

Performance analysis and terrain classification for a legged robot over rough terrain

Fernando L. Garcia Bermudez, Ryan C. Julian, Duncan W. Haldane, Pieter Abbeel, and Ronald S. Fearing

Abstract—Minimally actuated millirobotic crawlers navigate unreliably over uneven terrain—even when designed with inherent stability—mostly because of manufacturing variabilities and a lack of good models for ground interaction. In this paper, we investigate the performance of a legged robot as it traverses three distinct rough terrains: tile, carpet, and gravel. Furthermore, we present an accurate, robust, low-lag, and efficient algorithm for terrain classification that uses vibration data from the on-board inertial measurement unit and motor control data from back-EMF sensing and magnetic encoders.

I. INTRODUCTION

Millirobotic crawling platforms have in general been optimized from a design perspective at the mechanical level. The optimization goal was to achieve a certain behavior optimally, such as sprinting [4][17], climbing [3], walking [18][21], or turning [28]. What was not often reported, though, is that although the design may be optimal for a specific task, manufacturing variability can result in poor performance. This is even more significant at the millirobot scale and was a particularly crucial limitation of the actuation mechanism of the Micromechanical Flying Insect [1].

This has prompted several groups to work on adapting the robot’s gait to achieve better performance at tasks the robot was not designed for (e.g. walking on granular media instead of hard ground [26]) or to achieve better performance than that achieved solely through hand tuning [34]. Although the above adaptation strategies implied changing the leg motion profile, more passive solutions, such as tuning the leg stiffness, have been attempted with moderate success [16][17].

Minimal actuation presents yet another obstacle for gait adaptation, since the robot can have a number of uncontrollable dynamic modes. In contrast, many large robots depend on either extensive actuation or precise sensors in order to traverse their environment [13][22].

In this paper, we first focus on performance analysis of the legged robot, similar to what was done for RHex [30], Sprawlita [6], and both DASH and DynaRoACH [25]. Without adapting gait, we investigate how varying the stride frequency has an effect on locomotion over three distinct

This material is based upon work supported by the National Science Foundation under Grants No. CNS-0931463 and IIS-0931463. D.W. Haldane is funded by the National Science Foundation IGERT Grant No. DGE-0903711.

F.L. Garcia Bermudez, R.C. Julian, P. Abbeel, and R.S. Fearing are with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720 USA {fgb@eecs., ryanjulian@, pabbeel@cs., ronf@eecs.}berkeley.edu

D.W. Haldane is with the Department of Mechanical Engineering, University of California, Berkeley, CA 94720 USA dhaldane@berkeley.edu



Fig. 1: Robotic platform, outfitted with motion capture tracking dots, positioned on top of the gravel used for experimentation.

rough terrains. Using telemetry gathered from the on-board sensors, we then focus on terrain classification.

Previous authors have focused on vibrational terrain classification for wheeled robots, but little work has been done for legged robots [12][20][35]. Vail and Veloso [33] and Larson, et al. [24] both recognized that complex walking dynamics present both distinct challenges and opportunities to the terrain classification problem. In particular, a robot which recognizes its terrain can adapt its control policy to traverse that terrain most effectively [10][23]. Furthermore, a walking robot has an advantage in its ability to change its interaction with the terrain to optimize its ability to classify that terrain, making a walking robot an *active sensor* of its environment [24].

We develop a terrain classification algorithm, based on Support Vector Machines (SVM), and a robust feature set, which is well-suited for on-line implementation on small resource constrained robots. We explore how the dynamics of the legged robot affect its ability to classify the terrain it traverses. We then characterize the classifier’s performance, and compare it to similar terrain classifiers for both legged and wheeled robots.

II. ROBOTIC PLATFORM AND CONTROL POLICY

For this work, we used a newly designed six-legged dynamic running robot, pictured in Fig. 1, which was based on the OctoRoACH robot [28]. It features an inboard dual drive-train for independent control of the right and left set of legs. Its SCM fabricated [18] transmission is highly rigid



Fig. 2: Experimental test bed with interchangeable surfaces: (left) laminate tile, (middle) gravel, and (right) medium-pile carpet.

due to planarizing parallel four-bar linkages which couple the drive-train output crank to kinematic linkages that govern the motion of the legs. We expect this increased rigidity to isolate the compliance of the system to the deformable polymeric legs, allowing for more predictable and repeatable dynamic terrain response. The overall mass of the robot, with battery, is 30.4 grams and it can attain a maximum stride frequency of 20 Hz.

The robot is driven by embedded sensing and control hardware [2], which includes an inertial measurement unit and an 802.15.4 wireless radio for issuing control commands and downloading telemetry data¹. Unlike in previous work [28], where the on-board controller was mainly dependent on the motors’ back-electromotive force (back-EMF), the current robot adds magnetic encoders to each motor, enabling control of the relative phasing between the two alternating tripods².

The kinematic linkages described previously effectively define a hardware-based Buehler clock [31], in which the touchdown and liftoff angles ($\phi_{TD,LO} = \pm 42^\circ$) and duty cycle ($\delta = 0.5$) are rigidly defined by the geometric parameters of the fourbar linkages [17]. Therefore, from the controller side, the Buehler clock becomes defined by one parameter: the stride period, t_c , or inversely the stride frequency, f_c .

To assess the robot’s performance, we varied the stride frequency when traversing laminate tile, gravel, and medium-pile carpet (see Fig. 2). The dimensions of the experimental test bed are $1.14 \text{ m} \times 1.75 \text{ m}$. Ground truth data was collected using an OptiTrack³ motion capture system comprised of eight V100:R2 cameras located .45 m above the terrain and capable of sub-millimeter accuracy as specified by the manufacturer.

III. ROBOT PERFORMANCE ON ROUGH TERRAIN

The performance metric used for the experiments in this section is the robot’s maximum speed as it traverses each

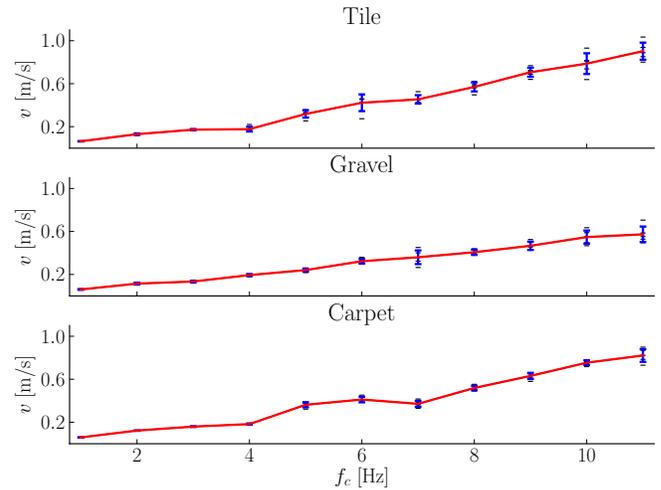


Fig. 3: Maximum speed achieved by the robot on three rough terrains as a function of the stride frequency ($n = 5$), along with a one standard deviation error bar. The raw data is plotted in black in the background. Note that the best performance is achieved on tile and is closely followed by that on carpet. The robot runs the worst on gravel, yet note the dips in performance for the other terrains at 4 and 7 Hz. These correspond to larger roll oscillations visible in Fig. 4.

terrain, which was calculated from the 3-D position data that the OptiTrack provides. Note that since the robot was not controlled for heading, the trajectories it took need not have been straight, particularly given that leg-ground interactions can substantially alter its direction. Nonetheless, the generalized traversal speed was computed using Python⁴ from the ℓ^2 -norm of the instantaneous position changes.

For each individual run, the robot’s legs were always started from the same position within their cycle. The position and orientation of the robot were not reset for each individual run. We always carefully monitored that the legs were indeed successfully reset before a new run was started, if not the case, we moved them into place by hand.

A. Performance as a Function of Stride Frequency

To gain a better understanding of how stride frequency variation affects the robot’s performance on the three terrains, we ran a parameter sweep on f_c of $1 \rightarrow 11$ Hz in steps of 1 Hz. Fig. 3 shows a plot of the maximum speed achieved at each stride frequency, along with a one standard deviation error bar. Each f_c setting was repeated five times ($n = 5$) and the resulting maximum speeds are shown in black as a scatter plot.

The deformable polymeric legs of the robot are tuned for best performance on hard ground, which explain why it achieves the greatest speeds on tile. This comes at a cost for our experiments, though, given the limited size of our test bed. Starting at about 6 Hz, the increased variance corresponds to the robot not always reaching its maximum speed while inside the tracking volume. The reported speed

¹Embedded board: https://github.com/biomimetics/imageproc_pcb

²Embedded code: <https://github.com/biomimetics/octoroch>

³NaturalPoint, Inc. OptiTrack: <http://www.naturalpoint.com/optitrack/>

⁴Scientific Tools for Python: <http://www.scipy.org/>

is thus a lower bound on the maximum achievable at those stride frequencies.

Apart from two dips in performance at 4 and 7 Hz, the robot’s speed increased with stride frequency on all terrains. The performance on carpet follows that of tile, but at a reduced overall speed. Since the carpet’s surface is more compliant than that of tile, the polymeric legs should require stiffening so as to achieve the same overall compliance between the robot and the terrain [14], thus improving its performance.

The lower speeds on gravel can be attributed to its fluidity. The gravel rocks are roughly the size of the robot’s feet (see Fig. 1) and as it moves over the terrain, the legs will either scoop rocks around or the robot’s feet will yield into the surface, impeding traction. As shown in previous studies of robot locomotion on fluidizing ground, leg dynamics tuned for hard ground generally need to be modified to achieve decent performance on flowing terrains [25][26][29].

Fig. 4 is a plot of the frequency spectra of the robot’s oscillations in the body frame as a function of the stride frequency. The spectra are averaged across the five sample runs taken at each stride frequency.

Note that the larger oscillatory peaks correspond to the robot’s roll dynamics. In particular, the roll frequency spectra on tile has distinct peaks at 4 and 7 Hz, which correspond to the dips in performance visible in Fig. 3. Since these frequencies excite large roll oscillations on the robot, some of the locomotive energy will be expended oscillating instead of running forward and the performance is thus reduced [8].

On carpet, the reduced performance at higher stride frequencies can now be understood in relation to the larger roll oscillations above 6 Hz, probably caused by the excessive compliance in the leg-ground interactions. The flowing gravel, on the other hand, might be helping dampen oscillations, but the corresponding loss of traction guarantees reduced locomotive speed.

IV. TERRAIN CLASSIFICATION

Terrain-sensitive gait adaptation is only effective if robots can identify terrains while they traverse them. We develop an accurate and robust terrain classification algorithm, using Support Vector Machines (SVM), that uses only data from the robot’s existing on-board IMU, back-EMF sensors, and magnetic encoders. We present a robust feature set, based on statistical moments, that minimizes computation and lag, and meets or exceeds the performance and robustness of larger and more computationally intensive feature sets from the literature. We discuss the relationship between stride frequency and terrain separability in feature space, and its implications to gait adaptation and active sensing. Finally, we analyze the algorithm’s performance using data from three terrains on our test bed, discuss the effect of episode length on classifier performance, and evaluate the performance of our features using principal component analysis.

A. Classification Algorithm Overview

Stiffness, damping, cohesion, friction, and surface irregularity comprise the terrain interface presented to a robot. As

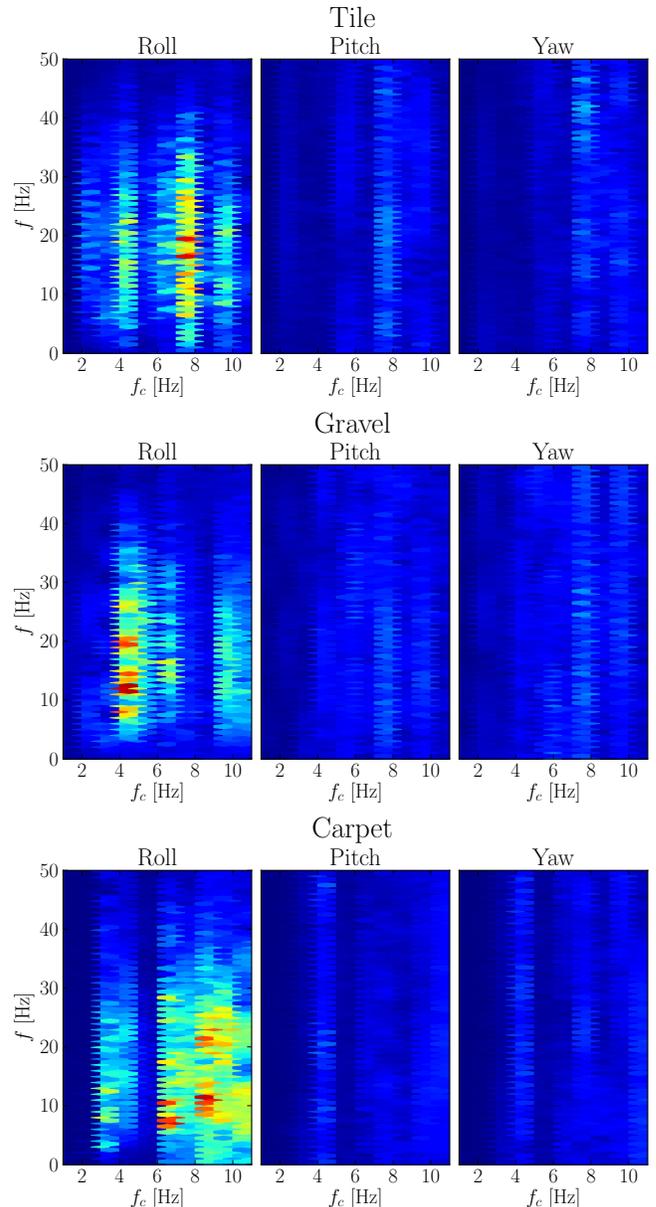


Fig. 4: Average frequency spectra of the robot’s oscillations on three rough terrains as a function of the stride frequency ($n = 5$). Note how roll oscillations dominate the dynamics and, in particular, how larger oscillations generally correlate to the dips in performance seen in Fig. 3.

a robot traverses the terrain, these terrain properties couple with the robot dynamics to produce vibrational signatures in body motion [20]. These vibrational signatures are especially significant in walking robots, whose periodic gait patterns exert a large range of contact forces and angles on the terrain, and excite terrain-robot vibration modes not seen by less dynamic locomotion strategies.

We measure these vibrational signatures using the robot’s on-board IMU, which senses translational accelerations and angular rates in three dimensions: x , y , z , yaw (ψ), $pitch$ (θ), and $roll$ (ϕ). At each of two motor shafts, on-board

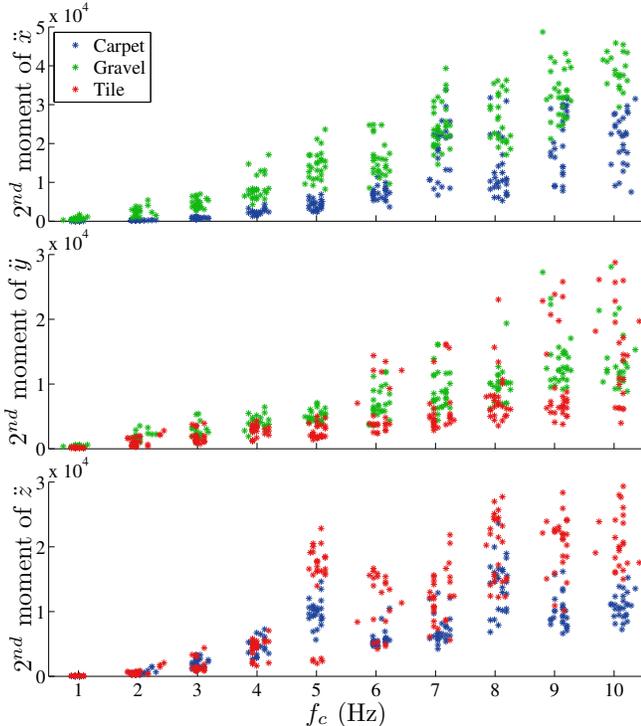


Fig. 5: Scatter plots of stride frequency (f_c) and the second moments of x , y , and z acceleration for selected terrains. Note how separability increases with stride frequency, as higher energy robot-terrain dynamics produce stronger vibrational terrain signatures.

magnetic encoders and back-EMF sensing (\mathcal{E}_L and \mathcal{E}_R) augment vibration measurements by sensing gross and instantaneous shaft angular velocity, respectively. These allow the algorithm to incorporate information about stride frequency—approximating velocity—and slippage, two parameters which presented challenges to previous terrain classification efforts. These sensors evolve ten time series, from which we extract features (see Section IV-B) that parameterize the vibrational signatures of terrains.

We then use an SVM to discriminate these terrains in feature space. Our implementation uses a soft-margin SVM [9] with a Gaussian radial basis function (RBF) kernel, and a one-against-one multi-class decision algorithm [19]. We chose SVMs for their versatility and modest computational requirements at classification time. The latter is important to our application, as we intend to implement this algorithm on the robot’s low-power microcontroller [2].

B. Feature Extraction

Our features operate on time series data that capture the vibrational signatures of terrains. Before performing feature extraction, we split the sensor data time series into sequential, disjoint episodes. There is a trade-off between episode length and classification accuracy, which is discussed in Section IV-C.

The dynamic interactions between our walking robots and the terrain make feature extraction challenging. Ter-

rain interaction signals combine with gait oscillations and the chaotic vibrations they cause in the robot’s compliant body to produce the vibration signatures seen at the IMU, confounding simple frequency response-based methods. The most effective of the many features we evaluated in initial experimentation were mean and variance, the first and second statistical moments of a sample set.

We followed this intuition and introduced higher statistical moments to the feature set. Statistical moments are used as features for classification of sound [15], radio signals [32], and other oscillating time series, as well as for pattern recognition in computer vision [27]. Our work shows that they are an effective and generalizable feature set for terrain classification in the face of chaotic robot-terrain interactions. Equation 1 gives the k^{th} statistical moment μ_k of a sample x with n members x_1, \dots, x_n .

$$\mu_k = \frac{1}{n} \sum_{i=1}^n x_i^k \quad (1)$$

Purely deterministic or functional decision methods often break down in the face of chaotic and noisy information generated by stochastic systems, such as that of a legged robot coupled with terrain. Our feature set characterizes the stochastic system probabilistically, thus avoiding this brittleness. The statistical moments of a sample are unbiased estimators of the moments of that sample’s underlying PDF [36]. Using relatively few samples, moment statistics succinctly capture information about the conditional probability distribution over the measurements in a time series, for a given terrain. Additionally, once the kernel transformation has been applied, SVMs are minimum ℓ^2 -norm linear discriminators [9]. Our classifier is, in effect, a minimum ℓ^2 -norm *method of moments* discriminator. It classifies a given terrain traversal episode by comparing that episode’s estimated PDF, parameterized by sample moments measured during a terrain interaction episode, with those of the representative PDFs for each terrain-IMU signal pair, parameterized by the training sample moments.

The only non-moment feature we introduced to the classifier is stride frequency, which we estimated using the robot’s on-board magnetic encoders. Weiss, et al. [35], Komma, et al. [23], and DuPont, et al. [12] all discuss variable speed as interfering with otherwise-effective terrain classifiers. As discussed above, the locomotion speed for walking robots can be parameterized by stride frequency. We found that robot dynamics vary significantly with stride frequency and increase the complexity of the classification problem. We were able to classify training data collected at only two values of stride frequency with 95% accuracy using only two features, but require at least fifteen (per the analysis in Section IV-D) to classify over the entire range of stride frequencies in our data set. Coyle, et al. addressed this problem in wheeled vehicles by introducing speed to the feature set [11]. We find that introducing stride frequency as a feature is analogously successful, and allows our algorithm to adapt to different dynamic regimes.

TABLE I: Terrain Classifier Confusion Matrix

		Predicted		
		Carpet	Gravel	Tile
True	Carpet	97.44%	0.00%	2.56%
	Gravel	0.00%	90.36%	9.64%
	Tile	3.80%	2.53%	93.67%

Although the complex dynamics of legged robots present a challenge to classification, we also find they have significant advantages. As portrayed in Fig. 5, the separability of the terrains in feature space varies markedly with stride frequency. At very low stride frequencies (below 2 Hz), robot locomotion does not excite terrain-robot dynamics interaction to a degree sufficient to separate the terrains for accurate classification. Above this threshold and in the typical operating range of the robot (above 2 Hz), these dynamic interactions produce distinct terrain vibration signatures, which allow for very accurate classification. Our robots are typically operated at a stride frequency above 2 Hz, so we trained our classifiers over this range.

Our final feature set comprises just twenty-five features: the second, third, and fourth moments (variance, skew, and kurtosis) of the accelerometer, gyroscope, and back-EMF time series in each episode, and average stride frequency over the episode as calculated from the magnetic encoder.

These features are computationally simpler than previous terrain classifiers, most of which use many Fast Fourier Transform or Power Spectral Density coefficients for feature vectors [5][11][12][23][24][35]. The moment features can be computed easily and efficiently on-line by sliding window methods, which were successfully used by Coyle, et al. to improve the responsiveness and efficiency of terrain classifiers on wheeled vehicles [10]. These methods update the current moment statistic by adding residuals, rather than carrying out the computation over the entire window for each new sample.

C. Experimental Results and Performance

We trained the classifier to distinguish between the three terrains in our test bed: laminate tile, gravel, and medium-pile carpet. We collected 165 three-second locomotion trials on the test bed, covering all three terrains, sweeping the stride frequencies between 1 and 11 Hz. Our classifier training software divided these trials into 1155×350 millisecond labeled training episodes, which we used to train and test the classifier for the results presented here.

We used MATLAB⁵ for data processing and presentation and the LIBSVM [7] library for classifier training and testing. We reserved 25% of the collected episodes for final testing, and used the remaining 75% for training and cross-validation. Before training, we normalized all features to the range $[0, 1]$ to ensure the best performance of the LIBSVM training algorithm.

⁵The MathWorks, Inc. Matlab: <http://www.mathworks.com/products/matlab/>

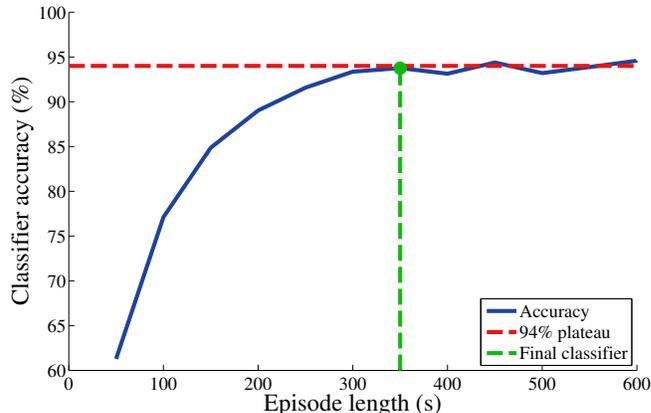


Fig. 6: Effect of episode length on classifier accuracy. The accuracy plateaus to 94% at around 350 milliseconds, our chosen episode length. Episode lengths above 550 milliseconds do not provide enough training data to prevent overfitting, and were neglected. Each classifier was trained and tuned using the same algorithms as the final classifier, which is a single point on this line at an episode length of 350 milliseconds.

To maximize classification performance, we performed a grid search over the SVM parameters: the Gaussian RBF shape parameter γ and the soft-margin c . The grid search iterates over a coarse 2D grid of parameters, then over a fine 2D grid within a tight range of the best parameters found in the first pass. For every pair of parameters, it trains the SVM using the training episodes with 10-fold cross-validation, and uses the cross-validation accuracy as the final score.

Finally, we used the 25% held-out test episodes to characterize the classifier with the highest overall accuracy found by the grid search. This resulted in a terrain classifier with 93.8% overall accuracy. We present detailed accuracy results for this classifier in Table I.

We found that episode length affects classifier accuracy dramatically. This is an important trade-off, because episode length determines the number of observations necessary to make a successful classification, and thus the lag between terrain transition and successful classification for the on-line algorithm. Fig. 6 shows the relationship between episode length and classifier accuracy. Classifier accuracy plateaus at around 94%. We tuned our classifier to 350 milliseconds, the smallest lag at the plateau. This is significantly shorter than the best classification times reported for similar algorithms on both legged and non-legged robots, which range from 500 milliseconds to 2.2 seconds [10][24].

D. Feature Performance Analysis

We used principal component analysis (PCA) to evaluate the effectiveness of our features. Fig. 7 summarizes our findings. Seven principal components explain 90% of the observed variance, ten explain 95%, and fifteen explain 99%, suggesting that the data has an effective rank of no more than 15. This exceeds the information density of previous

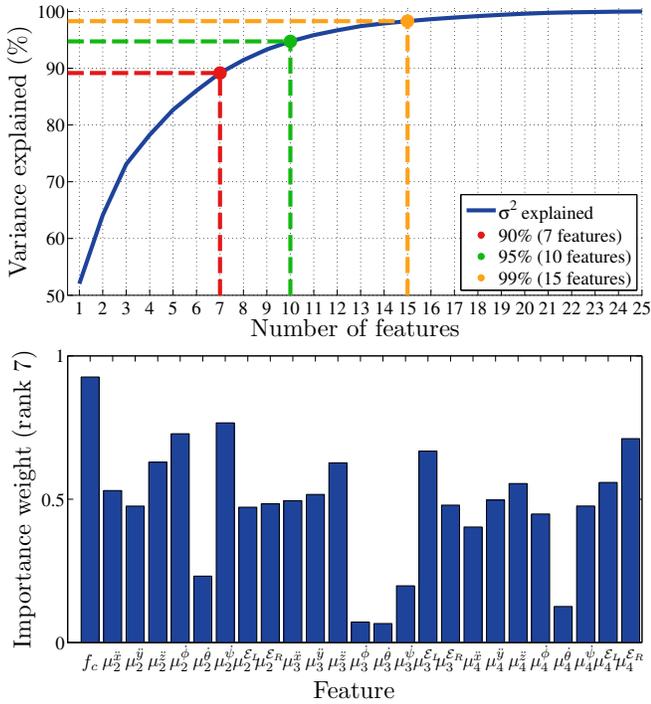


Fig. 7: Explained variance plot and normalized feature weights for a rank-7 approximation of the training features. Seven principal components explain 90% of the observed variance, ten explain 95%, and fifteen explain 99%, suggesting that the features have an effective rank of no more than 15.

pure FFT-based approaches by at least a factor of two, and is consistent the previous success of Weiss, et al. with non-FFT features. Brooks, et al. found that 15 FFT components were necessary to explain 90% of the variance for FFT features [5], and Weiss, et al. required 52 FFT components to explain 95% [35].

The high importance weight assigned to f_c confirms a strong relationship between separability and stride frequency (as discussed in Section IV-B). The highly-correlated nature of our time series partially explains the relatively low effective rank of the feature set. For example, all three moments of the pitch rate signal (μ_2^ϕ , μ_3^ϕ , and μ_4^ϕ) have near-zero importance weight in the rank-7 approximation. The pitch and roll rates of an ambulating robot are highly correlated, so little information is gained from the addition of the pitch rate features given the roll rate features, and their importance is devalued in the low-rank approximation. This analysis could form the foundation of a PCA-based feature selection and dimensionality reduction, such as those used by Komma, et al. [23], Brooks, et al. [5], and Weiss, et al. [35] for terrain classification.

V. CONCLUSIONS AND FUTURE WORK

The performance analysis of the robot traversing different rough terrains has highlighted the importance of leg tuning and its relationship to terrain characteristics, such as its

compliance. We also presented a terrain classifier and feature set which is accurate, robust to variations in stride frequency, low-lag, and efficient enough to be implemented on-line using the robot’s low-power microcontroller.

In future work, we will investigate how varying the intra-stride motion profile can improve the robot’s performance for terrains for which its legs weren’t tuned. We will also implement terrain classification on-line and use it to switch to the most efficient gait for those conditions.

VI. ACKNOWLEDGMENTS

The authors would like to thank Paul Birkmeyer and David Zarrouk for insightful discussions leading to the work in this paper, Humphrey Hu for help with data collection and set up of the motion capture arena, Andrew Pullin for help with the OctoRoACH embedded code and for photographing the robot, and members of the Biomimetic Millisystems Lab and the EECS community at the University of California, Berkeley for their advice and support.

REFERENCES

- [1] S. Avadhanula, R. J. Wood, E. Steltz, J. Yan, and R. S. Fearing, “Lift force improvements for the micromechanical flying insect,” in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, Oct. 2003, pp. 1350–1356. [Online]. Available: <http://dx.doi.org/10.1109/IROS.2003.1248832>
- [2] S. S. Baek, F. L. Garcia Bermudez, and R. S. Fearing, “Flight control for target seeking by 13 gram ornithopter,” in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, Sept. 2011, pp. 2674–2681. [Online]. Available: <http://dx.doi.org/10.1109/IROS.2011.6094581>
- [3] P. Birkmeyer, A. G. Gillies, and R. S. Fearing, “CLASH: Climbing vertical loose cloth,” in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, Sept. 2011, pp. 5087–5093. [Online]. Available: <http://dx.doi.org/10.1109/IROS.2011.6094905>
- [4] P. Birkmeyer, K. Peterson, and R. S. Fearing, “DASH: A dynamic 16g hexapedal robot,” in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, Oct. 2009, pp. 2683–2689. [Online]. Available: <http://dx.doi.org/10.1109/IROS.2009.5354561>
- [5] C. A. Brooks and K. Iagnemma, “Vibration-based terrain classification for planetary exploration rovers,” *IEEE Transactions on Robotics*, vol. 21, no. 6, pp. 1185–1191, Dec. 2005. [Online]. Available: <http://dx.doi.org/10.1109/TRO.2005.855994>
- [6] J. G. Cham, J. K. Karpick, and M. R. Cutkosky, “Stride period adaptation of a biomimetic running hexapod,” *The International Journal of Robotics Research*, vol. 23, no. 2, pp. 141–153, Feb. 2004. [Online]. Available: <http://dx.doi.org/10.1177/0278364904041323>
- [7] C. C. Chang and C. J. Lin, “LIBSVM: A library for support vector machines,” *ACM Trans. Intell. Syst. Technol.*, vol. 2, no. 3, May 2011. [Online]. Available: <http://portal.acm.org/citation.cfm?id=1961199>
- [8] J. E. Clark and M. R. Cutkosky, “The effect of leg specialization in a biomimetic hexapedal running robot,” *Journal of Dynamic Systems, Measurement, and Control*, vol. 128, no. 1, pp. 26+, March 2006. [Online]. Available: <http://dx.doi.org/10.1115/1.2168477>
- [9] C. Cortes and V. Vapnik, “Support-vector networks,” *Mach. Learn.*, vol. 20, no. 3, pp. 273–297, Sept. 1995. [Online]. Available: <http://dx.doi.org/10.1023/A:1022627411411>
- [10] E. Coyle, E. G. Collins, and L. Lu, “Updating control modes based on terrain classification,” in *IEEE Int. Conf. on Robotics and Automation*, May 2010, pp. 4417–4423. [Online]. Available: <http://dx.doi.org/10.1109/ROBOT.2010.5509870>
- [11] E. Coyle and E. G. Collins, Jr, “A comparison of classifier performance for vibration-based terrain classification,” in *26th Army Science Conference*, Orlando, Florida, Dec. 2008. [Online]. Available: <http://oai.dtic.mil/oai/oai?verb=getRecord&metadataPrefix=html&identifier=ADA505696>
- [12] E. M. DuPont, R. G. Roberts, M. F. Selekwa, C. A. Moore, and E. G. Collins, “Online terrain classification for mobile robots,” in *ASME Conference*, no. 42169, Nov. 2005, pp. 1643–1648. [Online]. Available: <http://dx.doi.org/10.1115/IMECE2005-81659>

- [13] K. Espenschied, R. Quinn, R. Beer, and H. Chiel, "Biologically based distributed control and local reflexes improve rough terrain locomotion in a hexapod robot," *Robotics and Autonomous Systems*, vol. 18, no. 1-2, pp. 59–64, July 1996. [Online]. Available: [http://dx.doi.org/10.1016/0921-8890\(96\)00003-6](http://dx.doi.org/10.1016/0921-8890(96)00003-6)
- [14] D. P. Ferris, M. Louie, and C. T. Farley, "Running in the real world: adjusting leg stiffness for different surfaces," *Proceedings of the Royal Society of London. Series B: Biological Sciences*, vol. 265, no. 1400, pp. 989–994, June 1998. [Online]. Available: <http://dx.doi.org/10.1098/rspb.1998.0388>
- [15] J.-W. Fu, Y.-C. Qi, and J.-S. Yuan, "Wavelet domain audio steganalysis based on statistical moments and PCA," in *IEEE Int. Conf. on Wavelet Analysis and Pattern Recognition*, vol. 4, Nov. 2007, pp. 1619–1623. [Online]. Available: <http://dx.doi.org/10.1109/ICWAPR.2007.4421711>
- [16] K. C. Galloway, J. E. Clark, M. Yim, and D. E. Koditschek, "Experimental investigations into the role of passive variable compliant legs for dynamic robotic locomotion," in *IEEE/RSJ Int. Conf. on Robotics and Automation*, Shanghai, China, 2011. [Online]. Available: <http://dx.doi.org/10.1109/ICRA.2011.5979941>
- [17] A. M. Hoover, S. Burden, X.-Y. Fu, S. Shankar Sastry, and R. S. Fearing, "Bio-inspired design and dynamic maneuverability of a minimally actuated six-legged robot," in *IEEE RAS/EMBS Int. Conf. on Biomedical Robotics and Biomechatronics*, Sept. 2010, pp. 869–876. [Online]. Available: <http://dx.doi.org/10.1109/BIOROB.2010.5626034>
- [18] A. M. Hoover, E. E. Steltz, and R. S. Fearing, "RoACH: An autonomous 2.4g crawling hexapod robot," in *IEEE Int. Conf. on Intelligent Robots and Systems*, Nice, France, Sept. 2008, pp. 26–33. [Online]. Available: <http://dx.doi.org/10.1109/IROS.2008.4651149>
- [19] C.-W. Hsu and C.-J. Lin, "A comparison of methods for multiclass support vector machines," *IEEE Transactions on Neural Networks*, vol. 13, no. 2, pp. 415–425, Mar. 2002. [Online]. Available: <http://dx.doi.org/10.1109/72.991427>
- [20] K. Iagnemma, H. Shibly, and S. Dubowsky, "On-line terrain parameter estimation for planetary rovers," in *IEEE Int. Conf. on Robotics and Automation*, May 2002, pp. 3142–3147. [Online]. Available: <http://dx.doi.org/10.1109/ROBOT.2002.1013710>
- [21] N. J. Kohut, A. M. Hoover, K. Y. Ma, S. S. Baek, and R. S. Fearing, "MEDIC: A legged millirobot utilizing novel obstacle traversal," in *IEEE Int. Conf. on Robotics and Automation*, May 2011, pp. 802–808. [Online]. Available: <http://dx.doi.org/10.1109/ICRA.2011.5980360>
- [22] J. Z. Kolter, M. P. Rodgers, and A. Y. Ng, "A control architecture for quadruped locomotion over rough terrain," in *IEEE Int. Conf. on Robotics and Automation*, May 2008, pp. 811–818. [Online]. Available: <http://dx.doi.org/10.1109/ROBOT.2008.4543305>
- [23] P. Kamma, C. Weiss, and A. Zell, "Adaptive bayesian filtering for vibration-based terrain classification," in *IEEE Int. Conf. on Robotics and Automation*, May 2009, pp. 3307–3313. [Online]. Available: <http://dx.doi.org/10.1109/ROBOT.2009.5152327>
- [24] A. C. Larson, R. M. Voyles, J. Bae, and R. Godzdzank, "Evolving gaits for increased discriminability in terrain classification," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, Oct. 2007, pp. 3691–3696. [Online]. Available: <http://dx.doi.org/10.1109/IROS.2007.4399500>
- [25] C. Li, A. M. Hoover, P. Birkmeyer, P. B. Umbanhowar, R. S. Fearing, and D. I. Goldman, "Systematic study of the performance of small robots on controlled laboratory substrates," in *SPIE Conf. on Micro- and Nanotechnology Sensors, Systems, and Applications*, Apr. 2010, pp. 76790Z–76790Z–13. [Online]. Available: <http://dx.doi.org/10.1117/12.851047>
- [26] C. Li, P. B. Umbanhowar, H. Komsuoglu, D. E. Koditschek, and D. I. Goldman, "Sensitive dependence of the motion of a legged robot on granular media," *Proceedings of the National Academy of Sciences*, Feb. 2009. [Online]. Available: <http://dx.doi.org/10.1073/pnas.0809095106>
- [27] R. J. Prokop and A. P. Reeves, "A survey of moment-based techniques for unoccluded object representation and recognition," *CVGIP: Graphical Models and Image Processing*, vol. 54, no. 5, pp. 438–460, Sept. 1992. [Online]. Available: [http://dx.doi.org/10.1016/1049-9652\(92\)90027-U](http://dx.doi.org/10.1016/1049-9652(92)90027-U)
- [28] A. O. Pullin, N. J. Kohut, D. Zarrouk, and R. S. Fearing, "Dynamic turning of 13 cm robot comparing tail and differential drive," in *IEEE Int. Conf. on Robotics and Automation*, May 2012. [Online]. Available: <http://dx.doi.org/10.1109/ICRA.2012.6225261>
- [29] F. Qian, T. Zhang, C. Li, P. Masarati, A. M. Hoover, P. Birkmeyer, A. Pullin, R. S. Fearing, and D. I. Goldman, "Walking and running on yielding and fluidizing ground," in *Robotics: Science and Systems (RSS) 2012*, Sidney, Australia, July 2012.
- [30] U. Saranlı, M. Buehler, and D. E. Koditschek, "RHex: A simple and highly mobile hexapod robot," *The International Journal of Robotics Research*, vol. 20, no. 7, pp. 616–631, July 2001. [Online]. Available: <http://dx.doi.org/10.1177/02783640122067570>
- [31] J. Seipel and P. Holmes, "A simple model for clock-actuated legged locomotion," *Regular and Chaotic Dynamics*, vol. 12, no. 5, pp. 502–520, Oct. 2007. [Online]. Available: <http://dx.doi.org/10.1134/S1560354707050048>
- [32] S. S. Soliman and S. Z. Hsue, "Signal classification using statistical moments," *IEEE Transactions on Communications*, vol. 40, no. 5, pp. 908–916, May 1992. [Online]. Available: <http://dx.doi.org/10.1109/26.141456>
- [33] D. Vail and M. Veloso, "Learning from accelerometer data on a legged robot," in *IFAC/EURON Symposium on Intelligent Autonomous Vehicles*, 2004. [Online]. Available: <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.1.9900>
- [34] J. D. Weingarten, G. A. D. Lopes, M. Buehler, R. E. Groff, and D. E. Koditschek, "Automated gait adaptation for legged robots," in *IEEE Int. Conf. on Robotics and Automation*, vol. 3, Apr. 2004, pp. 2153–2158 Vol.3. [Online]. Available: <http://dx.doi.org/10.1109/ROBOT.2004.1307381>
- [35] C. Weiss, H. Frohlich, and A. Zell, "Vibration-based terrain classification using support vector machines," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, Oct. 2006, pp. 4429–4434. [Online]. Available: <http://dx.doi.org/10.1109/IROS.2006.282076>
- [36] J. M. Wooldridge, "Applications of generalized method of moments estimation," *The Journal of Economic Perspectives*, vol. 15, no. 4, pp. 87–100, 2001. [Online]. Available: <http://www.jstor.org/stable/2696518>