Robotics and Autonomous Systems 62 (2014) 1790–1798

Contents lists available at ScienceDirect

Robotics and Autonomous Systems

journal homepage: www.elsevier.com/locate/robot

The effect of motor action and different sensory modalities on terrain classification in a quadruped robot running with multiple gaits

M. Hoffmann [a,b,*], K. Štěpánová [c], M. Reinstein [d]

[a] iCub Facility, Istituto Italiano di Tecnologia, Via Morego 30, 16123 Genova, Italy

[b] Artificial Intelligence Laboratory, Department of Informatics, University of Zurich, Andreasstrasse 15, 8050 Zurich, Switzerland

[c] Department of Cybernetics, Faculty of Electrical Engineering, Czech Technical University in Prague, Prague, 166 27, Czech Republic

[d] Center for Machine Perception, Department of Cybernetics, Faculty of Electrical Engineering, Czech Technical University in Prague, Prague, 166 27, Czech Republic

HIGHLIGHTS

- Preconditioning on gait used boosts terrain discrimination in a legged robot.
- Specific gaits are particularly suited for terrain perception.
- Inertial, tactile, and proprioceptive sensors are a robust terrain sensing set.
- Encoders in passive compliant joints performed best from the sensory set.

ARTICLE INFO

Article history:
Received 26 February 2014
Received in revised form
9 July 2014
Accepted 16 July 2014
Available online 24 July 2014

Keywords:
Legged robots
Terrain classification
Active perception
Gait change
Haptic terrain classification
Underactuated robots

ABSTRACT

Discriminating or classifying different terrains is an important ability for every autonomous mobile robot. A variety of sensors, preprocessing techniques, and algorithms in different robots were applied. However, little attention was paid to the way sensory data was generated and to the contribution of different sensory modalities. In this work, a quadruped robot traversing different grounds using a variety of gaits is used, equipped with a collection of proprioceptive (encoders on active, and passive compliant joints), inertial, and foot pressure sensors. The effect of different gaits on classification performance is assessed and it is demonstrated that separate terrain classifiers for each motor program should be employed. Furthermore, poor performance of randomly generated motor commands confirms the importance of coordinated behavior on sensory information structuring. The collection of sensors sensitive to active, "tactile", terrain exploration proved effective. Among the individual modalities, encoders on passive compliant joints delivered best results.

1. Introduction

The movement of every mobile robot, both legged and wheeled, is strongly affected by the interaction with the environment it is traversing. Successful perception of the terrain is a key ability that impacts the decision-making of the robot – whether to enter a particular area or which speed or gait to choose – and hence its performance in different situations. Furthermore, being able to perceive the terrain properties can be an important precondition for successful navigation performed by the robot (see e.g., [1]). Terrain can be discriminated in a supervised manner – matching the environment with predefined categories like tar or sand – or in an unsupervised manner. The former is known as *robotic terrain classification* (e.g., [2]). If the terrain classes are not available to the robot, different unsupervised methods can be employed to discover the different terrain types (e.g., [3]). A related notion has been put forth by Ojeda et al. [4]: terrain characterization, which aims at determining key parameters of the terrain that are relevant to the traversability by the robot. We will use *terrain discrimination* to encompass all of the above.

A large variety of methods is applied to terrain discrimination: these include different sensors, different methods to preprocess them, and finally different algorithms to discriminate between the terrains. Often, sensors that perceive at a distance like cameras,

* Corresponding author at: iCub Facility, Istituto Italiano di Tecnologia, Via Morego 30, 16123 Genova, Italy.
E-mail addresses: matej.hoffmann@iit.it (M. Hoffmann), stepakar@fel.cvut.cz (K. Štěpánová), reinstein.michal@fel.cvut.cz (M. Reinstein).

http://dx.doi.org/10.1016/j.robot.2014.07.006
0921-8890/

laser scanners or radars are used, sometimes supplemented by inertial sensors. Sensors that directly perceive the robot's physical interaction with the environment, like tactile sensors, or proprioceptive sensors, which reflect how the robot's body negotiates different grounds, are used less often. Furthermore, the type of motor action the robot is using while performing terrain discrimination is typically not considered. In this work, we investigate the possibilities of distinguishing terrain types using a collection of inertial, tactile and proprioceptive sensors and we specifically address the consequences of using different motor actions. We choose a platform with significantly different "motor regimes": a quadruped robot running with multiple gaits. In order to lend more credence to this investigation, we employ two different classification algorithms (naïve Bayes and Support Vector Machines—SVM) and a large pool of time and frequency domain features.

This article is structured as follows. Section 2 covers related work on this topic. In Section 3, the experimental setup and methods for data collection, feature computation and classification are sketched. Section 4 provides a mathematical formulation of the problem. In Section 5, results and analysis are presented. Finally, Section 6 concludes the paper with a discussion and directions of future research.

2. Related work

Typical examples of robot terrain classification involve wheeled platforms equipped with a camera, accompanied by a vibration sensor, like an accelerometer [5], or a microphone in the case of the Mars rovers [6]. Giguere and Dudek [7] introduced a metallic rod with an accelerometer dragged behind the robot; Ojeda et al. [4] experimented with a richer collection of sensors. The motor action of the robot during the classification task is typically not considered; occasionally, different speeds [8] or speed and load were considered [9].

There is much less work on legged robots. Filitchkin and Byl [10] used a compact camera mounted on the Little Dog robot for terrain classification; Belter and Skrzypczyński [11] used a laser scanner on a hexapod to map the terrain around the robot and choose footholds. However, legged robots are particularly suited for tactile exploration of the terrain being traversed. Hoepflinger et al. [12] estimated surface properties from joint motor currents and force sensing resistors while the robot's leg was probing the surface; Kim et al. [13] varied the gaits in a monopod platform and used ground reaction force and torque sensing for classification. Giguere et al. [3], Garcia Bermudez et al. [14], and Manjanna et al. [15] used hexapod robots with semi-circular legs and combined several sensory modalities: inertial sensors, encoders, and back-EMF sensors or motor current estimators. Regarding the effect of the motor signal, the effect of gait frequency on classification performance was studied in [14,15]. Finally, Larson et al. [16] used a simulated robot and a limb/terrain interaction model to evolve gaits that facilitate terrain classification based on gait bounce—estimation of the loping motion of the robot as it locomotes.

A variety of features is used to preprocess raw sensory streams. Leaving visual features aside, the most common is the combination of statistical moments (mean, variance, skewness, kurtosis etc.) in the time domain with frequency domain features [7,3,17,12]. Some authors rely on the frequency domain only [18,9,4]. Various algorithms are employed to perform the classification: Support Vector Machines (SVM) [8,5,6,13], other physical probabilistic models [19], neural networks [9,20,17,21,4], linear discriminant analysis [22,23], principal component analysis (PCA) [9,24,13], or adaptive Bayesian filtering [25]. In [17,3] comparison of the traditional classification methods is provided, furthermore, in [7] both supervised and unsupervised techniques are successfully applied showing that PCA can be used to reduce the data dimensionality without impacting on the classification results.

In this work, we will employ two standard classification algorithms (naïve Bayes and SVM) and a feature set comprising standard time and frequency domain features. The focus, however, will be on two aspects that have been largely overlooked so far: the effect of motor action and the collective as well as individual performance of inertial, tactile, and joint angle sensors to discriminate the properties of different terrains. Our work has a similar flavor to some of the work in terrain classification in legged robots that we reviewed above [14,3,12,15], but goes further in that very different dynamic locomotion patterns – gaits – are considered. We follow up on previous work on the Puppy quadruped robot [20,26]. In [20], a dataset comprising three gaits was used and the performance of different sensor sets was addressed, however, the data was collapsed across the gaits. Preconditioning on the gaits was first used in [26]; in this work, it is extended in the following ways: (i) A new, much more comprehensive feature set, and an additional classifier are introduced; (ii) A more elaborate analysis is performed and specific effects of individual gaits are studied; (iii) The effect of different sensory modalities is analyzed; (iv) The interplay of the four different factors impacting on the classification performance (gaits, sensory modalities, classifier, features) is investigated; (v) An additional dataset is added to study the contribution of additional sensors and the effect of random motor signals on classification.

3. Experimental setup

3.1. Robots and ground materials

The basic experimental setup is identical to [26]—the same dataset was used in this work. We recapitulate it here briefly for the reader's convenience. The Puppy robot (see Fig. 1 left) has four identical legs with two revolute joints per leg: the hip driven by servomotors and the knee passive compliant—with a spring attached across it. Four potentiometers measured the joint angles on the passive knee joints and four pressure sensors recorded forces at the robot's feet. Linear accelerations were measured by an onboard 3-axis accelerometer. All sensory channels were sampled at 50 Hz.

Five sets of position control commands for the servomotors were prepared, resulting in 5 distinct gaits (bound forwards, bound left/right, crawl, trot backwards). The target position trajectory for every motor was a sine wave at 1 Hz with a specific amplitude, offset, and phase lag. Gaits were chosen randomly and exercised in 2-second-intervals during which the sensory data were collected, forming what we call epochs. At the end of each epoch the robot could change the gait. There were two locomotion periods in every epoch. For analysis, only data from the second period is used; the first one – immediately after the gait transition – is discarded.

A small wall-enclosed arena of 2×1 m was prepared and subsequently covered with four different ground substrates—data was collected on each of them in turn. The materials were plastic foil, cardboard, Styrofoam and rubber. They differed in friction[1] and also in structure: cardboard and rubber had ridges that – especially in the case of cardboard – rendered the terrain non-flat relative to the robot size (about 1 cm high ridges vs. 6 cm leg length). When the robot was approaching the wall of the arena, this was detected by an infra-red distance sensor and the "trot back" gait was triggered until a safe distance was established and random gait selection resumed. In what follows, we will refer to the dataset coming from this robot as "Real robot".

[1] We estimated static friction coefficients between the ground materials and robot's feet by putting a block covered with the same adhesive skin as on the robot's feet on inclined planes covered with the different ground materials. As the adhesive skin has asymmetrical properties, two values were obtained for each material. The low/high values were: plastic foil: 0.39/0.39, cardboard: 0.64/1.10, Styrofoam: 0.74/1.06, rubber: 0.76/0.91.

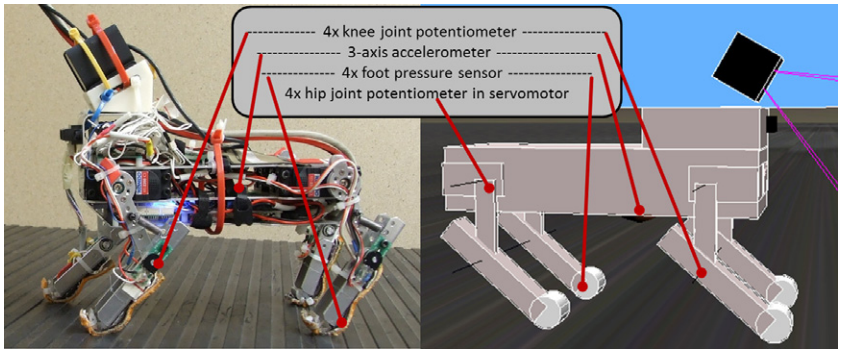

Fig. 1. Real and simulated robot and the sensor suite. The robot is 20 cm long. The camera that was also mounted on the robot was not used in the experiments. An infra-red distance sensor was used to trigger wall avoidance behavior, but was otherwise not used for terrain discrimination.

In order to be able to record larger datasets with a bigger gait repertoire, we additionally designed a model of Puppy in Webots [27], a physics-based simulator (see Fig. 1 right). For this model, we used the same gait repertoire (2 gaits had to be adapted) plus 4 additional gaits (turn left/right, pace, walk), obtaining a repertoire of 9 gaits. In addition, the actual angles of the hip joints were available as well. The dataset from this robot will be labeled "Simulated robot".

In the simulator, the arena was much bigger in size (25×25 m), so encounters with the walls were much less frequent. The "foil", "cardboard", and "rubber" were flat but differed in Coulomb friction coefficients that were a result of experimentation ($\mu = 2$, 11, and 20 respectively). To increase the differences between the substrates in the simulator, the "Styrofoam" ground ($\mu = 9$) was made uneven with randomly placed smooth bumps of up to 3 cm height.

Finally, to enrich the dataset compared to [26], we have added a third dataset, which comes from [28] and was also used in [20]. A twin version of the Puppy robot is used, which, in addition, has an Inertial Measurement Unit (IMU) with 3-axis accelerometer and gyroscope attached on its back; it also allows for reading the hip potentiometer values that are used by the servomotor internal PID controller. Experiments with the bound right and turn left gaits were conducted. Furthermore, a "random" gait was designed, which consisted of 4 uncorrelated, low-passed random signals. Since the IMU shifted the robot's Center of Mass upwards and made it less stable and also due to the inherent vulnerability of the "random" gait, the sample of terrains was changed to plastic foil, linoleum[2] and Styrofoam. Finally, the experimental protocol was different: The robot was not allowed to change gaits every 2 s. Instead, it continued with the same gait for individual runs of approximately 1 min duration. The dataset coming from this setup will be referred to as "Real robot 2", which is primarily used for control experiments (see Section 5.4).[3]

3.2. Datasets and feature computation

The three robot platforms described in the previous section gave rise to 3 datasets. Table 1 summarizes the properties of the datasets. In total, there were 1203 samples (locomotion periods) for the "Real robot", 18 002 for the "Simulated robot", and 2493 for the "Real robot 2", split between the different gait and ground combinations.

The raw data was compressed by extracting the most commonly used time-domain and frequency-domain features. Following a similar strategy as in [26], we took advantage of the periodic nature of the locomotion and created period-based features (where the period – our time window – had a duration of 1 s). However, as the feature set in [26] was very restricted (only 10 aggregate features in total), following an analogous approach to [29], the feature set was enriched as follows.

Table 1
Datasets.

Sensory channels	Real robot	Simulated robot	Real robot 2
Accelerometer	3 axes	3 axes	3 axes (IMU)
Gyroscope			3 axes (IMU)
Hip joint sensors		4	4
Knee joint sensors	4	4	4
Foot pressure sensors	4	4	4
Total nr. sensors	11	15	18
Total nr. features	209	285	342
Nr. gaits used	5	9	3
Nr. terrains	4	4	3

For each period and every sensor a set of 19 features was extracted. The time-domain features were: (1) Minimum value (*Min*), (2) Maximum value (*Max*), (3) *Mean*, (4) *Kurtosis*, (5) *Skewness*, (6) *Median*, (7) Standard deviation (*Std*), (8–11) Maximum and Mean values of the first and second differences (*Diff* 1 − *Max*, *Diff* 2 − *Max*, *Diff* 1 − *Mean*, *Diff* 2 − *Mean*—see formulas below), (12) Approximation of the integral of values via the trapezoidal method (*Trapz*), and (13) Amplitude of the Hilbert transform of the values (*Hilbert*). Features *Diff* 1 − *Max*, *Diff* 2 − *Max*, *Diff* 1 − *Mean*, *Diff* 2 − *Mean* were computed as follows:

$$\mathit{Diff}\,1 - \mathit{Max} = \max|\mathrm{diff}(\mathbf{x})|$$

$$\mathit{Diff}\,2 - \mathit{Max} = \max|\mathrm{diff}(\mathrm{diff}(\mathbf{x}))|$$

$$\mathit{Diff}\,1 - \mathit{Mean} = \mathrm{mean}\,|\mathrm{diff}(\mathbf{x})|$$

$$\mathit{Diff}\,2 - \mathit{Mean} = \mathrm{mean}\,|\mathrm{diff}(\mathrm{diff}(\mathbf{x}))|,$$

where $\mathbf{x}$ is the vector of data points of a particular sensory channel in the locomotion period (50 data points—sampling was at 50 Hz) and diff($\mathbf{x}$) transforms it into a vector of differences between adjacent elements of $\mathbf{x}$.

The frequency-domain features were: (14) Frequency with the highest amplitude (*FFT* 1*st*), (15) Magnitude of freq. with the highest amp. (*FFT* 1*st* − *Mag*), and similarly for frequencies with the second and third highest amplitude, (16) *FFT* 2*nd*, (17) *FFT* 2*nd* − *Mag*, (18) *FFT* 3*rd*, (19) *FFT* 3*rd* − *Mag*.

Applying all the features for every sensory channel of the 11/15/18 sensors in the three datasets gave rise to a pool of 209/285/342 features respectively (see Table 1).

3.3. Feature selection and classification

The aim of automatic feature selection is to decrease the number of features used (for real-time on-board processing for example) without significantly decreasing the classification accuracy.

[2] Low/high friction values in contact with the adhesive skin on Puppy's feet: 0.31/0.40.

[3] In this robot, the accelerometer readings from the IMU are used instead of the on-board accelerometer.

We used forward feature selection which is a greedy hill climbing method that starts with an empty subset and adds one by one the features that improve the classification accuracy the most. However, to keep the computational requirements within bounds, a stopping criterion to find the optimal subset of features is typically applied. Another possibility is to pick a fixed number of features that will be selected. We stopped the algorithm when 10 best features were selected or when the accuracy of classification on the validation data ceased to improve (the number 10 resulted from inspection of the graphs of accuracy development).

For classification, we used 10-fold cross-validation using the naïve Bayes classifier (MATLAB implementation) and SVM (implemented using the standard LIBSVM[4] library for MATLAB [30]), with radial basis function as kernel.[5] The resulting classification accuracies were averaged over 10 runs of the feature selection process.

4. Mathematical problem formulation

Taking the "Real robot" dataset as an example, the problem can be formalized as follows. The robot was running using 5 different gaits, G, on 4 different terrains, T:

$G = \{bound, boundleft, boundright, crawl, trotback\}$

$T = \{foil, Styrofoam, cardboard, rubber\}$.

The sensory data was collected from 11 sensory channels from 3 different sensor types or modalities. The complete feature set, F, thus looked like this:

$F = \{F^{ACC}, F^{KNEE}, F^{PRES}\}$, where

$ACC = \{acc\_x, acc\_y, acc\_z\}$

$KNEE = \{knee\_FL, knee\_FR, knee\_HL, knee\_HR\}$

$PRES = \{pres\_FL, pres\_FR, pres\_HL, pres\_HR\}$.

ACC denotes accelerometer (the modality) and *acc_x*, *acc_y*, *acc_z* its respective axes (individual sensory channels). Then, *KNEE* stands for the knee potentiometers, *PRES* for pressure sensors on the feet, with *FL*, *FR*, *HL*, *HR* denoting front left, front right, hind left, and hind right leg respectively. For each channel, 19 features from every second of the time series – one locomotion period – were computed. Thus, F^{ACC} can be expanded as follows:

$$F^{ACC} = (f_1^{acc\_x}, \ldots, f_{19}^{acc\_x}, f_1^{acc\_y}, \ldots, f_{19}^{acc\_y}, f_1^{acc\_z}, \ldots, f_{19}^{acc\_z}).$$

The expansion for the knees, F^{KNEE}, and pressure sensors, F^{PRES}, would be analogous with 4 times 19 features.

4.1. Classification problem and preconditioning on gait used

The basic classification problem is simply finding out T given F. This can be approached in a probabilistic fashion, $P(T|F)$, like in the naïve Bayes classifier, or by searching for a hyperplane separating the classes after a transformation of the input data in case of SVM. However, our main focus is two-fold. First, we want to investigate the impact of different gaits, $g \in G$, on this problem. That is, looking at the feature space F, the different terrains, $t \in T$, are not the only source of variability in the data, but are mixed with the effect of the different gaits. We will inspect their relative influence by applying Principal Component Analysis (PCA) as well as a nonlinear technique for visualizing multidimensional data in Section 5.1. Since the effect of the individual gaits on the feature space is crucial – as we will demonstrate – this should directly flow into the classification process.

There are several ways of incorporating the gait information. One could simply encode the gait as a single discrete feature, F^{GAIT}, which would transform the probabilistic formulation of the problem into $P(T|F^{SENSORS}, F^{GAIT})$. Another option would be to incorporate the additional motor information into the feature space by devising some motor features (e.g., amplitudes, offsets, and phase relationships of the sine waves that act as target motor signals), i.e. F^{MOTOR_PARAMS}. However, this is not the most effective solution. In particular, this is bound to fail for the naïve Bayes classifier due to its "naïve" assumption regarding conditional independence of the features. With this assumption and using Bayes theorem, the conditional distribution over the terrain class variable, T, becomes:

$$p(T|F_1, \ldots, F_n) = \frac{1}{Z} p(T) \prod_{i=1}^{n} p(F_i|T), \tag{1}$$

where the evidence $Z = p(F_1, \ldots, F_n)$ is a scaling factor dependent only on $F_1, \ldots, F_n$, that is, a constant if the values of the feature variables are known. From the term that remains on the right hand side, it is evident that the motor information will not improve the classification in this case: the gait or motor parameter features will become separate terms in $\prod_{i=1}^{n} p(F_i|T)$, e.g. $p(F^{GAIT}|T)$ will represent the probability of individual gaits being applied on different terrains (which carries no information about the terrain in our particular case with random choice of gait), while the joint probability distributions (sensory features attaining specific values under the application of specific gaits), which would be informative, are ignored. In the case of SVM, adding motor features would correspond to adding useful information into additional dimensions of the space, which may facilitate the placement of the separating hyperplane after the transformation of the space.

However, the solution that we pursue here is simpler and more effective: taking advantage of the fact that there is only a small number of gaits, the sensory data generated when applying these gaits will be kept separate, i.e. we will have for example

$$F^g = \{F^{ACC}, F^{KNEE}, F^{PRES}\}; \quad g = bound \tag{2}$$

and similarly for the other gaits. That is, we precondition on the gait used and train separate classifiers on the data generated as a result of application of different gaits. This will be compared with the "baseline" where sensory data from all gaits is put together in Section 5.2.

4.2. Effect of different sensory modalities on classification

Our second main objective is to study specifically the effect of different sensory modalities. The results are reported in Section 5.3. Two complementary approaches are employed. First, only features from a single modality are allowed. That is, $F = F^{ACC}$ only, for example. Second, the full feature set is used, but the data from the feature selection process is aggregated per modality and analyzed.

5. Results

5.1. Variability due to gaits vs. terrains

The interaction of a robot with different grounds gives rise to different patterns of sensory stimulation and can be exploited for terrain classification. However, the way in which this sensory stimulation is actively generated leads to significant transformations of the sensory patterns as well. In this section, we will look at how much of the variability in the data collected by a robot employing different gaits on different grounds can be attributed to the type of action used by the robot (the gait) and how much to the environment (the terrain).

[4] Available at http://www.csie.ntu.edu.tw/~cjlin/libsvm.

[5] Concretely, we used default settings where the radial basis function had the following form: $\exp(-\gamma * |u - v|^2)$, with u, v the feature vectors, and $\gamma = \frac{1}{nf}$, where nf was the number of features. The C-SVC (C-support vector classification) method with $C = 1$ was used.

First, we applied Principal Component Analysis (PCA). Two datasets were used: from real and simulated robot. For each dataset, the data collected from all gaits and grounds was put together and PCA was applied to the space of sensory features. Fig. 2(A, B) shows a projection onto the first two principal components (PCs): different markers correspond to different gaits, whereas different colors to different grounds the robot was running on (a random subset of the data was plotted). In the real robot (A), the trot backwards gait is nicely separated on the right; the bounding gait variants (straight, left and right) and the crawling gait are spread out in the left half of the graph. At the same time, irrespective of the gait used, the plastic foil ground (blue) is nicely separated by the 2nd PC at the top, with rubber occupying the other end of the figure – resembling the friction gradient of the terrains. In the simulated robot (right), the gaits are the dominant factor and determine the principal clusters in the data: bounding gaits at the bottom left, trot backwards in the bottom right, walk and crawl in the middle top, pace at the top right, and turn left and turn right in the right part of the figure. Some sub-clusters within these due to different terrains are visible.

In order to get a more quantitative understanding of the problem, we have also plotted contributions of individual PCs to the overall variance in the data (based on the eigenvalues of the principal vectors). This is depicted in Fig. 2(C, D). First, we see that in both cases, the first two PCs capture only about 40% of the variance and about 50 PCs are needed to capture 90% of variance (53 in the real, 49 in the simulated robot). In contrast, when data generated by individual gaits separately was taken, the variability dropped significantly: Between 24 and 35 PCs were needed in the real robot (PCA was rerun for data from every gait separately; plots not shown); in the simulated robot, it was only between 12 and 27, i.e. less than half compared to the case where data was collapsed across gaits.[6]

Second, Fig. 2(A, B) and (C, D) show that full variability of the data cannot be explained by a low-dimensional projection using a linear method. Therefore, in addition, we employed a non-linear method for visualizing high-dimensional data: Gaussian Process Latent Variable Models (GPLVM)—essentially a probabilistic dual version of PCA with a non-linear kernel [31]. A mapping between a latent space and an observed data space is used. In our case, in order to obtain a 2-dimensional projection of the data for visualization, the latent space dimensionality was set to 2. We used more recent extensions of the original method that employ sparse approximations to the full Gaussian process [32] and back constraints [33].[7] The results are depicted in Fig. 2(E, F). In the real robot, a mixed influence of the gaits and grounds can be seen. The trot gait is separated best in the bottom of the plot, split into three sub-clusters. The effect of terrain is pronounced too; in particular the foil ground data points are pulled to the left; Styrofoam and cardboard largely overlap. In the simulated robot, the gait variable is clearly dominating—resulting essentially in as many principal clusters as there were gaits.

In summary, this section demonstrated the difficulty of classifying terrains across different gaits. Therefore, our proposal is that considering sensory data generated by different motor patterns separately is a key prerequisite for successful terrain classification.

[6] The more dramatic drop in the simulator is at least partly to be explained by the fact that the differences between the grounds in the simulator are smaller (only friction differences, no texture, with the exception of the artificial "rough Styrofoam" with bumps).

[7] Specifically, we applied the fgplvm method from the GP lab MATLAB toolbox, https://github.com/SheffieldML/GPmat.git, with different settings: without and with back constraints and with different sparsification methods. Best results were obtained with the deterministic training conditional and back constraints – these are shown in the figure –, with 1000 iterations of the algorithm for the real and 500 for the simulated robot data.

Table 2
Classification accuracy per gait and per sensory modality—Real robot (%).

	SVM/(naïve Bayes)			
	All modalities	Acc	Knee	Pressure
Gaits together	95.2 (84.8)	68.6 (52.4)	95.1 (83.2)	75.0 (62.6)
Bound	98.3 (91.0)	65.6 (65.6)	96.1 (89.5)	92.6 (85.3)
Bound left	98.7 (96.3)	82.0 (80.1)	96.9 (95.0)	82.9 (78.9)
Bound right	98.3 (92.6)	71.6 (67.6)	96.2 (88.6)	90.3 (85.3)
Crawl	95.4 (96.3)	72.9 (66.9)	96.0 (95.8)	76.2 (74.5)
Trot back	99.0 (96.9)	83.5 (79.9)	98.7 (96.6)	71.0 (69.9)
Average	97.5 (93.0)	74.0 (68.8)	96.5 (91.4)	81.3 (76.1)

5.2. Preconditioning on gait for terrain classification

We have performed terrain classification on both the real and the simulated robot using the naïve Bayes and SVM classifier. The results are shown in Tables 2 (real robot) and 3 (simulated robot). In the real robot, when the data is collapsed across gaits and using all sensory modalities, an accuracy of 84.8% is obtained for the naïve Bayes, 95.2% for the SVM classifier (first cell in Table 2; means from 10 runs of 10-fold cross-validation). Performing the terrain classification for every gait separately – i.e. preconditioning on the gait used – results in significantly better results, attacking 100% for some gaits. The difference is more dramatic for the naïve Bayes classifier. A similar pattern is observed in the simulated robot (Table 3). For the gaits shared between the platforms, the sequence of the most potent gaits in terrain discrimination is the same: trot back, followed by the turning variants of the bounding gaits, indicating that certain coordination patterns facilitate terrain perception. The poor performance of the "walk" gait in the simulator is due to the fact that this gait was very unstable, resulting in many tumbling effects of the robot and a reduced number of samples collected.

Regarding the classification of individual terrains, we have computed confusion matrices (not shown here) for all the combinations displayed in Tables 2 and 3. They complement our intuition from the analysis of variability (Section 5.1, Fig. 2). In both datasets, the low-friction terrain, plastic foil, was the easiest ground to classify. In the real robot, the Styrofoam and cardboard were the most "confused" pair of terrains. In the simulated robot, this was the case for cardboard and rubber to some extent, while the rough Styrofoam was sometimes misclassified for either of the remaining three grounds.

5.3. Effect of different sensory modalities

In this section, we want to look deeper into the contribution of individual sensory channels. For analysis, we will group the sensory channels by their type and location in the robot – what we will call *sensory modality* – into accelerometers, knee joints (ang. pos. sensors in passive compliant knees), and feet pressure sensors.[8]

5.3.1. Classification using single modalities only

We studied the overall classification performance while the classifier could use features from individual sensory modalities only. In the real robot (Table 2), we see that the angular position sensors in the passive compliant knee joints were by far the most effective sensors. This can be at least partly explained by the fact that the interaction of the passive compliant knees with the ground

[8] Please note that in the following analysis we do not compensate for the fact that the accelerometers as a sensory modality consist of 3 sensory channels (and hence 3×19 features in the pool), while the other modalities (knee, pressure) comprise 4 channels.

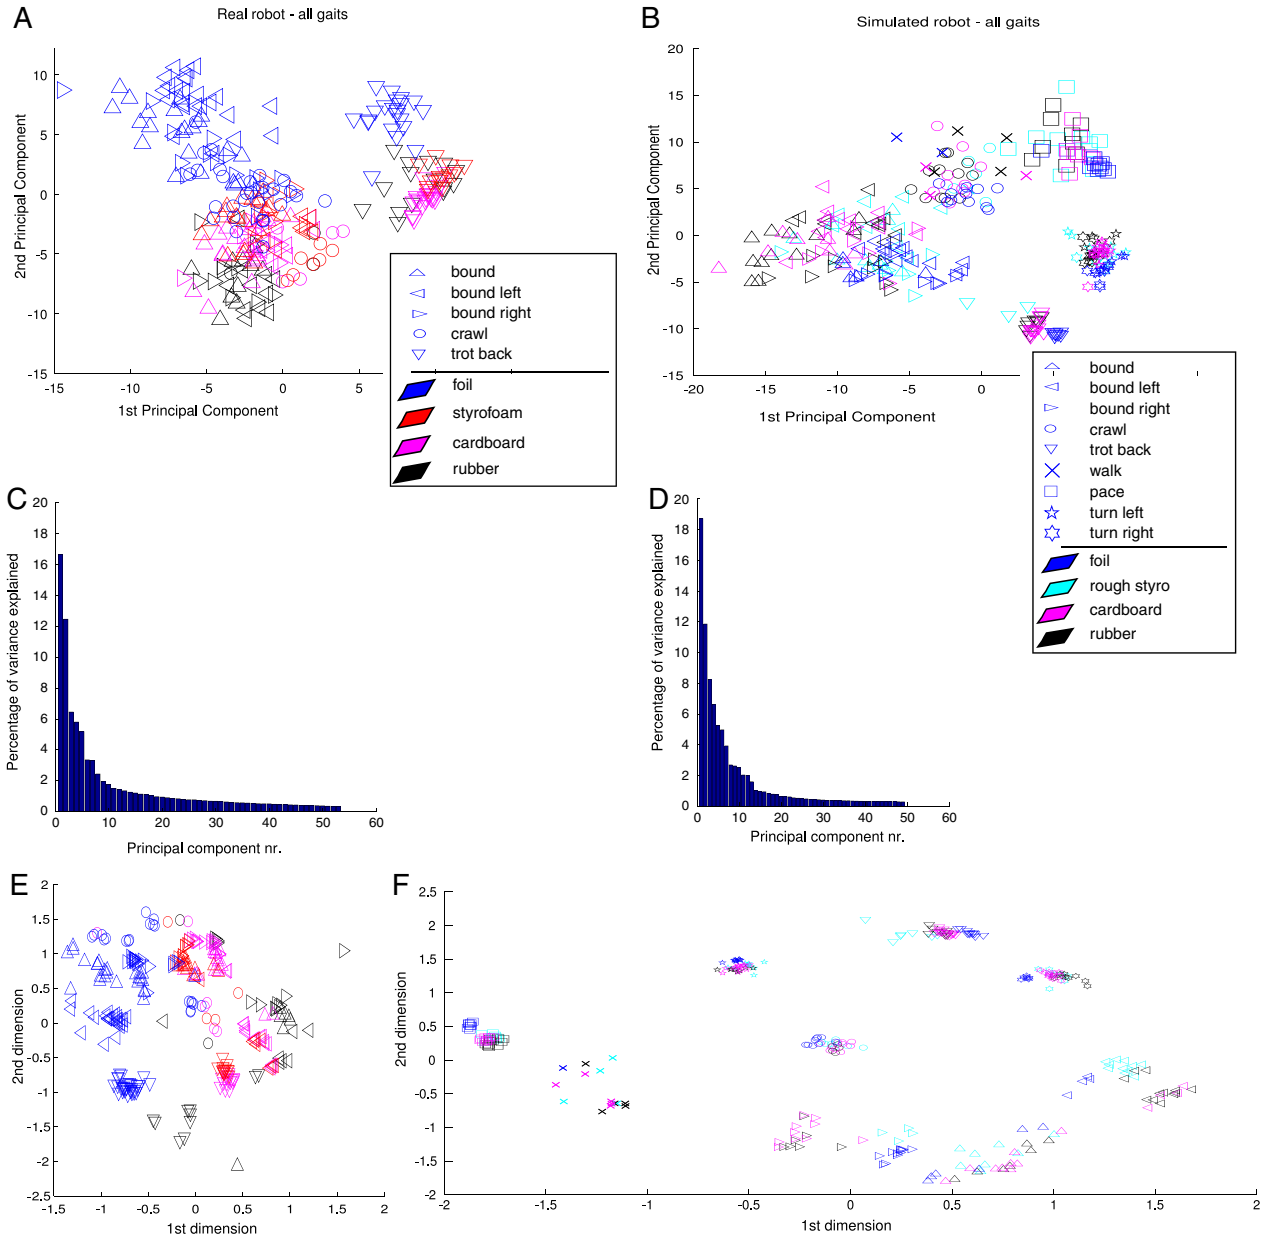

Fig. 2. Feature space visualization—variability due to gaits vs. terrains. Real robot data are on the left (A, C, E): robot was running with 5 different gaits (different markers in A, E) on 4 different grounds (foil—blue, Styrofoam—red, cardboard—magenta, rubber—black). Simulated robot data on the right (B, D, F). Robot was running with 9 different gaits on 4 different grounds (foil—blue, cardboard—magenta, rubber—black, rough Styrofoam—cyan in B and F). (A–B) Data in the space of first two principal components. (C–D) Contributions of first n principal components to explain 90% of variance in the data. (E–F) Data in the space of Gaussian Process Latent Variable Model with two latent variables (legend shared with A, B). Note: For visualization purposes, 1/4 of the data points were plotted in the real robot (A, E), 1/60 in the simulated (B, F). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 3
Classification accuracy per gait and per sensory modality—Simulated robot (%).

	SVM/(naïve Bayes)				
	All modalities	Acc	Hip	Knee	Pressure
Gaits together	81.2 (57.5)	77.9 (49.1)	67.8 (46.2)	77.1 (43.1)	62.6 (52.3)
Bound	81.6 (77.0)	79.1 (72.5)	76.9 (66.5)	71.9 (58.8)	63.6 (59.2)
Bound left	89.4 (84.4)	79.9 (71.7)	79.3 (75.2)	79.9 (69.1)	82.9 (67.7)
Bound right	87.8 (83.1)	77.9 (71.4)	77.4 (74.3)	75.2 (66.6)	81.4 (67.5)
Crawl	87.6 (86.4)	80.1 (72.1)	73.9 (74.6)	82.9 (79.2)	70.6 (61.4)
Trot back	91.1 (89.8)	90.6 (88.4)	39.6 (45.2)	87.7 (85.1)	73.3 (61.6)
Walk	70.8 (59.2)	67.2 (57.7)	57.7 (58.5)	61.5 (47.1)	61.7 (52.1)
Pace	94.6 (91.5)	93.2 (82.5)	82.5 (80.2)	92.1 (88.7)	79.5 (70.9)
Turn left	96.6 (90.9)	86.9 (79.3)	76.4 (74.3)	94.2 (91.0)	72.7 (67.2)
Turn right	95.3 (90.4)	87.7 (80.2)	76.8 (72.1)	94.0 (89.7)	73.5 (67.7)
Average	87.6 (81.0)	82.0 (72.5)	70.8 (66.7)	81.6 (71.8)	72.2 (62.7)

reflected the friction of the terrain: higher frictional forces resulted in greater bending of the joints. In the simulated robot (Table 3), the most effective were the knees and the accelerometers.

5.3.2. Aggregating data from feature selection

Whereas in the previous section the importance of every sensory modality was assessed based on the classification performance using only this particular set of sensory channels, here we perform a complementary analysis. From now on, only data from the real robot will be analyzed[9] and we will report results from SVM, the better performing classifier, only. As we described in Section 3.3, a forward feature selection algorithm was used to perform terrain classification. In 10 runs of the algorithm with 10-fold cross-validation, we summed the number of times every feature was selected. All the sensory channels and features were thus available (corresponding to the "All modalities" column of Table 2), but only a small portion of them (max. 10 out of 209 in every run) was actually picked. Running the algorithm first for all the gaits together and then for data from individual gaits, aggregating the counts of every feature, and finally grouping the "winning" features by sensory modality, gives rise to Fig. 3. We see a clear dominance of the angular position sensors in the knee joints—independent of the gait used, reconfirming our results from the previous section.

5.4. Control experiments with additional modalities and random motor commands

In order to further validate our findings from previous sections – regarding the effect of gaits and sensory modalities on terrain classification –, we performed additional experiments using the real robot equipped with additional sensors. This was the IMU (3-axis accelerometer, 3-axis gyroscope) and hip joint potentiometers (angular position sensors in the servo-controlled joints). As explained in Section 3.1, the addition of the IMU on the robot's back constrained the repertoire of gaits as well as terrains (difficult ridged grounds – cardboard and rubber – were left out and linoleum was added) and no gait transitions were induced.

The results of terrain classification are depicted in Table 4. The overall accuracies are higher, reflecting that the task was easier. However, the random "gait" – random low-passed motor commands – performed significantly worse, suggesting that coordinated behavior is a necessary basis for successful perception. Second, the performance of passive knee joint sensors is reconfirmed, but, the pressure sensors score even higher. This could be attributed to the fact that in this dataset all the grounds were flat (no ridges), hence better contact of the feet pressure sensors with the ground was guaranteed. Third, the data from the hip joint potentiometers and their poor performance highlights the difference in utility between the joint position sensing in active vs. passive compliant joints. At the same time, the drop in performance of the knees with the random gait indicates that a coordinated action of the leg is necessary to perceive the ground properties. Finally, the IMU data (accelerometers and gyroscopes) provides useful features as well.[10]

5.5. Effect of different features

We conclude our analysis by investigating the contributions of individual features in our set. Aggregate statistics across all sensory channels and gaits is shown in Fig. 4(left). The most

Table 4
Classification accuracy per gait and per sensory modality—Real robot 2 (%).

	SVM/(naïve Bayes)		
	Bound right	Turn left	Random
All modalities	100.0 (99.7)	100.0 (100.0)	94.4 (85.4)
Acc	97.5 (95.1)	96.8 (95.6)	76.3 (70.5)
Gyro	97.7 (95.7)	98.8 (95.5)	83.7 (77.5)
Hip	64.1 (63.0)	62.1 (64.7)	48.0 (44.4)
Knee	99.7 (98.9)	99.9 (98.9)	58.9 (56.3)
Pressure	99.9 (96.4)	100.0 (100.0)	70.1 (68.1)

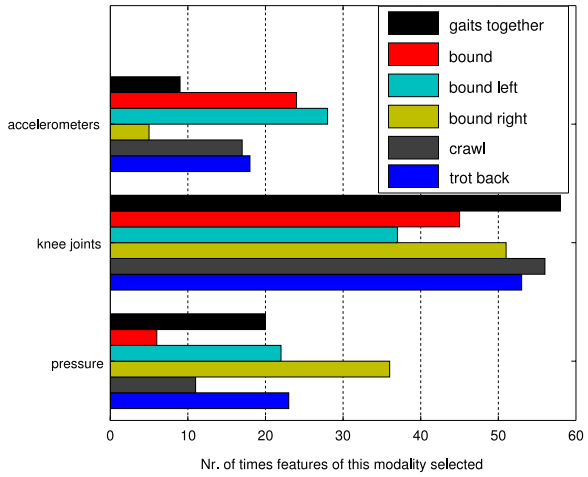

Fig. 3. Features of individual modalities in different gaits. (Real robot) In every run of the algorithm, at most 10 of the 209 features were selected. These were then summed up in 10 runs of the algorithm (separately for each gait) by the sensory modality. For example, if for the bound gait, in every of the 10 runs a front left knee joint skewness and a hind right knee joint median feature were selected, the knee modality would have a score of 20 for this gait.

frequently selected features are simple *Max* and then means of the 2nd differences (*Diff* 2 — *mean*) of every sensory channel per locomotion period. Furthermore, we studied the effectivity of different features in specific gaits and modalities. Individual gaits were not found to favor specific feature types (data not shown). The relative counts of every feature in different sensory modalities are shown in Fig. 4(right). For the knee sensors, it is the *Max* and *Diff* 2 — *mean* features that stand out; in the pressure sensors, it is also *Diff* 2 — *mean*, followed by frequency with highest amplitude, *FFT* 1*st*. For the accelerometers, the frequency features are slightly favored over the time domain features. Overall, unlike in the case of modalities, no very significant effect of specific features can be seen.

6. Discussion, conclusion, future work

In this work, we took a road less traveled in robot terrain discrimination and investigated the effect of action – gaits of a quadruped robot – on perception and studied how different gait-sensor-feature combinations impact on discrimination performance. A supervised learning scenario – classification – was used, since it allowed for easier quantification of the results. However, the findings should be equally applicable to unsupervised settings as well.

First, we will discuss the influence of the motor regimes. Using PCA as well nonlinear techniques to visualize high-dimensional data (GPLVM), we assessed the variability in the sensory data due to different motor patterns vs. due to different terrains. While the effect of these is intertwined, the gaits were found to account for a bigger part of the variability. Therefore, since the different motor regimes constitute information that is easily accessible to the robot, it is desirable to take it explicitly into account in the classification process. The gaits are substantially different coordination

[9] The main merit of the simulator was to investigate the effects of a larger gait repertoire. In order to analyze the properties of different sensory channels, however, only results from physical sensors will be reported.

[10] Note that the Real robot in Table 2 had a noisy on-board accelerometer only, whereas here (Table 4), the accelerometers from the IMU were used.

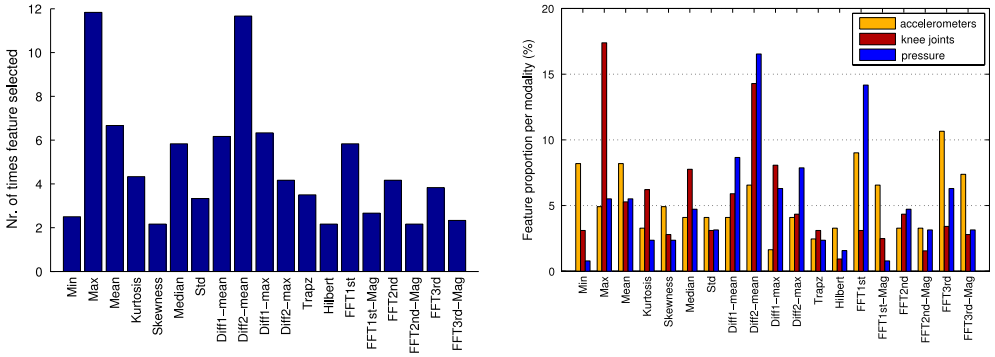

Fig. 4. Contribution of individual features. (Real robot) (*Left*) Overall counts of every feature. In every run of the algorithm, at most 10 of the 209 features were selected. These were then summed up in 10 runs of the algorithm (forward feature selection, SVM classifier) by the feature type across all sensory channels and finally averaged over results from different gaits. (*Right*) Relative counts of every feature per sensory modality. The count of every feature per sensory modality – sum over the sensory channels of the same type – is normalized by the overall count of the modality (accelerometers: 122, knee joints: 322, pressure 127).

patterns that induce distinct nonlinear transformations of the sensory flow and how the terrain is reflected in them. At the same time, the gaits are discrete entities and the number of them is limited. In light of these facts, rather than starting from the complex feature space where the effects of gaits and terrains are mixed and adding motor information to additional dimensions of the space, we suggest that the most effective solution is to precondition on the gait. That is, sensory data that was induced by the application of different gaits should be handled by different, gait-specific, classifiers. This was confirmed by our results: In all cases, significantly better classification performance was achieved on data coming from individual gaits, as opposed to the case where the data was collapsed across the different motor regimes. This effect was more pronounced in the case of the naïve Bayes classifier; SVM could deal with the variability coming from different motor regimes better.

Furthermore, different gaits lead to different overall performance — some turn out to be more suited for perception. This goes in the active perception tradition [34] which is popular in cognitive science, but has great potential in robotics too. Ugur et al. [35], for example, employ a similar approach by clustering data resulting from different actions of a robotic hand separately. Every coordinated motor pattern results in specific interaction with the environment and, consequently, in a specific spatio-temporal transformation of the incoming sensory data, which can be advantageous for subsequent perceptual tasks. This has been demonstrated by Lungarella and Sporns [36], for example, in different robots. We have confirmed the effect of coordinated vs. uncoordinated behavior on terrain perception by the drop in performance when random motor commands were applied.[11] Finally, the action context provided us with a natural way to parse the sensory data—into whole locomotion periods of the robot.

Second, we studied the role of different sensory channels – grouped into different sensory modalities – in the classification task. Overall, the combination of inertial, tactile and proprioceptive (joint angle) sensors turned out to be a powerful sensory suite for the task (in agreement with [20] where no features were used and different classifiers were applied). In particular, the angular position sensors in the passive compliant knee joints were found to be the most efficient sensory channel. This is, in our view, an interesting finding, as this type of sensor is often neglected in perception, although its variants are present in all robots that possess some compliant elements whose deflection can be measured (series elastic actuators for example). This further complements our previous study on the same platform, in which we showed that the knee joint sensors contribute the most to position increments estimation in a legged odometer [29]. This is also in line with recent work [37,38] where sensing through compliant artificial pneumatic muscles is studied. More generally, these results highlight the potential of contact-based sensing or "embodied information" [39] for robotics.

Third, the particular morphology of the robot and the properties of its sensors – or its embodiment – is the last component that is no less important for perceptual tasks. In our case, the nonlinear dynamics of the robot's interaction with the environment generates a rich stream of information that is preprocessed by the body and then picked up by the multimodal sensor collection. Iida and Pfeifer [40] have demonstrated this in a predecessor of our platform. An interesting future research direction would be to modify the robot's morphology, like the weight distribution or the stiffness of the springs in the knee joints, and study its effects on terrain discrimination.

Finally, while precise modeling and analytical understanding of the interaction of a legged robot with the ground and its projection into different sensory channels is practically unattainable, in this work, we have shown that a data-driven approach relying on a rich collection of sensors of different types – none of which was specifically designed to perceive terrain – and paying special attention to the active generation of the sensory data can lead to a close to perfect classification performance. The choice of particular sensory features was found to be less critical. In the future, we would like to test the generality of our approach by taking several different robot platforms – wheeled and legged with a similarly rich sensory set – and compare the relative importance of action and individual sensory modalities in terrain discrimination tasks.

Acknowledgments

M.H. was supported by the project FP7-ICT-270212 eSMCs. K.S. acknowledges the support of SGS grant No. 13/203/OHK3/3T/13 sponsored by the CTU in Prague. M.R. was supported by the Czech Science Foundation (Project Registration No. 14-13876S). M.H. would like to thank Rolf Pfeifer and Nico Schmidt from University of Zurich for support during this research. We are also indebted to Michal Vavrecka for helpful comments on this work.

References

[1] M. Reinstein, V. Kubelka, K. Zimmermann, Terrain adaptive odometry for mobile skid-steer robots, in: Proc. IEEE Int. Robotics and Automation, ICRA, Conf., IEEE, 2013, pp. 4706–4711.
[2] D. Tick, T. Rahman, C. Busso, N. Gans, Indoor robotic terrain classification via angular velocity based hierarchical classifier selection, in: Proc. IEEE Int. Robotics and Automation, ICRA, Conf., 2012, pp. 3594–3600.

[11] A detailed analysis of the information flows as a consequence of different gaits and grounds in our platform can be found in [28].

[3] P. Giguere, G. Dudek, Clustering sensor data for autonomous terrain identification using time-dependency, Auton. Robots 26 (2009) 171–186.
[4] L. Ojeda, J. Borenstein, G. Witus, R. Karlsen, Terrain characterization and classification with a mobile robot, J. Field Robot. 23 (2006) 103–122.
[5] C. Weiss, H. Tamimi, A. Zell, A combination of vision- and vibration-based terrain classification, in: Proc. IEEE/RSJ Int. Intelligent Robots and Systems, IROS, Conf., 2008, pp. 2204–2209.
[6] I. Halatci, C.A. Brooks, K. Iagnemma, Terrain classification and classifier fusion for planetary exploration rovers, in: Proc. IEEE Aerospace Conf., 2007, pp. 1–11.
[7] P. Giguere, G. Dudek, Surface identification using simple contact dynamics for mobile robots, in: Proc. IEEE Int. Robotics and Automation, ICRA, Conf., 2009, pp. 3301–3306.
[8] C. Weiss, H. Frohlich, A. Zell, Vibration-based terrain classification using support vector machines, in: Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems, IROS, 2006, pp. 4429–4434.
[9] E. Coyle, E.G. Collins, R.G. Roberts, Speed independent terrain classification using singular value decomposition interpolation, in: Proc. IEEE Int. Robotics and Automation, ICRA, Conf., 2011, pp. 4014–4019.
[10] P. Filitchkin, K. Byl, Feature-based terrain classification for littledog, in: Proc. IEEE/RSJ Int. Intelligent Robots and Systems, IROS, Conf., 2012, pp. 1387–1392.
[11] D. Belter, P. Skrzypczyński, Rough terrain mapping and classification for foothold selection in a walking robot, J. Field Robot. 28 (4) (2011) 497–528.
[12] M.A. Hoepflinger, C.D. Remy, M. Hutter, L. Spinello, R. Siegwart, Haptic terrain classification for legged robots, in: Proc. IEEE Int. Robotics and Automation, ICRA, Conf., 2010, pp. 2828–2833.
[13] K. Kim, K. Ko, W. Kim, S. Yu, C. Han, Performance comparison between neural network and SVM for terrain classification of legged robot, in: Proc. SICE Annual Conf. 2010, 2010, pp. 1343–1348.
[14] F.L. Garcia Bermudez, R.C. Julian, D.W. Haldane, P. Abbeel, R.S. Fearing, Performance analysis and terrain classification for a legged robot over rough terrain, in: Proc. IEEE/RSJ Int. Intelligent Robots and Systems, IROS, Conf., 2012, pp. 513–519.
[15] S. Manjanna, G. Dudek, P. Giguere, Using gait change for terrain sensing by robots, in: Tenth Conference on Computer and Robot Vision, CRV, 2013.
[16] A.C. Larson, R.M. Voyles, J. Bae, R. Godzdanker, Evolving gaits for increased discriminability in terrain classification, in: Proc. IEEE/RSJ Int. Intelligent Robots and Systems, IROS, Conf., IEEE, 2007, pp. 3691–3696.
[17] P. Giguere, G. Dudek, A simple tactile probe for surface identification by mobile robots, IEEE Trans. Robot. 27 (3) (2011) 534–544.
[18] E.G. Collins, E.J. Coyle, Vibration-based terrain classification using surface profile input frequency responses, in: Proc. IEEE Int. Robotics and Automation, ICRA, Conf., 2008, pp. 3276–3283.
[19] K. Iagnemma, F. Genot, S. Dubowsky, Rapid physics-based rough-terrain rover planning with sensor and control uncertainty, in: Proc. IEEE Int. Robotics and Automation, ICRA, Conf., vol. 3, 1999, pp. 2286–2291.
[20] J. Degrave, R. Cauwenbergh, F. wyffels, T. Waegeman, B. Schrauwen, Terrain classification for a quadruped robot, in: Proc. IEEE Int. Conf. Machine Learning and Applications, ICMLA, 2013.
[21] R. Jitpakdee, T. Maneewarn, Neural networks terrain classification using inertial measurement unit for an autonomous vehicle, in: Proc. SICE Annual Conf., 2008, pp. 554–558.
[22] C.A. Brooks, K. Iagnemma, Vibration-based terrain classification for planetary exploration rovers, IEEE Trans. Robot. 21 (6) (2005) 1185–1191.
[23] C. Brooks, K. Iagnemma, S. Dubowsky, Vibration-based terrain analysis for mobile robots, in: Proc. IEEE Int. Robotics and Automation, ICRA, Conf., 2005, pp. 3415–3420.
[24] E.M. DuPont, C.A. Moore, R.G. Roberts, Terrain classification for mobile robots traveling at various speeds: an eigenspace manifold approach, in: Proc. IEEE Int. Robotics and Automation, ICRA, Conf., 2008, pp. 3284–3289.
[25] P. Komma, C. Weiss, A. Zell, Adaptive Bayesian filtering for vibration-based terrain classification, in: Proc. IEEE Int. Robotics and Automation, ICRA, Conf., 2009, pp. 3307–3313.
[26] M. Hoffmann, N. Schmidt, R. Pfeifer, A. Engel, A. Maye, Using sensorimotor contingencies for terrain discrimination and adaptive walking behavior in the quadruped robot puppy, in: T. Ziemke, C. Balkenius, J. Hallam (Eds.), From Animals to Animats 12: Proc. Int. Conf. Simulation of Adaptive Behaviour (SAB), in: LNAI, vol. 7426, Springer, 2012, pp. 54–64.
[27] Webots, www.cyberbotics.com, commercial Mobile Robot Simulation Software. URL: http://www.cyberbotics.com.
[28] N. Schmidt, M. Hoffmann, K. Nakajima, R. Pfeifer, Bootstrapping perception using information theory: case studies in a quadruped robot running on different grounds, Adv. Complex Syst. J. 16 (2–3) (2013).
[29] M. Reinstein, M. Hoffmann, Dead reckoning in a dynamic quadruped robot based on multimodal proprioceptive information, IEEE Trans. Robot. 29 (2) (2013) 563–571.
[30] C.-C. Chang, C.-J. Lin, LIBSVM: a library for support vector machines, ACM Trans. Intell. Syst. Technol. (TIST) 2 (3) (2011) 27.
[31] N. Lawrence, Gaussian process latent variable models for visualisation of high dimensional data, Adv. Neural Inf. Process. Syst. 16 (2004) 329–336.
[32] N. Lawrence, Learning for larger datasets with the Gaussian process latent variable model, in: Proceedings of the Eleventh International Workshop on Artificial Intelligence and Statistics, 2007, pp. 243–250.
[33] N. Lawrence, J. Quinonero-Candela, Local distance preservation in the GP-LVM through back constraints, in: Proceedings of the 23rd International Conference on Machine Learning, ACM, 2006, pp. 513–520.
[34] R. Bajcsy, Active perception, Proc. IEEE 76 (8) (1988) 966–1005.
[35] E. Ugur, E. Oztop, E. Sahin, Goal emulation and planning in perceptual space using learned affordances, Robot. Auton. Syst. 59 (7–8) (2011) 580–595.
[36] M. Lungarella, O. Sporns, Mapping information flow in sensorimotor networks, PLoS Comput. Biol. 2 (2006) 1301–1312.
[37] S. Ikemoto, Y. Nishigori, K. Hosoda, Advantages of flexible musculoskeletal robot structure in sensory acquisition, Artif. Life Robot. 17 (1) (2012) 63–69.
[38] N. Sornkarn, M. Howard, T. Nanayakkara, Internal impedance control helps information gain in embodied perception, in: Proc. IEEE Int Robotics and Automation, ICRA, Conf., 2014.
[39] J. Schwendner, S. Joyeux, F. Kirchner, Using embodied data for localization and mapping, J. Field Robot. 31 (2) (2014) 263–295.
[40] F. Iida, R. Pfeifer, Sensing through body dynamics, Robot. Auton. Syst. 54 (2006) 631–640.

Matej Hoffmann received his Mgr. (M.Sc.) degree in Computer Science, Artificial Intelligence at Faculty of Mathematics and Physics, Charles University in Prague, Czech Republic, in 2006. Between 2006 and 2012 he completed his Ph.D. degree at the AI Lab, University of Zurich, Switzerland (Prof. Rolf Pfeifer). Since 2013, he is a Postdoctoral Researcher at iCub Facility, Italian Institute of Technology, with Prof. Giorgio Metta. His research interest is embodied cognition, in particular, the mechanisms underlying body representations and sensorimotor contingencies in humans and their implications for increasing the autonomy, resilience and robustness of robots.

Karla Štěpánová is a Ph.D. student at FEE CTU in Prague, Department of Cybernetics. She received her Mgr. (M.S.) degree in Physics of Condensed Matter at Faculty of Mathematics and Physics, Charles University in Prague, Czech Republic, in 2010. Her research interests are probabilistic models of cognition and EEG signal processing and analysis (with focus on higher cognitive functions).

Michal Reinštein, received the Ing. and Ph.D. degrees in Aircraft Information and Control Systems from the Czech Technical University in Prague in 2007 and 2011, where he is working as researcher at the Center for Machine Perception. His research interests concern state estimation and machine learning in robotic systems.