

Robot Adaptation to Unstructured Terrains by Joint Representation and Apprenticeship Learning

Sriram Siva¹, Maggie Wigness², John Rogers², and Hao Zhang¹

¹Colorado School of Mines, ²U.S. Army Research Laboratory

sivasriram@mines.edu, maggiee.b.wigness.civ@mail.mil, john.g.rogers59.civ@mail.mil, hzhang@mines.edu

Abstract—When a mobile robot is deployed in a field environment, e.g., during a disaster response application, the capability of adapting its navigational behaviors to unstructured terrains is essential for effective and safe robot navigation. In this paper, we introduce a novel joint terrain representation and apprenticeship learning approach to implement robot adaptation to unstructured terrains. Different from conventional learning-based adaptation techniques, our approach provides a unified problem formulation that integrates representation and apprenticeship learning under a unified regularized optimization framework, instead of treating them as separate and independent procedures. Our approach also has the capability to automatically identify discriminative feature modalities, which can improve the robustness of robot adaptation. In addition, we implement a new optimization algorithm to solve the formulated problem, which provides a theoretical guarantee to converge to the global optimal solution. In the experiments, we extensively evaluate the proposed approach in real-world scenarios, in which a mobile robot navigates on familiar and unfamiliar unstructured terrains. Experimental results have shown that the proposed approach is able to transfer human expertise to robots with small errors, achieve superior performance compared with previous and baseline methods, and provide intuitive insights on the importance of terrain feature modalities.

I. INTRODUCTION

Over the past several years, autonomous mobile robots have been more commonly used in unstructured field environments to address real-world applications, including disaster response, infrastructure inspection, homeland defense, and subterranean and planetary exploration [1, 2]. When deployed in an outdoor field environment, mobile robots need the essential capability of efficiently navigating through different types of terrains with a wide variety of characteristics. For example, as illustrated in Figure 1, when mobile robots perform disaster response tasks in a post-disaster situation, they are likely to encounter terrains such as grass, rock, pebble, mud, snow, and a mixture of these terrain types (e.g., a rocky terrain partially covered by mud and snow). In addition, the same type of terrains typically exhibits a variety of characteristics (e.g., different slope and softness). Such terrains whose types and characteristics cannot be fully modeled or determined before robot deployment are referred to as *unstructured terrains* [3, 4]. The capability of generating desired navigation behaviors (e.g., velocity control) that adapt to unstructured terrains is essential for an autonomous mobile robot to operate in complex field environments.

Given the importance of robot navigation over unstructured terrains, a large number of terrain classification and adaptation techniques were implemented [5, 6, 7, 8]. Especially, learning-based techniques have been attracting increasing attention over



Figure 1. Motivating scenarios of robot adaptation to unstructured terrains. When mobile robots operate in a field environment during disaster response, the robots need the essential capability of navigating over a variety of terrains, such as grass, rock, snow, pebble, and mud, which often cannot be accurately structured before onsite robot deployment. In this paper, we propose the TRAL method to integrate representation learning and apprenticeship learning under a unified mathematical framework and estimate terrain feature importance to enable robust robot adaptation to unstructured terrains.

the last few years. For example, unsupervised learning-based techniques were designed to perform semantic segmentation of terrains, and then navigational decisions were constructed by a robot through estimating traversability of unstructured terrains [9, 10, 11]. In addition, methods based upon terrain classification were implemented for mobile robots to reason about the possibility of traversing uneven terrains [12, 13]. Moreover, to teach navigational behaviors to robots, apprenticeship learning (a.k.a., learning from demonstration or imitation learning) was widely investigated [3, 11, 14] to transfer human expertise in navigation control to autonomous robots and generate adaptive behaviors according to given terrains (i.e., terrain adaptation).

Despite of promising performance of the previous methods, several challenges in robot adaptation to unstructured terrains have not yet been well addressed. Most existing learning-based methods focused on either terrain classification or navigation behavior generation only [15, 16, 17, 18, 19], or treated them as two separate components in the processing pipeline without a unified integration [14]. Moreover, previous adaptation techniques either used a single feature to represent terrains [20], or combined a number of features by simple concatenation [21]. The question of how to automatically estimate the importance of heterogeneous terrain features (e.g., from exteroceptive and proprioceptive sensors) and fuse them together for robust adaptation has not been well answered.

In this paper, we propose a novel joint *terrain representation and apprenticeship learning* (TRAL) approach to enable robot

adaptation to unstructured terrains. Instead of treating terrain classification and navigational behavior generation as separate tasks, our approach integrates them into a unified mathematical formulation under the regularized optimization paradigm. Our approach learns representations of unstructured terrains, which are jointly used to classify terrains and perform apprenticeship learning to generate navigational behaviors (e.g. velocity control) in the unified formulation. Furthermore, our approach is able to automatically learn weights of terrain features to create a multi-modal representation of terrains, by designing sparsity-induced norms as regularizers in the regularized optimization paradigm. These weights encode the importance of the terrain features, thus offering the insight of which terrain features are most critical for terrain classification and navigation behavior generation.

The key theoretical novelty of the paper is twofold:

- We propose a novel TRAL method that is able to integrate representation learning and apprenticeship learning under a unified mathematical framework, automatically estimate terrain feature importance, and incorporate heterogeneous terrain features to enable robot adaptation to unstructured terrains.
- We develop a new optimization algorithm that efficiently solves the formulated joint representation and apprenticeship learning problem, which holds a theoretical guarantee to converge to the optimal solution.

The remaining of the paper is organized as follows. Related work is reviewed in Section II. We discuss the formulation of TRAL for robot adaptation to unstructured terrains in Section III. The new optimization algorithm is described in Section IV. Experiments are presented in Section V. Finally, we conclude the paper in Section VI.

II. RELATED WORK

This section provides a review of related research on terrain classification and robot adaptation, and identifies key research challenges to be addressed in this paper.

A. Terrain Classification and Characterization

Terrain classification and characterization aims at extracting terrain features and associating them with terrain categories. Much of earlier methods was designed for larger vehicles and the characterization process was typically manual. There was an amount of research on Terramechanics [22], the guidance of autonomous vehicles through rough terrains. Similarly, Pre-planning was used for speed selection [23], e.g., for vehicles used in the DARPA Grand Challenge, to address the problem of unstructured terrain navigation through fixing the optimal speed for each terrain type by collecting data for months. [24] generated digital terrain maps and then fitted a Gaussian model to classify terrains. [25] employed terrain cohesion and internal frictional angle obtained from a variety of terrains to perform terrain characterization.

In learning-based methods, terrain classification was usually performed by learning from terrain data in a regressive manner. For example, [26] designed a feed forward neural network on

terrain images to classify terrain types. Motivated by the need to recognize visual terrain features that impact mobility, [27] used Hidden Markov Models to learn terrain features and an SVM classifier to identify terrain types. [15] used a series of sensors to measure robot vibrations on terrains to successfully classify terrains. [28] used a color-based classification to label obstacles based on terrain classes and carried out appropriate navigational behaviors.

Previous methods either used a single feature modality or a combination of features to represent terrains, and the problem of estimating the importance of feature modalities has not been well addressed yet, especially in the context of apprenticeship learning for robot terrain adaptation. Also, previous methods generally considered terrain classification and characterization as a preprocessing procedure for subsequential reasoning such as navigation planning.

B. Robot Adaptation

In the area of robot adaptation, high-level behavioral models were implemented to address the problem of robot adaptation in general [29, 30, 31]. Similarly, case-based reasoning [32] techniques were developed to make robots to adapt in dynamic environments [33, 34]. Many model-based approaches aimed at learning a single global model using function approximators, such as the Gaussian processes [35, 36] and neural networks [37]. Probabilistic Gaussian methods [38] were implemented to learn terrain models and update the models efficiently using sparse approximation. One of the key challenges of deploying such models is due to the difficulty of learning a global model that is accurate for the entire state space. Several approaches [39, 40] were developed to learn accurate local models instead of learning a global model. These techniques have acceptable local performance and allow for iterative local improvements of the policy, but they may not be feasible to be applied to an unexperienced domain. Recently, online learning approaches [41, 42] were also used to iteratively update model parameters during task execution. Although this online learning paradigm is widely used with robots operating in evolving environments, the effectiveness of the methods in the case of sudden change of terrain features is questionable.

For the research problem of robot adaptation to terrains, several approaches attempted to generate autonomously-adaptable dynamic maneuvering using a neural system model in simulations [43] and operational robots [44, 45]. Recently, methods based upon self-supervised learning [46, 9] were implemented to navigate in a minimal collision environment by maintaining minimal traversal and bumper history evaluations. [47] applied an array of inertial and ultrasonic sensors to calculate features of soil properties and carry out terrain adaptation. In the need for human-like operation and navigation [48], [11] used visual perception data and inverse optimal control trained with human supervision to learn to imitate expert navigation behaviors. By considering speed as an important factor for robot navigation, [49] presented a method for trading progress and velocity with respect to environment characteristics.

Given the promise of the previous learning-based techniques for robot adaptation to terrains, almost all existing approaches treated behavior generation and terrain classification as separate procedures, with no principled integration. Different from previous methods, our proposed approach explicitly addresses this challenge through integrating representation learning and apprenticeship learning under the same theoretical framework to perform joint terrain classification and behavior generation for terrain adaptation.

III. APPROACH

In this section, we introduce the novel TRAL approach that formulates robot adaptation to unstructured terrains as a joint representation and apprenticeship learning problem, from the unified regularized optimization perspective. An illustration of the TRAL approach is shown in Figure 2.

Notations: Matrices are denoted as boldface capital letters, and vectors as boldface lowercase letters. Given a matrix $\mathbf{M} = \{m_{ij}\} \in \mathbb{R}^{m \times n}$, we refer to its i -th row and j -th column as \mathbf{m}^i and \mathbf{m}_j , respectively. The Frobenius norm of \mathbf{M} is defined as $\|\mathbf{M}\|_F = \sqrt{\sum_{i=1}^m \sum_{j=1}^n m_{ij}^2}$. The ℓ_1 -norm of a vector $\mathbf{v} \in \mathbb{R}^n$ is computed by $\|\mathbf{v}\|_1 = \sum_{i=1}^n |v_i|$. The ℓ_2 -norm of \mathbf{v} is computed by $\|\mathbf{v}\|_2 = \sqrt{\mathbf{v}^\top \mathbf{v}}$.

A. Representation Learning for Terrain Classification

We denote a collection of n data instances of terrain features that are acquired by a robot when traversing a terrain as $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n] \in \mathbb{R}^{d \times n}$, where $\mathbf{x}_i \in \mathbb{R}^d$ denotes a feature vector of length d . We represent terrain types associated with \mathbf{X} as $\mathbf{Z} = [\mathbf{z}_1, \dots, \mathbf{z}_n] \in \mathbb{R}^{c \times n}$, where $\mathbf{z}_i \in \mathbb{Z}^c$ is an indicator vector with each element $z_{ij} \in \{0, 1\}$ indicating whether the i -th data instance of features \mathbf{x}_i has the j -th terrain type, and c is the number of terrain types. Given \mathbf{X} and \mathbf{Z} , representation learning for terrain classification can be simply formulated as an optimization problem [50, 51]:

$$\min_{\mathbf{V}} \|\mathbf{V}^\top \mathbf{X} - \mathbf{Z}\|_F^2 \quad (1)$$

where $\mathbf{V} \in \mathbb{R}^{d \times c}$ is a weight matrix with v_{ij} representing the importance of the i -th element in a feature vector with respect to the j -th terrain type. The loss function parameterized by \mathbf{V} in Eq. (1) models the squared error of utilizing the weighted features to represent the terrain types.

When the terrain features are computed by different feature extraction techniques and/or from different sensors, the feature vector contains multiple feature modalities, with each modality including features computed by one technique from one sensor. Then, each feature vector \mathbf{x}_i exhibits a multi-modal structure $\mathbf{x}_i = [(\mathbf{x}_i^1)^\top, \dots, (\mathbf{x}_i^m)^\top]^\top \in \mathbb{R}^d$, where m is the number of feature modalities in the vector. If d_j is the length of the j -th feature modality, we have $d = \sum_{j=1}^m d_j$.

Because different terrain features typically capture different characteristics of unstructured terrains (e.g., color, smoothness, and roughness) and often contribute differently toward terrain classification, it is essential to identify most descriptive feature modalities when multiple modalities are used together. In order

to enable this capability, we develop a norm to regularize the weight matrix \mathbf{V} [52]:

$$\|\mathbf{V}\|_R = \sum_{i=1}^c \sum_{j=1}^m \|\mathbf{v}_i^j\|_2 = \sum_{i=1}^c \sum_{j=1}^m \sqrt{\sum_{k=1}^{d_j} (v_{ik}^j)^2} \quad (2)$$

where $\mathbf{v}_i^j \in \mathbb{R}^{d_j}$ represents the weight vector of terrain features in the i -th modality with respect to the j -th terrain type. Since $\|\mathbf{V}\|_R$ uses the ℓ_2 -norm within each modality and the ℓ_1 -norm among modalities with respect to each terrain type, it enforces sparsity between modalities to identify the most discriminative feature modalities. That is, if a modality is not discriminative to represent the terrains for terrain classification, the weights of terrain features in this modality are assigned with zeros (in ideal case, usually they are small values close to 0); otherwise, the weights have big values. The effect of the norm $\|\mathbf{V}\|_R$ is illustrated in Figure 2.

Then, representation learning from multi-modal features for terrain classification is formulated as the following regularized optimization problem:

$$\min_{\mathbf{V}} \|\mathbf{V}^\top \mathbf{X} - \mathbf{Z}\|_F^2 + \lambda_1 \|\mathbf{V}\|_R \quad (3)$$

where $\lambda_1 \geq 0$ is a trade-off hyper-parameter applied to balance the loss function and the regularization term.

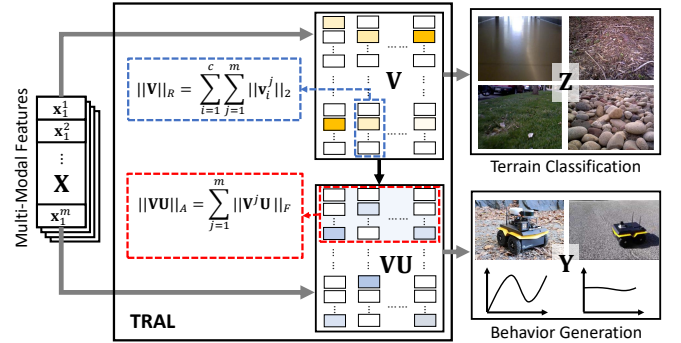


Figure 2. Illustration of the proposed TRAL approach for joint representation and apprenticeship learning under the unified regularized optimization framework, with the regularization terms $\|\mathbf{V}\|_R$ and $\|\mathbf{V}\mathbf{U}\|_A$ to identify the discriminative feature modalities.

B. Formulation of Joint Representation and Apprenticeship Learning for Unstructured Terrain Adaptation

Apprenticeship learning [48] is one of the most widely used learning-based paradigms to transfer human expertise to robots through directly expert demonstration, which can be viewed as the learning of a projection from observation or feature space into behavior space, e.g., by probabilistic [53] or optimization-based [14, 54] approaches.

Our goal in this paper is to achieve apprenticeship learning while retaining the advantage of discriminative features offered by representation learning. The theoretical novelty of the paper is that we introduce a principled integration of representation and apprenticeship learning in a unified formulation under the

mathematical framework of regularized optimization, which enables joint terrain classification and behavior generation.

Formally, given expert demonstrations on behavior controls $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_n] \in \mathbb{R}^{r \times n}$ (e.g., velocity, motor torque, and consumed power) that are associated with terrain features \mathbf{X} during robot navigation, where r is the number of controls we want to learn for behavior generation, we formulate joint representation and apprenticeship learning for terrain adaptation as a unified regularized optimization problem:

$$\min_{\mathbf{V}, \mathbf{U}} \|\mathbf{V}^\top \mathbf{X} - \mathbf{Z}\|_F^2 + \|(\mathbf{V}\mathbf{U})^\top \mathbf{X} - \mathbf{Y}\|_F^2 + \lambda_1 \|\mathbf{V}\|_C + \lambda_2 \|\mathbf{V}\mathbf{U}\|_A \quad (4)$$

where λ_2 is a tradeoff hyper-parameter. The new loss function (i.e., the 2nd term in Eq. (4)) encodes the difference between the learning model and expert demonstrations, in order to learn the projection from multi-modal feature space to robot behavior space. This loss function is parameterized by $\mathbf{V}\mathbf{U} \in \mathbb{R}^{d \times r}$, which first projects input multi-modal terrain features \mathbf{X} into a more discriminative space $\mathbf{V}^\top \mathbf{X}$ (parameterized by \mathbf{V} from representation learning), and further projects the representation space into the robot behavior space $\mathbf{U}^\top \mathbf{V}^\top \mathbf{X}$ (parameterized by $\mathbf{U} \in \mathbb{R}^{c \times r}$ for apprenticeship learning).

Because the output of the proposed joint learning approach is to generate robot behaviors, we also want our representation learning to learn discriminative terrain features that can better decide robot navigation behaviors. Therefore, we design a new norm for $\mathbf{V}\mathbf{U}$, which is mathematically defined as:

$$\|\mathbf{V}\mathbf{U}\|_A = \sum_{j=1}^m \|\mathbf{V}^j \mathbf{U}\|_F \quad (5)$$

where $\mathbf{V}^j \in \mathbb{R}^{d_j \times c}$ is the block of \mathbf{V} that includes weights of terrain features belonging to the j -th modality. This norm can be directly integrated into our objective function in Eq. (4) as a regularization term over $\mathbf{V}\mathbf{U}$, as demonstrated in Figure 2. This regularization enforces sparsity among feature modalities with respect to all robot behaviors, thus facilitating identifying discriminative features for robot behavior generation.

After solving the formulated regularized optimization problem in Eq. (4) based on Algorithm 1, we can obtain the optimal \mathbf{V}^* and \mathbf{U}^* . When a new multi-modal terrain feature vector $\mathbf{x}^o \in \mathbb{R}^d$ is acquired by a robot during execution, the robot's navigation behaviors can be determined by:

$$\mathbf{y}^o = (\mathbf{V}^* \mathbf{U}^*)^\top \mathbf{x}^o \quad (6)$$

which incorporates considerations of both representation learning and apprenticeship learning.

Different from conventional learning-based robot adaptation techniques [11, 12, 14], TRAL provides a unified formulation that integrates representation and apprenticeship learning under the same theoretical framework, instead of treating them as separate and independent procedures. Moreover, different from most existing apprenticeship learning methods that apply a single feature modality only [54] or simple feature concatenation [14], one of the TRAL's advantages is its capability to

Algorithm 1: An iterative algorithm to solve the formulated regularized optimization problem in Eq. (4)

Input : $\mathbf{X}, \mathbf{Y}, \mathbf{Z}$

- 1 Let $t = 1$. Initialize $\mathbf{V}(t)$ by solving $\min_{\mathbf{V}} \|\mathbf{V}^\top \mathbf{X} - \mathbf{Z}\|_F^2$
- 2 Initialize $\mathbf{U}(t)$ by solving $\min_{\mathbf{U}} \|(\mathbf{V}\mathbf{U})^\top \mathbf{X} - \mathbf{Y}\|_F^2$
- 3 **while** *not converge* **do**
- 4 Calculate the block diagonal matrix $\mathbf{D}(t+1)$, where the i -th diagonal element of $\mathbf{D}(t+1)$ is $\frac{\mathbf{1}}{2\|\mathbf{v}_i(t)\|_2} \mathbf{I}_i$
- 5 Calculate the block diagonal matrix $\widehat{\mathbf{D}}(t+1)$, where the j -th diagonal element is $\frac{1}{2\|\mathbf{V}^j \mathbf{U}(t)\|_F} \mathbf{I}_j$
- 6 For each
- 7 $\mathbf{v}_i(1 \leq i \leq c) = (\mathbf{X}\mathbf{X}^\top + \lambda_1 \mathbf{D} + \lambda_2 \widehat{\mathbf{D}})^{-1} \mathbf{X}(\mathbf{z}_i)^\top$
- 7 For each
- 8 $\mathbf{u}_i(1 \leq i \leq c) = (\mathbf{V}^\top \mathbf{X}\mathbf{X}^\top \mathbf{V} + \lambda_2 \mathbf{V}^\top \mathbf{V})^{-1} \mathbf{X}\mathbf{y}_i$
- 8 $t = t + 1$

Output: $\mathbf{V} = \mathbf{V}(t) \in \mathbb{R}^{d \times c}$ and $\mathbf{U} = \mathbf{U}(t) \in \mathbb{R}^{c \times r}$

automatically identify discriminative feature modalities, which can improve the robustness of robot adaptation to unstructured and unfamiliar terrains. Also, the weights learned by TRAL provide insights for intuitive analysis and future integration of heterogeneous terrain features for robot adaptation to unstructured terrains.

IV. OPTIMIZATION ALGORITHM

The second theoretical novelty of this paper is the design of an optimization algorithm to solve the formulated regularized optimization problem in Eq. (4). Although the problem in Eq. (4) is convex, it is challenging to solve using traditional gradient descent based solvers due to the non-smooth regularization terms and due to the dependency of the two loss functions.

To solve the formulated problem, we compute the derivative of Eq. (4) with respect to the columns of \mathbf{V} and set it to zero, while assuming \mathbf{U} to have a fixed value:

$$\mathbf{X}\mathbf{X}^\top \mathbf{v}_i - \mathbf{X}(\mathbf{z}_i)^\top + \lambda_1 \mathbf{D} \mathbf{v}_i + \lambda_2 \widehat{\mathbf{D}} \mathbf{v}_i = 0 \quad (7)$$

where \mathbf{D} is a diagonal matrix with the i -th diagonal element as $\frac{\mathbf{1}}{2\|\mathbf{v}_i\|_2}$, and $\widehat{\mathbf{D}}$ represents a block diagonal matrix with the j -th block computed by $\frac{\mathbf{1}}{2\|\mathbf{V}^j \mathbf{U}\|_F}$. Then, we obtain:

$$\mathbf{v}_i = (\mathbf{X}\mathbf{X}^\top + \lambda_1 \mathbf{D} + \lambda_2 \widehat{\mathbf{D}})^{-1} \mathbf{X}(\mathbf{z}_i)^\top \quad (8)$$

Fixing the value of \mathbf{V} , to compute \mathbf{U} , we take the derivative of Eq. (4) with respect to the columns of \mathbf{U} and set it to zero:

$$\mathbf{V}^\top \mathbf{X}\mathbf{X}^\top \mathbf{V} \mathbf{u}_i - \mathbf{V}^\top \mathbf{X}(\mathbf{y}_i)^\top + \lambda_2 \mathbf{V}^\top \mathbf{V} \mathbf{u}_i = 0 \quad (9)$$

Then we obtain:

$$\mathbf{u}_i = (\mathbf{V}^\top \mathbf{X}\mathbf{X}^\top \mathbf{V} + \lambda_2 \mathbf{V}^\top \mathbf{V})^{-1} \mathbf{X}\mathbf{y}_i \quad (10)$$

Since both \mathbf{D} and $\widehat{\mathbf{D}}$ are dependent on \mathbf{V} , they are unknown parameters that we need to estimate. Accordingly, we develop

an iterative algorithm to solve the formulated problem, which is described in Algorithm 1.

As follows, we provide a theoretical analysis of Algorithm 1. We first show a lemma cited from Nie et al. [55]:

Lemma 1: For any two given vectors \mathbf{v} and $\tilde{\mathbf{v}}$, the following inequality relation holds:

$$\|\tilde{\mathbf{v}}\|_2 - \frac{\|\tilde{\mathbf{v}}\|_2^2}{2\|\mathbf{v}\|_2} \leq \|\mathbf{v}\|_2 - \frac{\|\mathbf{v}\|_2^2}{2\|\mathbf{v}\|_2}$$

From Lemma 1, we can derive the following corollary:

Corollary 1: For given matrices \mathbf{M} and $\tilde{\mathbf{M}}$, the following inequality relation holds:

$$\|\tilde{\mathbf{M}}\|_F - \frac{\|\tilde{\mathbf{M}}\|_F^2}{2\|\mathbf{M}\|_F} \leq \|\mathbf{M}\|_F - \frac{\|\mathbf{M}\|_F^2}{2\|\mathbf{M}\|_F}$$

Then, we prove our algorithm's convergence as follows:

Theorem 1: Algorithm 1 iteratively decreases the value of the objective function in Eq. (4) and converges to the global optimal solution.

Proof: According to Step 4 of Algorithm 1, we obtain:

$$\mathbf{V}(t+1) = \min_{\mathbf{V}} \|\mathbf{V}^\top \mathbf{X} - \mathbf{Z}\|_F^2 \quad (11)$$

$$+\lambda_1 \text{Tr} \mathbf{V}^\top \mathbf{D}(t+1) \mathbf{V} + \lambda_2 \text{Tr} \mathbf{V}^\top \hat{\mathbf{D}}(t+1) \mathbf{V}$$

also from Step 5 of Algorithm 1, we obtain:

$$\mathbf{U}(t+1) = \min_{\mathbf{U}} \|(\mathbf{V}\mathbf{U})^\top \mathbf{X} - \mathbf{Y}\|_F^2 \quad (12)$$

$$+\lambda_2 \text{Tr} \mathbf{U}^\top \mathbf{V}^\top(t+1) \mathbf{V}(t+1) \mathbf{U}$$

Then, we can derive that

$$\begin{aligned} & \mathcal{J}(t+1) + \mathcal{F}(t+1) \\ & + \lambda_1 \text{Tr} \mathbf{V}^\top(t+1) \mathbf{D}(t+1) \mathbf{V}(t+1) \\ & + \lambda_2 \text{Tr} \mathbf{U}^\top(t+1) \hat{\mathbf{D}}(t+1) \mathbf{U}(t+1) \\ \leq & \mathcal{J}(t) + \mathcal{F}(t) + \lambda_1 \text{Tr} \mathbf{V}^\top(t) \mathbf{D}(t+1) \mathbf{V}(t) \\ & + \lambda_2 \text{Tr} \mathbf{U}^\top(t) \hat{\mathbf{D}}(t+1) \mathbf{U}(t), \end{aligned} \quad (13)$$

where $\mathcal{J}(t)$ is given as $\|\mathbf{V}^\top \mathbf{X}(t) - \mathbf{Z}\|_F^2$ and $\mathcal{F}(t)$ as $\|(\mathbf{V}\mathbf{U})^\top \mathbf{X}(t) - \mathbf{Y}\|_F^2 + \lambda_2 \text{Tr} \mathbf{U}^\top \mathbf{V}^\top(t) \mathbf{V}(t) \mathbf{U}$.

After substituting the definition of \mathbf{D} and $\hat{\mathbf{D}}$, we obtain

$$\begin{aligned} & \mathcal{J}(t+1) + \mathcal{F}(t+1) + \lambda_1 \sum_{i=1}^c \sum_{j=1}^m \frac{\|\mathbf{v}_i^j(t+1)\|_2^2}{2\|\mathbf{v}_i^j(t)\|_2} \\ & + \lambda_2 \frac{\|\mathbf{V}(t+1) \mathbf{U}(t+1)\|_F^2}{2\|\mathbf{V}(t) \mathbf{U}(t)\|_F} \\ \leq & \mathcal{J}(t) + \mathcal{F}(t) + \lambda_1 \sum_{i=1}^c \sum_{j=1}^m \frac{\|\mathbf{v}_i^j(t)\|_2^2}{2\|\mathbf{v}_i^j(t)\|_2} \\ & + \lambda_2 \frac{\|\mathbf{V}(t) \mathbf{U}(t)\|_F^2}{2\|\mathbf{V}(t) \mathbf{U}(t)\|_F} \end{aligned} \quad (14)$$

From Lemma 1, we can derive:

$$\begin{aligned} & \sum_{i=1}^c \sum_{j=1}^m \|\mathbf{v}_i^j(t+1)\|_2 - \sum_{i=1}^c \sum_{j=1}^m \frac{\|\mathbf{v}_i^j(t+1)\|_2^2}{2\|\mathbf{v}_i^j(t)\|_2} \\ \leq & \sum_{i=1}^c \sum_{j=1}^m \|\mathbf{v}_i^j(t)\|_2 - \sum_{i=1}^c \sum_{j=1}^m \frac{\|\mathbf{v}_i^j(t)\|_2^2}{2\|\mathbf{v}_i^j(t)\|_2}, \end{aligned} \quad (15)$$

and from Corollary 1, we can derive:

$$\begin{aligned} & \|\mathbf{V}(t+1) \mathbf{U}(t+1)\|_F - \frac{\|\mathbf{V}(t+1) \mathbf{U}(t+1)\|_F^2}{2\|\mathbf{V}(t) \mathbf{U}(t)\|_F} \\ \leq & \|\mathbf{V}(t) \mathbf{U}(t)\|_F - \frac{\|\mathbf{V}(t) \mathbf{U}(t)\|_F^2}{2\|\mathbf{V}(t) \mathbf{U}(t)\|_F}. \end{aligned} \quad (16)$$

Adding Eq. (13)-(16) on both sides, we obtain:

$$\begin{aligned} & \mathcal{J}(t+1) + \mathcal{F}(t+1) + \lambda_1 \sum_{i=1}^c \sum_{j=1}^m \|\mathbf{v}_i^j(t+1)\|_2 \\ & + \lambda_2 \|\mathbf{V}(t+1) \mathbf{U}(t+1)\|_A \\ \leq & \mathcal{J}(t) + \mathcal{F}(t) + \lambda_1 \sum_{i=1}^c \sum_{j=1}^m \|\mathbf{v}_i^j(t)\|_2 + \lambda_2 \|\mathbf{V}(t) \mathbf{U}(t)\|_A \end{aligned} \quad (17)$$

Eq. (17) decreases the value of the objective function in each iteration. Because our objective function is convex, Algorithm 1 converges to the global optimal solution of the formulated regularized optimization problem in Eq. (4). ■

Complexity analysis. Since the objective function in Eq. (4) is convex, Algorithm 1 converges very fast. In each iteration of Algorithm 1, it is trivial to compute Steps 4 and 5. We can compute Steps 6 and 7 by solving a system of linear equations, with a quadratic complexity.

V. EXPERIMENTS

This section discusses the experiment setup, present implementation details, and provides an analysis of our experimental results and comparison with other methods.

A. Experiment Setup and Dataset

In the experiments, a Clearpath Jackal mobile robot is used to navigate over a variety of unstructured terrains. In order to collect heterogeneous multi-sensory data that can be utilized to improve terrain representations, the mobile robot is equipped with multiple diverse proprioceptive and exteroceptive sensors. A structured-light camera is applied as an exteroceptive sensor to capture 3D colored point clouds of unstructured terrains in front of the robot with a frame rate of 30 Hz. Proprioceptive sensors are also installed on the robot to measure its internal information, including wheel odometry, motor speed, inertial measurement unit (IMU) measurement, power consumption, and battery status. The data frame rate from the proprioceptive sensors varied from 5–100 Hz. Linear interpolation based on data time-stamps is performed to ensure data frame rate to be consistent at 30 Hz. All of these data sources are used together as the input to terrain adaptation approaches for evaluation and comparison.

In addition, expert controls of the mobile robot are recorded as demonstrations of robot behaviors with the goal to navigate the mobile robot over unstructured terrains in straight lines as quickly as possible while maintaining safety (e.g., no flipping or crushing). The recorded demonstrations are utilized as the input to apprenticeship learning-based methods for training.

In the experiments, five types of unstructured terrains with various characteristics are used (shown in Figure 3), including:

- **Concrete:** The concrete surface exhibits smooth and hard characteristics. The mobile robot navigates over concrete floors in an indoor corridor.
- **Grass:** The grass terrain is less smooth compared to the concrete terrain, and is characterized by low ruggedness. The grass terrain often causes small wheel slips when the mobile robot navigates in places with a greater slope.
- **Mud:** This type of unstructured terrains has special characteristics in terms of roughness. Muddy ground surface is wet and typically soft, and can produce a lot of slips in robot motion. Mud terrains may also consist of weathered leaves and wood chips.
- **Pebble:** The pebble terrain is rough and characterized by small rocks, which we refer to as pebbles. The roughness requires a robot to navigate slowly and also results in relatively high power consumption. During demonstration, the robot is operated by human experts in a way to avoid sudden jerks, which is desirable in scenarios when a robot operates with fragile equipment.
- **Rock:** The rock terrain consists of larger rocks. This is the most challenging terrain type used for robot navigation in our experiments. The experts need to carefully control the robot and make detours to ensure safety of the mobile robot.



Figure 3. Instances of the five terrain types used in the experiments. The color images are observed by the camera installed in front of the mobile robot when navigating over these unstructured terrains.

Experiments of robot navigation are performed twenty times on each of the five unstructured terrain types. Each experiment lasts for a period of 10-25 seconds of an expert operating the mobile robot to navigate over a particular terrain type. In total, the experiments consist of 4000 data instances. All instances are organized into a dataset, which contains the heterogeneous multi-sensory data obtained by the robot during navigation on unstructured terrains, as well as the associated demonstrations.

B. Implementation

From the multi-sensory data obtained by robot sensors, commonly adopted terrain features are extracted. For exteroceptive data from depth and color images, real-time feature extraction techniques, including Histogram of Oriented Gradients (HOG) [56], color descriptors, and Local Binary Patterns (LBP) [57], are used to compute visual features of the terrains from both color and depth images. The extracted HOG, Color and LBP

features capture different visual characteristics of the terrains. Additionally, the point cloud data collected by the color-depth sensor is used to compute Elevation Maps based on grids [58] to discretely represent the grid-wise elevation and variance of the unstructured terrain. 3D point clouds with transformation and wheel odometry data are utilized to form the grid-based local elevation map. Normal vector maps and variance maps are also calculated based upon the generated elevation map of unstructured terrains. The proprioceptive sensor data including IMU and wheel odometry, are used together with the visual and elevation features to form a feature vector that is used as the input to our TRAL approach.

We implement four versions of the TRAL approach. First, we set λ_1 and λ_2 in Eq. (4) to 0, which performs joint terrain classification and robot behavior generation but cannot identify discriminative features. Then, we set $\lambda_1 = 0$ and $\lambda_2 = 0.1$ to identify discriminative features only useful for terrain classification. In addition, we assign $\lambda_1 = 0.1$ and $\lambda_2 = 0$ to identify discriminative features only useful for behavior generation. Finally, the full version of the proposed TRAL approach is implemented with $\lambda_1 = \lambda_2 = 0.1$, which identifies features that contribute to the joint terrain classification and behavior generation. We observe that TRAL shows stable performance when $\lambda_1, \lambda_2 \in (0.05, 5)$. The TRAL approach is implemented in Linux on the Jackal robot’s onboard computer with an i5 2.5 GHz CPU and 8G memory.

We also compare the TRAL approach with previous apprenticeship learning methods used for robot adaptation, including Learning from Demonstration (LfD) based upon probabilistic graphical models [59], Sequence-based Multi-modal Apprenticeship Learning (SMAL) [14] that implements robot adaptation based upon optimization-based state perception and MDP-based robot decision making, and an imitation learning method using optimization-based multiclass classification [54] with no representation learning (NoChar).

C. Result on Familiar Terrains

In this set of experiments, we evaluate the proposed TRAL approach’s performance when the mobile robot navigates over familiar unstructured terrains, in which the terrain types have been experienced by the robot during the training period. The robot is trained using terrain features and expert demonstrations collected in all five types of unstructured terrains used, and then tested using terrain features only from the five types of terrains. In order to evaluate how well human expertise is transferred to a mobile robot in terms of ground speed, which is viewed as one of the most critical performance metrics that indicates how fast a navigation task can be completed by the robot [23, 60], we use the Root Mean Square Error (RMSE) as the metric to measure the magnitude of deviation of navigation behaviors generated by apprenticeship learning methods from the expert controls. Although RMSE provides the magnitude of deviation, it cannot present the ratio of deviation, comparing with the expert control. Therefore, we also use error ratio (ER), defined as the average of the ratio of the deviation magnitude over the control magnitude, in order to measure the percentage

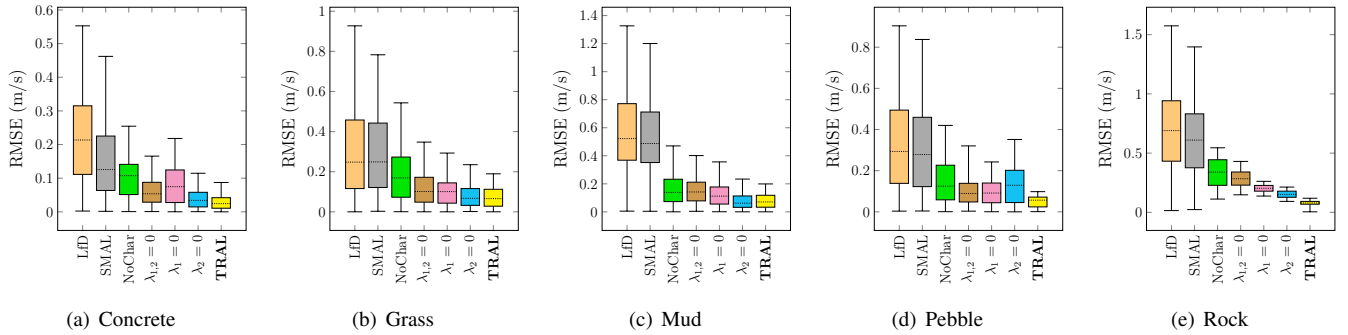


Figure 4. Quantitative results of RMSE on ground velocity obtained by TRAL in the scenarios of navigation over *familiar* unstructured terrains. Comparison with previous apprenticeship learning methods and our baseline methods is also presented. The results of each method over each type of terrains are presented as a boxplot that depicts the distribution of the results from all testing instances based on a five number summary (i.e., minimum, first quartile, median, third quartile, and maximum).

Table I
MEAN AND VARIANCE OF RMSE AND ER ON GROUND VELOCITY OBTAINED BY TRAL AND PREVIOUS/BASELINE METHODS OVER FAMILIAR UNSTRUCTURED TERRAINS.

| Approach | RMSE (m/s) | ER (%) |
|---------------------|-------------------------------------|-------------------------------------|
| LfD [59] | 0.39 (± 0.20) | 21.35 (± 10.16) |
| SMAL [14] | 0.34 (± 0.18) | 12.5 (± 10.10) |
| NoChar [54] | 0.17 (± 0.08) | 10.75 (± 4.12) |
| $\lambda_{1,2} = 0$ | 0.13 (± 0.05) | 7.48 (± 4.96) |
| $\lambda_1 = 0$ | 0.11 (± 0.04) | 5.33 (± 3.40) |
| $\lambda_2 = 0$ | 0.08 (± 0.04) | 3.40 (± 2.41) |
| TRAL | 0.05 (± 0.03) | 2.41 (± 1.76) |

of deviation from the expert controls. For both RMSE and ER, a smaller value indicates a better performance.

The quantitative result obtained by our proposed approach is demonstrated in Figure 4 in the scenario of navigation over each of the familiar unstructured terrains, using a boxplot that displays the distribution of the results from all testing instances based on a five number summary (i.e., minimum, first quartile, median, third quartile, and maximum). The overall mean and variance of RMSE and ER on ground velocity obtained by our TRAL approach are summarized in Table II. TRAL achieves the RMSE of 0.05 ± 0.03 m/s, which shows the magnitude of deviation is 5 cm/s from expert control on average. The error rate ER demonstrates that the difference between the generated control and expert control is 2.41% with a variance of 1.67% for robots navigating over familiar terrains.

Quantitative comparisons with several previous approaches and baseline methods are demonstrated in Figure 4 using the boxplot graph and comparisons of the overall performance are summarized in Table I. It is observed that the previous apprenticeship learning methods tested in our experiment, including LfD, SMAL, and NoChar obtain a larger deviation magnitude and percentage. A likely reason that explains this phenomenon is that they cannot perform representation learning to identify discriminative terrain features. The raw feature vectors directly computed from the sensor measurements are noisy, which can decrease the performance of navigational behavior generation. Since our approach combines two loss functions to perform a joint optimization for both representation and apprenticeship

learning, even without applying the two regularization norms, the version of our approach with $\lambda_{1,2} = 0$ is able to build good representations to improve the performance of behavior generation. When regularization terms are used to promote modality sparsity (i.e., with non-zero λ_1 or λ_2), these versions of our approach can find more discriminative features toward terrain classification or behavior generation, which can further reduce the deviation. When both regularization norms are used to learn terrain representations toward both terrain classification and navigational behavior generation, the complete TRAL approach obtains superior performance on both RMSE and ER over the previous methods.

D. Result on Unfamiliar Unstructured Terrains

In this set of experiments, we evaluate TRAL’s performance when the mobile robot navigates over unfamiliar unstructured terrains, in which the terrain type that the robot navigates over has not been experienced during training. The robot is trained using terrain features and expert demonstrations collected in four categories of the unstructured terrains, and then evaluated using the fifth terrain type that has not been used in the training period. For example, Figure 5(c) demonstrates the results from the scenario that the robot is trained using data from concrete, grass, pebble and rock terrain types, and evaluated on the mud terrain what is not experienced by the robot in training.

The RMSE results obtained by TRAL and comparisons with previous and baseline techniques on each type of terrains that are not familiar to the robot are demonstrated in Figure 5. The overall performance of TRAL and other methods are presented in Table II. When the robot navigates over unfamiliar terrains, we observe phenomena that are similar to the results obtained when it navigates on familiar terrains: apprenticeship learning approaches without using representation learning do not work well; integrating the representation and apprenticeship learning decreases the errors; and the proposed regularization terms can further increase the performance. The complete TRAL method still obtains the best performance over other methods, mostly because of its capability of identifying the most discriminative terrain features that can generalize well to unfamiliar scenarios. On the other hand, since unfamiliar terrains are not seen in

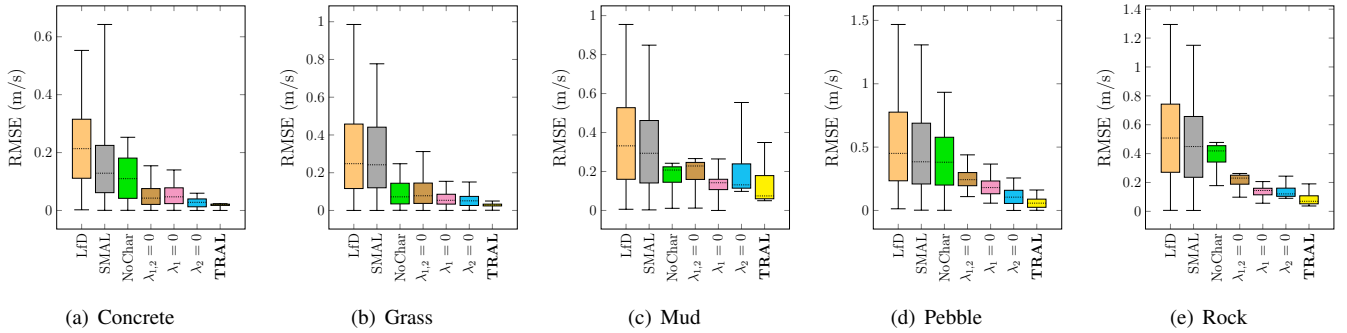


Figure 5. Quantitative results of RMSE on ground velocity obtained by TRAL in the scenarios of navigation over *unfamiliar* unstructured terrains. Comparison with previous apprenticeship learning methods and our baseline methods is also presented.

Table II
MEAN AND VARIANCE OF RMSE AND ER ON GROUND VELOCITY OBTAINED BY TRAL AND PREVIOUS/BASELINE METHODS OVER UNFAMILIAR UNSTRUCTURED TERRAINS.

| Approach | RMSE (m/s) | ER (%) |
|---------------------|---------------------|-----------------------|
| LfD [59] | 0.35 (± 0.21) | 33.15 (± 19.59) |
| SMAL [14] | 0.29 (± 0.19) | 29.35 (± 16.80) |
| NoChar [54] | 0.23 (± 0.07) | 20.73 (± 6.26) |
| $\lambda_{1,2} = 0$ | 0.16 (± 0.05) | 22.82 (± 6.97) |
| $\lambda_1 = 0$ | 0.11 (± 0.03) | 14.22 (± 7.52) |
| $\lambda_2 = 0$ | 0.08 (± 0.04) | 13.14 (± 4.56) |
| TRAL | 0.06 (± 0.03) | 7.45 (± 3.89) |

the training period, all approaches evaluated in this case exhibit higher errors (i.e., higher RMSE and ER values) comparing to the results from the situations when robots navigate on familiar unstructured terrains.

E. Discussion

The TRAL approach offers the desirable ability to automatically estimate the importance of feature modalities for terrain classification and behavior generation. The results of modality importance averaged from both familiar and unfamiliar experimental scenarios are demonstrated in Figure 6. It is observed that two feature modalities, i.e., IMU and terrain variance, are most discriminative in determining both terrain types and robot navigational behaviors. Intuitively, both feature modalities can be considered to contain complementary information because the IMU data is greatly affected by and correlated with terrain variances. Moreover, it is interesting to note that HOG features are discriminative for classifying the terrain types, but not for learning robot navigation behaviors. LBP features on the other hand exhibit the opposite impact. When comparing all visual features, we observe that features extracted from color images are generally more discriminative than the features from depth images. Due to TRAL’s ability to integrate representation and apprenticeship learning in the unified framework, and due to the efficiency of our convex optimization formulation, TRAL is able to achieve high-speed processing on the robot’s onboard computer. Including the time spent on feature extraction, our approach obtains a processing rate of 20 Hz. Without feature extraction, our approach itself achieves a processing speed of 80 Hz on average. This result indicates the runtime advantage

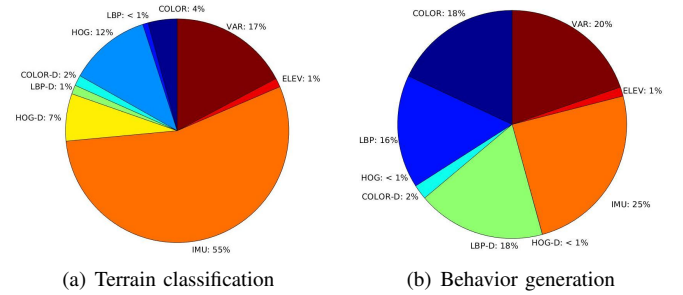


Figure 6. Normalized importance of terrain feature modalities with respect to terrain classification and behavior generation.

of our TRAL approach for real-time robotics applications and its potential to facilitate high-speed mobile robot navigation.

VI. CONCLUSION

In this paper, we present the TRAL approach to address the research problem of robot adaptation to unstructured terrains. As a theoretical novelty, TRAL formulates representation and apprenticeship learning in the unified regularized optimization framework to perform joint robot learning for terrain classification and behavior generation, which also automatically estimates the importance of terrain feature modalities. Our second theoretical novelty is the design of an optimization algorithm to solve the formulated problem, which possesses a theoretical convergence guarantee. The proposed approach is extensively evaluated using two real-world scenarios in the experiments, in which a mobile robot navigates over familiar or unfamiliar unstructured terrains. Experimental results have demonstrated that TRAL is able to transfer human expertise to mobile robots with small errors, obtain superior performance compared with previous and baseline methods, and offer intuitive insights on the importance of terrain feature modalities.

ACKNOWLEDGMENTS

This research was supported by NSF awards IIS-1849348, IIS-1849359, CNS-1823245, DOE grant DE-FE0031650, and ARO grant W911NF-17-1-0447. The authors also would like to thank Christopher Reardon for valuable discussions and the anonymous reviewers for their constructive comments.

REFERENCES

- [1] Shinji Kawatsuma, Mineo Fukushima, and Takashi Okada. Emergency response by robots to Fukushima-Daiichi accident: summary and lessons learned. *Industrial Robot: An International Journal*, 39(5):428–435, 2012.
- [2] Kristopher Toussaint, Nicolas Pouliot, and Serge Montambault. Transmission line maintenance robots capable of crossing obstacles: State-of-the-art review and challenges ahead. *JFR*, 26(5):477–499, 2009.
- [3] David Silver, J Andrew Bagnell, and Anthony Stentz. Learning from demonstration for autonomous navigation in complex unstructured terrain. *IJRR*, 29(12):1565–1592, 2010.
- [4] Ayse Naz Erkan, Raia Hadsell, Pierre Sermanet, Jan Ben, Urs Muller, and Yann LeCun. Adaptive long range vision in unstructured terrain. In *IROS*, 2007.
- [5] Christopher A Brooks and Karl Iagnemma. Vibration-based terrain classification for planetary exploration rovers. *TRO*, 21(6):1185–1191, 2005.
- [6] Ayanna Howard and Homayoun Seraji. Vision-based terrain characterization and traversability assessment. *Journal of Robotic Systems*, 18(10):577–587, 2001.
- [7] Sang-Ho Hyon. Compliant terrain adaptation for biped humanoids without measuring ground surface and contact forces. *TRO*, 25(1):171–178, 2009.
- [8] Yoon-Gu Kim, Jeong-Hwan Kwak, Dae-Han Hong, In-Huck Kim, Dong-Hwan Shin, and Jinung An. Autonomous terrain adaptation and user-friendly teleoperation of wheel-track hybrid mobile robot. *International Journal of Precision Engineering and Manufacturing*, 13(10):1781–1788, 2012.
- [9] Michael Shneier, Tommy Chang, Tsai Hong, Will Shackelford, Roger Bostelman, and James S Albus. Learning traversability models for autonomous mobile vehicles. *AuRo*, 24(1):69–86, 2008.
- [10] Panagiotis Papadakis. Terrain traversability analysis methods for unmanned ground vehicles: A survey. *Engineering Applications of Artificial Intelligence*, 26(4):1373–1385, 2013.
- [11] Maggie Wigness, John G Rogers, and Luis E Navarro-Serment. Robot navigation from human demonstration: learning control behaviors. In *ICRA*, 2018.
- [12] Fernando L Garcia Bermudez, Ryan C Julian, Duncan W Haldane, Pieter Abbeel, and Ronald S Fearing. Performance analysis and terrain classification for a legged robot over rough terrain. In *IROS*, 2012.
- [13] Dylan Williamson, Navinda Kottege, and Peyman Moghadam. Terrain characterisation and gait adaptation by a hexapod robot. In *Australasian Conference on Robotics and Automation (ACRA)*, 2016.
- [14] Fei Han, Xue Yang, Yu Zhang, and Hao Zhang. Sequence-based multimodal apprenticeship learning for robot perception and decision making. In *ICRA*, 2017.
- [15] Christian Weiss, Holger Frohlich, and Andreas Zell. Vibration-based terrain classification using support vector machines. In *IROS*, 2006.
- [16] Michael Happold, Mark Ollis, and Nikolas Johnson. Enhancing supervised terrain classification with predictive unsupervised learning. In *RSS*, 2006.
- [17] Thomas M Howard and Alonzo Kelly. Optimal rough terrain trajectory generation for wheeled mobile robots. *IJRR*, 26(2):141–166, 2007.
- [18] Anelia Angelova, Larry Matthies, Daniel Helmick, Gabe Sibley, and Pietro Perona. Learning to predict slip for ground robots. In *ICRA*, 2006.
- [19] Sebastian Thrun, Michael Montemerlo, and Andrei Aron. Probabilistic terrain analysis for high-speed desert driving. In *RSS*, 2006.
- [20] Paul Filitchkin and Katie Byl. Feature-based terrain classification for littledog. In *IROS*, 2012.
- [21] Stefan Laible, Yasir Niaz Khan, and Andreas Zell. Terrain classification with conditional random fields on fused 3D LIDAR and camera data. In *European Conference on Mobile Robots*, 2013.
- [22] Mieczyslaw Gregory Bekker. *Introduction to terrain-vehicle systems*. University of Michigan Press, 1969.
- [23] Chris Urmson, Charlie Ragusa, David Ray, Joshua Anhalt, Daniel Bartz, Tugrul Galatali, Alexander Gutierrez, Josh Johnston, Sam Harbaugh, Hiroki “Yu” Kato, et al. A robust approach to high-speed navigation for unrehearsed desert terrain. *JFR*, 23(8):467–508, 2006.
- [24] Cheng Wang and Nancy F Glenn. Integrating LiDAR intensity and elevation data for terrain characterization in a forested area. *IEEE Geoscience and Remote Sensing Letters*, 6(3):463–466, 2009.
- [25] Karl Iagnemma, Shinwoo Kang, Hassan Shibly, and Steven Dubowsky. Online terrain parameter estimation for wheeled mobile robots with application to planetary rovers. *TRO*, 20(5):921–927, 2004.
- [26] Lauro Ojeda, Johann Borenstein, Gary Witus, and Robert Karlsen. Terrain characterization and classification with a mobile robot. *JFR*, 23(2):103–122, 2006.
- [27] Eric Trautmann and Laura Ray. Mobility characterization for autonomous mobile robots using machine learning. *AuRo*, 30(4):369–383, 2011.
- [28] Roberto Manduchi, Andres Castano, Ashit Talukder, and Larry Matthies. Obstacle detection and terrain classification for autonomous off-road navigation. *AuRo*, 18(1):81–102, 2005.
- [29] Lynne E Parker. Task-oriented multi-robot learning in behavior-based systems. In *IROS*, 1996.
- [30] Lynne E Parker. Case study for life-long learning and adaptation in cooperative robot teams. In *Sensor Fusion and Decentralized Control in Robotic Systems*, volume 3839, pages 92–102, 1999.
- [31] Lynne E Parker. Lifelong adaptation in heterogeneous multi-robot teams: Response to continual variation in individual robot performance. *AuRo*, 8(3):239–267, 2000.
- [32] Ian Watson and Farhi Marir. Case-based reasoning: A review. *Knowledge Engineering Review*, 9(4):327–354,

- 1994.
- [33] Michael W Floyd, Michael Drinkwater, and David W Aha. Trust-guided behavior adaptation using case-based reasoning. Technical report, 2015.
- [34] Qiqian Zhang, Hui Qian, and Miaoliang Zhu. Parameter adaptation by case-based mission-planning of outdoor autonomous mobile robot. In *Intelligent Vehicles Symposium*, 2005.
- [35] Jonathan Ko and Dieter Fox. GP-BayesFilters: Bayesian filtering using Gaussian process prediction and observation models. *AuRo*, 27(1):75–90, 2009.
- [36] Marc Deisenroth and Carl E Rasmussen. PILCO: A model-based and data-efficient approach to policy search. In *ICML*, pages 465–472, 2011.
- [37] Ian Lenz, Ross A Knepper, and Ashutosh Saxena. DeepMPC: Learning deep latent features for model predictive control. In *RSS*, 2015.
- [38] Christian Plagemann, Sebastian Mischke, Sam Prentice, Kristian Kersting, Nicholas Roy, and Wolfram Burgard. Learning predictive terrain models for legged robot locomotion. In *IROS*, 2008.
- [39] Sergey Levine and Pieter Abbeel. Learning neural network policies with guided policy search under unknown dynamics. In *NIPS*, 2014.
- [40] Austin D Buchan, Duncan W Haldane, and Ronald S Fearing. Automatic identification of dynamic piecewise affine models for a running robot. In *IROS*, 2013.
- [41] Alexander Kleiner, Markus Dietl, and Bernhard Nebel. Towards a life-long learning soccer agent. In *Robot Soccer World Cup*, pages 126–134, 2002.
- [42] Sebastian Thrun and Tom M Mitchell. Lifelong robot learning. In *The Biology and Technology of Intelligent Autonomous Agents*, pages 165–196. 1995.
- [43] Gentaro Taga. A model of the neuro-musculo-skeletal system for human locomotion. *Biological Cybernetics*, 73(2):97–111, 1995.
- [44] Hiroshi Kimura, Yasuhiro Fukuoka, and Hiroyuki Nakamura. Biologically inspired adaptive dynamic walking of the quadruped on irregular terrain. In *Robotics Research*, pages 329–336. 2000.
- [45] Hiroshi Kimura, Seiichi Akiyama, and Kazuaki Sakurama. Realization of dynamic walking and running of the quadruped using neural oscillator. *AuRo*, 7(3):247–258, 1999.
- [46] Andrew Lookingbill, John Rogers, David Lieb, J Curry, and Sebastian Thrun. Reverse optical flow for self-supervised adaptive autonomous robot navigation. *IJCV*, 74(3):287–302, 2007.
- [47] Samir Nabulsi, M Armada, and Héctor Montes. Multiple terrain adaptation approach using ultrasonic sensors for legged robots. In *Climbing and Walking Robots*. 2006.
- [48] Brenna D Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. A survey of robot learning from demonstration. *Robotics and Autonomous Systems*, 57(5):469–483, 2009.
- [49] Dieter Fox, Wolfram Burgard, and Sebastian Thrun. The dynamic window approach to collision avoidance. *RAM*, 4(1):23–33, 1997.
- [50] Zheng Zhang, Yong Xu, Jian Yang, Xuelong Li, and David Zhang. A survey of sparse representation: algorithms and applications. *IEEE Access*, 3:490–530, 2015.
- [51] Hao Zhang, Fei Han, and Hua Wang. Robust multimodal sequence-based loop closure detection via structured sparsity. In *RSS*, 2016.
- [52] Fei Han, Xue Yang, Yiming Deng, Mark Rentschler, Dejun Yang, and Hao Zhang. SRAL: Shared representative appearance learning for long-term visual place recognition. *RA-L*, 2(2):1172–1179, 2017.
- [53] Sylvain Calinon and Aude Billard. Incremental learning of gestures by imitation in a humanoid robot. In *HRI*, 2007.
- [54] Nathan Ratliff, J Andrew Bagnell, and Siddhartha S Srinivasa. Imitation learning for locomotion and manipulation. In *Humanoids*, 2007.
- [55] Feiping Nie, Heng Huang, Xiao Cai, and Chris H Ding. Efficient and robust feature selection via joint ℓ_2 , ℓ_1 -norms minimization. In *NIPS*, 2010.
- [56] Navneet Dalal and Bill Triggs. Histograms of oriented gradients for human detection. In *CVPR*, 2005.
- [57] Timo Ahonen, Abdenour Hadid, and Matti Pietikainen. Face description with local binary patterns: Application to face recognition. *PAMI*, (12):2037–2041, 2006.
- [58] Péter Fankhauser, Michael Bloesch, Christian Gehring, Marco Hutter, and Roland Siegwart. Robot-centric elevation mapping with uncertainty estimates. In *Mobile Service Robotics*, pages 433–440. 2014.
- [59] Fenglu Ge, Wayne Moore, and Michael Antolovich. Learning from demonstration using GMM, CHMM and DHMM: A comparison. In *Australasian Joint Conference on Artificial Intelligence*, 2015.
- [60] Shingo Shimoda, Yoji Kuroda, and Karl Iagnemma. Potential field navigation of high speed unmanned ground vehicles on uneven terrain. In *ICRA*, 2005.