Traversability Transfer Learning Between Robots with Different Cost Assessment Policies

Josef Zelinka[✉] , Miloš Prágr , Rudolf Szadkowski , Jan Bayer ,
and Jan Faigl

Faculty of Electrical Engineering, Czech Technical University in Prague,
Technická 2, 166 27 Prague, Czech Republic
{zelinjo1,pragrmi1,szadkrud,bayerja1,faiglj}@fel.cvut.cz
https://comrob.fel.cvut.cz

Abstract. Predicting mobile robots' traversability over terrains is crucial to select safe and efficient paths through rough and unstructured environments. In multi-robot missions, knowledge transfer techniques can enable learning terrain traversability assessment the robots did not experience individually. The knowledge can be incrementally aggregated for homogeneous robots since they can treat foreign knowledge as their own. However, robots with different perceptions might experience the same terrain differently, so it is impossible to aggregate the shared knowledge directly. In this paper, we show how to learn a model that transfers the experience between heterogeneous robots, enabling each robot to use the whole sum of the experience of the multi-robot team. The proposed approach uses correlation to combine individual neural networks that assess the traversability of individual robots. The presented method has been verified in a real-world deployment of multi-legged walking robots with different cost assessment policies.

Keywords: Transfer learning · Traversability assessment · Mobile robot

1 Introduction

The studied transfer learning is motivated by deployments of a heterogeneous team of multi-legged walking robots, each exploring and perceiving various terrain types. It is desirable to explore as quickly and efficiently as possible during terrain exploration. As the robots are deployed, they collect a large amount of information about the environment and experience the traversability cost of the traversed terrain. The collected traversability information can be encoded into a traversability cost model that can assess the cost; however, such a quality assessment is limited to the experience collected by the individual. A group of cooperating robots can improve their performance on a given task by sharing their experiences. The knowledge of each robot can be enhanced by sharing the

obtained knowledge among the robots. Hence, it can enable the team to improve its overall performance. Motivated by groups of social animals learning experiences from one another [17], we aim to implement such transfer learning patterns into multi-robot systems.

Furthermore, missions such as exploration of unknown environments can be speeded up by parallelization of the exploration process using a large group of robots [23]. Thus, having multiple robots, we can explore transferring the collected knowledge between the robots. For homogeneous teams with a single robot type, the knowledge transfer that is called inductive transfer learning [14] can be utilized [21]. Such a transfer is possible only when all robots have the same morphology and sensor equipment; otherwise, the homogeneity of the team is lost. Changes in morphology or equipment can lead to variation in the terrain perception. However, changes of the identical robots can be caused by damage to the robot during the mission or hardware updates in later operational deployments. With inductive transfer learning, changed robots would not contribute to the shared knowledge nor benefit from it. In that perspective, heterogeneity seems to be natural, highlighting the importance of transfer learning in multi-robot heterogeneous team [13].

Fig. 1. The used hexapod walking robot in the experimental deployment of the proposed method in the Bull Rock Cave, where it builds elevation map and collects the dataset.

We propose a transfer learning method to share knowledge among heterogeneous robots to enhance their cost assessment capabilities. In the proposed approach, the robots transfer their individually learned models implemented as the Convolutional Neural Network (CNNs) regressor. For knowledge transfer between two different robots, correlation of predictions on terrains traversed by both robots is used to determine the relationship between the models, thus creating an augmented model. After the relation between models is learned, the robots are ready to exchange the traversal experience they have already collected. The proposed approach has been experimentally verified on data from the deployment of a real hexapod walking robot with adaptive locomotion control [5] in a natural cave system, see Fig. 1. Based on the achieved results, the proposed method allows two robots to share the knowledge and exploit the traversability cost models experienced by the other robot.

The paper is organized as follows. Section 2 summarizes related work with the emphasis on traversability assessment and transfer learning techniques deployed in robotics. The proposed transfer learning method is introduced in Sect. 3. In Sect. 4, we report on the experimental results of the proposed method using real hexapod walking robots. The paper is concluded in Sect. 5.

2   Related Work

Traversability assessment is studied in various fields such as planetary exploration [6,20], search and rescue missions [3], and agriculture, or off-road driving [8]. Two main classes of approaches can be identified in the literature: traversability classification and prediction of traversability cost as a continuous score. The simplest terrain classification can be a binary classifier to determine whether the terrain is traversable or not [10]. However, the authors of [6] report improved path planning results avoiding impassable terrain and also better-optimized paths using a continuous score. Therefore, in this paper, we follow the idea of traversability as a continuous score.

The traversability assessment can be based on proprioceptive and exteroceptive sensory signals, where the exteroceptive data processing approach can be further categorized into geometry- or appearance-based [15]. Nevertheless, hybrid methods might benefit from combined approaches. The rest of this section provides an overview of the most related traversability approaches to support our traversability assessment choices.

Proprioceptive traversability assessment uses information captured by sensors that measure the robot's internal properties during the robot's interaction with the terrain, e.g., speed, tilt, shakiness, energy, or vibration. Thus, the proprioceptive traversability assessment can estimate traversability only on currently traversed terrains. An example of traversability assessment based on the energy expenditure is reported in [12].

The exteroceptive, geometry-based approaches use range measurements such as LiDARs and RGB-D cameras to construct maps of the perceived environment. The maps are then used to examine terrain properties such as roughness, edges,

slope, or features the robot might not be able to traverse. Obstacle extractions from the maps using filtering and clustering are presented in [16]. On the other hand, the visual appearance of the terrain can be studied using image-processing and classification of terrain types into categories with defined properties [1]. Methods using appearance and geometrical properties might suffer from wrongly classified terrains in cases where range sensing is not sufficient, e.g., unexpected covered hole [22]. Therefore, we have chosen hybrid approaches to leverage the advantages of the individual methods.

In addition to traversability assessment, transfer learning is assumed to improve the assessment by exploiting individual experiences of the particular robots in a team. Transfer learning can be defined as a machine learning approach to boost the knowledge in the target domain by the transfer from the source domain [24]. Transfer learning is already established technique in the fields such as text [7] and image [18] classification. In [24], the authors combine text and image classification using the matrix factorization method to enhance image classification by information extracted from their annotations, thus merging the two tasks.

Similar to text and image classification, robotics is a domain where labeled data are costly to obtain. Besides, it is relatively hard to train robots to adapt to the demands of various environments. The knowledge transfer is a way to benefit from deployments of multi-robot teams. The authors of [4] adopt transfer learning to reduce the learning time of the particle swarm optimization for faster optimization of robot's gaits (walking patterns). Transfer learning applied in the learning of humanoid robots to solve tasks by observing human behavior is described in [11]. The idea is to transfer knowledge about a human motion to the robot that is requested to perform a similar motion. In [19], learned navigation patterns around obstacles are transferred into new environments to enhance planning capabilities.

The aforementioned transfer learning approaches in the robotics domain provide supportive evidence of successfully deployed techniques. Therefore, we focus on deploying transfer learning among heterogeneous robots that might yield different traversability assessments [9].

3   Method

The proposed method for transferring knowledge from one robot to another is motivated to improve cost assessment capabilities by learning from one another. Two roles of the robots can be distinguished: a provider of new information called teacher T; and receiving robot denoted student S. Each i-th robot collects a dataset $D^i = \{(t_1^i, c_1^i), (t_2^i, c_2^i), \dots\}$ that consists of features (t_j^i, c_j^i) describing the perceived terrain $t_j^i \in \mathcal{T}$ and labels describing the cost of traversing a particular segment of the terrain $c_j^i \in C$. Hence, observations of the i-th robot are stored in the dataset D^i that represents the robot's experience with the environment.

In this case, the two robots are in the roles of teacher and student, respectively, and the experience with the same terrain, i.e., $t^S = t^T$, the observed costs

might not necessarily be equal, $c^S \neq c^T$, because of different terrain perception. The proposed approach targets to extrapolate individual datasets on newly observed terrains. The extrapolation is realized through analyzing the relation between the cost assessment of the student and the teacher. The relation is then used to enrich the student's extrapolation capabilities by the teacher's cost assessment.

The remainder of this section describes the proposed method for transferring knowledge between the robots with heterogeneous terrain perceptions. First, the proposed cost assessment learning model is introduced in Sect. 3.1, together with the training procedure. The transfer learning framework is then introduced in Sect. 3.2.

3.1 Cost Assessment Learning Model

The individual robot's cost assessment model $M = (r, a)$ is trained using its dataset D. The model comprises the regressor $r : \mathcal{T} \rightarrow C$, which returns the cost estimation, and certainty evaluation $a : \mathcal{T} \rightarrow U$, where $u \in U = \mathbb{R}$ denotes the certainty of the model over the particular terrain segment $t \in \mathcal{T}$.
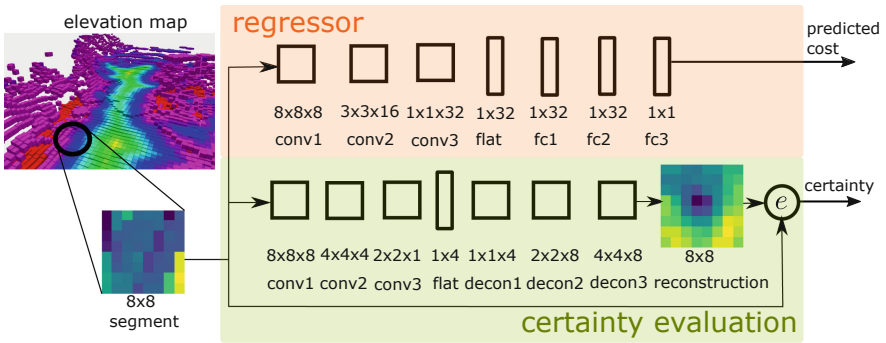
Fig. 2. The elevation map is segmented and sent into the model M. The model M is composed of a regressor (orange) providing the cost prediction r and autoencoder (olive) from which the certainty a is computed. The input of the model is an 8×8 segment of elevation map, which is processed by the neural networks. The depicted architecture of the neural networks shows the convolutional (conv), flattening (flat), deconvolution (decon), and fully connected (fc) layers with their dimensions. For certainty evaluation, the log of reconstruction error $\log(e(t; g))$ computation is indicated by the e-node. (Color figure online)

Both the functions r and a are implemented with separate neural networks, where the terrain t is represented as a set of elevation map segments, each encoded by $n \times n$ matrix of real values, $t \in \mathcal{T} = \mathbb{R}^{n \times n}$. Topological dependencies between the matrix values are the same as in images; therefore, convolutional layers are used similarly to image processing. Thus, for processing the elevation

map segments, convolutional layers are added to the neural networks for the regressor r and certainty evaluator a. The learning architecture is depicted in Fig. 2.

The certainty evaluation a is trained indirectly with a convolutional autoencoder. The autoencoder $g : \mathcal{T} \rightarrow \mathcal{T}$ maps given segments t to reconstructed segments $g(t)$, where the reconstruction error $e(t; g) = ||t - g(t)||$ is minimized during the training. Trivially, the reconstruction error would be zero for all $t \in \mathbb{R}^{n \times n}$ if the map g is an identity function. However, due to the bottleneck architecture of the autoencoder, the map g cannot be an identity function; so, the segments have different reconstruction errors. Here, we assume the trained autoencoder has a low reconstruction error on segments presented in the dataset and a higher reconstruction error for other segments. The certainty of the model is thus represented by the log of the reconstruction error $a(t) = \log(e(t; g))$, where higher values correspond to the terrain segments that are dissimilar to segments the model has been trained on.

3.2 Transfer Learning Framework

The proposed framework uses trained models M_S and M_T where the student uses the teacher's knowledge by considering the relation κ. Similar features are used to obtain comparable predictions to determine the relationship between the models. We assume that if the teacher and the student traversed through the same region, the features collected in that region are similar for the teacher and the student. Additionally, both models should be certain about the previous observation of the terrain with similar (if not equal) certainty. Hence, at least one similar terrain observation T_{sim}, where both robots are certain about the previous observation of the terrain, is necessary to learn the relation κ successfully.

Using a set of similar terrain observations T_{sim} containing n samples of terrain observations, predictions of M_S and M_T models about the cost $C = (c_i)_{i=1}^{n}$ and certainties $U = (u_i)_{i=1}^{n}$ are obtained. The indicator of the certainty $A = (\alpha_i)_{i=1}^{n}$ is created as

$$\alpha_i = \begin{cases} 0 & \text{if } u_i^S < \theta \vee u_i^T < \theta \\ 1 & \text{otherwise} \end{cases}. \tag{1}$$

The individual indicator is zero for samples where one robot's certainty about the terrain observation is less than the empirically set threshold θ. The relation κ between student's and teacher's cost assessment models is determined by the average of the oriented differences between corresponding cost predictions $c \in C$

$$\kappa = \frac{1}{\sum_{i=1}^{n} \alpha_i} \sum_{i=1}^{n} \alpha_i (c_i^S - c_i^T), \tag{2}$$

where the certainty indicator α_i is used to remove samples where the robot model is not certain enough.

The obtained relation κ is used to enhance student's future predictions about the newly observed terrain cost c^p to facilitate better path planning decisions. With the next terrain observation t_{new}, both models M_S and M_T of the student and teacher, respectively, are used to predict the cost and certainty $(c, u) = M(t_{new})$. Then, the certainties are compared and the prediction with the higher certainty is selected. If the teacher's prediction is selected, the cost c^T is corrected by the relation between the models κ as

$$c^p = \begin{cases} c^S, & if \ u^S > u^T \\ c^T + \kappa, & \text{otherwise} \end{cases}. \tag{3}$$

The feasibility of the proposed cost models and transfer learning framework has been empirically validated using real datasets. The achieved results are reported in the following section.

4   Results

The proposed method has been verified in an experimental deployment using a real hexapod walking robot shown in Fig. 1. The robot is equipped with the Intel RealSense D435 RGB-D and T265 tracking cameras, and terrain's features are stored into an elevation map [2]. During the deployment, the robot collects datasets further used in the model learning, knowledge transfer, and evaluation of the learned traversability cost assessment models. In particular, the used dataset has been collected in the Bull Rock Cave, Czech Republic, and the proposed method has been evaluated as follows.

Various cost perceptions are simulated using different cost calculation methods instead of deploying heterogeneous robots. The student's costs are computed as an angular distance of the pitch and roll from the leveled position of the robot. On the other hand, the teacher's costs are computed as the robot's relative slowdown compared with the commanded velocity v_{cmd}, which characterizes the difficulty of the terrain as a resistance difference from regular walking. An individual terrain segment i holds the information about the robot's state changes between two consecutive feature collection places. The cost of the i-th segment is defined as the median over multiple traversed consecutive segments

$$c^T = \text{median}\{c_{sk}\}_{k=1}^n \tag{4}$$

for $c_{si} = v_{cmd}\Delta t_i / s_i$ with the segment duration Δt_i and length s_i.

The student and teacher models are trained as described in Sect. 3.1. The cost regressor is trained using 2000 epochs, and the autoencoder is trained for 100 epochs. It is presumed that the autoencoder is uncertain in the previously unobserved terrains. Therefore, the model benefits from overfitted autoencoder. The architecture of the regressor and autoencoder neural networks is as in Fig. 2, and for each layer, ReLU activation functions are utilized.

Datasets have been collected in three parts of the Bull Rock Cave with different traversability properties. The student's model is trained on data collected

(a) Chiffon   (b) Room   (c) Hall
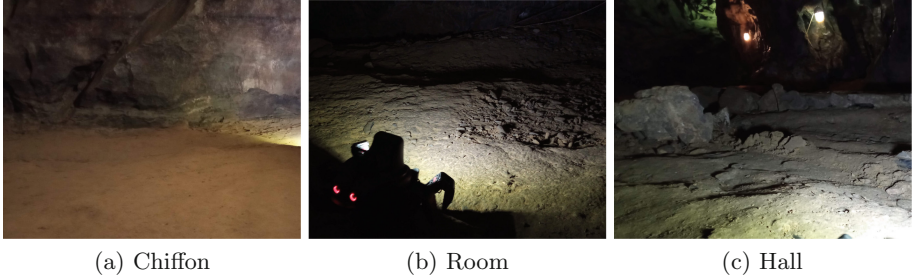
Fig. 3. Different terrains of the Bull Rock Cave used in the evaluation of the proposed method.

in the *Chiffon* and *Hall* parts of the cave, while the teacher's dataset is collected in the *Room* and *Hall* parts. Terrains in both *Room* and *Hall* consist of similar leveled, packed surfaces. In *Chiffon*, the robot has experienced a slightly sandy surface, which makes the robot's movement marginally slower due to its legs sinking into the sand during motion. The visual appearances of the terrains are displayed in Fig. 3.

Since both the teacher's and the student's models are trained on datasets collected on the *Hall* terrain, both models should have certain predictions for the respective terrain, which is ideal for demonstration of the transfer learning. Hence, the *Hall* dataset is selected as the training dataset for the transfer relation κ with the uncertainty threshold set to $\theta = 3$ that has been found empirically. The relation between the models is determined to be $\kappa = -0.52$, indicating that the teacher's cost assessments are, on average, by 0.52 higher than the student's assessments. Therefore, 0.52 is subtracted whenever the student uses the teacher's prediction.

The *Room* dataset is selected as the testing environment for the scenario, where the teacher's knowledge enhances student's predictions. In this setup, the teacher should be able to make better cost predictions than the student because the teacher previously observed the terrain in the *Room*. The results presented in Fig. 4 show that during the evaluation phase, the teacher's cost assessments are used more often because the teacher is more certain about the terrain sample. In most cases, transfer learning improved the cost assessment. The values of the *Mean Square Error* (MSE) using the transfer learning are depicted in Table 1, where we can also observe the negative transfer.

Only the student is trained in *Chiffon*, and therefore, the student accumulated better knowledge about the *Chiffon* terrain. From Fig. 4, we can observe that the teacher's and student's predictions are used almost equally by the transfer learning component. The comparable prediction usage might be caused by the fact that the flat terrain in *Chiffon* is partially similar to the teacher's training domain of *Hall* and *Room*. However, teacher's cost assessments rarely improve the student's knowledge, thus resulting in a decrease in cost assessment capabilities using the transfer learning in this scenario, which is further shown in Table 1.
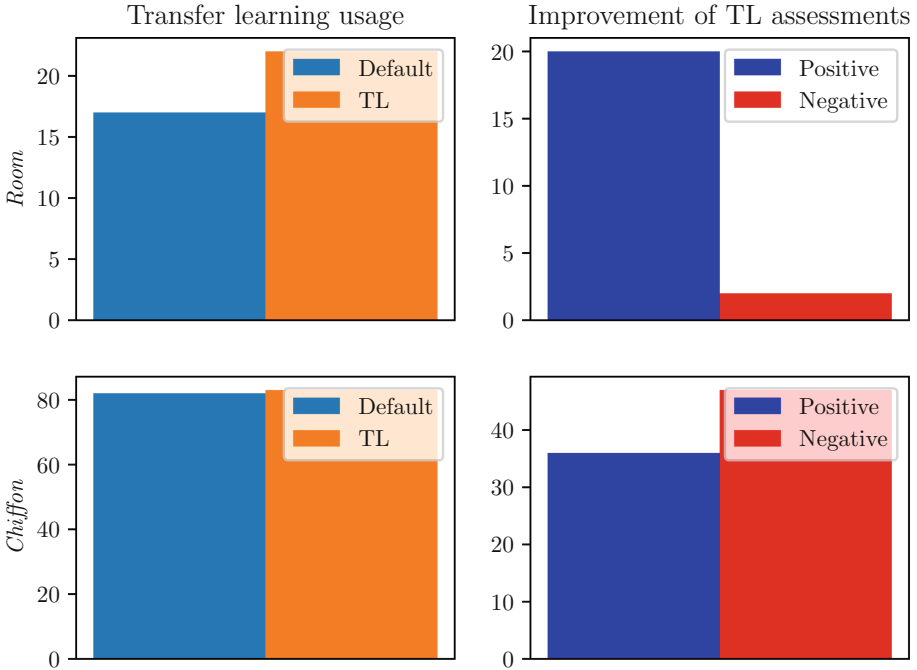
Fig. 4. Results of the transfer learning performed in *Room* (top row) and *Chiffon* (bottom row). The left column shows the amounts of usage of the transfer learning cost assessments versus the default estimations made by the student's model. The numbers of positive and negative improvements achieved by transfer learning assessments compared to the student's estimation are illustrated in the right column.

The sparse improvement of the cost assessment using the transfer learning is likely caused by the sandy surface being fairly similar to the packed surfaces in *Hall* and *Room*. Note that the real cost of moving over the sandy terrain is higher than on packed surfaces.

Discussion – The results show that the knowledge transfer from the teacher improved the student's ability to assess cost in the case of the terrain previously observed by the teacher. The positive transfer learning is represented by the transfer over the *Room* terrain. However, there can be confusion in the model selection during the transfer learning phase, which can be observed for the *Chiffon* scenario. Nevertheless, in *Chiffon*, both models performed similarly well, producing lower MSE than in *Room*. The negative transfer might be solved by making the transfer learning component more strict and enforcing stricter requirements on the teacher's estimation.

Table 1. Mean square errors of the predictions compared to the ground truth and percentage improvement of the transfer learning model.

Scenario	Positive transfer	Negative transfer
	Room	*Chiffon*
Default (MSE)	6.19	2.70
Transfer (MSE)	4.10	2.98
Percentage improvement	33.76%	-10.00%

5   Conclusion

In this paper, we show a method to accomplish transfer learning across robots with heterogeneous terrain perception. The proposed transfer learning approach is based on creating cost assessment models for the individual robots using convolutional neural networks. The cost assessment models are then used to estimate cost and uncertainty for terrains observed after the models are created. The transfer learning is addressed by an augmented model created using a correlation between the students and teachers models established on a training terrain dataset. The feasibility of the proposed approach is validated on experimental data collected in real cave terrains. The results indicate that the approach is viable, and terrain cost assessment can be improved by transfer learning. Different cost assessment models are used to simulate heterogeneous robots. In our future work, we plan to experimentally evaluate the method using robots with different morphology and sensory equipment. Besides, we also aim to deploy the learning method directly during the exploration task. It is expected to improve the individual robot's performance by avoiding hard to traverse areas experienced by the other robots.

References

1. Ball, D., et al.: Vision-based obstacle detection and navigation for an agricultural robot. J. Field Robot. **33**, 1107–1130 (2016). https://doi.org/10.1002/rob.21644
2. Bayer, J., Faigl, J.: Speeded up elevation map for exploration of large-scale subterranean environments. In: Mazal, J., Fagiolini, A., Vasik, P. (eds.) MESAS 2019. LNCS, vol. 11995, pp. 190–202. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-43890-6_15
3. Cafaro, B., Gianni, M., Pirri, F., Ruiz, M., Sinha, A.: Terrain traversability in rescue environment. In: IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR), pp. 1–8 (2013). https://doi.org/10.1109/SSRR.2013.6719358

4. Degrave, J., Burm, M., Kindermans, P.J., Dambre, J., Wyffels, F.: Transfer learning of gaits on a quadrupedal robot. Adapt. Behav. **23**(2), 69–82 (2015). https://doi.org/10.1177/1059712314563620

5. Faigl, J., Čížek, P.: Adaptive locomotion control of hexapod walking robot for traversing rough terrains with position feedback only. Robot. Auton. Syst. **116**, 136–147 (2019). https://doi.org/10.1016/j.robot.2019.03.008

6. Gennery, D.B.: Traversability analysis and path planning for a planetary rover. Auton. Robot. **6**, 131–146 (1999). https://doi.org/10.1023/A:1008831426966

7. Howard, J., Ruder, S.: Universal language model fine-tuning for text classification. In: 56th Annual Meeting of the Association for Computational Linguistics (ACL), pp. 328–339 (2018). https://doi.org/10.18653/v1/P18-1031

8. Huertas, A., Matthies, L., Rankin, A.: Stereo-based tree traversability analysis for autonomous off-road navigation. In: IEEE Workshops on Applications of Computer Vision (WACV/MOTION 2005), pp. 210–217 (2005). https://doi.org/10.1109/ACVMOT.2005.111

9. Kolvenbach, H., Bellicoso, D., Jenelten, F., Wellhausen, L., Hutter, M.: Efficient gait selection for quadrupedal robots on the moon and mars. In: 14th International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS) (2018). https://doi.org/10.3929/ethz-b-000261939

10. Langer, D., Rosenblatt, J.K., Hebert, M.: A behavior-based system for off-road navigation. IEEE Trans. Robot. Autom. **10**(6), 776–783 (1994). https://doi.org/10.1109/70.338532

11. Makondo, N., Hiratsuka, M., Rosman, B., Hasegawa, O.: A non-linear manifold alignment approach to robot learning from demonstrations. J. Robot. Mechatron. **30**(2), 265–281 (2018). https://doi.org/10.20965/JRM.2018.P0265

12. Martin, S., Corke, P.: Long-term exploration & tours for energy constrained robots with online proprioceptive traversability estimation. In: IEEE International Conference on Robotics and Automation (ICRA), pp. 5778–5785 (2014). https://doi.org/10.1109/ICRA.2014.6907708

13. Murphy, R.R., et al.: Use of remotely operated marine vehicles at Minamisanriku and Rikuzentakata Japan for disaster recovery. In: IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR), pp. 19–25 (2011). https://doi.org/10.1109/SSRR.2011.6106798

14. Pan, S.J., Yang, Q.: A survey on transfer learning. IEEE Trans. Knowl. Data Eng. **22**(10), 1345–1359 (2010). https://doi.org/10.1109/TKDE.2009.191

15. Papadakis, P.: Terrain traversability analysis methods for unmanned ground vehicles: a survey. Eng. Appl. Artif. Intell. **26**(4), 1373–1385 (2013). https://doi.org/10.1016/J.ENGAPPAI.2013.01.006

16. Peng, Y., Qu, D., Zhong, Y., Xie, S., Luo, J., Gu, J.: The obstacle detection and obstacle avoidance algorithm based on 2-D lidar. In: IEEE International Conference on Information and Automation, pp. 1648–1653 (2015). https://doi.org/10.1109/ICInfA.2015.7279550

17. Pongrácz, P., Miklósi, Á., Kubinyi, E., Gurobi, K., Topál, J., Csányi, V.: Social learning in dogs: the effect of a human demonstrator on the performance of dogs in a detour task. Anim. Behav. **62**(6), 1109–1117 (2001). https://doi.org/10.1006/ANBE.2001.1866

18. Quattoni, A., Collins, M., Darrell, T.: Transfer learning for image classification with sparse prototype representations. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1–8 (2008). https://doi.org/10.1109/CVPR.2008.4587637

19. Saha, O., Dasgupta, P., Woosley, B.: Real-time robot path planning from simple to complex obstacle patterns via transfer learning of options. Auton. Robot. **43**(8), 2071–2093 (2019). https://doi.org/10.1007/S10514-019-09852-5
20. Singh, S., et al.: Recent progress in local and global traversability for planetary rovers. In: IEEE International Conference on Robotics and Automation (ICRA), vol. 2, pp. 1194–1200 (2000). https://doi.org/10.1109/ROBOT.2000.844761
21. Stone, P., Veloso, M.: Multiagent systems: a survey from a machine learning perspective. Auton. Robot. **8**, 345–383 (2000). https://doi.org/10.1023/A:1008942012299
22. Tennakoon, E., Peynot, T., Roberts, J., Kottege, N.: Probe-before-step walking strategy for multi-legged robots on terrain with risk of collapse. In: IEEE International Conference on Robotics and Automation (ICRA), pp. 5530–5536 (2020). https://doi.org/10.1109/ICRA40945.2020.9197154
23. Tribelhorn, B., Dodds, Z.: Evaluating the Roomba: a low-cost, ubiquitous platform for robotics research and education. In: IEEE International Conference on Robotics and Automation (ICRA), pp. 1393–1399 (2007). https://doi.org/10.1109/ROBOT.2007.363179
24. Zhuang, F., et al.: A comprehensive survey on transfer learning. Proc. IEEE **109**(1), 43–76 (2021). https://doi.org/10.1109/JPROC.2020.3004555