

RCA: Ride Comfort-Aware First-Person Navigation via Self-Supervised Learning

Xinjie Yao

CMU-RI-TR-22-15

May, 2022



The Robotics Institute
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA

Thesis Committee:

Jean Oh, *Co-Chair*
Ji Zhang, *Co-Chair*
Stephen Smith
Tanmay Shankar

*Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Robotics.*

Copyright © 2022 Xinjie Yao. All rights reserved.

To my parents, for their love and supports.

Abstract

Under shared autonomy, wheelchair users expect vehicles to provide safe and comfortable rides while following users’ high-level navigation plans. To find such a path, vehicles negotiate with different terrains and assess their traversal difficulty. Most prior works model surroundings either through geometric representations or semantic classifications, which do not reflect perceived motion intensity and ride comfort in downstream navigation tasks. We propose to model ride comfort explicitly in traversability analysis using proprioceptive sensing. We develop a self-supervised learning framework to predict traversability costmap from first-person-view images by leveraging vehicle states as training signals. Our approach estimates how the vehicle would “feel” if traversing over based on terrain appearances. We then show our navigation system provides human-preferred ride comfort through robot experiments together with a human evaluation study.

Acknowledgments

I would like to express my deepest appreciation to my advisors Jean and Ji for their enormous support and guidance during my master study. It has been a great pleasure to work with them since the RISS program. I felt extremely grateful to them for allowing me to explore research ideas with great patience and offering resources to help achieve my goals. Their feedback and encouragements have guided me through the world of unknown.

Besides my advisors, I very much appreciate the rest of my thesis committees, for their insightful comments and uplifting feedback. Meanwhile, thanks should go to my colleagues from Robotics Institute as they have shared technical insights and offered assistance at different stages of my project. I gratefully acknowledge the help from Zhanxin, Peter, Tanmay, Hyeonwoo, Peng, Chao and other people from BIG.

I owe thanks to a number of people outside the lab for livening me up, Xuhua, Ziyang, Mononito, Shuqin, Zixuan, Carl, Fan, Xuanyi, Gaurav, and Jed. Thank you for spreading positivity and picking me up when I was down. You have pulled me out of my darkness and made me see light again.

Finally, I'd like to thank my parents for their unconditional love and relentless support.

Contents

1	Introduction	1
1.1	Motivation	1
1.2	Thesis Contributions	2
2	Related Work	5
2.1	Geometry-based	5
2.2	Semantic-based	5
2.3	RL-based	6
2.4	Multimodal-based	6
3	Problem Formulation	9
4	Approach	11
4.1	A Navigation System	11
4.2	An Overview of RCA Approach	13
4.3	Input Data Preprocessing	13
4.3.1	Vehicle State Data	13
4.3.2	Image Data	14
4.3.3	Data Augmentation	14
4.4	Estimating Traversability Cost from Vehicle States	15
4.4.1	Unsupervised Learning for Classifying Vehicle States	15
4.4.2	Traversal Cost Generation	16
4.5	Self-supervised Traversal Cost Prediction	16
5	Experiments	19
5.1	Training the Learning Modules	20
5.1.1	Dataset	20
5.1.2	Implementations	20
5.1.3	Evaluation Metrics for Learning Accuracy	20
5.2	Robot Experiments	21
5.2.1	Experiment Design	21
5.2.2	Perceived Vehicle Motion Intensity (PMI)	22
5.2.3	Human Evaluation	22

6	Results	25
6.1	Learning Verification	25
6.2	Robot Experiments	26
6.2.1	PMI Analysis	26
6.2.2	Human Evaluation Analysis	27
6.2.3	PMI vs. Human Evaluation	28
7	Conclusions	29
A	Appendix	31
A.1	Dataset Examples	31
	Bibliography	33

List of Figures

4.1	A navigation system with the terrain analysis module.	11
4.2	An overview of RCA. During training, the vehicle experiences physical vibrations by traversing various types of terrains. Using recorded vehicle states and first-person view images, the system learns to estimate traversability cost from images. Instead of manual labeling, we first use unsupervised learning to coarsely group vehicle states into clusters based on vehicle-terrain interactions. Next, with those clusters, we define a continuous cost function to reflect the traversability based on vehicle dynamics. Finally, we train a prediction model to associate terrain images directly to the traversability costs. During deployment, the learned model estimates a terrain costmap to support autonomous navigation.	12
5.1	(a) The tracked vehicle used for data collection. (b) The robotic wheelchair platform used in the experiments equipped with a 4.1GHz i7 computer and a NVIDIA GTX 1660Ti GPU card. Both platforms are equipped with a Velodyne Puck Lidar, a MEMS-based IMU, and a monocular camera.	19
5.2	Three approaches are tested in five scenes including asphalt roads, grass, curbs, tactile pavings, gravels (top row). The wheelchair starts at the same initial pose and tend to reach the desired goal points (red cross). The bottom row shows recorded wheelchair trajectories in a top-down view (our RCA in blue; semantic-based in red; and 3D-based in green). In (1), RCA detours away from the curb and uneven surfaces until it reaches the ramp. In (2), RCA takes a side way while avoiding the turbulence brought by tactile pavings. In (3)&(4), RCA foresees the effect from sinking in the soft terrain and avoids the vegetation although going through it is the shortest path. In (5), RCA keeps the wheelchair away from the vegetation and gravels.	22
6.1	Visualization of vehicle state feature space using t-SNE projection. Samples are colored by semantic classes of recorded scenes.	25

6.2	Prediction results of normalized amplitude spectra along roll, pitch and Z axis. It shows relative amounts of amplitude for each frequency component. Ground truth signals are marked in blue, while the estimates ones are shown in orange.	26
6.3	Normalized Perceived Vehicle Motion Intensity Scores[6] for Z (upward pointing), yaw, roll and pitch axes. Our approach (blue) shows consistently smaller scores for all axes, demonstrating less disturbances. Semantic-based baseline (orange) exhibits a higher value along Yaw axis, implying its frequent and aggressive turning motions. 3D-based baseline (green) tends to have larger values along Z, Roll, and Pitch axes, indicating its frequent shakiness over uneven terrains.	27
6.4	Percentage of samples ranking three approaches among five dimensions. Each vertical bar consists of three segments representing the percentage of samples voting an approach as the first (bottom, darker color), the second (middle, moderate color), the third (top, lighter color). Our approach marks as the blue bar, while the semantic-based baseline is in red and the 3D-based baseline is in green. Across five evaluation criteria over five scenes, our approach is most likely to be ranked as the top. We fail to show a clear winner between semantic-based approach and 3D-based approach.	28
A.1	Our dataset includes various types of soft and hard terrains.	31

List of Tables

6.1 Spearman Correlation Scores	28
---	----

Chapter 1

Introduction

1.1 Motivation

There are more than 2.7 million people in the United States who primarily use wheelchairs for their mobility on a daily basis, and the number of wheelchair users is also expected to grow [12]. With the advancement of the autonomous vehicle technology, autonomous wheelchairs are gradually becoming attainable [14]. Previous studies show that having the sense of control is important to human users, *i.e.*, they prefer systems that provide high-level controls over fully autonomous ones, promoting the concept of shared autonomy [17, 23]. In this context, our work addresses low-level robotic vehicle control given human users' high-level control inputs. Specifically, we focus on finding safe and comfortable paths that generally follow a user's high-level navigation plan.

Finding a collision-free, smooth path has been a central challenge in robot navigation. 3D-based approaches, as in [2], interpret a surrounding environment as rigid 3D objects and estimate the traversability based on their geometric properties. While these approaches can detect obstacles and the general traversability of terrain well for large-sized vehicles with robust mobility, they lack semantic understanding of social norms as well as subtle differences in traversability that are important for wheelchair applications.

The other stream of prior works relies on semantic classification of environments, associating each predefined semantic class with a traversal difficulty. Existing works

on this end require fine-grained terrain semantic classes and merge them into levels of traversability[25]. In these approaches, the entire region with the same class label is assigned the same traversability cost, ignoring the variations within each region. Moreover, such a terrain analysis depends heavily on human assessment of traversing difficulty which should be customized for various purposes of navigation, including wheelchair applications.

To address these shortcomings of existing approaches, we propose a first-person image-based approach where traversability is defined as riding comfort and can be measured in terms of the vehicle state. As in the semantic classification methods, the main idea is to use vision to estimate the traversability. Instead of requiring human-defined semantic labels, our framework learns to associate the visual input with how the vehicle state changes when the robot is actually traversing over the spot. Because our approach’s learning signal comes directly from the vehicle state, we propose a self-supervised learning approach for learning this mapping between vision and internal vehicle state and, eventually, traversability cost. To enable a direct mapping from vision to cost, we use unsupervised learning as in [33] to first classify vehicle states into coarse categories. Whereas [33] uses this idea for semantic classification and image segmentation, we propose a method for generating a continuous traversability cost given an image input.

We fully integrate the proposed approach (RCA) on a robotic wheelchair platform for our experiments and evaluation. Based on human evaluation, the proposed approach is strongly preferred to the 3D-based and semantic classification methods according to stability, path normality, safety, trustworthiness, and preference.

1.2 Thesis Contributions

Our contributions are summarized as follows:

1. we propose to define ride comfort in terms of proprioceptive sensing and include that in traversability analysis to improve the general quality of navigation paths under shared autonomy;
2. we develop a self-supervised learning framework (RCA) for predicting a traversability costmap based on first-person-view image inputs;

3. we show that human assessment agrees with the normalized Perceived Vehicle Motion Intensity (PMI) score, supporting the intuition that proprioceptive data can capture the traversability and ride comfort; and
4. we share our findings from a human evaluation study on how humans determine the ride quality according to various aspects including vehicle dynamics and perceived motion intensity.

1. Introduction

Chapter 2

Related Work

Off-road navigation for unmanned ground vehicles (UGV) has been studied extensively over the past years. Under unstructured environments, UGV observes terrains and plans a path based on its interpretation of the surroundings.

2.1 Geometry-based

Conventional approaches of vehicle navigation model the environments through geometric expressions and obtain outstanding results with rigid objects. Representations of the environments involve 3D pointcloud [2], signed distance field [18], and 2.5D elevation map[22][3]. They do not take into account the complexity of vehicle-terrain interactions in which deformable and nonrigid terrains bring additional ambiguity into the measurement of traversability.

2.2 Semantic-based

With the advancement of the urban-scene image segmentation task[5], semantic-based approaches demonstrate promising results due to the rich texture information from vision sensors. Traversal difficulty is typically modeled by the visual appearance of terrains, namely via either mapping semantic terrain classes to traversability [25] [19] [8] or defining a texture-dependent score function to reflect difficulties [11] [20]. The

association between visual distinctness and traversability is constrained to the vision domain and is commonly assessed by human experts, whose subjective interpretations could be erratic given the complex dynamics.

2.3 RL-based

Another stream of emerging works deploy end-to-end learning systems through Inverse Reinforcement Learning[32], Reinforcement Learning applied to model predictive controller[15], and Imitation Learning from experts using additional sensory inputs[21]. These approaches generally require a large amount of vehicle driving data in real world or in high fidelity simulators and may suffer from the vehicle hazards caused by the aggressive exploration behaviors[10].

2.4 Multimodal-based

Apart from leveraging the single modality of the sensory inputs, self-supervised learning from cross-modality sensor inputs shows successful practices to supplement the complementary information across modalities[16]. It investigates the learning of associations between one modality and labels generated from other modalities. During deployment, the predictive model estimates the associate labels at the absence of other modalities. Proprioceptive data are typically generated from traversing actual terrains and are linked to its exteroceptive data for future prediction under this reciprocal manner. Approaches using force-torque signals [26], vehicle vibrations [1][29], and acoustic signals [33][13][24] have been proposed to construct this cross-modal association with terrain visual inputs. Vehicle dynamics based approaches fail to leverage the texture information provided from vision images to create labels, and instead only translate proprioceptive signals to the vision domain[26][1][29]. Approaches using acoustic waveform [33][13][24] regard terrain properties as semantic classes and formulate the navigation task as semantic classification[25]. They fail to illustrate vehicle dynamics in the feature space and are categorized as a semantic-based approach. One recent approach [33] selects acoustic signal triplets based on the Euclidean distance between visual embeddings. With a Siamese encoder, it

obtains a distinct feature space using the acoustic signal reconstruction loss and triplet loss. Those features are further clustered into semantic classes for downstream segmentation task.

While prior works mostly represent terrain traversability with semantic labels or geometry properties, the proposed RCA approach includes ride comfort in assessing traversal cost and determines vehicle dynamics and perceived motion intensity by complimenting information from the exteroceptive and proprioceptive sensing via self-supervised learning.

2. Related Work

Chapter 3

Problem Formulation

The target problem is predicting the navigation costs using first-person-view monocular camera images. Let $\mathcal{I}_t \in \mathcal{I}$ denote an image input at time step t where \mathcal{I} denotes the projection space of the monocular camera images. We aim to find a function that maps this projection space to a 2D cost map, denoted by $\Phi(\mathcal{I}) \rightarrow \mathbb{R}^2$.

Based on the intuition that traversability directly affects the vehicle state, we formulate this problem as finding two mapping functions. Here, we introduce the vehicle state as an internal representation of the cost as follows. Let \mathbf{s}_t denote the vehicle state at time t , consisting of the 3D robot pose including position, orientation, angular velocity, and linear acceleration, denoted by $\mathbf{s}_t = [\mathbf{p}_t \in \mathbb{R}^3, \mathbf{q}_t \in \mathbf{SO}(3), \dot{\mathbf{q}}_t, \ddot{\mathbf{p}}_t]$, respectively. The first subproblem is to find a mapping function $\Phi_{\mathcal{I} \rightarrow \mathcal{S}}(\mathcal{I}) \rightarrow \mathcal{S}$ that can predict the vehicle states from input images, minimizing the error between the predicted and the true states. The second mapping function, $\Phi_{\mathcal{S} \rightarrow \mathbb{R}^2}(\mathcal{S}) \rightarrow \mathbb{R}^2$, maps the vehicle states to traversability costs, quantifying the traversability based on vehicle states. Using this formulation, our self-supervised approach automatically generates training data in the form of an image and the corresponding traversability costmap. This training data can then be used to learn the mapping function from image to cost $\Phi(\mathcal{I}) \rightarrow \mathbb{R}^2$.

3. Problem Formulation

Chapter 4

Approach

We first describe the overall navigation system and an overview of the proposed approach in [Section 4.1](#) and [Section 4.2](#), followed by technical details on key modules of the proposed approach.

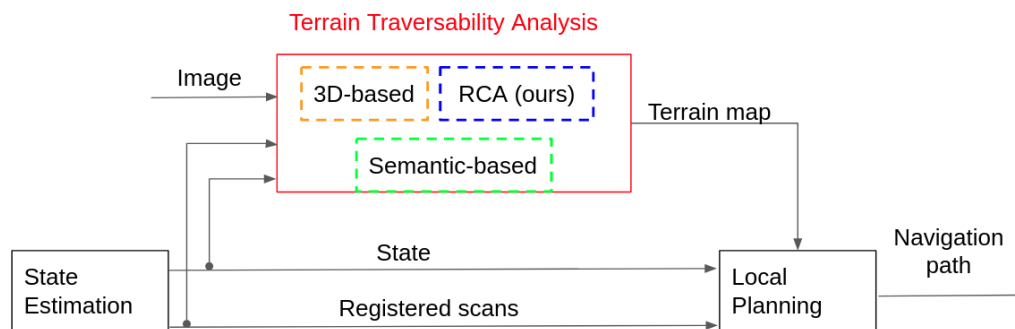


Figure 4.1: A navigation system with the terrain analysis module.

4.1 A Navigation System

[Fig. 4.1](#) shows an overall navigation system where the proposed approach contributes to the terrain traversability analysis module.

The State Estimation Module consists of a data pipeline[30] that processes 3D lidar scans, images, and inertial measurements. It generates vehicle motions and registers laser scans based on the estimated pose. We note that only the vehicle state from this module constitutes the vehicle state used in our approach as defined in

4. Approach

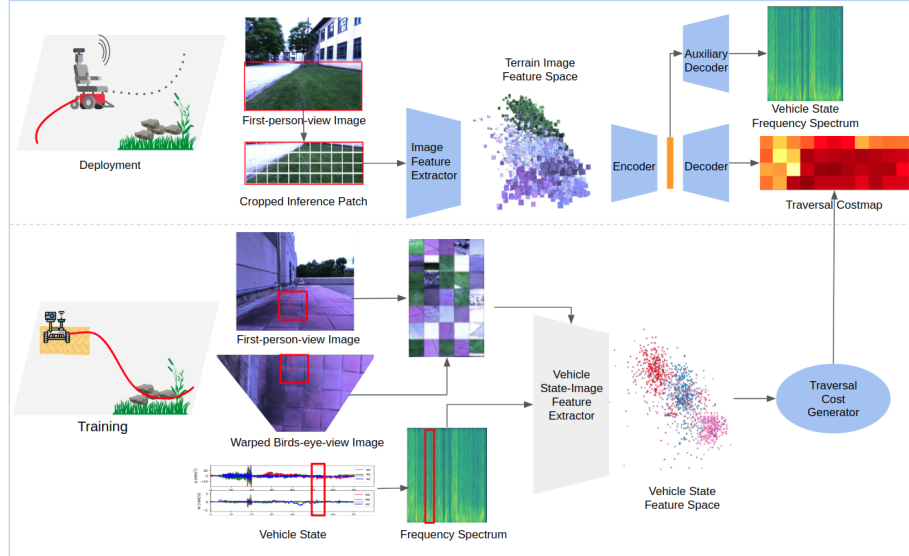


Figure 4.2: An overview of RCA. During training, the vehicle experiences physical vibrations by traversing various types of terrains. Using recorded vehicle states and first-person view images, the system learns to estimate traversability cost from images. Instead of manual labeling, we first use unsupervised learning to coarsely group vehicle states into clusters based on vehicle-terrain interactions. Next, with those clusters, we define a continuous cost function to reflect the traversability based on vehicle dynamics. Finally, we train a prediction model to associate terrain images directly to the traversability costs. During deployment, the learned model estimates a terrain costmap to support autonomous navigation.

Chapter 3 and the rest is included for supporting the baselines.

The Terrain Traversability Analysis Module takes data from onboard sensors, *e.g.*, camera or 3D sensor, to produce a navigation costmap. Our RCA approach is specifically designed for a monocular camera only setting, but we include a 3D-based approach [2] as well as a semantic classifier approach [4] as the baselines for evaluation. We also assume that standard vehicle state information such as 6-DOF pose of the vehicle is available. The Terrain Traversability Analysis Module outputs a local terrain map indicating how traversable local region is either in a discrete or continuous fashion.

The Local Planning Module takes charge of avoiding obstacles in the vicinity of the vehicle, providing low-level planning. This module 3D projects first-person view image-based terrain traversability analysis output onto the registered lidar scans and uses it for collision avoidance planning. The planning algorithm generates

collision-free paths based on a trajectory library[31].

4.2 An Overview of RCA Approach

Fig. 4.2 describes RCA framework for predicting a traversability costmap using a monocular camera only. Here, we propose a self-supervised learning approach for the terrain traversability analysis module by exploiting a robot’s vehicle state as an additional input to the learning algorithm in addition to the image input. Intuitively, our approach aims to learn to predict how a robot would “feel” based on the first-person view of the terrain ahead of the robot. We then translate the vehicle state to a continuous numerical cost function using self-supervised learning. By using proprioceptive data as the learning signal for traversability, the model learns to map visual inputs to traversability costs.

4.3 Input Data Preprocessing

Our approach takes the following two types of inputs: the vehicle states and the first-person view camera images. This pair of data can be collected in an unsupervised manner, *e.g.*, by having a robot explore an environment. Such pairs, however, generally result in a data imbalance issue between these two types of data. More specifically, perspective projection distorts further regions where vehicle-terrain interactions happen and thus are discarded. To address this issue we augment the training data.

4.3.1 Vehicle State Data

To model experienced disturbances from vehicle-terrain interaction, we take the angular velocity along the roll and pitch axes $\dot{\mathbf{q}}^r, \dot{\mathbf{q}}^p$, linear acceleration along the Z axis (pointing downward) $\ddot{\mathbf{p}}^z$ as inputs. Signals from three axes are truncated into a window of n samples, and their frequency components are analyzed by the Fast Fourier Transform (FFT) to obtain amplitude spectra \mathbf{A}^d respectively. Finally, three amplitude spectra are concatenated into a single vector.

4.3.2 Image Data

As the vehicle traverses only sub-regions captured by the image, it is essential to obtain the image of those sub-regions. Thus, we project the vehicle trajectory into the distortion-corrected image, and keep footprints with the similar heading direction as the viewpoint. To account for scale ambiguity caused by perspective projection, we keep the closest m footprints. Around each footprint, we crop a box sizing $w \times h$ to indicate areas traversed by the robot and associate them with corresponding vehicle states. For vehicle's footprint i , denote ψ_i as the yaw angle, $\mathbf{p}_i^w \in \mathbb{R}^3$ as the position of footprint in the world frame, and $\mathbf{p}_i^c \in \mathbb{R}^2$ as the projected footprint in the image frame. Let \mathbf{K} be the camera intrinsic matrix, $\mathbf{R}_w^c \in \mathbf{SO}(3)$ be the rotation matrix, and $\mathbf{t}_w^c \in \mathbb{R}^3$ be the translation vector from the world frame to the camera frame. The projected vehicle footprints are,

$$\mathbf{p}_i^c = \mathbf{K}[\mathbf{R}_w^c | \mathbf{t}_w^c] \mathbf{p}_i^w \quad \forall i \in \{j : |\psi_j - \psi_0| \leq \theta\}$$

4.3.3 Data Augmentation

To remedy the data imbalance issue caused by discarding a subset of footprint image patches, we propose to augment the data by generating imaginary images for flat surfaces. We first generate a birds-eye-view (BE-view) image from the first-person-view (FP-view) input, and then take the cropped samples from the BE-view images. The imaginary BE-view images are warped using the Euclidean homography \mathbf{H}_{bc} from the FP-view frame c to the BE-view frame b as follows. For any point in the FP-view image, \mathbf{p}^c , the corresponding point in the BE-view image is,

$$\mathbf{p}^b = \lambda \mathbf{K} \cdot \mathbf{H}_{bc} \cdot \mathbf{K}^{-1} \cdot \mathbf{p}^c \quad (4.1)$$

where λ is a scale factor. Denote \mathbf{n} as the normal of the ground plane, and d as the distance from origin to the plane. We can write the Euclidean homography \mathbf{H}_{bc} as,

$$\mathbf{H}_{bc} = \mathbf{R}_b^c - \frac{\mathbf{t}_b^c \mathbf{n}^T}{d}$$

Corresponding projected footprints are warped using Eq. (4.1). Then image patches for data augmentation are cropped from BE-view images around these transformed footprints.

4.4 Estimating Traversability Cost from Vehicle States

Our RCA framework only utilizes vehicle states during training as intermediate learning signals. These learning signals are eventually turned into numerical costs. This way, the main learning algorithm uses these predicted costs, along with the images, as self-supervision, rather than utilizing vehicle states directly. This section describes how we estimate traversability costs from vehicle states for that training. The idea is composed of two parts: first, we use an unsupervised learning method as in [33] to generate coarse labels for vehicle state clusters; second, we use the learned clusters to estimate a continuous cost function given a vehicle state input.

4.4.1 Unsupervised Learning for Classifying Vehicle States

After input data preprocessing, we have image patches indicating traversal regions and frequency spectra describing experienced motions. As opposed to previous self-supervised vibration-based approaches which handcraft features[1][29], we exploit the complementary information shared by the visual and vehicle state domains to extract latent features.

To learn this association, we generally follow the unsupervised feature learning framework proposed in[33]. In essence, this framework brings visually-similar samples closer and visually-distinct samples further away in the target space, *e.g.*, acoustic data in their work. We customize the framework for our interpretation on riding comfort as follows. First, we extract image features using the Deep Encoding Pooling (DEP) network [28], which is specialized for ground terrain recognition. Next, in order to ensure that negative samples are selected from different ground truth classes, we leverage semantic classes clustered in the visual feature space. Such prior knowledge serves as a reference for computing Euclidean distance and selecting negative sample.

4. Approach

Based on the PCA projection of the feature space, we perform k-Means clustering to obtain coarse labels for terrain classes.

4.4.2 Traversal Cost Generation

Wheelchair passengers are exposed to greater physiological risks[7] and psychological barriers[9] with drastic changes in the vehicle motions. Thus we hypothesize that larger and more frequent movements along three dimensions should be avoided, and assigned with higher cost. The amplitude spectra show amounts of motion variations at different frequencies, and vehicle state clusters obtained in [Section 4.4.1](#) serve as a prior to show distinct vehicle dynamics.

The traversal cost function is considered as a weighted average of roll, pitch, and Z dimension measuring similarity to the average amplitude spectrum of its vehicle state class. Different vehicle state classes are also magnified with a hyperparameter to offset from each other. Let $\mathbf{A}_i^{\mathbf{d},\mathbf{k}} \in \mathbb{R}^m$ denote the amplitude spectrum along dimension d from the vehicle state class k for the sample i . N^k describes the number of samples within the vehicle state class k . A traversal cost for sample i from vehicle state class k is,

$$T_i^k = \omega_k \cdot \sum_{d=1}^3 \mathbf{M}^{\mathbf{d},\mathbf{k}} \cdot \mathbf{A}_i^{\mathbf{d},\mathbf{k}} \quad (4.2)$$

where ω_k is a weight parameter of vehicle state class k , and where $\mathbf{M}^{\mathbf{d},\mathbf{k}}$ is the mean amplitude spectrum along dimension d from the vehicle state class k ,

$$\mathbf{M}^{\mathbf{d},\mathbf{k}} = \frac{1}{N^k} \sum_{i=1}^{N^k} \mathbf{A}_i^{\mathbf{d},\mathbf{k}}$$

4.5 Self-supervised Traversal Cost Prediction

With the traversal cost function, we associate terrain visual appearances with motion measurements. We implement a encoder-decoder network, to estimate amplitude spectra of vehicle states and to predict traversal costs labeled by [Eq. \(4.2\)](#).

To align the visual feature space with the vehicle state feature space, terrain patches are first extracted by the same feature extractor in [Section 4.4.1](#). Then

terrain visual features are further encoded. During training, two decoders start back propagation at different epochs. In early epochs, the encoded space is translated toward vehicle states as only weights from the auxiliary vehicle state decoder get updated. Then later, the traversal cost decoder joins and learns from the partially translated feature space. We use \mathcal{L}_2 loss and smooth- \mathcal{L}_1 loss for regressing amplitude spectrum and traversal cost, respectively.

$$\mathcal{L}_2 = \|g(f(\mathbf{I}_i)) - \mathbf{A}_i\|^2$$

Define $x = h(f(\mathbf{I}_i)) - T_i$ as the difference between the traversal cost estimation and the computed traversal cost,

$$smooth_{\mathcal{L}_1} = \begin{cases} 0.5x^2 & \text{if } |x| < 1 \\ |x| - 0.5 & \text{otherwise} \end{cases}$$

The regression loss is written as,

$$\mathcal{L} = \beta\mathcal{L}_2 + (1 - \beta)smooth_{\mathcal{L}_1}$$

Given a camera image, it infers the traversal cost for each cropped patch and aggregates them as 2D costmaps.

4. Approach

Chapter 5

Experiments

We first describe the training dataset and implementation details, followed by descriptions of the main robot experiments.



Figure 5.1: (a) The tracked vehicle used for data collection. (b) The robotic wheelchair platform used in the experiments equipped with a 4.1GHz i7 computer and a NVIDIA GTX 1660Ti GPU card. Both platforms are equipped with a Velodyne Puck Lidar, a MEMS-based IMU, and a monocular camera.

5.1 Training the Learning Modules

5.1.1 Dataset

We collected a training dataset in an urban area using a tracked vehicle (Fig. 5.1a) that competently covered a variety of challenging terrains without posing risks to human operators. The camera captures frontal images at 5 Hz with 1280×1024 resolution. The state estimation module provides motion measurements at 200 Hz. Our dataset includes 2 hours of operated driving at 1 m s^{-1} including variations of soft and hard surfaces with different elevations to ensure a broad spectrum of vehicle motion profiles.

5.1.2 Implementations

Estimating Traversability Cost from Vehicle States: The vehicle state is split into windows of 256 samples and a stride size of 128 samples. The size of the hidden space vector is 32. The image is cropped into patches of 256×256 . We adopt Adam Optimizer with a learning rate as 0.001, a weight decay as 0.0001, and $\alpha = 0.4, \gamma = 20$. We trained for 120 epochs using a batch size of 64.

Self-supervised Traversal Cost Prediction: The lower half of the input image is cropped into small patches of 256×256 pixels. We adopt Adam Optimizer with a learning rate as 0.0001 and a weight decay as 0.00001, training in 150 epochs and batches of 64 each. We set $\beta = 1$ during the first 40 epochs and adjust to $\beta = 0.4$ for the remaining. As for the k-MEANS clustering, $k = 3$ was chosen empirically.

5.1.3 Evaluation Metrics for Learning Accuracy

Following previous works [33][13], we use two metrics to evaluate classification of vehicle state unsupervised learning using weakly supervised labels. One is normalized mutual information (NMI) defined as:

$$\text{NMI}(Y, C) = \frac{2I(Y, C)}{H(Y) + H(C)}$$

where $I(Y, C)$ denotes the mutual information between cluster sets Y and C , and $H(Y)$, $H(C)$ denote the entropy of cluster sets Y and C respectively. We define Y as the cluster set of the ground truth labels. The clustering accuracy can be defined as:

$$\text{Accuracy}(Y, C) = \frac{1}{N} \sum_k \max_j |y_k \cap c_j|$$

For evaluating self-supervised learning for cost prediction, we also measure how well vehicle states are reconstructed given the visual input to indicate how much the feature domain is translated from image to vehicle states. We use Mean Square Error (MSE) average over three axes defined as:

$$\text{MSE} = \frac{1}{3} \sum_{d=1}^3 \left(\frac{1}{n} \sum_{i=1}^n \left(\mathbf{A}_i^d - \hat{\mathbf{A}}_i^d \right)^2 \right)$$

5.2 Robot Experiments

To evaluate the quality of RCA, we compare the performance with two baselines: a 3D-based approach [2] and a semantic classification-based one [27]. To assess the performance of three terrain analysis approaches in actual navigation tasks, we conduct robot experiments with a wheelchair-based vehicle (Fig. 5.1b) navigating various terrain conditions. To investigate the performance according to various factors determining ride comfort and to evaluate perceived motion profiles, we design and conduct a human evaluation study.

5.2.1 Experiment Design

For evaluation, we record a set of robot videos in a controlled setting such that, for each scenario, the wheelchair robot repeats the same task using the proposed RCA and the two baseline approaches, *i.e.*, starting at the same pose towards a short-term goal point specified by a human controller. Each run is repeated three times. The scenarios consists of commonly seen objects and terrains in the urban area.

5. Experiments




Figure 5.2: Three approaches are tested in five scenes including asphalt roads, grass, curbs, tactile pavings, gravels (top row). The wheelchair starts at the same initial pose and tend to reach the desired goal points (red cross). The bottom row shows recorded wheelchair trajectories in a top-down view (our RCA in blue; semantic-based in red; and 3D-based in green). In (1), RCA detours away from the curb and uneven surfaces until it reaches the ramp. In (2), RCA takes a side way while avoiding the turbulence brought by tactile pavings. In (3)&(4), RCA foresees the effect from sinking in the soft terrain and avoids the vegetation although going through it is the shortest path. In (5), RCA keeps the wheelchair away from the vegetation and gravels.

5.2.2 Perceived Vehicle Motion Intensity (PMI)

The PMI Score [6] provides a quantitative measurement on the perception of inertial of motions. It hypothesizes the perception of motions is affected by acceleration, jerk, and their interaction, following a normal distribution with mean as,

$$\mu_{\psi} = \omega_A A_{\max} + \omega_J J_{\max} + \omega_{AJ} (A_{\max} \times J_{\max}) \quad (5.1)$$

where A_{\max} , J_{\max} are the maximum of acceleration and jerk for a given motion respectively, and where ω_A , ω_J , and ω_{AJ} are weights on acceleration, jerk and their interactions.

5.2.3 Human Evaluation

We designed a Amazon Mechanical Turk (AMT) study to assess human evaluation of the performances on five dimensions of ride comfort. During a session, workers are expected to watch three videos showing different robot behaviors in the same

scenario; the order of three videos is randomized. The participants are asked to rank the three videos from the most preferred (1) to the least preferred (3) according to the following five criteria:

- Stability: A ride is considered stable when noticing few frame-to-frame jitters and gentle orientation changes.
- Path normality: A ride is considered normal if the path matches with a human's expectation.
- Safety: A ride is considered as safe when observing continuously smooth motions, predictable paths, and avoiding obstacles at proper distance.
- Trustworthiness: A ride is considered as trustworthy if a passenger could rely on the wheelchair to complete trips independently.
- Overall preference: A ride is preferred when you would like to have a same wheelchair in your community.

5. Experiments

Chapter 6

Results

6.1 Learning Verification

For our verification purpose, we first visualize the vehicle state feature space colored by semantic classes of recording scenes (Fig. 6.1). We note that during the unsupervised learning process, those semantics classes serve as a prior for triplet selection. The visualization validates our hypothesis that semantic classes do not fully represent the inherent structure of vehicle dynamics features, and thus our RCA approach based on vehicle states is more targeted toward safe and comfortable navigation.

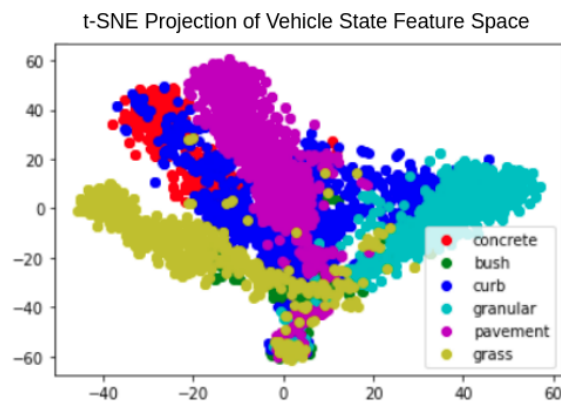


Figure 6.1: Visualization of vehicle state feature space using t-SNE projection. Samples are colored by semantic classes of recorded scenes.

We then measure the performance of unsupervised vehicle state clustering in

6. Results

terms of NMI and Accuracy, 0.56 and 86.84, respectively. Next, the performance of the vehicle state prediction, measured by MSE, was 0.53. Fig. 6.2 qualitatively shows the prediction results against the ground truth. The main evaluation was verified through the downstream task of robot navigation.

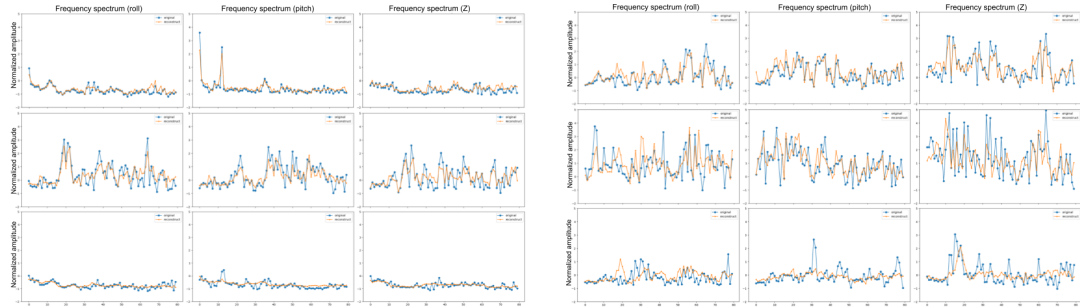


Figure 6.2: Prediction results of normalized amplitude spectra along roll, pitch and Z axis. It shows relative amounts of amplitude for each frequency component. Ground truth signals are marked in blue, while the estimates ones are shown in orange.

6.2 Robot Experiments

We used 45 robot videos for 5 scenarios shown in Fig. 5.2 for each of the three approaches with 3 repeated trials. After briefing on the learning performance, we report the PMI results to provide quantitative analysis on our robot experiments. We then report on the full human evaluation results based on 55 MTurk workers' responses.

6.2.1 PMI Analysis

Fig. 6.3 shows the aggregated PMI (Eq. (5.1)) results from averaging each run of the robot experiments. Our approach brought less turbulence and unstable motions to the vehicle, especially in the vertical, roll, and pitch axes. As traversal cost is defined based on the amount of variations in the vehicle states, it inherently switches to regions with less intense displacements along three axes. Limited by the angle of view from the front-facing camera, the wheelchair displays occasional turnings along long trips. Although it introduces a fair amount of yaw axis changes to avoid uneven terrains, the PMI originated from heading changes is relatively comparable

as the 3D-based approach, which tends to generate a straight-line path wherever terrain heights permit. The semantic-based method avoids moderately rough terrains; however, it fails to provide fine-grid estimation within a semantic class, and thus results in more frequent turnings to search for alternative passable semantic classes.

6.2.2 Human Evaluation Analysis

Fig. 6.4 presents a vote breakdown in percentage received by each approach among all participants. Based on these collective responses, we found that our approach is substantially better than two baseline approaches on average. The results show consistent top rankings across five dimensions with moderate variations across scenarios. It does not show a clear first runner up approach. The results also cast light on the positive correlation between objective factors (i.e. stability, safety) to subjective factors (i.e. trustworthy, preference).

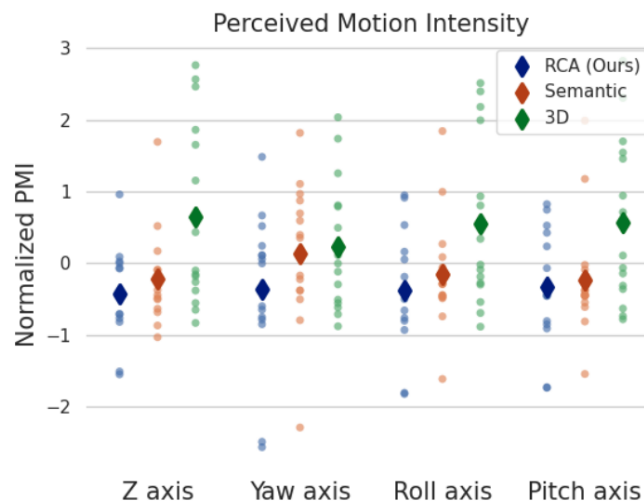


Figure 6.3: Normalized Perceived Vehicle Motion Intensity Scores[6] for Z (upward pointing), yaw, roll and pitch axes. Our approach (blue) shows consistently smaller scores for all axes, demonstrating less disturbances. Semantic-based baseline (orange) exhibits a higher value along Yaw axis, implying its frequent and aggressive turning motions. 3D-based baseline (green) tends to have larger values along Z, Roll, and Pitch axes, indicating its frequent shakiness over uneven terrains.

6.2.3 PMI vs. Human Evaluation

Tab. 6.1 reports the Spearman Correlation scores between stability rankings and Perceived Motion Intensity ranks along Z, Roll, Pitch, Yaw axes among all experiment runs with $p < 0.01$. It reveals strong positive correlations between human perceived stability scores and Perceived Motion Intensity computed from acceleration and jerk along Z, Roll, Pitch, Yaw axes. It further validates that our approach improves the perceived stability relatively by reducing motion intensity along Z, Roll, Pitch, Yaw axes.

Table 6.1: Spearman Correlation Scores

	PMI Z	PMI Roll	PMI Pitch	PMI Yaw
Stability	0.69	0.70	0.61	0.80




Figure 6.4: Percentage of samples ranking three approaches among five dimensions. Each vertical bar consists of three segments representing the percentage of samples voting an approach as the first (bottom, darker color), the second (middle, moderate color), the third (top, lighter color). Our approach marks as the blue bar, while the semantic-based baseline is in red and the 3D-based baseline is in green. Across five evaluation criteria over five scenes, our approach is most likely to be ranked as the top. We fail to show a clear winner between semantic-based approach and 3D-based approach.

Chapter 7

Conclusions

In conclusion, we introduce a self-supervised ride comfort-aware approach for terrain traversability analysis. It is motivated by curtailing disturbances caused by traversing uneven or soft surfaces to enhance wheelchair ride comfort. Our RCA approach uses internal vehicle states to self-supervise the process of traversal cost estimation from visual images. According to Perceived Motion Intensity Scores, the results of robot experiments strongly support the effectiveness of our approach on reducing disturbances. Based on human evaluation of 55 MTurk workers, RCA is ranked the highest when compared to the baseline approaches in terms of stability, path normality, safety, trustworthiness, and overall preference. We also recognize that RCA does not perform consistently well across all scenarios. Whereas RCA generally performs well by estimating the disturbing motions brought by uneven or soft terrains, the semantic-based approach provides human-interpretive estimation in the pixel space and the 3D-based approach provides robust and accurate geometric measurements of rigid objects. These findings lead to our future direction for unifying comfort-aware, semantic-based, and 3D-based approaches towards the goal of supporting safe and comfortable rides for the wheelchair users.

7. Conclusions

Appendix A

Appendix

We also include additional information for reference.

A.1 Dataset Examples

Example scenes of our dataset is shown in [Fig. A.1](#).



Figure A.1: Our dataset includes various types of soft and hard terrains.

A. Appendix

Bibliography

- [1] Christopher A. Brooks and Karl Iagnemma. Self-supervised terrain classification for planetary surface exploration rovers. *J. Field Robotics*, 29:445–468, 2012. [2.4](#), [4.4.1](#)
- [2] Chao Cao, Hongbiao Zhu, Fan Yang, Yukun Xia, Howie Choset, Jean Oh, and Ji Zhang. Autonomous exploration development environment and the planning algorithms. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2021. [1.1](#), [2.1](#), [4.1](#), [5.2](#)
- [3] R. Omar Chavez-Garcia, Jérôme Guzzi, Luca Maria Gambardella, and Alessandro Giusti. Learning ground traversability from simulations. *IEEE Robotics and Automation Letters*, 3:1695–1702, 2018. [2.1](#)
- [4] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *ECCV*, 2018. [4.1](#)
- [5] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. [2.2](#)
- [6] Ksander N. de Winkel, Florian Soyka, and Heinrich H. Bühlhoff. The role of acceleration and jerk in perception of above-threshold surge motion. March 2020. doi: 10.1007/s00221-020-05745-7. ([document](#)), [5.2.2](#), [6.3](#)
- [7] R P Gaal, N Rebholtz, Ralf D. Hotchkiss, and P. Pfaelzer. Wheelchair rider injuries: causes and consequences for wheelchair design and selection. *Journal of rehabilitation research and development*, 1997. [4.4.2](#)
- [8] Biao Gao, Shaochi Hu, Xijun Zhao, and Huijing Zhao. Fine-grained off-road semantic segmentation and mapping via contrastive learning. *CoRR*, abs/2103.03651, 2021. [2.2](#)
- [9] Takashi Isezaki, Arinobu Nijjima, Akihiro Miyata, Tomoki Watanabe, and Osamu Mizuno. Wheelchair users’ psychological barrier estimation based on inertial

- and vital data. In Margherita Antona and Constantine Stephanidis, editors, *Universal Access in Human-Computer Interaction. Interaction Techniques and Environments*, pages 403–413. Springer International Publishing, 2016. ISBN 978-3-319-40244-4. [4.4.2](#)
- [10] Gregory Kahn, Pieter Abbeel, and Sergey Levine. BADGR: an autonomous self-supervised learning-based navigation system. *CoRR*, abs/2002.05700, 2020. [2.3](#)
- [11] Dongshin Kim, Jie Sun, Sangmin Oh, James M. Rehg, and Aaron F. Bobick. Traversability classification using unsupervised online visual learning for outdoor robot navigation. *Proceedings 2006 IEEE International Conference on Robotics and Automation*, pages 518–525, 2006. [2.2](#)
- [12] Alicia M. Koontz, Dan Ding, Yih-Kuen Jan, Sonja de Groot, and Andrew Hansen. Wheeled mobility. *BioMed research international*, 2015. [1.1](#)
- [13] Akiyoshi Kurobe, Yoshikatsu Nakajima, Kris Kitani, and Hideo Saito. Audio-visual self-supervised terrain type recognition for ground mobile platforms. *IEEE Access*, 9:29970–29979, 2021. doi: 10.1109/ACCESS.2021.3059620. [2.4](#), [5.1.3](#)
- [14] Jesse Leaman and Hung Manh La. A comprehensive review of smart wheelchairs: Past, present, and future. *IEEE Transactions on Human-Machine Systems*, 47(4):486–499, 2017. doi: 10.1109/THMS.2017.2706727. [1.1](#)
- [15] Joonho Lee, Jemin Hwangbo, Lorenz Wellhausen, Vladlen Koltun, and Marco Hutter. Learning quadrupedal locomotion over challenging terrain. *CoRR*, abs/2010.11251, 2020. [2.3](#)
- [16] Michelle A. Lee, Yuke Zhu, Peter Zachares, Matthew Tan, Krishnan Srinivasan, Silvio Savarese, Li Fei-Fei, Animesh Garg, and Jeannette Bohg. Making sense of vision and touch: Learning multimodal representations for contact-rich tasks. *IEEE Transactions on Robotics*, 36(3):582–596, 2020. doi: 10.1109/TRO.2019.2959445. [2.4](#)
- [17] Katharina Muelling, Arun Venkatraman, Jean-Sebastien Valois, John Downey, Jeffrey Weiss, Shervin Javdani, Martial Hebert, Andrew B Schwartz, Jennifer L Collinger, and J Andrew Bagnell. Autonomy infused teleoperation with application to bci manipulation. *arXiv preprint arXiv:1503.05451*, 2015. [1.1](#)
- [18] Helen Oleynikova, Alexander Millane, Zachary Taylor, Enric Galceran, Juan Nieto, and Roland Siegwart. Signed distance fields: A natural representation for both mapping and planning. In *RSS 2016 Workshop: Geometry and Beyond-Representations, Physics, and Scene Understanding for Robotics*. University of Michigan, 2016. [2.1](#)
- [19] Kyohei Otsu, Masahiro Ono, Thomas J. Fuchs, Ian A. Baldwin, and Takashi Kubota. Autonomous terrain classification with co- and self-training approach.

- IEEE Robotics and Automation Letters*, 1:814–819, 2016. [2.2](#)
- [20] Simone Palazzo, Dario C. Guastella, Luciano Cantelli, Paolo Spadaro, Francesco Rundo, Giovanni Muscato, Daniela Giordano, and Concetto Spampinato. Domain adaptation for outdoor robot traversability estimation from rgb data with safety-preserving loss. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020. [2.2](#)
- [21] Yunpeng Pan, Ching-An Cheng, Kamil Saigol, Keuntaek Lee, Xinyan Yan, Evangelos A. Theodorou, and Byron Boots. Agile autonomous driving using end-to-end deep imitation learning. In *Robotics: Science and Systems*, 2018. [2.3](#)
- [22] Milos Prágr, Petr íek, Jan Bayer, and Jan Faigl. Online incremental learning of the terrain traversal cost in autonomous exploration. *Robotics: Science and Systems XV*, 2019. [2.1](#)
- [23] Mario Selvaggio, Marco Cagnetti, Stefanos Nikolaidis, Serena Ivaldi, and Bruno Siciliano. Autonomy in physical human-robot interaction: A brief survey. *IEEE Robotics and Automation Letters*, 2021. [1.1](#)
- [24] Abhinav Valada and Wolfram Burgard. Deep spatiotemporal models for robust proprioceptive terrain classification. *CoRR*, abs/1804.00736, 2018. [2.4](#)
- [25] K. S. Viswanath, Kartikeya Singh, Peng Jiang, P. B. Sujit, and Srikanth Saripalli. Offseg: A semantic segmentation framework for off-road driving. *2021 IEEE 17th International Conference on Automation Science and Engineering (CASE)*, pages 354–359, 2021. [1.1](#), [2.2](#), [2.4](#)
- [26] Lorenz Wellhausen, Alexey Dosovitskiy, René Ranftl, Krzysztof Walas, César Cadena, and Marco Hutter. Where should i walk? predicting terrain properties from images via self-supervised learning. *IEEE Robotics and Automation Letters*, 4:1509–1516, 2019. [2.4](#)
- [27] Zhanxin Wu, Xinjie Yao, and Jean Oh. Semantic segmentation in complex scenes for robotics navigation. In *Carnegie Mellon Robotics Institute Summer Scholars Working Papers Journal*, 2021. [5.2](#)
- [28] Jia Xue, Hang Zhang, and Kristin J. Dana. Deep texture manifold for ground terrain recognition. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 558–567, 2018. [4.4.1](#)
- [29] H. W. Yu and B. H. Lee. A bayesian approach to terrain map inference based on vibration features. In *2017 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI)*, pages 272–277, 2017. doi: 10.1109/MFI.2017.8170440. [2.4](#), [4.4.1](#)
- [30] Ji Zhang and Sanjiv Singh. Loam: Lidar odometry and mapping in real-time. In *Proceedings of Robotics: Science and Systems (RSS '14)*, July 2014. [4.1](#)

Bibliography

- [31] Ji Zhang, Chen Hu, Rushat Gupta Chadha, and Sanjiv Singh. Falco: Fast likelihood-based collision avoidance with extension to human-guided navigation. *Journal of Field Robotics*, 37(8):1300 – 1313, December 2020. [4.1](#)
- [32] Zeyu Zhu, Nan Li, Ruoyu Sun, Huijing Zhao, and Donghao Xu. Off-road autonomous vehicles traversability analysis and trajectory planning based on deep inverse reinforcement learning. *2020 IEEE Intelligent Vehicles Symposium (IV)*, pages 971–977, 2020. [2.3](#)
- [33] Jannik Zurn, Wolfram Burgard, and Abhinav Valada. Self-supervised visual terrain classification from unsupervised acoustic feature learning. *IEEE Transactions on Robotics*, 37:466–481, 2021. [1.1](#), [2.4](#), [4.4](#), [4.4.1](#), [5.1.3](#)