Noname manuscript No.
(will be inserted by the editor)

Autonomous exploration with online learning of traversable yet visually rigid obstacles

Miloš Prágr · Jan Bayer · Jan Faigl

Received: date / Accepted: date

Abstract This paper concerns online learning of terrain properties combining haptic perception with exteroceptive sensing to reason about forces needed to pass through terrains that visually appear as untraversable obstacles. Terrain learning is studied within the context of autonomous exploration. We propose predicting the traversability of potentially obstructing terrains by active perception to establish a connection between the observed geometric environment model and deliberately sampled forces to pass through the terrain using a haptic sensor that probes the terrain in front of the robot. The developed solution uses a Gaussian Process regressor in online learning and force prediction. The robot is navigated following the information gain to improve traversability and spatial models. The proposed approach has been experimentally verified in fully autonomous exploration with a multi-legged walking robot. The robot is navigated through visually looking obstacles and explores "hidden" areas while following the expected information gain to explore the terrain properties of the mission area.

Keywords Mobile robot · Exploration · Active perception · Haptic sensing · Gaussian Process regression

The presented work has been supported by the Czech Science Foundation (GAČR) under research project No. 18-18858S and project No. 19-20238S.

Miloš Prágr · Jan Bayer · Jan Faigl
Department of Computer Science, Faculty of Electrical Engineering, Czech Technical University in Prague,
E-mail: {pragrmi1,bayerja1,faiglj}@fel.cvut.cz

1 Introduction

Terrain properties like appearance and geometry can be used to reason about the traversability of mobile robots by assigning terrain classes (Bradley et al., 2015), computing a terrain traversal cost function (Sofman et al., 2006), or discriminating untraversable terrains (Stelzer et al., 2012). Further, we reason about terrains that appear untraversable due to their geometrical properties. These visually appearing obstacles can be traversable, such as sparse vegetation or a curtain-covered doorway, which appear as a wall when presumed to be rigid. Assuming a non-rigid terrain is a rigid obstacle to avoid leads to safe behavior. However, in autonomous exploration, such behavior might be overly cautious and results in an incomplete terrain model if the robot is prevented from visiting areas separated by terrain that appears untraversable yet can be traversed by the robot.

In this paper, we propose to use both exteroceptive and haptic sensing to actively learn to predict the traversability of potentially obstructing obstacles. A haptic sensor (3D-printed bumper) is used to sample the force needed to pass through these visually appearing obstacles. A Gaussian Process (GP) regressor (Rasmussen and Williams, 2006) exploits the obstacle appearance in online learning and predicts the forces required to be generated by the robot to pass through an obstacle. The proposed method is demonstrated in a real-world autonomous exploration scenario where a multi-legged walking robot actively learns terrains that can be passed through. Besides exploring the unobserved areas, the robot actively collects information about the force prediction model driven by the expected information gain from interaction with potential obstacles. An example of the robot decision-making using the proposed model is visualized in Fig. 1.

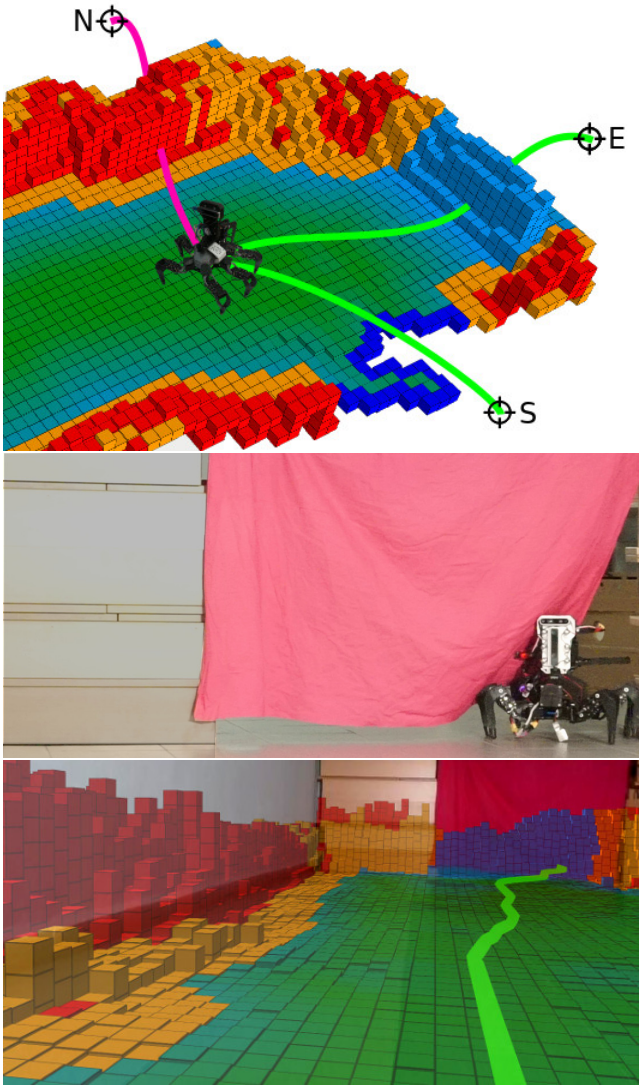

Fig. 1 A visualization of autonomous decision-making in the proposed terrain learning approach; untraversable terrains are in the red and orange, while the green and blue areas can be traversed. For the "blue obstacles", the robot has already learned that such visually appearing obstacles can be traversed through. (top) The circular black targets represent possible exploration goal locations. The southern (S) goal is reachable over terrain that appears traversable, while both the northern (N) and eastern (E) goals are located behind obstacles. However, the eastern goal is considered reachable, visualized by the green path from the current robot's location to the possible goal location. (middle) Robot walking through a traversable obstacle in the form of purple fabric. (bottom) Projection of the traversability on robot vision a few moments before traversing the fabric. Since the robot has learned that the purple fabric (shown at the northeast of the view) is traversable, the robot plans a path through it.

Regarding the existing work, including our previous work on terrain learning (Prágr et al., 2018b, 2019), the main contributions of the presented work are considered as follows.

- Model characterizing the force needed to pass through terrains (obstacles), incrementally learnable using exteroceptive and haptic measurements.
- Proposed model's capability to exploit observations about terrain traversability that visually appears to be rigid, but for which the experienced haptic interaction provides evidence of its traversability.
- A robotic system with active haptic perception using information-theoretic estimation of the expected information gain of the object touching to ensure traversability of terrains that visually appear impassable.
- Experimental evaluation of the developed incremental model learning and online prediction in an autonomous robotic exploration scenario with a real hexapod walking robot. The model's incremental learning capacity is exploited in online decision-making based on the expected information gain from visiting and interacting with terrain obstacles.
- Experimental evaluation of the developed incremental model learning with real outdoor vegetation.

The rest of the paper is structured as follows. Section 2 provides a brief review of the most related approaches that concern terrain traversal by unmanned ground vehicles and mobile robot exploration. Section 3 specifies the problem of mobile robot exploration of the environment with terrains that appears untraversable yet can be traversed. The proposed learning and modeling of such environments and the autonomous exploration framework to build a complete traversability map of the unknown environment with areas "hidden" by non-rigid terrains visually appearing as obstacles are proposed in Section 4. The employed haptic sensor is overviewed in Section 5, together with the reported evaluation results from the real experimental deployment of the proposed system. Section 6 concludes the paper.

2 Related Work

In this section, a short survey of the related work is provided. First, we focus on approaches concerned with describing terrain traversability, and the herein presented work is placed in their context. Second, we provide a short overview of robotic exploration.

2.1 Terrain Traversability

It is not desirable to enter areas that would bar mobile robots from continuing their missions by damaging the robot, such as by falling off a high cliff or impeding further motion as it happened by getting stuck in soft sand (Brown and Webster, 2010). Besides, avoiding terrains that do not pose an immediate danger to the robot but are hard to traverse is desirable. Such terrains may cause unnecessary energy consumption or slowly wear the robot body. Thus, autonomous mobile robots have to consider local terrain properties during navigation through the operational environment. A rich body of literature reports on systems concerning terrain traversability. An extensive review and taxonomy of such approaches can be found in Papadakis (2013).

The terrain traversability can be described either by classifying the terrains into a set of terrain classes (Belter et al., 2019; Rothrock et al., 2016; Kragh et al., 2015; Giguere and Dudek, 2008), or by assigning a continuous traversability score to the observed terrain properties (Kottege et al., 2015), such as terrain appearance and geometry (Prágr et al., 2018a; Prágr and Faigl, 2019). A common yet straightforward approach is to classify terrain either as an untraversable obstacle or a free traversable space. Stelzer et al. (2012) use geometric terrain properties to classify terrain as untraversable or free and compute a continuous index to describe the traversability of the latter class. In Kragh et al. (2015), the Support Vector Machine (SVM) classifier learns three classes (ground, vegetation, and object) in an agricultural environment; while 40 hand-labeled terrain classes are used in Bradley et al. (2015), where some of them are denoted as obstacles.

Terrain traversability scores are computed directly from remotely observed terrain appearance and geometry or describe the difficulty or energy consumption previously experienced by the robot when traversing over the respective terrain. Sofman et al. (2006) use overhead imagery to learn traversability log-scale score based on ground LiDAR data. Overhead features are utilized to predict the energy required to traverse various outdoor terrains in Prágr et al. (2020). The cost of transport, an energy-over-velocity cost originating in biology (Tucker, 1975), is modified for use with battery-powered robots in Kottege et al. (2015). McGhee and Frank (1968) propose to measure the stability of a multi-legged robot in terms of its foothold support polygon. Furthermore, Prágr et al. (2019) learn to predict stability based on inertial measurements of the robot shaking in an active perception scenario.

Terrain geometry, which can serve as a traversability indicator, can be characterized in terms of its slope (Gu et al., 2008; Brunner et al., 2013), step height (Homberger et al., 2016), or roughness (Krüsi et al., 2016; Belter et al., 2019). The terrain shape is described based on the Eigen-statistics of the point cloud covariance matrix in Lalonde et al. (2006) and Kragh et al. (2015). Approaches that consider terrain color use the HSV (Sofman et al., 2006) or Lab (Otsu et al., 2016) color space to avoid illumination sensitivity of the RGB color space. Cunningham et al. (2019) propose to use thermal imagery to predict slip during Mars rover missions.

Autonomous robots operating in outdoor environments might encounter hard-to-traverse-vegetation, and thus attention is given to such terrains (Sofman et al., 2006; Bradley et al., 2015). Ünsalan and Boyer (2004) compared indices that characterize vegetation using a LiDAR sensor. In Petrou et al. (2015), the vegetation height is classified using overhead imagery. The elevation of the supporting terrain occluded by the vegetation is estimated using a GP model of the vegetation height, and foothold supports in Homberger et al. (2019).

Furthermore, mobile robots may also be deployed in environments with dynamic obstacles such as closing doors or moving people. Approaches to handle such dynamic environments may either filter out dynamic objects and extract a static map (Burgard et al., 1999), or use spatial-temporal maps to represent and predict changes in the environment such as a door being closed or opened (Biber and Duckett, 2005; Halodová et al., 2019). However, to the best of the authors' knowledge, none of the existing dynamic environment approaches is designed to handle terrain that appears like a non-moving obstacle for the whole time while it can be passed through with sufficient forward force.

In Baleia et al. (2015), a haptic antenna is used to classify the traversability of visually untraversable yet possibly traversable objects. The therein proposed traversability predictor is used in a self-supervised manner. Upon encountering a potential obstacle during its mission execution, the robot recalls the k-nearest appearing obstacles using a feature similarity metric. The robot computes its confidence levels regarding the obstacle being traversable and untraversable and decides whether to move forward or avoid the obstacle. If neither the obstacle traversability nor non-traversability can be observed with sufficient confidence, the robot uses the antenna to assess the obstacle and expand its memory.

Kahn et al. (2021) present a self-supervised, end-to-end learning system to navigate potentially traversable terrains that appear untraversable without relying on a Simultaneous Localization and Mapping (SLAM) sys-

tem. The robot uses a random walk policy to collect a dataset, where it identifies collision, bumpiness, and position events using its Inertial Measurement Unit (IMU) and wheel odometry. The robot learns to predict the events given input image and action. The learned models can be exploited in navigation with respect to (w.r.t.) an arbitrary reward function that considers the three event types.

In the herein presented work, we aim to build both the observed geometric model of the environment and a model that predicts the traversability of potentially obstructing terrains by a small amount of data provided by a haptic sensor correlated with exteroceptive sensing. Compared to Baleia et al. (2015) and Kahn et al. (2021), we characterize the force needed to pass through potential obstacles, and hence the proposed approach is robot agnostic. Unlike the method proposed by Baleia et al. (2015), which focuses on classifying objects encountered during the robot mission that can be avoided when deemed untraversable, we address relatively large, obstacle-like terrains that might block access to additional sections of the environment.

Besides, considering the previous work on learning using collected data (Kahn et al., 2021), the proposed system is employed in autonomous exploration with the additional complexity of incremental learning of the force to pass through obstacles on the robot during the deployment. Based on the visual appearance, the force predictions are utilized in online decision-making to discriminate the objects the robot cannot traverse and identify unknown obstacles the robot should sample next. Since our approach is focused on learning the force, it samples and learns only on terrains that are unknown and thus informative. We use a learning method that requires tens of samples and can be used online, directly on the robot during the exploration. Thus, the proposed approach uses much less data than general approaches such as Kahn et al. (2021), which rely on long-term data collection and offline processing.

Because we employ the proposed approach in the exploration context, an overview of mobile robot exploration approaches is presented in the following paragraphs.

2.2 Mobile robot exploration

Mobile robot exploration is an active perception scenario where one or a group of mobile robots build a model of the mission environment. In frontier-based exploration (Faigl and Kulich, 2015), the robot follows frontiers, the boundaries between the observed traversable and not yet observed areas (Yamauchi, 1997). Alternatively, the probabilistic representation of the cell occupancy in the occupancy grids (Moravec and Elfes, 1985) can be used in the exploration strategy that maximizes the information gain (Bourgault et al., 2002; Makarenko et al., 2002). Beside grid maps, Gaussian Processes (GPs) (Vasudevan et al., 2009; Ruiz and Olariu, 2015), Gaussian Mixture Models (GMMs) (O'Meadhra et al., 2019), or Hilbert maps (Ramos and Ott, 2016) can be used to create continuous maps that are not resolution-dependant. Since the GP regressors provide predictive variance for their queries, they are particularly suited for active perception scenarios. Jadidi et al. (2018) use a GP-based representation to construct frontier maps, while the GMM is used in Tabib et al. (2019).

Exploration is not limited to building maps and geometric models but may also concern modeling a phenomenon underlying the spatial environment such as temperature (Luo and Sycara, 2018). In informative path planning (Singh et al., 2007), the goal is to find the most informative path subject to a particular constraint, such as the robot energy budget. Hence, the robot explores as much of the environment as possible while avoiding battery depletion that would lead to its immobilization, as noted in Tiwari et al. (2019), where a framework for operation range estimation is presented to support robots ranging from multi-rotor fliers to ground vehicles.

When the goal is to find extrema of the modeled phenomena, exploration-exploitation tradeoff-based approaches such as Gaussian Process Upper Confidence Bound (Srinivas et al., 2010) can be utilized. Furthermore, the active learning of the underlying model can be combined with the traditional geometric exploration (Prágr et al., 2019). For example, the robot localization model can be incorporated into information-based exploration approaches, such as the localization uncertainty represented using the differential entropy of the robot position distribution by directly adding it to the mapping uncertainty (Bourgault et al., 2002; Stachniss et al., 2005). However, since the differential entropy differs from the Shannon entropy of the binary cell occupancy distribution in scale, particularly when considering dynamic environment size, Carrillo et al. (2018) argue that it is not desirable to combine them directly and employ Rényi entropy (Rényi, 1961) to create an uncertainty utility function.

In this paper, we present a combination of the spatial map exploration with the active building of the obstacle traversability model characterized as the force to pass through. The proposed approach is demonstrated within autonomous robotic exploration, in an escape-like scenario, where the robot first explores all areas accessible without interacting with obstacles. Only when no such areas are available does it actively learn the

obstacle traversability. Consequently, the robot selects the exploration goals independently for the respective models. Hence, even though the spatial and prediction models yield information gains in Shannon's discrete and differential entropy, respectively, we circumvent the need to combine these two quantities.

3 Problem Specification

We address mobile robot exploration in environments where obstacles can be non-rigid and passable by the robot. The robot is tasked to explore an environment modeled as the grid map $\mathcal{M}_{2.5D}$, where each cell ν corresponds to a foothold of the hexapod walking robot used in the experimental verification. Hence, the cell size d_ν corresponds to the foothold size. The robot moves through the environment along a path ψ that can be expressed as

$$\begin{aligned} &\psi = (\nu_1, \nu_2, \cdots, \nu_n), \\ &\text{s.t.} \\ &\forall i \in 1, \cdots, n: \ p(\nu_i) = 1, \\ &\forall i \in 1, \cdots, n-1: \ \nu_{i+1} \in 8nb(\nu_i), \end{aligned} \tag{1}$$

where $8nb(\nu)$ is the 8-neighborhood function on the grid, and $p(\nu)$ returns the probability that the cell ν is passable, denoted as the robot's traversability.

Flat areas are considered traversable, and the environment geometry is used to determine areas that appear as obstacles. Since obstacles may be non-rigid, the robot's traversability through such areas cannot be determined only by geometry. Rather, an obstacle is traversable if the robot can exert force sufficient to pass through. Hence, the robot's traversability through a grid cell ν that appears as an obstacle is

$$p(\nu) = \begin{cases} 0 & \text{if } F(\nu) > \mathrm{F}_{\mathrm{trav}}, \\ 1 & \text{otherwise}, \end{cases} \tag{2}$$

where $F(\nu)$ is the force needed to pass through the cell ν, and the threshold $\mathrm{F}_{\mathrm{trav}}$ is the maximum force that can be exerted by the robot when trying to pass through the obstacle.

In the explored environment, the force to pass through the obstacles is not known for the individual obstacles. However, it is assumed the force is similar for similar-appearing obstacles, and thus the robot can predict the force needed to pass through the obstacles described by their respective appearance descriptors A as

$$f_{\mathrm{predict}} : A \rightarrow \hat{F}. \tag{3}$$

Besides, while we assume that the appearance description is sufficiently discriminative to distinguish the obstacles in each individual explored environment, similar-appearing obstacles in different deployments may have different rigidity. For example, dry summer grass is easier to traverse than wet grass prevalent during spring, even though they appear similar. Hence, the robot learns the rigidity predictor f_{predict} online during the exploration, starting from scratch for each deployment.

Since the task of the robot is to explore the environment where some areas may be reached only by traversing through the non-rigid obstacles, the portion of the environment that is explored is the benchmark value. The proposed method is thus evaluated and compared to a baseline model that considers the obstacles untraversable.

4 Proposed Traversability Model

The proposed method to characterize the traversability of apparent yet potentially traversable obstacles is presented as a part of the autonomous exploration. The robot is equipped with a haptic bumper sensor to experience the possible traversability of particular terrain areas. The traversability of obstructing obstacles the robot can walk through is characterized by sampling the force needed to pass through the obstacle. These haptic measurements are considered the traversability ground truth. The robot incrementally learns a Gaussian Process (GP) regressor (Rasmussen and Williams, 2006) employed to predict the force to pass through from the appearance of the apparent obstacles and thus to assess the traversability of the obstacles.

The idea of the proposed traversability model is demonstrated in an exploration-exploitation scenario set up as a robotic *escape mission*. The robot first explores the areas observed by the exteroception that appear traversable without interacting with apparent obstacles. After all such reachable areas are explored, the robot actively uses its haptic sensor to learn a model of obstacle traversability (force model exploration). Furthermore, when the robot learns that some apparent obstacle is traversable, it reverts to exploring the area that may lie behind such obstacles (force model exploitation).

The relation of the proposed terrain model, its learning based on the measured sensory input, and the decision-making in the exploration setup is depicted in Fig. 2. It principally works as follows. The *exteroceptive* part is responsible for continuously building an elevation map of the robot's surroundings using RGB-D sensory input. Further, the *exteroceptive* model identifies the areas that appear untraversable by the robot from

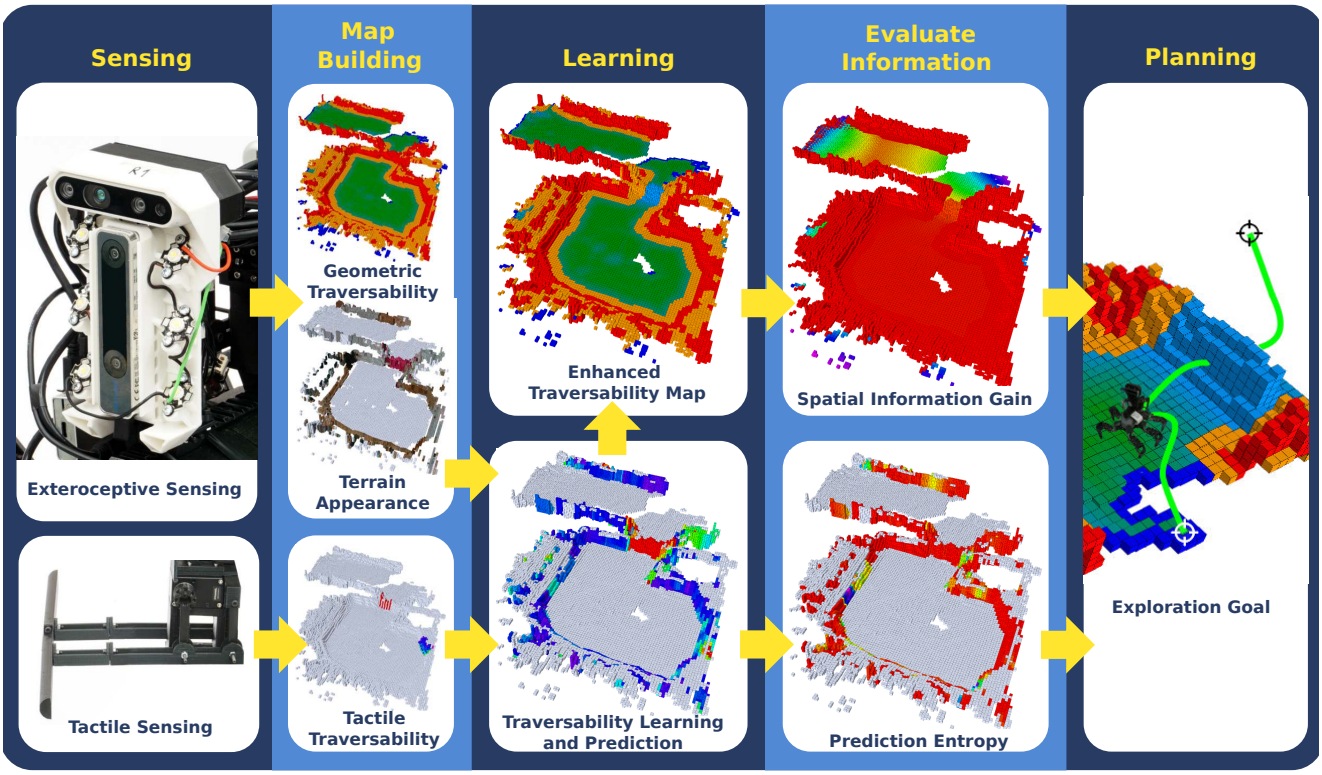

Fig. 2 Individual modules employed in building the traversability model and its usage in the autonomous exploration.

terrain geometry. The geometric properties of the terrain are then passed to the *learning* module. The *haptic* module accumulates measurements of the force needed to pass through the apparent obstacles as the force experienced by the haptic sensor during the mission. These ground-truth force measurements are also passed to the *learning* module, which pairs the respective appearance characterization from the *exteroceptive* model with these experienced haptic observations. The module realizes online learning and prediction of the force, and thus the traversability for the observed yet untraversed terrains. The *exteroceptive* model is updated with these traversability predictions to allow traversal through non-rigid obstacles. In *exploration*, the robot uses information gain predictions provided by the *exteroceptive* and *learning* parts to select the next exploration goal.

The cell traversability $\hat{p}(\nu)$ is reported based on the inputs from the *haptic*, *exteroceptive*, and *learning* modules. The cells with traversability measurements $p_{haptic}(\nu)$ are considered ground truth and reported as $\hat{p}(\nu) = p_{haptic}(\nu)$, regardless of the other modules. If there is no ground truth traversability reported for the cell by the *haptic* model, its geometric traversability $p_{geom}(\nu)$ provided by the *exteroceptive* module is used. If such cell appears traversable, the traversability is set to $\hat{p}(\nu) = p_{geom}(\nu) = 1$. For cells marked as potential obstacles with $p_{geom}(\nu) = 0$, the traversability is assessed by the *learning* module, reporting the traversability prediction $p_{predict}(\nu)$. Finally, unobserved cells are reported as traversable $\hat{p}(\nu) = p_{unobserved} = 1$ to allow traversal of such areas, which is desirable since cells hidden directly behind traversable obstacles cannot be observed before traversal. The traversability assesment can be summarized as

$$\hat{p}(\nu) = \begin{cases} p_{haptic}(\nu) & \text{if } p_{haptic}(\nu) \text{ is known} \\ p_{predict}(\nu, \mathcal{P}) & \text{if } p_{geom}(\nu) \text{ is known and } p_{geom}(\nu) = 0 \\ 1 & \text{otherwise} \end{cases}. \tag{4}$$

The individual exteroceptive and haptic models, traversability predictions, and the terrain learning process are detailed in the following sections. The symbols used in the description are overviewed in Table 1.

4.1 Exteroceptive Model

Robot's visual and depth perceptions are utilized to construct a colored elevation map $\mathcal{M}_{2.5D}$, see Fig. 3a. The elevation map is a grid map with the squared cell of the size d_ν, and its underlying representation is based on a memory-efficient quadtree data structure (Bayer and Faigl, 2020). For each cell $\nu \in \mathcal{M}_{2.5D}$, the geometrical traversability model (visualized in Fig. 3b) provides

Table 1 Used symbols.

Description	Symbol	Description	Symbol
First introduced in Section 3			
environment gridmap	$\mathcal{M}_{2.5D}$	gridmap cell	ν
cell size	d_ν	path	ψ
cell traversability	p	cell grid 8-neighborhood	$8nb$
force to pass through	F	maximum force exerted by robot	$\mathrm{F}_{\mathrm{trav}}$
terrain appearance	A	force prediction function	f_{predict}
First introduced in Section 4			
reported traversability	$\hat{p}$	ground-truth haptic traversability	p_{haptic}
apparent geometric prediction	p_{geom}	predicted traversability	$p_{predict}$
unobserved-cell traversability	$p_{unobserved}$		
First introduced in Section 4.1			
step height	Δ	maximum allowed step height	$\mathrm{t}_{\mathrm{trav}}$
information gained by observing a cell	I^{cell}_{geom}	information gained by observing from a cell	I^{model}_{geom}
cell δ neighborhood	δ	sensor range	δ_{sensor}
spatial goal cluster radius	d_{cl}	spatial goal minimum cluster size	$\mathrm{n}^{\mathrm{thr}}_{\mathrm{cl}}$
spatial goal set	$\mathcal{G}_{geom}$		
First introduced in Section 4.2			
force measurement	z^{force}	k-th force measurement at cell	z^{force}_k
force reported at cell after k measurements	F_k	force uncertainty at cell after k measurements	σ^2_k
bumper sensor measurement uncertainty	$\sigma^2_{\mathrm{sensor}}$		
First introduced in Section 4.3			
traversability predictor	$\mathcal{P}$	mono-color appearance descriptor	A^{AB}
force prediction mean	μ_F	force prediction variance	σ^2_F
approximated predictor information gain	$I_{predict}$		
First introduced in Section 4.4			
GP model noise variance	σ^2_n	exponential kernel output variance	σ_{exp}
exponential kernel lengthscale	l_{exp}		
First introduced in Section 4.5			
planning cost	c	the cheapest path	ψ^*
exploration goal	$\nu^*_{explore}$	spatial exploration goal	ν^*_{geom}
predictor exploration goal	$\nu^*_{predict}$	predictor initialization goal	$\nu^{*-\Psi}_{predict}$
predictor learning goal	$\nu^{*-I}_{predict}$	minimum prediction goal information	$\mathrm{I}^{\mathrm{thr}}_{\mathrm{predict}}$
First introduced in Section 5.3			
histogram appearance descriptor	A^{hist}	histogram descriptor radius	r_{hist}

Some single-use symbols are omitted for simplicity.

$p_{geom}(\nu)$, the probability that the robot can traverse the cell ν, by comparing local differences of height to the threshold $\mathrm{t}_{\mathrm{trav}}$

$$p_{geom}(\nu) \begin{cases} 0 & \text{if } \max\limits_{\xi \in 8nb(\nu)} \Delta(\nu, \xi) > \mathrm{t}_{\mathrm{trav}} \\ 1 & \text{otherwise} \end{cases}, \tag{5}$$

where $8nb(\nu)$ is the 8-neighborhood of the cell ν, the particular value of $\mathrm{t}_{\mathrm{trav}}$ depends on the kinematics of the used robot. The step height $\Delta(\nu_a, \nu_b)$ is defined as

$$\Delta(\nu_a, \nu_b) = |elevation(\nu_a) - elevation(\nu_b)|, \tag{6}$$

with $elevation(\nu)$ denoting the estimated height of the terrain at the cell ν. Note that all the results presented in this paper are for $\mathrm{t}_{\mathrm{trav}} = 12\,\mathrm{cm}$ based on the kinematics constraints and motion gait of the utilized hexapod walking robot.

The information about the geometric traversability model gained by observing an unknown cell ξ is approximated as the entropy of the binary distribution p_{geom} that depends on the 8-neighborhood of the cell. Since the knowledge whether one cell is traversable corresponds to one bit, the information gained by observing ξ with unknown height is approximated as

$$I^{cell}_{geom}(\xi) \approx \frac{k(\xi) + 1}{9}, \tag{7}$$

where $k(\xi)$ is the number of the unknown cells in the neighborhood of ξ. Thus, the expected information gained by perceiving the terrain from the position of the cell ν is

$$I^{model}_{geom}(\nu) = \sum_{\xi \in \delta(\nu, \delta_{\mathrm{sensor}})} \begin{cases} I^{cell}_{geom}(\xi) & \text{if } observable(\nu, \xi) \\ 0 & \text{otherwise} \end{cases}, \tag{8}$$

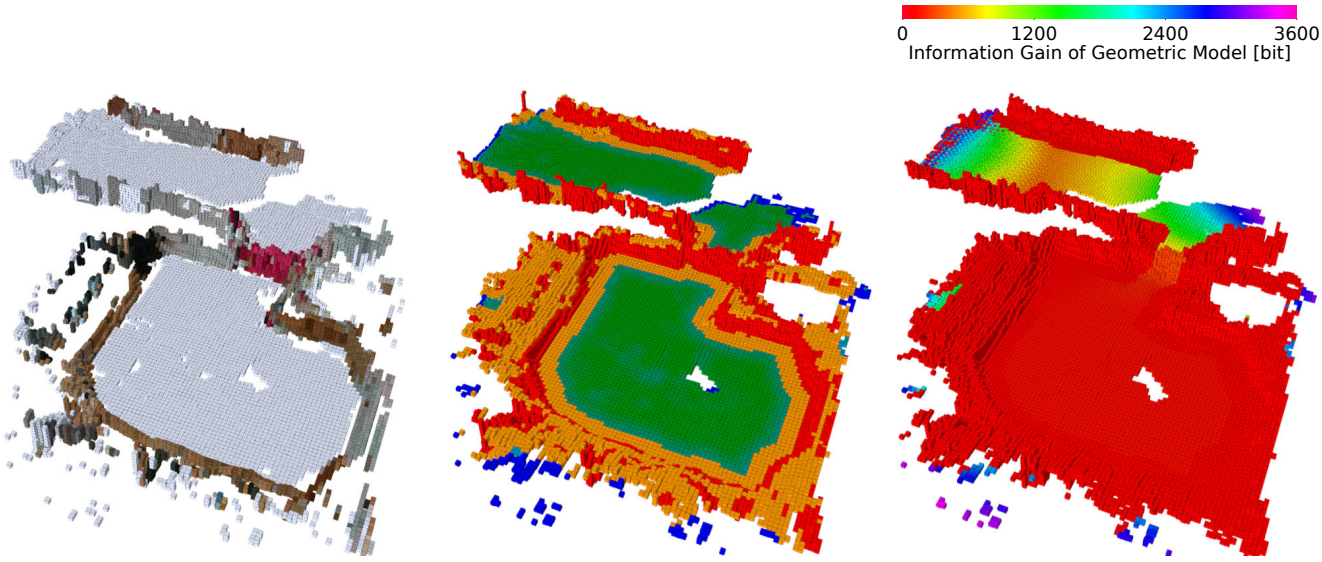

(a) (b) (c)

Fig. 3 Example of (a) the colored elevation grid map (obstacles colored, the ground is in the light blue); (b) the potential obstacles determined solely on their geometric properties (obstacles are in the red, traversable terrain in the green, and cells too close to obstacles in the yellow); and (c) the information gain of the geometric model indicating terrain areas with unsure spatial traversability because of lack of strong evidence from exteroceptive measurements.

where $\delta(\nu, \delta_{\text{sensor}})$ is the neighborhood of ν defined by the sensor range δ_{sensor}, which value depends on the used sensor, e.g., $\delta_{\text{sensor}} = 2\,\text{m}$. The function $observable(\nu, \xi)$ returns `true` if the cell ξ is observable from ν, which is determined by casting a ray from ν to ξ in the current elevation map $\mathcal{M}_{2.5D}$. An example of the information gain of the geometric traversability model is depicted in Fig. 3c.

In an active perception scenario, the goal locations $\mathcal{G}_{geom}$ to improve the geometric traversability model are selected as a subset of cells with non-zero information gain according to Algorithm 1.

4.2 Haptic Model

The *haptic* model uses measurements provided by the bumper sensor to predict the traversability of the observed obstacles that can be non-rigid. In particular, the sensor measures the force $F(\nu)$ needed to pass through the obstacle at cell ν. Since the haptic sensor utilizes a high sampling frequency, multiple different measurements are acquired for a single cell ν. Therefore, each cell uses a Kalman filter that fuses the measurements as

$$\begin{aligned} F_k(\nu) &= \frac{\sigma^2_{\text{sensor}} F_{k-1}(\nu) + \sigma^2_{k-1}(\nu) z^{\text{force}}_k(\nu)}{\sigma^2_{\text{sensor}} + \sigma^2_{k-1}(\nu)}, \\ \sigma^2_k(\nu) &= \frac{\sigma^2_{\text{sensor}} \sigma^2_{k-1}(\nu)}{\sigma^2_{\text{sensor}} + \sigma^2_{k-1}(\nu)}, \end{aligned} \tag{9}$$

where $F_k(\nu)$ is the value reported by the *haptic* model for the cell ν after k measurements were assigned to the cell, $z^{\text{force}}_k(\nu)$ is the k-th measurement assigned to the cell, σ^2_{sensor} is the bumper sensor measurement uncertainty, the initial force at the cell $F_0(\nu)$ equals the first measurement assigned to the cell, and the initial filter variance is $\sigma^2_0(\nu) = 1$. An example of the acquired traversability experience projected onto the elevation map (visualized in Fig. 3) is shown in Fig. 4a.

The haptic sensor is considered to provide the ground truth traversability measurements. Therefore, a binary value of the traversability p_{haptic} is utilized for cells where the ground truth measurements are available, with the traversability of the cell ν traversed by the haptic sensor being determined w.r.t. (2).

4.3 Traversability Prediction

The ground-truth force measurements reported by the *haptic* module are limited to the particular obstacles the robot has interacted with. Hence, the ground truth traversability is relatively sparse. Therefore, the traversability *learning* module combines the exteroceptive and haptic information about the environment and determines the traversability for each observed cell ν where the ground truth is unavailable. The traversability predictor $\mathcal{P}$ is learned from the haptic experience that is extrapolated using the appearance A perceived by exteroceptive sensing.

Algorithm 1: Goal locations clustering.

Input: $\mathcal{M}_{2.5D}$– 2.5D grid map with assessed frontiers,
$\mathrm{d_{cl}}$– Cluster radius, $\mathrm{n_{cl}^{thr}}$– Minimal cluster size.
Output: $\mathcal{G}_{geom}$– Clustered cells with non-zero entropy.

$A \leftarrow \emptyset$ ▷ Init. set of clusters.
for $\nu \in \mathcal{M}_{2.5D} : I^{model}_{geom}(\nu) > 0$ **do** ▷ For each map cell with non-zero entropy.
  if $A = \emptyset$ **then** ▷ If no clusters in set.
    $A \leftarrow \{\{\nu\}\}$ ▷ Create a new cluster.
  else
    $d \leftarrow$ `distanceToClosestCluster`(ν, A)
    if $d < d_{cl}$ **then**
      `addToClosestCluster`(ν, A) ▷ Add point to existing cluster.
    else
      $A \leftarrow A \cup \{\{\nu\}\}$ ▷ Create new cluster.

$\mathcal{G}_{geom} \leftarrow \emptyset$ ▷ Init. cluster representants.
for $A_i \in A$ **do** ▷ For each clusters.
  if $|A_i| > n^{thr}_{cl}$ **then**
    $\mathcal{G}_{geom} = \mathcal{G}_{geom} \cup \{$`cellClosestToAverageCoordinates`$(A_i)\}$ ▷ Create new representants.
return $\mathcal{G}_{geom}$

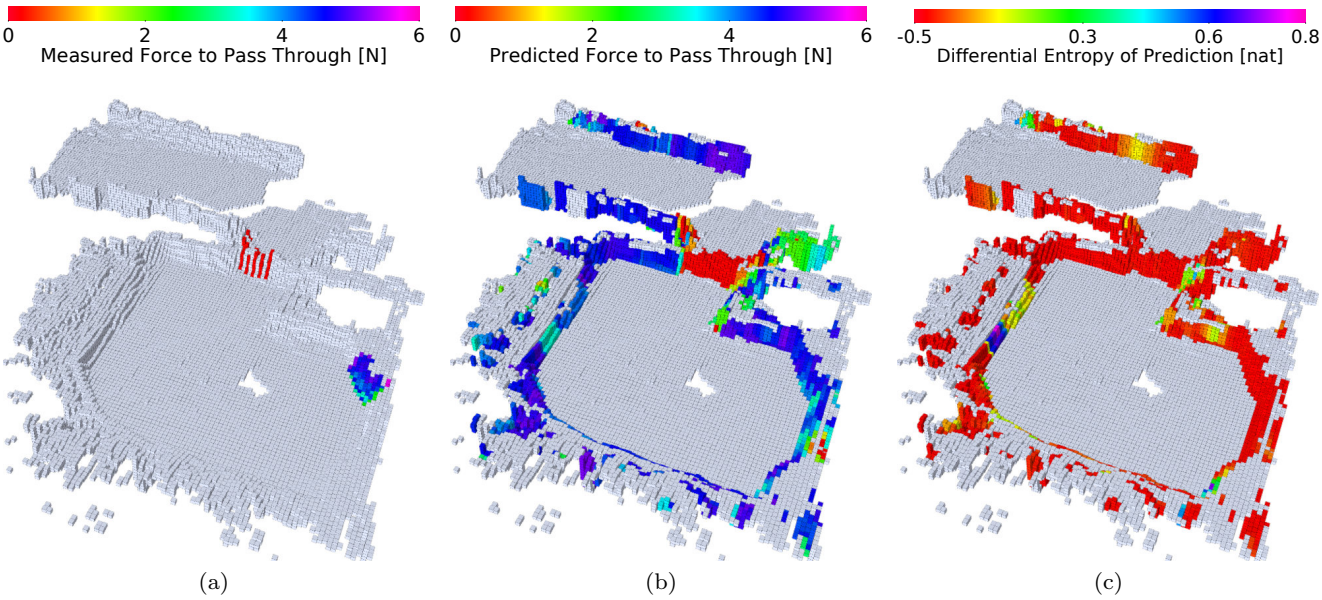

(a) (b) (c)

Fig. 4 (a) Example of haptic based traversability ground truth perceived by the robot as the measured force to pass through the obstacle. (b) The traversability prediction of the potential obstacles as the predicted force to pass through the obstacle and (c) the respective prediction of the differential entropy.

The *learning* module learns to predict the force to pass through the observed potentially untraversable cell ν using its appearance. The terrain description $A(\nu)$ of each potentially untraversable cell ν is based on its color appearance. We advocate the use of simple descriptors, which are easy to compute and read. In scenarios with large, mono-color obstacles, we use a pair of the cell's a and b colors in the Lab color space computed from the RGB-colored elevation grid, as shown in Fig. 3a,

$$A^{\mathrm{AB}}(\nu) = (a, b). \tag{10}$$

The traversability predictor $\mathcal{P}$ is learned from each observed potentially untraversable cell ν that carriers the haptic force measurement $F(\nu)$ as

$$\mathcal{P} \leftarrow \mathrm{learn}(\{(A(\nu), F(\nu)\}_{\{\nu\}}). \tag{11}$$

It predicts the force to pass through the potentially untraversable cell ν distributed as

$$\mathcal{N}(\mu_F(\nu,\mathcal{P}),\sigma_F^2(\nu,\mathcal{P})) \leftarrow \text{predict}(\nu,\mathcal{P}), \tag{12}$$

where the force prediction mean $\mu_F(\nu,\mathcal{P})$ determines the traversability prediction $p_{predict}(\nu,\mathcal{P})$ considering the force threshold $\mathrm{F}_{\mathrm{trav}}$ as in (2). Besides, if the predictor $\mathcal{P}$ has not yet been learned, it assesses cells as untraversable. An example of the traversability prediction over the elevation grid map of the environment is shown in Fig. 4b.

The predicted distribution is utilized to estimate the expected information gain associated with sampling cells of unknown haptic traversability ground truth to steer the robotic exploration towards collecting the required information to improve the traversability model. The information expected from the haptic measurement at the cell ν with unknown haptic traversability $p_{haptic}(\nu)$ is approximated as the differential entropy of the respective predicted distribution (see Fig. 4c)

$$I_{predict}(\nu,\mathcal{P}) \approx H(\mathcal{N}(\mu_F,\sigma_F^2)) = \frac{1}{2}\log(2\pi e\sigma_F^2(\nu,\mathcal{P})). \tag{13}$$

Notice that for cell ν with the known haptic traversability ground truth $p_{haptic}(\nu)$, the ground truth with $\sigma_F^2 = 0$ is considered instead of the prediction, and thus no additional information can be gained, and the differential entropy at any such cell is undefined.

4.4 Gaussian Process Regressor

The traversability prediction is based on a GP regressor, briefly described here to make the paper self-contained. Given an observed function $f(x)$ with the noise ϵ

$$y = f(x) + \epsilon, \quad \epsilon \in \mathcal{N}(0,\sigma_n^2), \tag{14}$$

GP is a distribution over the functions (Rasmussen and Williams, 2006)

$$f(x) \sim \mathcal{GP}(m(x),K(x,x')) \tag{15}$$

where $m(x)$ and $K(x,x')$ are mean and covariance, respectively, defined as

$$m(x) = E\left[f(x)\right], \tag{16}$$

$$K(x,x') = E\left[\left(f(x)-m(x)\right)\left(f(x')-m(x')\right)\right]. \tag{17}$$

Given the train data X and the test data X_*, the latent values f_* at X_* are

$$\begin{aligned}\mu(X_*) &= K(X,X_*)\left[K(X,X)+\sigma_n^2 I\right]^{-1} y,\\ (\sigma(X_*))^2 &= K(X_*,X_*)\\ &\quad - K(X,X_*)^T\left[K(X,X)+\sigma_n^2 I\right]^{-1} K(X,X_*),\end{aligned} \tag{18}$$

where $K(X,X')$ is the covariance function. In this work, the used covariance function is the exponential kernel

$$K(\boldsymbol{x},\boldsymbol{x'}) = \sigma_{\exp}^2 \exp(-\frac{1}{l_{\exp}^2}\|\boldsymbol{x}-\boldsymbol{x'}\|). \tag{19}$$

In the considered exploration scenario, the model is learned online using only the available onboard computational resources of the robot. Hence, it is necessary to consider the computational requirements as the computation of GPs can generally be demanding. Therefore, the GP regressor is relearned from the accumulated traversability observations and terrain appearance with a fixed rate of 0.03 Hz. The real performance of the proposed terrain learning model is reported in Section 5 within the autonomous exploration that is briefly described in the following section.

4.5 Exploration Scenario

The proposed terrain traversability approach with haptic and exteroceptive sensing is intended to model the robot's operational environment where some parts can look like obstacles in exteroceptive data but can be traversed. Since the robot environment is represented by the grid map $\mathcal{M}_{2.5D}$, the robot plans its paths through the environment w.r.t. the cost

$$c(\nu_a,\nu_b) = \|(\nu_a,\nu_b)\| + c_d(\mathcal{M}_{2.5D}), \tag{20}$$

where ν_a and ν_b are two cells that are 8-neighbors, the norm $\|(\nu_a,\nu_b)\|$ is the respective Euclidean distance between the cells' centers, and c_d is a non-negative cost. The cost c_d decreases with the distance from the closest untraversable cell to penalize robot presence close to the obstacles as in (Bayer and Faigl, 2019). The path cost c (20) is used to assess a cost of path ψ, and select the shortest path in the environment as

$$\begin{aligned}&\psi^*(\nu_{start},\nu_{goal}) = \text{argmin}_{\psi\in\Psi(\nu_{start},\nu_{goal})}\, c(\psi),\\ &c(\psi) = \textstyle\sum_{i=0}^{|\psi|-1} c(\nu_i,\nu_{i+1}),\end{aligned} \tag{21}$$

where $\Psi(\nu_{start},\nu_{goal})$ is the set of all possible paths from ν_{start} to ν_{goal}, and the path $\psi \in \Psi(\nu_{start},\nu_{goal})$ is a path starting at ν_{start} and ending at ν_{end}. The cheapest path ψ^* is determined using the A* algorithm.

The exploration mission is considered as an escape-like scenario. The exploration procedure is overviewed in Algorithm 2, and it works as follows. The robot first explores the areas accessible without interacting with apparent obstacles. Only after all reachable areas are explored, the haptic sensor is actively used to learn the

Algorithm 2: Learning the proposed traversability model in a robotic escape mission.

Input: $z^{\text{range}}_{1,\dots,m}$ – Colored range measurements; $z^{\text{force}}_{1,\dots,n}$ – Force measurements.
Output: $\mathcal{M}_{2.5D}$ – Colored elevation gridmap including traversability assesments for apparent obstacles; $\mathcal{P}$ – Traversability predictor for apparent obstacles.

$\mathcal{M}_{2.5D} \leftarrow \varnothing$ ▷ Initialize gridmap.
$\mathcal{P} \leftarrow \varnothing$ ▷ Initialize predictor.
while exploration is running **do**
  $\mathcal{M}_{2.5D} \leftarrow$ `updateModel`$(\mathcal{M}_{2.5D}, z^{range}_{i,\dots,i'})$ ▷ Update the gridmap using the latest range measurements.
  $\mathcal{P} \leftarrow$ `updateModel`$(\mathcal{P}, z^{force}_{j,\dots,j'}, \mathcal{M}_{2.5D})$ ▷ Update the predictor using the force measurements and terrain descriptors.
  $\mathcal{M}_{2.5D} \leftarrow$ `applyPredictions`$(\mathcal{P})$ ▷ Apply traversability predictions to the gridmap.
  for $\nu \in \mathcal{M}_{2.5D} : \hat{p}(\nu) = 1$ **do** ▷ For each traversable cell ν (including predicted traversability).
    compute $I^{model}_{geom}(\nu)$ ▷ Compute the spatial information gained by observing from the cell w.r.t. (8).
  if $\exists \nu \in \mathcal{M}_{2.5D} : \hat{p}(\nu) = 1, I^{model}_{geom}(\nu) > 0$ **then** ▷ Explore the spatial model if any information can be gained.
    $\mathcal{G}_{geom} \leftarrow$ `cluster`$(\forall \nu \in \mathcal{M}_{2.5D} : I^{model}_{geom}(\nu) > 0)$ ▷ Create the spatial exploration goals.
    $\nu^* \leftarrow \text{argmin}_{\nu_{geom} \in \mathcal{G}_{geom}} c(\psi^*(\nu_{robot}, \nu_{geom}))$ ▷ Select the closest goal.
  else if $\mathcal{P}$ *is not learned* **then** ▷ Start exploring the predictor if it does not exist.
    $\nu^* \leftarrow \text{argmin}_{\nu \in \mathcal{M}_{2.5D} | p_{geom}(\nu)=0} c(\psi^*(\nu_{robot}, \nu))$ ▷ Select the closest apparent obstacle as goal to start learning.
  else ▷ Otherwise, continue exploring the predictor.
    for $\nu \in \mathcal{M}_{2.5D} : p_{geom}(\nu) = 0$ **do** ▷ For each potential obstacle ν.
      compute $I_{predict}(\nu, \mathcal{P})$ ▷ Compute the predictor information gained w.r.t. (13).
    if $\max_{\nu \in \mathcal{M}_{2.5D} | p_{geom}(\nu)=0} I_{predict}(\nu, \mathcal{P}) > I^{thr}_{predict}$ **then** ▷ If enough information can be gained.
      $\nu^* \leftarrow \text{argmax}_{\nu \in \mathcal{M}_{2.5D} | p_{geom}(\nu)=0} I_{predict}(\nu, \mathcal{P})$ ▷ Select the most informative cell as goal.
    else
      $\nu^* \leftarrow \varnothing$ ▷ Otherwise, select no goal.
  if $\nu^* \neq \varnothing$ **then**
    `navigateTo`(ν^*) ▷ Navigate to goal if it exists.
  else
    `finishExploration()` ▷ Otherwise, end exploration.
return $\mathcal{M}_{2.5D}$,$\mathcal{P}$

model of obstacle traversability. The exploration strategy selects the goal as

$$\nu^*_{explore} = \begin{cases} \nu^*_{geom} & \text{if } \nu^*_{geom} \text{ exists} \\ \nu^*_{predict} & \text{otherwise} \end{cases}, \tag{22}$$

where the geometric exploration goal ν^*_{geom} is selected if it is possible to gain any additional information about the geometric model. The prediction model improvement goal $\nu^*_{predict}$ is selected otherwise. In particular, the robot selects the closest geometry exploration goal as

$$\nu^*_{geom} = \text{argmin}_{\nu_{geom} \in \mathcal{G}_{geom}} c(\psi^*(\nu_{robot}, \nu_{geom})), \tag{23}$$

where ν_{robot} is the cell corresponding to the current robot position. The goal to improve the prediction is selected either as the cell with the highest potential information gain about the prediction model, or the closest potentially untraversable cell if the prediction model is not yet learned

$$\nu^*_{predict} = \begin{cases} \nu^{*-I}_{predict} & \text{if } \mathcal{P} \text{ is learned} \\ \nu^{*-\Psi}_{predict} & \text{otherwise} \end{cases}, \tag{24}$$

where the cell with the highest information gain potential is

$$\nu^{*-I}_{predict} = \text{argmax}_{\nu \in \mathcal{M}_{2.5D} | p_{geom}(\nu)=0} I_{predict}(\nu, \mathcal{P}), \tag{25}$$

and the closest potentially untraversable cell is

$$\nu^{*-\Psi}_{predict} = \text{argmin}_{\nu \in \mathcal{M}_{2.5D} | p_{geom}(\nu)=0} c(\psi^*(\nu_{robot}, \nu)). \tag{26}$$

Note that the area in the vicinity of the prediction goal is temporarily cleared as traversable in $\mathcal{M}_{2.5D}$ to allow the robot to approach the sampling location. Finally, the robot does not pursue prediction goals associated with less than $\text{I}^{\text{thr}}_{\text{predict}}$ information gained. Therefore, when there are no geometric exploration goals and sampling, no potentially untraversable cell is associated with more than $\text{I}^{\text{thr}}_{\text{predict}}$ information gained, the exploration stops.

5 Experimental Results and Discussion

The proposed system for active terrain traversability learning using visual and haptic cues has been experi-

mentally verified in two scenarios. First, the robot is deployed in an escape-like exploration scenario to demonstrate active learning of passing through obstacles and thus explore areas that would remain inaccessible if only visual sensing would be used. The robot is deployed in an indoor office arena containing rigid obstacles the robot cannot traverse and non-rigid obstacles the robot can pass through. In the second scenario, we showcase the predictor in an outdoor setting with realistic vegetation.

Prior to the results from each scenario, a brief description of the utilized robot and its sensors are presented in Section 5.1. The two deployment scenarios are individually presented in Section 5.2 and Section 5.3, respectively. The results are further discussed in Section 5.4.

5.1 Robot and Sensory Equipment

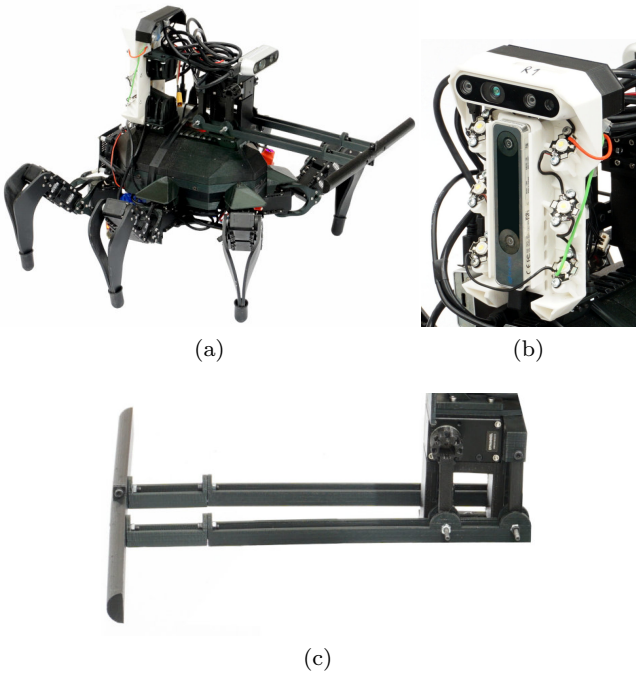

(a) (b)

(c)

Fig. 5 (a) The hexapod walking robot used in the experimental deployment, (b) its rear-facing sensor rig, (c) and the hinge mechanism of its haptic bumper sensor.

The proposed system is deployed on the hexapod walking robot (Faigl and Čížek, 2019) shown in Fig. 5a. The robot has six legs attached to its trunk, each comprising three Dynamixel AX12 servomotors. The robot, including its legs, can fit into a square with the side length of about 40 cm. The robot is equipped with exteroceptive sensor rigs to localize the robot and build the colored elevation map. The forward-facing rig comprises the Intel RealSense D435 RGB-D camera (D435 for short). The rear-facing camera rig holds another D435 and the Intel RealSense T265 tracking camera (T265 for short), see Fig. 5b. The localization using the rear-facing T265 is selected to avoid losing tracked features when the robot approaches obstacles.

Table 2 System performance.

Component	Update Rate [Hz]	CPU usage [%]
Exteroception (each D435)	1	25
Tactile sensing	300	4
Localization	200	7
Map building	2	7
Path planning	5	
Feature description	1	6
Learning and prediction	0.03	13
Locomotion	10	3

The robot carries a haptic sensor designed as a bumper mounted on a parallelogram hinge, see Fig. 5c. The sensor is based on the Dynamixel XM430 servomotor that is actuated and set to return to a pre-set position. The servomotor provides torque measurements that are paired with the tabulated force values obtained by letting the bumper push on a force sensor prior to the deployment. The measured force values allow us to transform the bumper into a sensor measuring the force to pass through obstacles. The sensor has been individually calibrated for each presented experiment to avoid eventual bumper changes. The traversability threshold has been set to $\mathrm{F}_{\mathrm{trav}} = 2\,\mathrm{N}$.

Furthermore, the robot is equipped with two simple reflexes that help to sample the obstacle rigidity and traverse through non-rigid obstacles that can impede its visual sensors. First, when the robot gets close, within 0.25 m, to an obstacle, it tries to pass through by walking forward. The behavior is stopped either 10 s after leaving the obstacle vicinity or by triggering the second reflex when the bumper observes values higher than $\mathrm{F}_{\mathrm{trav}}$. In such a case, the robot samples the obstacle for 2 s to ensure that the observed value is not an outlier caused by bumper motion. Then, the robot engages a backward motion for 8 s to clear the obstacle.

The onboard computational resources are the Nvidia Jetson TX2 with 8 GB RAM running Robot Operating System (ROS) Melodic (Quigley et al., 2009) that demonstrated sufficient computational power according to the real computational requirements overviewed in Table 2. The proposed system is parametrized, as in Table 3. Similar to the exploration strategy proposed in Karolj et al. (2020), we use fixed values of the kernel

Table 3 System parametrization.

Symbol	Description	Value	Reasoning
t_{trav}	Obstacle detection step height	0.12 m	Robot step height
δ_{sensor}	RGB-D sensor range	2 m	Properties of the sensor
d_ν	Size of the squared grid cell of $\mathcal{M}_{2.5D}$	0.05 m	Size of the robot foothold
d_{cl}	Geometric goal cluster radius	1 m	Twice size of the robot
n_{cl}^{thr}	Geometric goal minimum cluster size	10	Set empirically
σ^2_{sensor}	Bumper sensor uncertainty in Kalman fuser	0.01	Bumper sensor calibration
σ_n^2	Gaussian Process noise variance σ_{exp}	0.1	Set empirically
σ_{exp}	Gaussian Process exponential kernel σ_{exp}	1	Set empirically
l_{exp}	Gaussian Process exponential kernel (simple features) l_{exp}	1	Set empirically
	Gaussian Process exponential kernel (histogram features) l_{exp}	0.4	Set empirically
F_{trav}	Maximum force to push through obstacle	2 N	Properties of the robot
$I_{predict}^{thr}$	Miminum prediction model MI	0.26 nat	Set empirically ($\sigma_F^2 \approx 0.1$)
r_{hist}	Color histogram descriptor radius	0.15 m	Half of the bumper width

hyper-parameters instead of optimizing them when recomputing the GP because the optimization process is computationally costly and thus not suitable for online deployments. We exploit that the ranges of the feature descriptors values and measured forces remain the same between the individual deployments and set the hyper-parameters empirically. In particular, the feature sensitivity can be adjusted via the kernel lengthscale with the intuition that it is possible to extrapolate roughly within the lengthscale distance of the known data. The kernel output variance is set so that the already sampled terrains report prediction entropies that are considered known w.r.t. $I_{predict}^{thr}$.

5.2 Exploration Scenario

In the exploration scenario, the robot has been deployed in an arena split into three sections, see Fig. 6. The approximate size of the arena is 35 m^2. The robot cannot see the other sections of the arena from each section. The robot starts in a small arena section located in an office room with rigid brown wooden obstacles, rigid white walls, and purple fabric. The fabric can be traversed to access a corridor containing the second and third arena sections. The two sections in the corridor contain similar obstacles as the first section. The corridor sections are divided by another purple fabric, enabling the robot which has already learned individual terrains' rigidity to identify the fabric as a traversable area. The goal of the experimental deployment is to explore all three sections and thus prove the ability of the proposed approach to sample and learn the force to pass through the individual obstacle types. The force sensor calibration in the exploration scenario is based on the values depicted in Table 4, which are used to compute a cubic spline to create a function of the force based on the torque measurements.

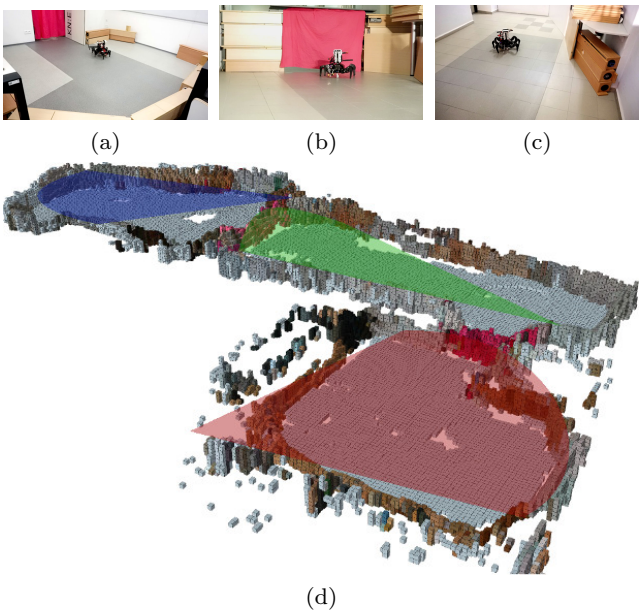

Fig. 6 The robot (a) in the first arena section, (b) leaving the second section, (c) exploring the third section; and (d) the map projection of the three respective image vision cones in the red, green and blue.

Table 4 Bumper Sensor Calibration Values

Torque [Nm]	0	0.17	0.34	0.51	0.60	0.79	10^*
Force [N]	0	1.05	2.02	3.10	4.09	6.34	10×10^4

*A limit value used in interpolation; not actually measured.

Particular maps showing the arena before the first sampling of the traversable fabric in the first section, during the exploration of the second section, and near the end of the experiment are shown in Fig. 7, 8, and 9, respectively. Note, a video of the therein depicted experimental run is also presented in Online Resource 1.

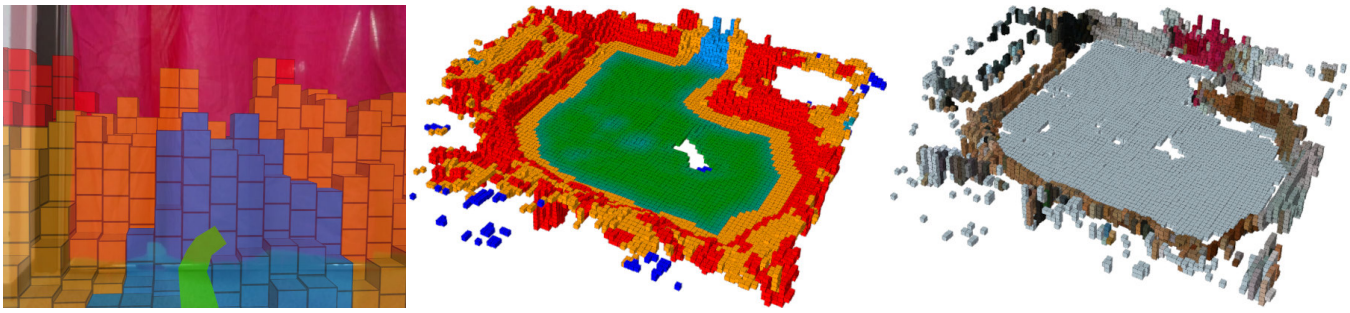

(a) (b) (c)

Fig. 7 The traversability (a) as seen from the robot and (b) its overview, and (c) color features in the arena before sampling the traversable fabric in the first section. Note, an obstacle is cleared as traversable even though the robot has not yet learned its rigidity because it needs to approach the obstacle to learn its traversability.

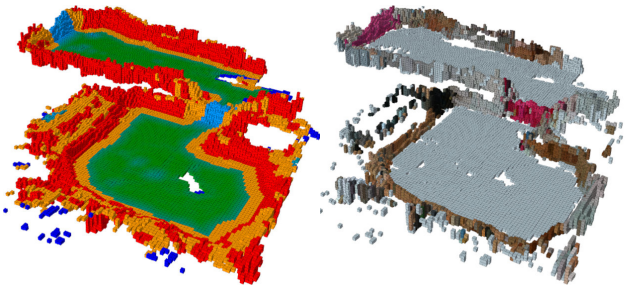

(a) (b)

Fig. 8 The (a) traversability and (b) color features in the arena during the exploration of the second section.

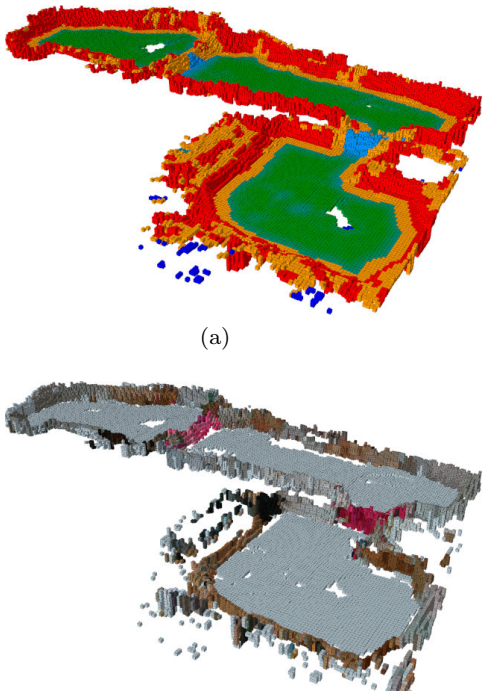

(a) (b)

Fig. 9 The (a) traversability and (b) color features near the end of the experiment.

The robot behaved similarly in five deployment runs, of which we choose two particular runs to report the robot's behavior here. After exploring the initial section, the robot samples the obstacles. In general, if the robot has sampled the wall or a wooden obstacle first, it chooses to sample the purple curtain second, since its color in the Lab descriptor space is distant from the colors of the rigid obstacles and thus it has a high prediction model uncertainty. Hence, the robot walks through the curtain and enters the second section of the arena. There, the robot resumes spatial exploration and may attempt to traverse obstacles it considers traversable based on its previous experience. Since the robot has not necessarily sampled all the obstacles available in the first area, its traversability predictions might be too optimistic. It is expected behavior as it is a result of the incremental nature of the learning process. Nevertheless, the robot obtains new force ground truth when it attempts to traverse obstacles erroneously considered as traversable, thus further aiding the learning process as demonstrated in the alternative experimental run where the robot sampled the traversable fabric first, see Fig. 10.

The differences between the two experimental runs in the proportional representation of predictions considered as sure with regards to the uncertainty threshold $\mathrm{I}^{\mathrm{thr}}_{\mathrm{predict}}$, and of the predictions thresholded as either traversable or untraversable w.r.t. $\mathrm{F}_{\mathrm{trav}}$, can be seen in Fig. 11a and Fig. 11b, respectively. The evolution of the prediction entropy distribution is depicted in Fig. 12. It can be seen that although the robot predicts a large portion of obstacles as traversable after only sampling the purple fabric in Fig. 11b, a large portion of these predictions is not sure. After the robot learns that some terrains are rigid by sampling them in

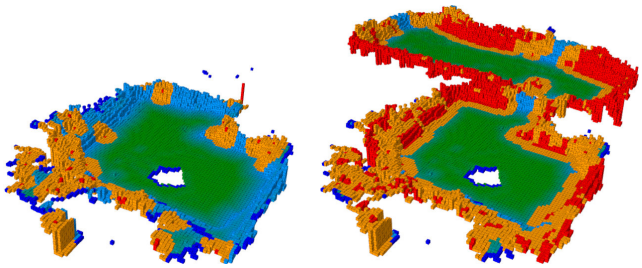

(a) (b)

Fig. 10 The predicted traversability in the experimental run where the robot has sampled the traversable fabric first: (a) overly optimistic predictions after sampling the fabric and (b) corrected predictions after the robot has sampled other obstacles due to trying to traverse through them while exploring the corridor section.

the second area, sure untraversable predictions emerge, even though some uncertain traversable predictions remain, particularly in the first section of the map.

5.3 Outdoor Scenario

In the outdoor scenario, the proposed system has been deployed in several locations in the Prokop Valley in Prague, Czech Republic. The bumper is used to collect the haptic data for several terrains in the area, which are paired with feature descriptors of the respective terrains. The bumper sensor calibration is depicted in Fig. 13. The terrains consists of several types of grass with varying density and appearance (some of them traversable), and a rigid tree trunk and rocks, see Fig. 14.

Since the terrains are more complex than in the indoor exploration, similarly to Belter et al. (2019), we utilize an alternative terrain descriptor in the outdoor scenario. The outdoor descriptors A^{hist} are based on color histograms, where each cell $\nu \in \mathcal{M}_{2.5D}$ is provided a 10-bit color by projecting the camera image to the $\mathcal{M}_{2.5D}$. Then, the color space is shrunk to 8 different colors: white, black, grey, blue, green, red, brown, and magenta. The relative amount of the cell color prototypes within the radius r_{hist} given by the half of the bumper size is used to build an 8-dimensional color histogram for each cell $\nu \in \mathcal{M}_{2.5D}$ as illustrated in Fig. 15. Besides, since the appearance descriptors differ from the indoor experiment in both scale and dimensionality, we adjust the kernel lengthscale, see Table 3.

The measurements from each terrain are split into testing and training sets. The algorithm is incrementally presented with training sets for the individual terrains, simulating a robot learning the terrains in a sequence. Table 5 shows the predictions on the respective testing sets in three alternative training sequences. The results suggest that after being presented with the particular terrain, the robot learns the terrain and reports entropy below the uncertainty threshold $\mathrm{I}^{\text{thr}}_{\text{predict}}$; hence, marking the terrain as known.

5.4 Discussion

Based on the experimental deployment, we can conclude that the proposed robotic system actively learns traversable obstacles and can explore areas hidden behind such obstacles and thus explore a greater portion of such environment than a system considering all obstacles as untraversable. Unlike the state-of-the-art approaches that consider mobile robot interaction with non-rigid obstacles (Baleia et al., 2015; Kahn et al., 2021), the herein proposed system concerns mobile robot exploration, and the reasoning about the traversability of non-rigid obstacles is done in the context of the environment geometric map, which is built online as a part of the exploration process.

The difference is further manifested in the selected learning approach, where Kahn et al. (2021) uses an end-to-end network learned from a large amount of data gathered using a time-correlated random walk policy. Besides, visual-tactile sensing considered by Pearson et al. (2021) combines tactile whisker sensors with visual perception for place recognition in Multimodal Predictive Coding Network (MultiPredNet), a bio-inspired approach that comprises visual, tactile, and multi-sensory modules. While such approaches provide significant advantages in the form of navigation policy learning and crossmodal reconstruction, respectively, the herein proposed approach is focused on the problem of apparent-yet-non-rigid obstacles in mobile robot exploration. Hence, it learns only in areas associated with the high information gain concerning the prediction model and can learn online during the mission itself.

6 Conclusion

We present a system for online learning of the force needed to pass through obstacles employed in autonomous exploration to assess traversability in an environment with terrain that appears untraversable yet can be traversed. The robot actively learns the geometric model of its surroundings with model learning to predict the traversability of potentially obstructing terrains using a haptic sensor. Gaussian Process regressor is utilized for the force prediction representing the traversability of the potentially obstructing terrains. The robot actively navigates based on expected information gain

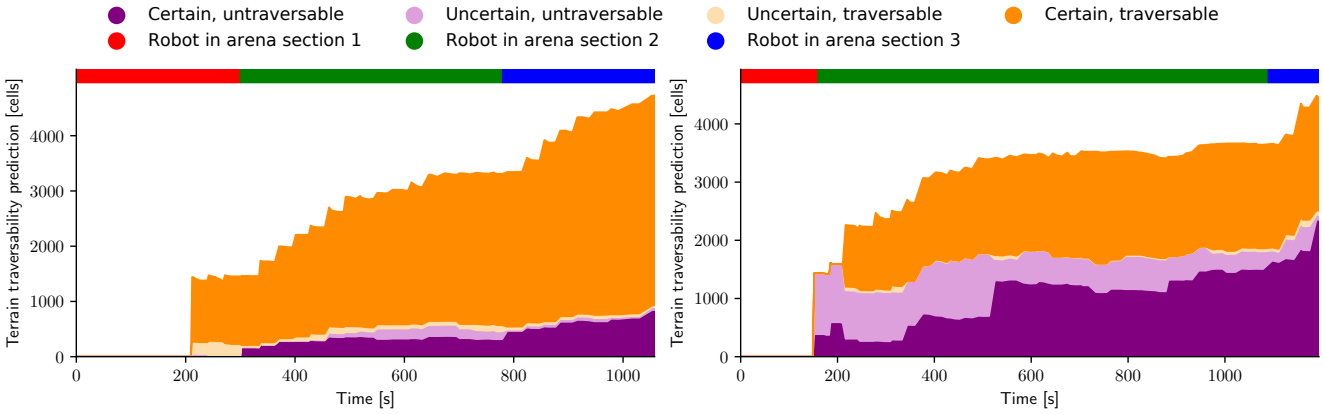

Fig. 11 Evolution of the predictions for potentially untraversable cells (a) in the main presented experimental run and (b) in the experimental run where the robot has sampled the traversable fabric first.

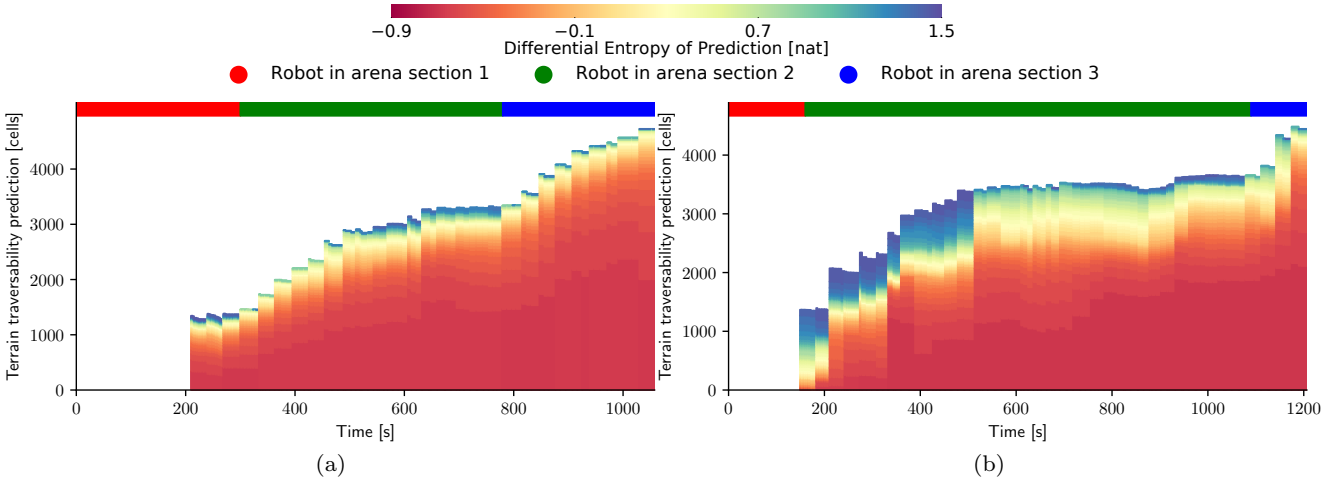

Fig. 12 Evolution of the predicted differential entropies for potentially untraversable cells in (a) the main presented experimental run; and in (b) the experimental run where the robot has sampled the traversable fabric first.

Table 5 Outdoor Scenario Predictions: set test mean **Prediction**, set test mean prediction **Entropy**, set test **RMSE**.

		Sparse Grass			**Dense Grass**			**Tree Trunk**			**Rock**		
Sequence 1	*Ground Truth*	1.106 ± 0.006			2.546 ± 0.377			13.869 ± 0.282			12.368 ± 0.480		
	Learning Step	**Prediction**	**Entropy**	**RMSE**	**Prediction**	**Entropy**	**RMSE**	**Prediction**	**Entropy**	**RMSE**	**Prediction**	**Entropy**	**RMSE**
	Sparse Grass	1.109	−0.552	0.007	1.111	0.902	1.484	1.107	1.418	12.765	1.109	1.398	11.270
	Dense Grass	1.224	−0.643	0.270	2.643	−0.534	0.423	1.837	1.410	12.035	2.268	1.350	10.108
	Tree Trunk	1.171	−0.643	0.196	2.711	−0.535	0.472	13.495	−0.484	0.749	7.394	1.334	4.999
	Rock	1.206	−0.647	0.277	2.707	−0.538	0.428	13.509	−0.488	0.779	12.316	−0.195	0.642
		Dense Grass			**Rock**			**Sparse Grass**			**Tree Trunk**		
Sequence 2	*Ground Truth*	2.543 ± 0.365			12.430 ± 0.510			1.110 ± 0.018			13.356 ± 0.714		
	Learning Step	**Prediction**	**Entropy**	**RMSE**	**Prediction**	**Entropy**	**RMSE**	**Prediction**	**Entropy**	**RMSE**	**Prediction**	**Entropy**	**RMSE**
	Dense Grass	2.612	−0.468	0.477	2.478	1.329	9.967	2.112	1.207	1.006	2.441	1.411	10.938
	Rock	2.641	−0.471	0.509	12.299	−0.216	0.671	3.937	1.201	2.983	7.499	1.389	5.893
	Sparse Grass	2.639	−0.517	0.509	12.245	−0.218	0.783	1.114	−0.653	0.105	4.994	1.388	8.400
	Tree Trunk	2.666	−0.517	0.542	12.402	−0.226	0.674	1.137	−0.653	0.228	13.566	−0.626	0.839
		Tree Trunk			**Rock**			**Sparse Grass**			**Dense Grass**		
Sequence 3	*Ground Truth*	13.789 ± 0.522			12.449 ± 0.534			1.106 ± 0.003			2.475 ± 0.360		
	Learning Step	**Prediction**	**Entropy**	**RMSE**	**Prediction**	**Entropy**	**RMSE**	**Prediction**	**Entropy**	**RMSE**	**Prediction**	**Entropy**	**RMSE**
	Tree Trunk	13.746	−0.306	0.742	14.026	1.393	1.641	14.013	1.413	12.909	14.238	1.393	11.769
	Rock	13.658	−0.325	0.652	12.844	−0.424	0.933	12.676	1.364	11.574	11.866	1.202	9.414
	Sparse Grass	13.400	−0.328	0.770	12.819	−0.432	0.889	1.265	−0.538	0.419	3.521	0.679	1.618
	Dense Grass	13.394	−0.333	0.789	12.827	−0.442	0.869	1.245	−0.548	0.375	2.534	−0.666	0.399

The ground truth values are reported including the standard deviation. Each sequence represents a particular order in which the individual training sets (terrains) are passed to the predictor, and the respective predictions represent values after learning the particular terrain in the sequence. The split into training and testing sets is done randomly and differs for each sequence. The highligted values signify terrains considered known since their prediction entropy is below $\mathrm{I}^{\mathrm{thr}}_{\mathrm{predict}}$.

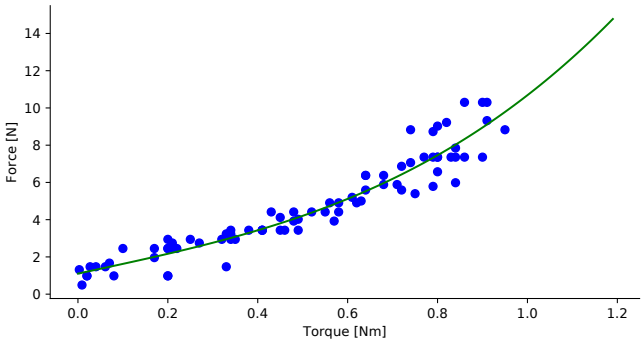

Fig. 13 Calibration data of the bumper sensor for the outdoor experiment. A cubic polynomial is fitted to the collected torque and force measurements.

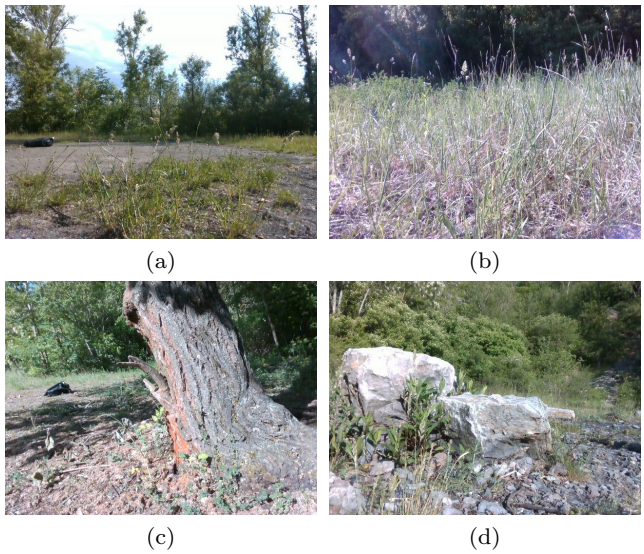

Fig. 14 The terrains used in the outdoor experiemnts: (a) sparse grass, (b) dense grass with hay, (c) tree trunk, (d) rocks.

from both the traversability predictor and the geometric model. The proposed system has been deployed in a fully autonomous experiment in an arena where the robot passed through an occluding non-rigid obstacle showing that the traversability properties have been successfully learned. Besides, we also show the performance of the predictor in an environment with real vegetation. The experimental results suggest that the robot successfully navigates an environment with non-rigid obstacles and chooses to explore areas that provide information for the rigidity and spatial models and can discriminate different natural terrains.

In the future, we aim to develop a unified framework to combine traversal costs of visually traversable terrains and apparent obstacles, thus adding a class of exploration goals. Furthermore, we intend to generalize the proposed robot-terrain interaction modeling to encompass traversable obstacles, rigid appearing terrains

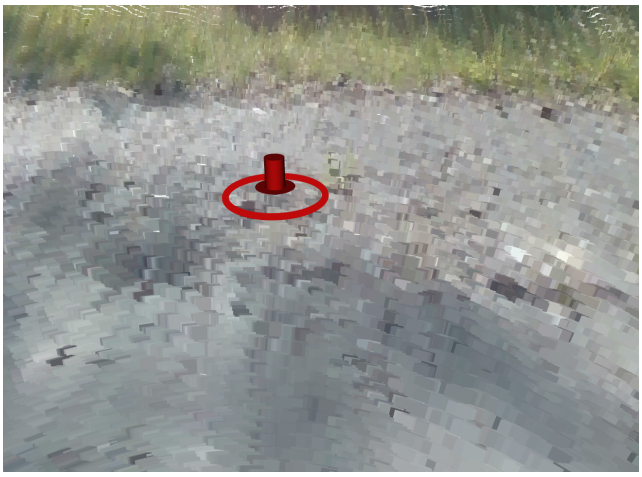

Fig. 15 The area (red circle) around a point of interest (red marker) signifying the cells used to compute the histogram terrain descriptor, projected on the elevation map with localized color measurements.

that do not support the robot, and terrains that change appearance after physical interaction. The exploration could also be focused on constraints such as the robot's battery capacity. Such extensions lead to deploying the proposed approach in non-myopic scenarios, where the robot considers its plans further than the immediate next navigation goal (waypoint). Hence, we aim to exploit multi-goal path planning (Faigl and Kulich, 2013) for such extended exploration scenarios.

Conflict of interest

The authors declare that they have no relevant financial or non-financial interests to disclose.

References

Baleia J, Santana P, Barata J (2015) On Exploiting Haptic Cues for Self-Supervised Learning of Depth-Based Robot Navigation Affordances. *Journal of Intelligent & Robotic Systems* 80(3–4):455–474, https://doi.org/10.1007/s10846-015-0184-4

Bayer J, Faigl J (2019) On Autonomous Spatial Exploration with Small Hexapod Walking Robot using Tracking Camera Intel RealSense T265. In: European Conference on Mobile Robots (ECMR), pp 1–6, https://doi.org/10.1109/ECMR.2019.8870968

Bayer J, Faigl J (2020) Speeded Up Elevation Map for Exploration of Large-Scale Subterranean Environments. In: 2019 Modelling and Simulation for Autonomous Systems (MESAS), pp 190–202, https://doi.org/10.1007/978-3-030-43890-6_15

Belter D, Wietrzykowski J, Skrzypczyński P (2019) Employing Natural Terrain Semantics in Motion Planning for a Multi-Legged Robot. *Journal of Intelligent & Robotic Systems* 93(3):723–743, https://doi.org/10.1007/s10846-018-0865-x

Biber P, Duckett T (2005) Dynamic Maps for Long-Term Operation of Mobile Service Robots. In: Robotics: Science and Systems (RSS), https://doi.org/10.15607/RSS.2005.I.003

Bourgault F, Makarenko AA, Williams SB, Grocholsky B, Durrant-Whyte HF (2002) Information based adaptive robotic exploration. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp 540–545, https://doi.org/10.1109/IRDS.2002.1041446

Bradley DM, Chang JK, Silver D, Powers M, Herman H, Rander P, Stentz A (2015) Scene Understanding for a High-mobility Walking Robot. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp 1144–1151, https://doi.org/10.1109/IROS.2015.7353514

Brown D, Webster G (2010) Now a stationary research platform, nasa's mars rover spirit starts a new chapter in red planet scientific studies. *NASA Press Release*

Brunner M, Brüggemann B, Schulz D (2013) Rough Terrain Motion Planning for Actuated, Tracked Robots. In: International Conference on Agents and Artificial Intelligence (ICAART), Springer, pp 40–61, https://doi.org/10.1007/978-3-662-44440-5_3

Burgard W, Cremers AB, Fox D, Hähnel D, Lakemeyer G, Schulz D, Steiner W, Thrun S (1999) Experiences with an interactive museum tour-guide robot. *Artificial Intelligence* 114(1):3 – 55, https://doi.org/10.1016/S0004-3702(99)00070-3

Carrillo H, Dames P, Kumar V, Castellanos JA (2018) Autonomous robotic exploration using a utility function based on Rényi's general theory of entropy. *Autonomous Robots* 42(2):235–256, https://doi.org/10.1007/s10514-017-9662-9

Cunningham C, Nesnas IA, Whittaker WL (2019) Improving slip prediction on Mars using thermal inertia measurements. *Autonomous Robots* 43(2):503–521, https://doi.org/10.1007/s10514-018-9796-4

Faigl J, Kulich M (2013) On Determination of Goal Candidates in Frontier-Based Multi-Robot Exploration. In: European Conference on Mobile Robots (ECMR), pp 210–215

Faigl J, Kulich M (2015) On Benchmarking of Frontier-Based Multi-Robot Exploration Strategies. In: European Conference on Mobile Robots (ECMR), pp 1–8, https://doi.org/10.1109/ECMR.2015.7324183

Faigl J, Čížek P (2019) Adaptive Locomotion Control of Hexapod Walking Robot for Traversing Rough Terrains with Position Feedback Only. *Robotics and Autonomous Systems* 116:136–147, https://doi.org/10.1016/j.robot.2019.03.008

Giguere P, Dudek G (2008) Clustering Sensor Data for Terrain Identification using a Windowless Algorithm. In: Robotics: Science and Systems (RSS), https://doi.org/10.15607/RSS.2008.IV.004

Gu J, Cao Q, Huang Y (2008) Rapid traversability assesment in 2.5 d grid based map on rough terrain. *International Journal of Advanced Robotic Systems* 5(4), https://doi.org/10.5772/6233

Halodová L, Dvořáková E, Majer F, Vintr T, Mozos OM, Dayoub F, Krajník T (2019) Predictive and adaptive maps for long-term visual navigation in changing environments. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp 7033–7039, https://doi.org/10.1109/IROS40897.2019.8967994

Homberger T, Bjelonic M, Kottege N, Borges PVK (2016) Terrain-dependant Control of Hexapod Robots using Vision. In: International Symposium on Experimental Robotics (ISER), Springer, pp 92–102, https://doi.org/10.1007/978-3-319-50115-4_9

Homberger T, Wellhausen L, Fankhauser P, Hutter M (2019) Support Surface Estimation for Legged Robots. In: IEEE International Conference on Robotics and Automation (ICRA), pp 8470–8476, https://doi.org/10.1109/ICRA.2019.8793646

Jadidi MG, Miro JV, Dissanayake G (2018) Gaussian processes autonomous mapping and exploration for range-sensing mobile robots. *Autonomous Robots* 42(2):273–290, https://doi.org/10.1007/s10514-017-9668-3

Kahn G, Abbeel P, Levine S (2021) BADGR: An Autonomous Self-Supervised Learning-Based Navigation System. *Robotics and Automation Letters* 6(2):1312–1319, https://doi.org/10.1109/LRA.2021.3057023

Karolj V, Viseras A, Merino L, Shutin D (2020) An Integrated Strategy for Autonomous Exploration of Spatial Processes in Unknown Environments. *Sensors* 20(13):3663

Kottege N, Parkinson C, Moghadam P, Elfes A, Singh SPN (2015) Energetics-informed hexapod gait transitions across terrains. In: IEEE International Conference on Robotics and Automation (ICRA), pp 5140–5147, https://doi.org/10.1109/ICRA.2015.7139915

Kragh M, Jørgensen RN, Pedersen H (2015) Object Detection and Terrain Classification in Agricultural Fields Using 3D Lidar Data. In: International Conference on Computer Vision Systems (ICVS), pp 188–

197, https://doi.org/10.1007/978-3-319-20904-3_18

Krüsi P, Bosse M, Siegwart R (2016) Driving on Point Clouds: Motion Planning, Trajectory Optimization, and Terrain Assessment in Generic Nonplanar Environments. *Journal of Field Robotics* 34(5):940–984, https://doi.org/10.1002/rob.21700

Lalonde JF, Vandapel N, Huber DF, Hebert M (2006) Natural terrain classification using three-dimensional ladar data for ground robot mobility. *Journal of Field Robotics* 23(10):839–861, https://doi.org/10.1002/rob.20134

Luo W, Sycara K (2018) Adaptive Sampling and Online Learning in Multi-Robot Sensor Coverage with Mixture of Gaussian Processes. In: IEEE International Conference on Robotics and Automation (ICRA), pp 6359–6364, https://doi.org/10.1109/ICRA.2018.8460473

Makarenko AA, Williams SB, Bourgault F, Durrant-Whyte HF (2002) An experiment in integrated exploration. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), vol 1, pp 534–539 vol.1, https://doi.org/10.1109/IRDS.2002.1041445

McGhee RB, Frank AA (1968) On the stability properties of quadruped creeping gaits. *Mathematical Biosciences* 3:331–351, https://doi.org/10.1016/0025-5564(68)90090-4

Moravec H, Elfes A (1985) High resolution maps from wide angle sonar. In: IEEE International Conference on Robotics and Automation (ICRA), vol 2, pp 116–121, https://doi.org/10.1109/ROBOT.1985.1087316

Otsu K, Ono M, Fuchs TJ, Baldwin I, Kubota T (2016) Autonomous Terrain Classification with Co- and Self-Training Approach. *Robotics and Automation Letters* 1(2):1–6, https://doi.org/10.1109/LRA.2016.2525040

O'Meadhra C, Tabib W, Michael N (2019) Variable Resolution Occupancy Mapping Using Gaussian Mixture Models. *Robotics and Automation Letters* 4(2):2015–2022, https://doi.org/10.1109/LRA.2018.2889348

Papadakis P (2013) Terrain traversability analysis methods for unmanned ground vehicles: A survey. *Engineering Applications of Artificial Intelligence* 26(4):1373–1385, https://doi.org/10.1016/j.engappai.2013.01.006

Pearson MJ, Dora S, Struckmeier O, Knowles TC, Mitchinson B, Tiwari K, Kyrki V, Bohte S, Pennartz CM (2021) Multimodal Representation Learning for Place Recognition Using Deep Hebbian Predictive Coding. *Frontiers in Robotics and AI* 8

Petrou Z, Manakos I, Stathaki T, Mücher C, Adamo M (2015) Discrimination of Vegetation Height Categories With Passive Satellite Sensor Imagery Using Texture Analysis. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 8(4):1442–1455, https://doi.org/10.1109/JSTARS.2015.2409131

Prágr M, Faigl J (2019) Benchmarking Incremental Regressors in Traversal Cost Assessment. In: International Conference on Artificial Neural Networks (ICANN), pp 685–697, https://doi.org/10.1007/978-3-030-30487-4_52

Prágr M, Čížek P, Faigl J (2018a) Incremental Learning of Traversability Cost for Aerial Reconnaissance Support to Ground Units. In: Modelling and Simulation for Autonomous Systems (MESAS), Springer, https://doi.org/10.1007/978-3-030-14984-0_30

Prágr M, Čížek P, Faigl J (2018b) Cost of Transport Estimation for Legged Robot Based on Terrain Features Inference from Aerial Scan. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp 1745–1750, https://doi.org/10.1109/IROS.2018.8593374

Prágr M, Čížek P, Bayer J, Faigl J (2019) Online Incremental Learning of the Terrain Traversal Cost in Autonomous Exploration. In: Robotics: Science and Systems (RSS), https://doi.org/10.15607/RSS.2019.XV.040

Prágr M, Váňa P, Faigl J (2020) Aerial Reconnaissance and Ground Robot Terrain Learning in Traversal Cost Assessment. In: 2019 Modelling and Simulation for Autonomous Systems (MESAS), pp 3–10, https://doi.org/10.1007/978-3-030-43890-6_1

Quigley M, Conley K, Gerkey BP, Faust J, Foote T, Leibs J, Wheeler R, Ng AY (2009) ROS: an open-source robot operating system. In: ICRA Workshop on Open Source Software, pp 1–6

Ramos F, Ott L (2016) Hilbert maps: Scalable continuous occupancy mapping with stochastic gradient descent. *International Journal of Robotics Research* 35(14):1717–1730, https://doi.org/10.1177/0278364916684382

Rasmussen CE, Williams CKI (2006) *Gaussian processes for machine learning*. Adaptive computation and machine learning, MIT Press, Cambridge, Mass

Rothrock B, Kennedy R, Cunningham C, Papon J, Heverly M, Ono M (2016) SPOC: Deep Learning-based Terrain Classification for Mars Rover Missions. In: AIAA SPACE 2016, American Institute of Aeronautics and Astronautics, https://doi.org/10.2514/6.2016-5539

Ruiz AV, Olariu C (2015) A general algorithm for exploration with Gaussian processes in complex, unknown environments. In: IEEE International Conference on Robotics and Automation (ICRA), pp 3388–

3393, https://doi.org/10.1109/ICRA.2015.7139667
Rényi A (1961) On Measures of Entropy and Information. In: Berkeley Symposium on Mathematical Statistics and Probability, vol 1, pp 547–561
Singh A, Krause A, Guestrin C, Kaiser W, Batalin M (2007) Efficient planning of informative paths for multiple robots. In: International Joint Conference on Artifical Intelligence, pp 2204–2211
Sofman B, Lin E, Bagnell JA, Cole J, Vandapel N, Stentz A (2006) Improving Robot Navigation Through Self-Supervised Online Learning. *Journal of Field Robotics* 23(11-12):1059–1075, https://doi.org/10.1002/rob.20169
Srinivas N, Krause A, Kakade S, Seeger M (2010) Gaussian process optimization in the bandit setting: no regret and experimental design. In: Intl. Conf. International Conference on Machine Learning (ICML), Haifa, Israel, pp 1015–1022
Stachniss C, Grisetti G, Burgard W (2005) Information Gain-based Exploration Using Rao-Blackwellized Particle Filters. In: Robotics: Science and Systems (RSS), Robotics: Science and Systems Foundation, https://doi.org/10.15607/RSS.2005.I.009
Stelzer A, Hirschmüller H, Görner M (2012) Stereo-vision-based navigation of a six-legged walking robot in unknown rough terrain. *International Journal of Robotics Research* 31(4):381–402, https://doi.org/10.1177/0278364911435161
Tabib W, Goel K, Yao J, Dabhi M, Boirum C, Michael N (2019) Real-Time Information-Theoretic Exploration with Gaussian Mixture Model Maps. In: Robotics: Science and Systems (RSS), vol 15, https://doi.org/10.15607/RSS.2019.XV.061
Tiwari K, Xiao X, Malik A, Chong NY (2019) A unified framework for operational range estimation of mobile robots operating on a single discharge to avoid complete immobilization. *Mechatronics* 57:173–187
Tucker VA (1975) The Energetic Cost of Moving About: Walking and running are extremely inefficient forms of locomotion. Much greater efficiency is achieved by birds, fish—and bicyclists. *American Scientist* 63(4):413–419
Vasudevan S, Ramos F, Nettleton E, Durrant-Whyte H, Blair A (2009) Gaussian Process modeling of large scale terrain. In: IEEE International Conference on Robotics and Automation (ICRA), IEEE, pp 1047–1053, https://doi.org/10.1002/rob.20309
Yamauchi B (1997) A frontier-based approach for autonomous exploration. In: CIRA, IEEE, pp 146–151, https://doi.org/10.1109/CIRA.1997.613851
Ünsalan C, Boyer KL (2004) Linearized vegetation indices based on a formal statistical framework. *Transactions on Geoscience and Remote Sensing* 42(7):1575–1585, https://doi.org/10.1109/TGRS.2004.826787