Terrain Traversal Cost Learning with Knowledge Transfer Between Multi-legged Walking Robot Gaits

Miloš Prágr, Rudolf Szadkowski, Jan Bayer, Josef Zelinka, Jan Faigl
Czech Technical University, Faculty of Electrical Engineering, Technicka 2, 166 27, Prague, Czech Republic
Email: pragrmi1@fel.cvut.cz

***Abstract*—The terrain traversal abilities of multi-legged walking robots are affected by gaits, the walking patterns that enable adaptation to various operational environments. Fast and low-set gaits are suited to flat ground, while cautious and high-set gaits enable traversing rough areas. A suitable gait can be selected using prior experience with a particular terrain type. However, experience alone is insufficient in practical setups, where the robot experiences each terrain with only one or just a few gaits and thus would infer novel gait-terrain interactions from insufficient data. Therefore, we use knowledge transfer to address unsampled gait-terrain interactions and infer the traversal cost for every gait. The proposed solution combines gait-terrain cost models using inferred gait-to-gait models projecting the robot experiences between different gaits. We implement the cost models as Gaussian Mixture regressors providing certainty to identify unknown terrains where knowledge transfer is desirable. The presented method has been verified in synthetic showcase scenarios and deployment with a real walking robot. The proposed knowledge transfer demonstrates improved cost prediction and selection of the appropriate gait for specific terrains.**

I. Introduction

Understanding mobile robot traversability over terrains is crucial for autonomous robot operations in challenging environments. Multi-legged walking robots benefit from many controllable degrees of freedom, allowing motion adaptation to a wide variety of rough terrains impassable for similar-sized wheeled or tracked platforms. The legs' motion can be designed almost arbitrarily, but rhythmic motion patterns called gaits can be observed in animal locomotion. Gaits provide different capabilities [1] over different terrains [2]. Assessing gaits' terrain traversability and selecting the most suitable one is important for operating in challenging terrains.

The terrain assessment can be based on appearance and geometry terrain features [4] or experienced energy consumption [5], stability [1], [6] or slippage [7]. Machine learning approaches have been proposed to combine sensed surroundings and expected locomotion performance in terrain-aware locomotion control [8], [10]. However, such methods rely on dense datasets unavailable in setups where robots can collect only sparse experience for each gait.

In an ideal setup, the terrain traversal experience is collected for each gait individually, best for each terrain type. From such abundant experience, a terrain-gait cost model can be trained directly, modeling all possible terrain-gait interactions. However, a robot experiences each terrain with only one or just a few gaits, for which the experience can be aggregated into specialized gait-terrain models. Using the specialized models constrains the gait selection because the performance prediction can be only for the few terrain-gait interactions observed so far. We assume that gaits with correlated common experiences behave similarly, as indicated in Fig. 1. Then, we can enhance the sparse experience by transferring the experience collected for different gaits.

This work was supported by the Czech Science Foundation (GAČR) under research projects No. 19-20238S and No. 21-33041J.

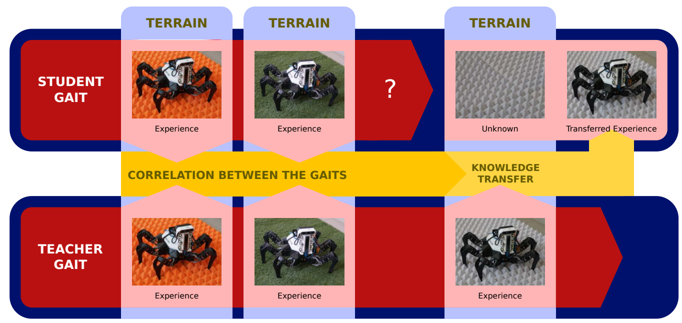

Fig. 1. The assumption on correlated traversal experience. If the robot has a similar experience when traversing two terrains using one gait, but its experience is limited to only one of the terrains with another gait, it assumes the experience with the second gait also matches between both terrains.

In this paper, we propose a terrain cost transfer learning algorithm that combines individual terrain-gait expert models, each trained in a different domain, into a general model that yields predictions about not yet experienced terrain-gait interactions. The model is based on Gaussian Mixture (GM) regressor experts [11] connected by GM regressors capturing the relations between the experience of each pair of gaits. The expert models' predictions are combined by weighting their respective prediction certainties. The feasibility of the proposed approach has been demonstrated in extreme scenarios that the learner might encounter and experimentally verified in deployment on a real hexapod walking robot.

The rest of the paper is organized as follows. Section II overviews existing methods of multi-legged robot locomotion with a focus on knowledge transfer. In Section III, a probabilistic representation of the knowledge transfer is derived. The proposed transfer learning algorithm is presented in Section IV, and experimental verification results are reported in Section V. The paper is concluded in Section VI.

II. Related Work

Mobile robot exteroception provides information about the surrounding terrain utilized in locomotion control. For different terrains, particular locomotion patterns can be more suitable than others. The energy consumption of a hexapod walking robot is minimized by gaits switching in [5]. The relation between the energy consumption and used gait is studied in [8], where the authors show that in different contexts, such as body configuration or required speed, different gaits are optimal. However, the body-terrain interaction model is apriori unknown and may change in time; therefore, machine learning is employed to obtain the knowledge.

In [12], enhanced knowledge is extracted from simulated ANYmal traversal to augment *Reinforcement Learning* (RL) algorithms. Robot exteroception is employed in [10], where ground and aerial robots are trained in simulators and real-world deployments similarly to [12]. However, RL requires a huge amount of real instances from different gaits and terrains without simulation. In contrast, the herein presented approach focuses on traversal experience gained continually.

The lack of training data can be addressed by transfer learning [13], where a robot collects local terrain-gait interactions models to infer the global model. In [14], quadrupedal gaits are trained for different terrain inclinations with training boosted by sharing parameters across multiple tasks. However, parameter sharing requires access to the model parameters that is impossible for black-box models, which are suitable for modeling complex body-terrain interaction.

The model-ensemble approaches combine local models (viewed as black boxes) into a global model. An ensemble is used in [15] where multiple specialized models trained by RL are used for unseen examples. The authors of [16] propose a transfer learning for homogeneous robots that inductively merge separately trained terrain classifier ensembles. However, such a direct inductive knowledge expansion is impossible in heterogeneous setups, where the experience is gathered under different dynamics and environments. Transfer for heterogeneous robots can be found in [17], where domains are reduced into lower-dimensional latent spaces and the transformation between these spaces is found. Still, such a transfer relies on assessing the received model quality before integration.

The authors of [18] hypothesize that uncertainty is a part of the experience utilized in further knowledge inference in animal cognition. Hence, model certainty and similarity to prior knowledge are suitable quality measures. The algorithm proposed in [19] builds Bayesian networks with the dependency between models measured before merging. A linear combination of classifiers where weights are calculated from experience similarity is proposed in [20]. However, both methods are limited to classification. In this paper, we address transfer learning in traversal cost regression, which can be considered a more complex task.

III. Problem Specification

In the studied problem, we address a robot that can use a discrete set of motion gaits. The robot walks over multiple terrains and measures different performances for each pair of a used gait and traversed terrain. Predicting performance for a given gait and terrain enhances the capability to assess the planned path feasibility. We focus on setups where the robot samples the terrain only with a few gaits, and thus, it lacks an exhaustive observation of all its gaits' performance for most of the terrains. For such a case, we propose extrapolating the gait performance from collected data using transfer learning.

The motion performance over the terrain $t \in \mathcal{T}$ (input) is evaluated by the traversal cost $c \in \mathcal{C}$ (output). For each gait, $g \in G$, the experience of the relation between the terrain and the cost is represented by the joint distribution $P(\mathcal{T}_g, \mathcal{C}_g)$ that describes which pairs of terrains and costs are likely. In terms of transfer learning [13], the marginal distribution $P(\mathcal{T}_g)$ represents the *domain*, while the task is to learn the terrain-to-cost *regressor* $r_g(t) = P(c|\mathcal{T} = t)$ from the observed *instances* $D_g = \{(c,t)|c \in \mathcal{C}_g, t \in \mathcal{T}_g\}$. We consider the transfer learning between *teacher* g_T and *student* g_S gait experiences, where the goal is to use the teacher's experience $P(\mathcal{T}_T, \mathcal{C}_T)$ for performance improving of the student's regressor r_S in the teacher's domain $P(\mathcal{T}_T)$.

For continually experienced terrain, the domain of each gait is generally different $P(\mathcal{T}_S) \neq P(\mathcal{T}_T)$. Moreover, since the gaits interact with the terrain differently [5], the conditional distributions are different as well, and $P(\mathcal{C}_S|\mathcal{T}_S) \neq P(\mathcal{C}_T|\mathcal{T}_T)$. The different domains and conditional probabilities classify the scenario as heterogeneous transfer learning.

Two heterogeneous experience sets can be harmonized by identifying similarities in both sets [19], [20]. If both gaits' performance (cost) correlates as they traverse similar terrains, we assume correlation on all terrains. The correlation between the student and teacher models can be expressed as

$$P'(\mathcal{C}_S|\mathcal{T}_T = t) = \int_{\mathcal{C}_T} P(\mathcal{C}_S|\mathcal{T}_T = t, \mathcal{C}_T) P(\mathcal{C}_T|\mathcal{T}_T = t), \quad (1)$$

where the *transfer probability* $P(\mathcal{C}_S|\mathcal{T}_T = t, \mathcal{C}_T)$ is the probability that the model contains information supporting similarity between the teacher's and student's costs. The resulting model $P'(\mathcal{C}_S|\mathcal{T}_T)$ is denoted as the *transfer model* since it is inferred from the *direct model* of the teacher, which is learned directly from *instances*.

In the following section, we propose an algorithm approximating the transfer probability $P(\mathcal{C}_S|\mathcal{T}_T, \mathcal{C}_T)$ to infer the transfer model $P'(\mathcal{C}_S|\mathcal{T}_T)$ that augments the student's model.

IV. Proposed Transfer of Traversal Experience

The proposed knowledge transfer algorithm combines the cost experiences of individual gaits. We assume that the functional relations between each pair of gaits are the same for all terrains. Thus, if the teacher and student gaits[1] experience the respective costs c_T^1 and c_S^1 on one shared terrain t_1, then the measured relation between the teacher and student costs $\gamma(c_T) = c_S$ can be extrapolated on the terrain t_2 unseen by the student. We propose to use a set of pairwise models that

[1] *"Teacher gait"* denotes any gait-terrain interaction model that provides some knowledge to another model, which is denoted as *"student gait."*

capture the relations between the costs experienced by the individual gaits, which are used to transfer the traversal cost predictions produced by the models learned for each of the gaits. In reference to (1), for each student-teacher gait pair $S \cap T$, we approximate $p_{S\cap T}(c_S|c_T, t)$ as $p_{S\cap T}(c_S|c_T)$.

The experience for each individual gait g is described by the direct regressor $r_g^{\text{direct}} : \boldsymbol{a} \to c_g$, where $\boldsymbol{a}$ is a descriptor of the terrain appearance, and c_g is the cost experienced when traversing the respective terrain with the gait g. Besides, the relationship between the costs experienced when using the gaits g_S and g_T is described by the *cost transfer model* $k_S^T : c_T \to c_S$, which is created from the costs experienced on terrains traversed using both the teacher gait g_T and student gait g_S. When predicting the cost $\hat{c}(\boldsymbol{a})$ of traversing a terrain appearing as $\boldsymbol{a}$ using the student gait g_S, the learner can use both the *direct* r_S^{direct} or *transfer* $r_{S,T}^{\text{transfer}} = k_S^T \circ r_T^{\text{direct}}$ regressor that uses the cost transfer models k_S^T to transform the teacher's experience from r_T^{direct}. Therefore, the final cost prediction is an output of the *combined regressor* r_S^{combined} that combines the direct and transfer regressors' outputs based on their respective certainties. In the rest of this section, we describe the building blocks of the proposed approach in detail.

A. Gait Cost Regressors

The direct cost regressors are implemented as Gaussian Mixture (GM) regressors [11]. For each gait g, the regressor r_g^{direct} is trained on the dataset $\{(\boldsymbol{a}^i, c_g^i)\}_{i=1}^{|D_g|}$, where the i-th instance contains the cost c_g^i measured during the gait's g locomotion on a terrain with the appearance $\boldsymbol{a}^i$, and $|D_g|$ is the number of instances collected for the gait g. For example, a robot pairs the terrain's color with the velocity experienced over the terrain, which yields a regressor that infers the expected velocity over terrains based on their color.

Since the number of traversed terrains is not known beforehand, the number of components of the GM is selected using the Bayesian Information Criterion (BIC) [21] as $q^* = \operatorname{argmin}_{q\in 1,\cdots,Q} BIC(q, D_g)$, where Q is the highest allowed number of components. The BIC is computed as $BIC(q, D_g) = \overline{q}(q, D_g)\ln(|D_g|) - 2\ln(\hat{L}(q, D_g))$, where the number of the estimated parameters $\overline{q}(q, D_g)$ depends on the number of components q and the dimensionality of the descriptors in the dataset D_g, and $\hat{L}(q, D_g)$ is the maximized likelihood of the GM with q components given the data D_g.

Given a query appearance $\boldsymbol{a}$, the GM regressor returns the cost mean $\mu_{c_g|\boldsymbol{a},D_g}$ and covariance (the variance $\sigma^2_{c_g|\boldsymbol{a},D_g}$ for 1D outputs), and the cost prediction is the expected value $r_g^{direct}(\boldsymbol{a}) = E(\mathcal{N}(\mu_{c_g|\boldsymbol{a},D_g}, \sigma^2_{c_g|\boldsymbol{a},D_g})) = \mu_{c_g|\boldsymbol{a},D_g}$. Therefore, the posterior probability of the prediction is

$$p_g(r_g^{direct}|\boldsymbol{a}, D_g) = p_{\mathcal{N}(\cdots)}(E(\cdots)) = (2\pi\sigma^2_{c_g|\boldsymbol{a},D_g})^{-\frac{1}{2}}. \quad (2)$$

B. Cost Transfer Model

The cost transfer is modeled using a GM regressor. Unlike the dataset for direct regressors, the dataset for the transfer model cannot be directly captured since it is learned from paired student and teacher costs, but the robot cannot traverse a terrain using two gaits simultaneously. Hence, we use the GMs learned by the student's and teacher's direct regressors to sample the training set in two steps.

The samples for the cost transfer model k_S^T are obtained using the respective GMs from r_S^{direct} and r_T^{direct} to generate familiar terrains appearances, drawing $n_{\text{gait}}^{\text{generate}} = 10000$ samples for each gait. Then, for each generated terrain appearance $\boldsymbol{a}$, we query student and teacher regressors obtaining the respective costs $r_S^{\text{direct}}(\boldsymbol{a})$ and $r_T^{\text{direct}}(\boldsymbol{a})$. Since the terrain appearances are drawn independently for each gait using the respective GM, many generated appearances by one GM can have a low prior probability for the other GM. In such a case, the terrain appearance with a low prior causes the GM to make arbitrary cost predictions, which is potentially confusing to the transfer function. For instance, consider a teacher that has learned the costs over sand, pavement, and turf terrains and a student that knows pavement, turf, and mud. Then, the relation between the teacher's and student's costs is learned only over pavement and turf, which are known to both. It is addressed by filtering the generated samples as follows.

Even though the transfer function models the student cost distribution as $p_{S\cap T}(c_S|c_T)p_T(c_T|\boldsymbol{a})$ with the student cost conditional only on the teacher's cost c_T, in (1), the student cost distribution is also conditional on the terrain that can be expressed as $p_{S\cap T}(c_S|c_T, \boldsymbol{a})p_T(c_T|\boldsymbol{a})$. We filter the samples by drawing $n_{\text{filter}}^{\text{generated}} = 2000$ from the discrete distribution of all samples weighted with their prior terrain probability $p_{S\cap T}(\boldsymbol{a})$. Thus, we create the dataset used to learn the transfer GM from the respective paired student and teacher costs[2]. However, since the prior probabilities are not known, they are approximated by multiplying the student's and teacher's prior probabilities as $p_{S\cap T}(\boldsymbol{a}) \approx p_T(\boldsymbol{a})p_S(\boldsymbol{a})$.

C. Combined Regressor

Let $\mathcal{G}_g = \mathcal{G} \setminus \{g\}$ denote the set of all gaits without the gait g. Given the student gait g_S and the set of teacher gaits $\mathcal{G}_S$, the combined regressor combines the direct and transfer regressors to select a likely cost prediction

$$r_S^{\text{combined}}(\boldsymbol{a}) = \lambda_S^S(\boldsymbol{a})r_S^{\text{direct}}(\boldsymbol{a}) + \sum_{g_T\in\mathcal{G}_S} \lambda_S^T(\boldsymbol{a})r_{S,T}^{\text{transfer}}(\boldsymbol{a}). \quad (3)$$

The weighting function $\lambda_S^g(\boldsymbol{a})$ is defined as

$$\lambda_S^g(\boldsymbol{a}) = \begin{cases} 1 & \text{if dominant}_S^g(\boldsymbol{a}), \\ 0 & \text{if } \exists h \in \mathcal{G}_g : \text{dominant}_S^h(\boldsymbol{a}), \\ \frac{p_S^g(\boldsymbol{a})}{\sum_{h\in\mathcal{G}} p_S^h(\boldsymbol{a})} & \text{otherwise,} \end{cases} \quad (4)$$

where a gait is dominant in the model if

$$\text{dominant}_S^g(\boldsymbol{a}) = \begin{cases} 1 & \text{if } \forall h \in \mathcal{G}_g : p_S^g(\boldsymbol{a}) > \alpha p_S^h(\boldsymbol{a}), \\ 0 & \text{otherwise,} \end{cases} \quad (5)$$

and $\alpha = 1.5$ is set empirically. For the student gait g_S, $p_S^S(\boldsymbol{a}) = p_S(r_S^{\text{direct}}|\boldsymbol{a})$ is the posterior probability of the re-

[2]The transfer model does not directly include the terrain appearance, and thus only costs are included in the transfer mixture.

gressor r_S^{direct} as defined in (2). For each teacher gait $g_T \in \mathcal{G}_S$, the transfer probability is

$$p_S^T(\boldsymbol{a}) = p_{S\cap T}(r_{S,T}^{\text{transfer}}|\boldsymbol{a}) = \\ = \frac{p_T(\boldsymbol{a})}{p_S(\boldsymbol{a})} p_{S\cap T}(k_S^T|r_T^{\text{direct}}(\boldsymbol{a})) p_T(r_T^{\text{direct}}|\boldsymbol{a}), \quad (6)$$

where $p_{S\cap T}(k_S^T|r_T^{\text{direct}}(\boldsymbol{a}))$ is the posterior probability of k_S^T computed similarly to the *direct* case in (2), $p_T(r_T^{\text{direct}}|\boldsymbol{a})$ is the posterior probability of r_T^{direct}, and $p_T(\boldsymbol{a})$ and $p_S(\boldsymbol{a})$ are the prior probabilities from r_T^{direct}'s and r_S^{direct}'s GMs.

V. Results

The proposed cost transfer learning has been validated in three scenarios. First, the predictor has been tested on synthetic data to show particular cases encountered by the learner. The approach has been then verified on a dataset captured with a real robot, reported in Section V-B. Finally, the learned predictor has been deployed in a real-time gait selection on the robot reported in Section V-C.

A. Evaluation on Synthetic Data

The synthetic data showcase the proposed approach in model scenarios. The data set represents a set of distinct gaits traversing a set of terrains, where each instance is a pair representing terrain appearance and the respective traversal cost. For each terrain-gait interaction, we generate 100 instances from the two-dimensional normal distribution shown in Table I. Each run is repeated ten times, and the samples generated in one of the trials are shown in Fig. 2.

TABLE I
Mean value $\boldsymbol{\mu}$ of distributions of generated synthetic data

Scenario / Terrain	t_1	t_2	t_3	t_4	t_2'	t_3'	t_4'
Terrain appearance	70	10	20	40	50	60	40
Cost in *Reference* scenario	30	10	30	*10	-	-	-
Cost in *Transfer* scenario	-	0	−20	0	-	-	-
Cost in *Flat Cost* scenario	-	20	20	20	-	-	-
Cost in *Unrelated* scenario	-	-	-	-	40	10	0

Appearance and cost values are drawn from $\mathcal{N}\left(\boldsymbol{\mu}, \begin{pmatrix} 1 & -0.1 \\ -0.1 & 1 \end{pmatrix}\right)$.
The symbol '-' represents that data has not been generated.
*Used only for testing.

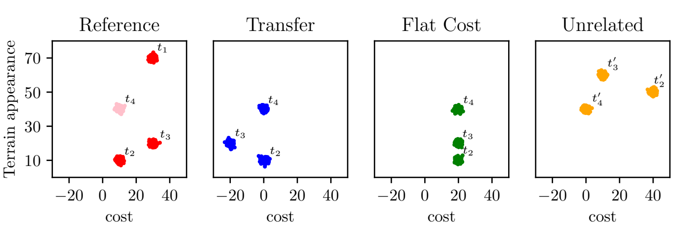

Fig. 2. The data points of the synthetic scenario drawn from the distributions defined as listed in Table I. The data points represent the union between the training and testing data used in one of ten trials. The data set is divided into four subsets; *Reference* consists of terrains experienced by the learner; and *Transfer*, *Flat Cost*, and *Unrelated* represent scenarios that might be encountered by the learner. For *Reference*, the dark data points are terrains where the direct predictor is learned, and light points represent a terrain where knowledge transfer is required and thus used only for testing.

There are four tasks, each corresponding to a hypothetical gait yielding a gait-specific model: *Reference, Transfer, Flat Cost*, and *Unrelated*. The *Reference* yields the direct cost regressor $r_{\text{reference}}$, whose performance is further enhanced by the other tasks. The *Reference* domain comprises terrains t_1, t_2, and t_3 accompanied by the target domain t_4 of an unobserved terrain where knowledge transfer is needed. The terrains are organized in two pairs $\langle t_1, t_3\rangle$, and $\langle t_2, t_4\rangle$, where the traversal cost is distributed similarly for each pair. *Transfer, Flat Cost*, and *Unrelated* demonstrate scenarios of teaching (providing experience to) the student *Reference*, which receives the knowledge about the terrain t_4.

Transfer, which is the ideal task for knowledge transfer, contains t_2 and t_3 terrains that overlap with the terrains from the *Reference* scenario, and t_4 corresponds to the test-only terrain from *Reference*. The overlapping terrains differ in traversal difficulty, giving the transfer model the information necessary to connect the costs between gaits. Hence, a performance increase is expected.

In the *Flat Cost* scenario, the teacher gait experiences the same terrains as *Transfer*. However, the cost distribution is the same for all the terrains. Therefore, knowledge transfer is not feasible since multiple *Reference* student gait traversal costs map on one *Flat Cost* teacher gait traversal cost. In the *Unrelated* scenario, the teacher does not traverse any terrains whose appearance would overlap with the *Reference* data set. Therefore, the traversed terrains are denoted t_2', t_3', and t_4', and knowledge transfer is impossible. For both scenarios, uncertain and arbitrary predictions are expected on t_4.

TABLE II
Prediction performance on synthetic data over 10 trials

Scenario / Terrain	t_1	t_2	t_3	t_4	all
Reference (baseline) scenario					
Direct	0.98 (0.07)	0.97 (0.04)	1.02 (0.07)	12.55 (1.21)	6.33 (0.60)
Transfer scenario					
Transfer	10.00 (0.80)	0.97 (0.04)	1.02 (0.06)	1.00 (0.08)	5.07 (0.40)
Combined	**0.99** (0.09)	**0.97** (0.03)	**1.02** (0.07)	**1.06** (0.19)	**1.02** (0.06)
Flat Cost scenario					
Transfer	8.68 (1.75)	9.94 (0.62)	9.91 (0.70)	10.08 (0.74)	9.71 (0.49)
Combined	**0.98** (0.07)	**2.39** (0.70)	**2.18** (0.64)	**10.08** (0.74)	**5.33** (0.41)
Unrelated scenario					
Transfer	1.10 (0.12)	19.94 (2.11)	1.76 (1.13)	20.94 (2.65)	14.59 (0.65)
Combined	**0.98** (0.07)	**1.03** (0.17)	**1.02** (0.07)	**20.94** (2.65)	**10.51** (1.32)

Mean values of the RMSE with the standard deviation in the brackets.

The learner's prediction Root Mean Square Error (RMSE) with respect to the *Reference* sampled costs is depicted in Table II. *Direct baseline predictor* performs well on the terrains t_1, t_2, and t_3 where it has been learned but cannot provide reasonable predictions on t_4 not observed in the learning. In the *Transfer* scenario, the *Transfer model* exhibits low error for the overlapping terrains t_2 and t_3 but increased error for the unknown (to the teacher) t_1, while the proposed *Combined model* selects whether to use the direct or transfer predictions and performs well for all the terrains. In the *Flat Cost* and *Unrelated* scenarios, the *Transfer model* is generally not capable of providing sufficient predictions on t_2 and t_3, which is the expected behavior. The *Flat Cost Combined model* error is high on t_2 and t_3, suggesting that violating the assumption on correlated traversal experience brings confusion into the model. The *Unrelated Combined model* performs close to the *Direct baseline* on t_2 and t_3, suggesting that the proposed approach is robust against unrelated but not deliberately confusing data. We conclude that the proposed approach performs as expected.

B. Evaluation using Real Robot Data

The proposed transfer learning has been further examined using data captured with the hexapod walking robot *SCARAB II* [22] that was deployed in a laboratory environment. The robot can locomote with three tripod gaits: (i) the *Fast* paced gait with a low stance position and short gait period; (ii) the *Tall* gait with a high stance and long gait period; and (iii) the *Basic* gait with parameters between the two former gaits. The gait parameters are listed in Table III. The robot with the *Fast* gait moves swiftly on flat terrains but struggles on rough terrains. The slower *Tall* gait traverses rough terrains with little effort. Finally, the *Basic* gait moves moderately fast on flat areas but is slowed down on rough terrains.

TABLE III
PARAMETRIZATION OF THE USED GAITS

Gait Parameter / Gait	*Fast*	*Basic*	*Tall*
Gait Cycle Duration [s]	1.10	2.00	2.90
Step Height [m]	0.04	0.05	0.07

The robot traversal cost is measured as the relative slow-down compared to the commanded velocity v_{cmd}. The cost c_w is continually captured over the time-window w as

$$c_w = c_{\max} \tanh\left(\frac{v_{\text{cmd}}}{c_{\max}} \frac{\Delta_t(w)}{\|w\|}\right), \tag{7}$$

where $\|w\|$ is the distance covered over the window and $\tanh()$ scaled by the maximum cost $c_{\max}$ smooths high cost variances observed when the traversed distance $\|w\|$ is short, e.g., when the robot is stuck. The window size w is considerably longer than the robot's gait period as $\Delta_t(w) \approx 20\,\text{s}$.

Fig. 3. (left) The hexapod walking robot with tracking and mapping cameras used in the real experimental validation. (right) The robot on the blue *Plate*, purple *Fabric*, and *Orange* and *Gray* compliant spikes.

The robot is equipped with the Intel RealSense T265 and RGB-D D435 cameras, see Fig. 3, used for localization and mapping, respectively. Mapping is performed as building colored elevation grid [23] with the cell size $d_\nu = 0.07\,\text{m}$. The terrain appearance is encoded as the $\texttt{LAB}$ color descriptor using the mean values of the a and b channels over the cells in the $d_{desc} = 0.22\,\text{m}$ radius around the evaluated cell. The robot captures its terrain traversing experience by pairing the measured cost with the descriptor corresponding to its position at the midpoint of the cost computation window.

The four distinct terrains utilized for collecting data are depicted in Fig. 3: the flat blue *Plate*; flat ground covered by purple *Fabric* that is somewhat slippery yet relatively easy to traverse; and *Orange* and *Gray* colored spikes made from compliant foam (sound-proofing material), which is not hard to traverse for the *Tall* gait, but slows down the *Basic* gait and immobilizes the *Fast* gait. The collected appearance-based descriptors and experienced costs are listed in Table IV.

TABLE IV
APPEARANCE AND COST OF THE EXPERIMENTAL TERRAINS

Gait / Terrain	*Plate*	*Orange*	*Gray*	*Fabric*
Fast gait - appearance a	-3.43 (0.72)	33.31 (6.74)	-1.93 (2.07)	49.05 (9.34)
Fast gait - appearance b	-13.87 (1.45)	41.39 (10.39)	5.73 (2.76)	-4.50 (3.57)
Fast gait - cost	1.14 (0.05)	9.88 (0.24)	9.93 (0.11)	1.10 (0.06)
Basic gait - appearance a	-3.42 (0.66)	36.59 (4.65)	-2.07 (1.29)	45.89 (9.22)
Basic gait - appearance b	-14.42 (1.76)	45.60 (8.28)	4.54 (2.33)	-4.43 (4.13)
Basic gait - cost	1.27 (0.09)	3.89 (1.16)	3.17 (0.97)	1.21 (0.04)
Tall gait - appearance a	-3.44 (0.72)	37.81 (4.80)	-1.04 (1.88)	47.85 (9.09)
Tall gait - appearance b	-15.37 (1.76)	46.76 (7.21)	3.50 (2.63)	-3.41 (4.12)
Tall gait - cost	1.44 (0.04)	1.75 (0.08)	1.70 (0.11)	1.38 (0.06)

Mean values with standard deviation (in the brackets) of the real captured data are listed.

In the prior models, each gait misses knowledge of one of the terrains for which it is expected to use the knowledge transfer. The *Fast* gait misses *Fabric*, while the *Tall* and *Basic* gaits miss *Gray*. The *Orange* and *Plate* terrains are known to all gaits and can be used to learn the cost transfer models.

The dataset contains three passes for each gait-terrain pair, where two passes are used for learning if applicable, and the remaining one is reserved for testing. We sample 100 points from each selected pass, and thus the learning sets contain 200 points for each used gait-terrain pair. Since the datasets are relatively sparse, especially for the gait-terrain pairs where the *Fast* gait is stuck on the spikes, the individual samples are augmented by adding $\mathcal{N}(0, 1)$ noise to each terrain appearance descriptor dimension.

TABLE V
PREDICTION PERFORMANCE OF THE ROBOT GAIT MODEL

	Regressor	*Plate*	*Orange*	*Gray*	*Fabric*	All
Fast	Direct	0.24 (2.80)	0.32 (2.85)	0.10 (3.63)	6.40 (3.42)	3.22 (3.17)
	Transfer from *Basic*	0.05 (4.68)	5.94 (1.70)	7.28 (-8.60)	0.11 (10.64)	4.95 (2.11)
	Transfer from *Tall*	0.79 (4.70)	1.55 (3.82)	1.41 (-6.14)	1.22 (11.86)	1.71 (3.56)
	Combined	**0.04** (–)	**1.55** (–)	**0.10** (–)	**1.31** (–)	**1.15** (–)
Basic	Direct	0.10 (1.92)	1.24 (-0.95)	2.02 (0.75)	0.10 (2.54)	1.20 (1.06)
	Transfer from *Fast*	0.14 (4.77)	2.30 (3.64)	2.11 (13.64)	2.15 (-5.65)	2.03 (4.10)
	Transfer from *Tall*	0.24 (4.46)	2.27 (0.72)	1.87 (0.22)	0.44 (4.54)	1.53 (2.48)
	Combined	**0.11** (–)	**2.20** (–)	**2.12** (–)	**0.11** (–)	**1.54** (–)
Tall	Direct	0.22 (2.60)	0.21 (2.81)	0.18 (3.16)	0.11 (2.78)	0.19 (2.84)
	Transferfrom *Fast*	0.07 (8.93)	0.10 (8.00)	0.77 (10.40)	1.00 (-4.86)	0.66 (5.62)
	Transfer from *Basic*	0.04 (7.07)	0.28 (5.61)	0.37 (-1.22)	0.11 (2.60)	0.25 (3.52)
	Combined	**0.07** (–)	**0.12** (–)	**0.49** (–)	**0.10** (–)	**0.27** (–)

Mean values of the RMSE with the mean component-model certainty in the brackets computed from 10 trials.
The *Fast* model is learned without *Fabric* terrain and *Basic* and *Tall* modes are learned without *Gray* spikes terrain, which is highlighted in gray.

The results in Table V suggest that the *Fast* gait receives the knowledge about the purple *Fabric*, and the combined model is the best performer for the *Fast* gait. The *Basic* and *Tall* gaits receive the knowledge about the *Gray* spikes, but it is not as pronounced in their respective prediction RMSE since, by chance, the uninformed direct predictor evaluates the *Gray* terrain correctly. However, as the *Tall* and *Basic* models have no experience with the terrain, the transfer model is preferred over the uncertain guess. Overall, the results suggest that the robot correctly combines the direct and transfer models. The proposed combination of direct and transfer models improves or provides similar results as the standalone models.

C. Experimental Deployment on the Hexapod Walking Robot

The proposed approach has been experimentally verified on the real hexapod walking robot, which uses the predictions to select a suitable gait when traversing a laboratory environment in real-time. While walking through the environment, the robot uses the models learned in Section V-B to select the lowest cost gait at its current location as

$$g^* = \mathrm{argmin}_{g \in \mathcal{G}} \max_{\boldsymbol{a} \in A(\mathbf{s}_{\mathrm{robot}})} r_g^{\mathrm{comb}}(\boldsymbol{a}), \tag{8}$$

where the cost is maximized over terrain appearance descriptors of the grid map cells within the $0.15\,\mathrm{m}$ radius area $A(\mathbf{s}_{\mathrm{robot}})$ around the robot's position $\mathbf{s}_{\mathrm{robot}}$.

During the deployment, the robot traversed three terrains shown in Fig.4. First, the robot traverses the blue *Plate* directly known to all of the gaits. Then the *Gray* spikes and *Purple* fabric where knowledge transfer is required for some gaits. The cost predictions and selected gaits (see Fig. 5) suggest the proposed approach is suitable for gait selection on the robot since the *Fast* gait is preferred over the flat *Plate* and *Fabric*, while the *Tall* gait is selected over the rough *Gray* spikes. Overall, a suitable gait is selected with only small irregularities on the terrain boundaries, and the robot behaves as expected.

Fig. 4. The robot using the *Fast* gait over the purple *Fabric*, *Tall* gait over the *Gray* spikes, and the *Fast* gait over the blue *Plate*.

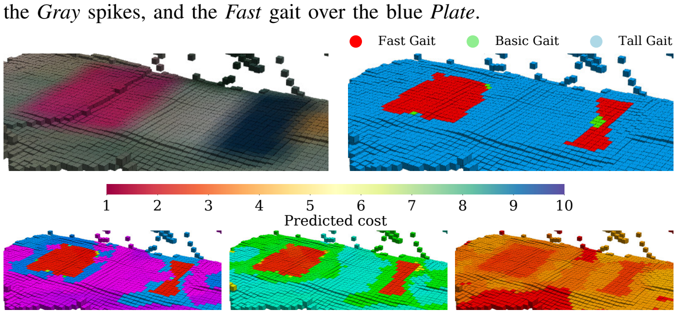

Fig. 5. Top, left to right: the colored elevation map of the deployment track, and the selected gaits. Bottom, left to right: the cost predictions (maximized over $A(\mathbf{s}_{\mathrm{robot}})$) for the *Fast*, *Basic*, and *Tall* gaits, respectively.

VI. Conclusion

In this paper, we present a method for transferring the traversal cost experience between gaits of a multi-legged walking robot that learns an independent model for each gait and then infers functional relations between models to transfer predictions between gaits. The prediction certainty of the individual models (implemented as Gaussian Mixture regressors) is used to identify terrains unknown to the student's gait, for which the student tries to transfer the prediction from other models. The method is verified both on synthetic data, designed to showcase specific learning scenarios, and on real data captured by a walking robot in a laboratory environment, where the proposed approach improves the regressor performance. Besides, the developed approach has been deployed in real-time robot locomotion gait selection. Next, we aim to investigate the gait model confidence for active learning.

References

[1] R. B. McGhee and A. A. Frank, "On the stability properties of quadruped creeping gaits," *Mathematical Biosciences*, vol. 3, pp. 331–351, 1968.

[2] H. Kolvenbach, D. Bellicoso, F. Jenelten, L. Wellhausen, and M. Hutter, "Efficient gait selection for quadrupedal robots on the moon and mars," in *International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*, 2018.

[3] P. Krüsi, M. Bosse, and R. Siegwart, "Driving on Point Clouds: Motion Planning, Trajectory Optimization, and Terrain Assessment in Generic Nonplanar Environments," *Journal of Field Robotics*, vol. 34, no. 5, pp. 940–984, 2016.

[4] D. Belter, J. Wietrzykowski, and P. Skrzypczyński, "Employing Natural Terrain Semantics in Motion Planning for a Multi-Legged Robot," *Journal of Intelligent & Robotic Systems*, vol. 93, no. 3, pp. 723–743, 2019.

[5] N. Kottege, C. Parkinson, P. Moghadam, A. Elfes, and S. P. N. Singh, "Energetics-informed hexapod gait transitions across terrains," in *IEEE Intl. Conf. Robotics and Automation (ICRA)*, 2015, pp. 5140–5147.

[6] B. Lin and S. Song, "Dynamic modeling, stability and energy efficiency of a quadrupedal walking machine," in *IEEE Intl. Conf. Robotics and Automation (ICRA)*, 1993, pp. 657–670.

[7] R. Gonzalez and K. Iagnemma, "Slippage estimation and compensation for planetary exploration rovers. State of the art and future challenges," *Journal of Field Robotics*, vol. 35, no. 4, pp. 564–577, 2018.

[8] W. Xi, Y. Yesilevskiy, and C. D. Remy, "Selecting gaits for economical locomotion of legged robots," *International Journal of Robotics Research*, vol. 35, no. 9, pp. 1140–1154, 2016.

[9] S. Gangapurwala, M. Geisert, R. Orsolino, M. Fallon, and I. Havoutis, "RLOC: Terrain-aware legged locomotion using reinforcement learning and optimal control," *arXiv*, vol. abs/2012.03094, 2020.

[10] M. L. Iuzzolino, M. E. Walker, and D. Szafir, "Virtual-to-real-world transfer learning for robots on wilderness trails," in *IEEE/RSJ Intl. Conf. Intelligent Robots and Systems (IROS)*, 2018, pp. 576–582.

[11] S. Calinon, F. Guenter, and A. Billard, "On learning, representing, and generalizing a task in a humanoid robot," *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, vol. 37, no. 2, pp. 286–298, 2007.

[12] J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning quadrupedal locomotion over challenging terrain," *Science Robotics*, vol. 5, no. 47, 2020.

[13] F. Zhuang, Z. Qi, K. Duan, D. Xi, Y. Zhu, H. Zhu, H. Xiong, and Q. He, "A comprehensive survey on transfer learning," *Proceedings of the IEEE*, vol. 109, no. 1, pp. 43–76, 2021.

[14] J. Degrave, M. Burm, P.-J. Kindermans, J. Dambre, and F. wyffels, "Transfer learning of gaits on a quadrupedal robot," *Adaptive Behavior*, vol. 23, no. 2, pp. 69–82, 2015.

[15] C. Devin, A. Gupta, T. Darrell, P. Abbeel, and S. Levine, "Learning modular neural network policies for multi-task and multi-robot transfer," in *IEEE Intl. Conf. Robotics and Automation (ICRA)*, 2017, pp. 2169–2176.

[16] R. Szadkowski, M. Prágr, and J. Faigl, "Transfer of inter-robotic inductive classifier," in *International Conference on Automation, Control and Robots (ICACR)*, 2020, pp. 32–36.

[17] B. Bócsi, L. Csató, and J. Peters, "Alignment-based transfer learning for robot models," in *International Joint Conference on Neural Networks (IJCNN)*, 2013, pp. 1–7.

[18] D. C. Knill and A. Pouget, "The Bayesian brain: the role of uncertainty in neural coding and computation," *Trends in Neurosciences*, vol. 27, no. 12, pp. 712–719, 2004.

[19] R. Luis, L. E. Sucar, and E. F. Morales, "Inductive transfer for learning bayesian networks," *Machine Learning*, vol. 79, pp. 227–255, 2010.

[20] J. Gao, W. Fan, J. Jiang, and J. Han, "Knowledge transfer via multiple model local structure mapping," in *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2008, p. 283–291.

[21] G. Schwarz, "Estimating the dimension of a model," *The Annals of Statistics*, vol. 6, no. 2, pp. 461–464, 1978.

[22] M. Forouhar, P. Čížek, and J. Faigl, “SCARAB II: A small versatile six-legged walking robot,” in *5th Full-Day Workshop on Legged Robots at IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 1–2.

[23] J. Bayer and J. Faigl, “On autonomous spatial exploration with small hexapod walking robot using tracking camera intel realsense T265,” in *European Conference on Mobile Robots (ECMR)*, 2019, pp. 1–6.