Learning ground traversability from simulations. *IEEE Robotics and Automation Letters*, 3(3):1695–1702, 2018. doi: 10.1109/LRA.2018.2801794.
- [8] Annett Chilian and Heiko Hirschmüller. Stereo camera based navigation of mobile robots on rough terrain. In *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4571–4576, 2009. doi: 10.1109/IROS.2009.5354535.
- [9] MMSegmentation Contributors. MMSegmentation: Openmmlab semantic segmentation toolbox and benchmark. <https://github.com/open-mmlab/mms Segmentation>, 2020.
- [10] Tiago Cortinhal, George Tzelepis, and Eren Erdal Aksoy. Salsanext: Fast, uncertainty-aware semantic segmentation of lidar point clouds for autonomous driving, 2020.
- [11] Hendrik Dahlkamp, Adrian Kaehler, David Stavens, Sebastian Thrun, and Gary R. Bradski. Self-supervised monocular road detection in desert terrain. In *Robotics: Science and Systems*, 2006.
- [12] Fucheng Deng, Xiaorui Zhu, and Chao He. Vision-based real-time traversable region detection for mobile robot in the outdoors. *Sensors*, 17(9), 2017. ISSN 1424-8220. doi: 10.3390/s17092101. URL <https://www.mdpi.com/1424-8220/17/9/2101>.
- [13] Péter Fankhauser and Marco Hutter. A universal grid map library: Implementation and use case for rough terrain navigation. In *Robot Operating System (ROS)*, pages 99–120. Springer, 2016.
- [14] A Geiger, P Lenz, C Stiller, and R Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013. doi: 10.1177/0278364913491297.
- [15] Tianrui Guan, Divya Kothandaraman, Rohan Chandra, and Dinesh Manocha. Ganav: Group-wise attention network for classifying navigable regions in unstructured outdoor environments, 2021.
- [16] Robert A Hewitt, Alex Ellery, and Anton de Ruiter. Training a terrain traversability classifier for a planetary rover through simulation. *International Journal of Advanced Robotic Systems*, 14(5):1729881417735401, 2017. doi: 10.1177/1729881417735401.
- [17] Noriaki Hirose, Amir Sadeghian, Marynel Vázquez, Patrick Goebel, and Silvio Savarese. Gonet: A semi-supervised deep learning approach for traversability estimation. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3044–3051, 2018. doi: 10.1109/IROS.2018.8594031.
- [18] Gabriel M Hoffmann, Claire J Tomlin, Michael Montemerlo, and Sebastian Thrun. Autonomous automobile trajectory tracking for off-road driving: Controller design, experimental validation and racing. In *2007 American Control Conference*, pages 2296–2301. IEEE, 2007.
- [19] Christopher J. Holder and Toby P. Breckon. Learning to drive: Using visual odometry to bootstrap deep learning for off-road path prediction. In *2018 IEEE Intelligent Vehicles Symposium (IV)*, pages 2104–2110, 2018. doi: 10.1109/IVS.2018.8500526.
- [20] Peng Jiang, Philip R. Osteen, Maggie Wigness, and Srikanth Saripalli. Rellis-3d dataset: Data, benchmarks and analysis. *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1110–1116, 2021.
- [21] Gregory Kahn, Pieter Abbeel, and Sergey Levine. Badgr: An autonomous self-supervised learning-based navigation system. *IEEE Robotics and Automation Letters*, 6(2):1312–1319, 2021.
- [22] Muhammad Khan, Karsten Berns, and Abubakr Muhammad. Vehicle specific robust traversability indices using roadmaps on 3d pointclouds. *International Journal of Intelligent Robotics and Applications*, 4:1–17, 12 2020. doi: 10.1007/s41315-020-00148-x.
- [23] Muhammad Mudassir Khan, Haider Ali, Karsten Berns, and Abubakr Muhammad. Road traversability analysis using network properties of roadmaps. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2960–2965, 2016. doi: 10.1109/IROS.2016.7759458.
- [24] Sung-Keun Kim and J. Russell. Framework for an intelligent earthwork system: Part i. system architecture. *Automation in Construction*, 12:1–13, 2003.
- [25] Nathaniel Kingry, Myungjin Jung, Evan Derse, and Ran Dai. Vision-based terrain classification and solar irradiance mapping for solar-powered robotics. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5834–5840, 2018. doi: 10.1109/IROS.2018.8593635.
- [26] Ashish Kumar, Zipeng Fu, Deepak Pathak, and Jitendra Malik. Rma: Rapid motor adaptation for legged robots, 2021.
- [27] Karl Kurzer. Path planning in unstructured environments : A real-time hybrid a* implementation for fast and deterministic path generation for the kth research concept vehicle. Master’s thesis, 2016.
- [28] Roberto Manduchi, A. Castano, Ashit Talukder, and L. Matthies. Obstacle detection and terrain classification for autonomous off-road navigation. *Autonomous Robots*, 18:81–102, 01 2005. doi: 10.1023/B:AURO.0000047286.

- 62481.1d.
- [29] Sango Matsuzaki, Kimitoshi Yamazaki, Yoshitaka Hara, and Takashi Tsubouchi. Traversable region estimation for mobile robots in an outdoor image. *J. Intell. Robot. Syst.*, 92(3-4):453–463, 2018. doi: 10.1007/s10846-017-0760-x.
- [30] Daniel Maturana, Po-Wei Chou, Masashi Uenoyama, and Sebastian Scherer. Real-time semantic mapping for autonomous off-road navigation. In Marco Hutter and Roland Siegwart, editors, *Field and Service Robotics*, pages 335–350, Cham, 2018. Springer International Publishing. ISBN 978-3-319-67361-5.
- [31] Nipun D. Nath and A. Behzadan. Deep convolutional networks for construction object detection under different visual conditions. In *Frontiers in Built Environment*, 2020.
- [32] Panagiotis Papadakis. Terrain traversability analysis methods for unmanned ground vehicles: A survey. *Engineering Applications of Artificial Intelligence*, 26(4):1373–1385, 2013. ISSN 0952-1976. doi: <https://doi.org/10.1016/j.engappai.2013.01.006>. URL <https://www.sciencedirect.com/science/article/pii/S095219761300016X>.
- [33] David Paz, Hengyuan Zhang, Qinru Li, Hao Xiang, and Henrik I. Christensen. Probabilistic semantic mapping for urban autonomous driving applications. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2059–2064, 2020. doi: 10.1109/IROS45743.2020.9341738.
- [34] Rudra P. K. Poudel, Stephan Liwicki, and R. Cipolla. Fast-scnn: Fast semantic segmentation network. In *BMVC*, 2019.
- [35] Michael J. Procopio, Jane Mulligan, and Greg Grudic. Learning terrain segmentation with classifier ensembles for autonomous robot navigation in unstructured environments. *Journal of Field Robotics*, 26(2):145–175, 2009. doi: <https://doi.org/10.1002/rob.20279>.
- [36] Redmond R Shamshiri, Cornelia Weltzien, Ibrahim A Hameed, Ian J Yule, Tony E Grift, Siva K Balasundram, Lenka Pitonakova, Desa Ahmad, and Girish Chowdhary. Research and development in agricultural robotics: A perspective of digital farming. 2018.
- [37] Rene Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *ICCV*, 2021.
- [38] Dominic Roberts and Mani Golparvar-Fard. End-to-end vision-based detection, tracking and activity analysis of earthmoving equipment filmed at ground level. *Automation in Construction*, 2019. ISSN 0926-5805. doi: <https://doi.org/10.1016/j.autcon.2019.04.006>.
- [39] Ryan D. Rosenfeld, Mark G. Restrepo, William H. Gerard, Walter E. Bruce, Atiena A. Branch, Gregory C. Lewin, and Nicola Bezzo. Unsupervised surface classification to enhance the control performance of a ugv. In *2018 Systems and Information Engineering Design Symposium (SIEDS)*, pages 225–230, 2018. doi: 10.1109/SIEDS.2018.8374741.
- [40] Brandon Rothrock, Ryan Kennedy, Christopher T. Cunningham, Jeremie Papon, Matthew Heverly, and Masahiro Ono. Spoc: Deep learning-based terrain classification for mars rover missions. 2016.
- [41] Fabian Schilling, Xi Chen, John Folkesson, and Patric Jensfelt. Geometric and visual terrain classification for autonomous mobile navigation. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2678–2684, 2017. doi: 10.1109/IROS.2017.8206092.
- [42] Jongwon Seo, Seungsoo Lee, Jeonghwan Kim, and Sung-Keun Kim. Task planner design for an automated excavation system. *Automation in Construction*, 20(7):954–966, 2011. ISSN 0926-5805. doi: <https://doi.org/10.1016/j.autcon.2011.03.013>.
- [43] H. Shariati, Anuar Yeraliyev, B. Terai, S. Tafazoli, and Mahdi Ramezani. Towards autonomous mining via intelligent excavators. In *CVPR Workshops*, 2019.
- [44] Anukriti Singh, Kartikeya Singh, and P. B. Sujit. Offroadtranseg: Semi-supervised segmentation using transformers on offroad environments, 2021.
- [45] Juil Sock, Jun Kim, Jihong Min, and Kiho Kwak. Probabilistic traversability map generation using 3d-lidar and camera. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5631–5637, 2016. doi: 10.1109/ICRA.2016.7487782.
- [46] Pei Sun, Henrik Kretschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2446–2454, 2020.
- [47] Vivekanandan Suryamurthy, Vignesh Sushrutha Raghavan, Arturo Laurenzi, Nikos G. Tsagarakis, and Dimitrios Kanoulas. Terrain segmentation and roughness estimation using rgb data: Path planning application on the centauro robot. In *2019 IEEE-RAS 19th International Conference on Humanoid Robots (Humanoids)*, pages 1–8, 2019. doi: 10.1109/Humanoids43949.2019.9035009.
- [48] Hugues Thomas, Charles R. Qi, Jean-Emmanuel Deschaud, Beatriz Marcotegui, François Goulette, and Leonidas J. Guibas. Kpconv: Flexible and deformable convolution for point clouds. *Proceedings of the IEEE International Conference on Computer Vision*, 2019.
- [49] Kasi Viswanath, Kartikeya Singh, Peng Jiang, P.B. Sujit, and Srikanth Saripalli. Offseg: A semantic segmentation framework for off-road driving. In *2021 IEEE 17th International Conference on Automation Science and Engineering (CASE)*, pages 354–359, 2021. doi: 10.1109/CASE49439.2021.9551643.
- [50] Martin Wermelinger, Péter Fankhauser, Remo Diethelm, Philipp Krüsi, Roland Siegwart, and Marco Hutter. Navigation planning for legged robots in challenging terrain. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1184–1189, 2016. doi: 10.1109/IROS.2016.7759199.
- [51] Maggie Wigness, Sungmin Eum, John G Rogers, David Han, and Heesung Kwon. A rugd dataset for autonomous navigation and visual perception in unstructured outdoor

- environments. In *International Conference on Intelligent Robots and Systems (IROS)*, 2019.
- [52] Huikai Wu, Junge Zhang, Kaiqi Huang, Kongming Liang, and Yu Yizhou. Fastfcn: Rethinking dilated convolution in the backbone for semantic segmentation, 2019.
- [53] Tianyi Wu, Sheng Tang, Rui Zhang, and Yongdong Zhang. Cgnet: A light-weight context guided network for semantic segmentation. *IEEE Transactions on Image Processing*, 30:1169–1179, 2021.
- [54] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M. Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021. URL <https://openreview.net/forum?id=OG18MI5TRL>.
- [55] Jia Xue, Hang Zhang, K. Dana, and K. Nishino. Differential angular imaging for material recognition. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6940–6949, 2017.
- [56] Changqian Yu, Changxin Gao, Jingbo Wang, Gang Yu, Chunhua Shen, and Nong Sang. Bisenet v2: Bilateral network with guided aggregation for real-time semantic segmentation. *International Journal of Computer Vision*, 129:1–18, 11 2021. doi: 10.1007/s11263-021-01515-2.
- [57] Liangjun Zhang, Jinxin Zhao, Pinxin Long, Liyang Wang, Lingfeng Qian, Feixiang Lu, Xibin Song, and Dinesh Manocha. An autonomous excavator system for material loading tasks. *Science Robotics*, 6(55), 2021. doi: 10.1126/scirobotics.abc3164. URL <https://robotics.sciencemag.org/content/6/55/eabc3164>.
- [58] Yimo Zhao, Peilin Liu, Wuyang Xue, Ruihang Miao, Zheng Gong, and Rendong Ying. Semantic probabilistic traversable map generation for robot path planning. In *2019 IEEE International Conference on Robotics and Biomimetics (ROBIO)*, pages 2576–2582, 2019. doi: 10.1109/ROBIO49542.2019.8961533.
- [59] Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip H.S. Torr, and Li Zhang. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In *CVPR*, 2021.
- [60] Yan Zhou, Ying Huang, and Zhenhua Xiong. 3d traversability map generation for mobile robots based on point cloud. In *2021 IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM)*, pages 836–841, 2021. doi: 10.1109/AIM46487.2021.9517463.