Environment Identification for a Running Robot Using Inertial and Actuator Cues

Philippe Giguere and Gregory Dudek and Chris Prahacs and Shane Saunderson
Centre for Intelligent Machines
McGill University
Montréal, QC, H3A 2A7, Canada.
Email: {philg,dudek,cprahacs,shane}@cim.mcgill.ca

Abstract— In this paper, we explore the idea of using inertial and actuator information to accurately identify the environment of an amphibious robot. In particular, in our work with a legged robot we use internal sensors to measure the dynamics and interaction forces experienced by the robot. From these measurements we use simple machine learning methods to probabilistically infer properties of the environment, and therefore identify it. The robot's gait can then be automatically selected in response to environmental changes. Experimental results show that for several environments (sand, water, snow, ice, etc.), the identification process is over 90 *per cent* accurate. The requisite data can be collected during a half-leg rotation (about 250 ms), making it one of the fastest and most economical environment identifiers for a dynamic robot. For the littoral setting, a gait-change experiment is done as a proof-of-concept of a robot automatically adapting its gait to suit the environment.

I. Introduction

In this paper we demonstrate adaptive gait control to a wide range of environments for a legged robot. In particular, we demonstrate that inertial sensors and actuator feedback are sufficient to leverage a Bayesian classifier that rapidly identifies the environment, despite large amounts of noise and intermittent contact. This information then allows the robot to chose its gait both qualitatively and quantitatively to adapt to the current environment. Furthermore, we believe this is the first work that demonstrates the efficiency of such methods over such a wide range of environmental contexts including swimming underwater, walking on slippery ice, and traversing the open spaces of a typical university office complex. Practical implications of this include, for example, the ability of the robot to switch from walking to swimming gaits as it moves from a sand beach or surf-zone to deep water.

Our experimental testbed, AQUA, is an amphibious hexapod with six independently-controlled leg actuators. The robot can negotiate rugged terrains, and with the use of amphibious legs, it can also swim in water to a depth of 10 m. Proper selection of gait for each type of environment is of crucial importance. Therefore any autonomous version of the robot would have to identify the environment in order to select the proper gait. Since for this robot the leg forces are by nature very impulse-like, the robot dynamics highly depend of the surface mechanical properties. Conveniently, this behavior can be viewed as a mechanism for probing and estimating the dynamic properties of the surface.

Fig. 1. The hexapod robot, shown equipped with the semi-circle legs. Communication with the operator station occurs over the fiber-optic tether.

Due to the complexity and variability of the robot-ground interactions, our approach will rely on a probabilistic framework that could be extended to become non-parametric. This has the further advantage of keeping the learning phase simple (merely recording a few segments of walking). We assume that any environment type we need to identify will persist for at least a few leg cycles.

II. Description of the Robot

Our vehicle (Fig. 1) [1] is a hexapod robot that was specifically designed for amphibious locomotion. This robot is based on the highly successful RHex [2] robot, and has been improved to have autonomous underwater behavior [3]. Various legs have been designed for the appropriate terrain: semi-circle compliant legs for rugged terrain, amphibious straight legs for beach and water, flippers for underwater swimming. The robot is equipped with a 3-axis Inertial Measurement Unit (3DM-GX1^{TM}) from Microstrain. Two PC/104 single-Board computers, one a 300 MHz Pentium-equivalent running QNX and the other a 1400 MHz Pentium-M running GNU/Linux, are used for on-board computation. Although it was not used in this experiment, the robot is also equipped with 3 cameras, one of which is connected directly through Firewire to the Pentium-M subsystem.

A. Robot Gait

Locomotion of the robot on land can be achieved by rotating the compliant legs in two groups of three legs. In this (simplest) walking mode, the three legs, two on one side and one on the other, form at low speed a statically stable tripod. At higher speed, this configuration is also dynamically stable. While one tripod formation is in contact with the ground and propelling the robot forward, the other tripod formation is circulated rapidly around to be ready for the next support phase. This quick alternation of support coupled with the compliant nature of the legs results in a complex dynamic interaction between the robot and the ground. This work will exploit this feature in order to try to identify the terrain.

III. RELATION TO OTHER WORK

Identifying the robot state or environment via actuator feedback is a problem that has previously received some attention. The groundbreaking work on the Ambler project suggested that for certain types of terrain, identification could be inferred from a vehicle's footfall data [4]. In subsequent work, Krotkov and Hoffman describe how terrain characteristics can be estimated for leg contact forces [5]. This work, however, never goes beyond a tentative proposal for measuring very local terrain properties (i.e. the terrain under a single foot by using deliberative placement). Suzuki et al. [6] discuss the potential for terrain identification during walking for legged robots, but do not implement it. Other close works are related to terrain identification through vibration analysis of wheeled vehicles. In the work by Sadhukhan et al. [7], a neural network is employed to classify terrain types by the spectrum of vertical vibration. Subsequently, Brooks et al. in [8] present a technique based on pairwise classifiers using vibration spectrum data that has been reduced in dimensionality via principal component analysis (PCA).

In work closer to what we report here, Skaff et al. [11] show how the state estimation of the robot can be improved by detecting the ground or in-flight phase for a similar robot. Their method relies on using leg contact forces or acceleration measurements. In their work, the emphasis is placed in estimating the robot state within a gait cycle, whereas here we are concerned on identifying the environment spanning gait cycles. Lenser and Veloso [12], [13] on the other hand is closer to our work, using signal statistics to classify and identify environments. In particular, [13] includes the topic of robot adaptation to its environment. Although the changes presented are for adapting thresholds used in vision, it is in spirit similar to our goal of gait adaptation. The problem of gait adaptation for a similar robot has been explored by Weingarten et al. [10].

This work also bears some link with the multi-modal sensor fusion field. The work of Smith and Cheeseman [9] formalized the notion of uncertainty in merging measurements for various locations. An important part of this work is related to Bayes techniques, particularly the naive Bayes classifier, and to a lesser extend Bayes networks. A good introduction to probabilistic reasoning can be found in Pearl [14] and early

work on this matter was done by Kortenkamp [15], where a Bayesian belief network was used to combine sonar and visual information.

Finally, some ties can be made with the field of fault detection algorithms. These systems often rely on identifying a fault condition by comparing it to the working state signature. In our system a mechanical failure, for example, would be easily seen by the large change in the robot's dynamic signature. Moreover, given a large and varied sensor collection, it becomes possible to tackle the difficult issue of differentiating between a faulty robot and a faulty sensor. For example, Carlson et al. [16] describe how the Dempster-Shafer conflict metric *Con* can be used as an interpretation of sensing quality.

IV. PROBABILISTIC FRAMEWORK FOR THE ENVIRONMENT IDENTIFICATION PROBLEM

We adopt a probabilistic approach to environment identification. Not only is this method mathematically simple, but it allows us to express the unavoidable uncertainties resulting from the mismatch between sample classes and real environments. The problem we are addressing in this paper can be formalized as follows.

Let C_{env} represent the set of n possible environments the robot is expected to encounter. The set of possible values for C_{env} is limited to:

$$C_{env} \in \{C_0, C_1, ..., C_n\} \qquad (1)$$

where C_0 represents the null hypothesis. For each of these environments, the terrain is expected to have different mechanical characteristics. In turn, the mechanical interaction between the robot's legs and the environment surface generates different reactions such as varying current consumption for the leg's motor, accelerations and rotation of the robot's body. As such, for each element in C_{env}, we expect a particular collection of sensor outputs in the sensory space S. For our particular vehicle, the number of dimensions for this space is determined by the number of internal sensors (18):

- 3 accelerometers (a_x, a_y and a_z)
- 3 rate gyroscopes ($\dot{\phi}_p$, $\dot{\phi}_r$ and $\dot{\phi}_y$ for pitch, roll and yaw rates)
- 6 leg angle encoders (θ_{leg_0} through θ_{leg_5})
- 6 motor current estimators (I_{leg_0} through I_{leg_5})

The sensory output S of the moving robot is a function of the current environment $C \in C_{env}$ and gait parameters G:

$$S = \phi(C, G) \qquad (2)$$

To simplify the problem, the gait parameters G was kept constant across all test cases[1]. Given this, Eq. 2 can be simplified to:

$$S = \phi(C) \qquad (3)$$

---

[1]A more complete approach would have to take the gait parameters into account. Given the limited gait parameter space for this robot (for the gait used in the tests, this correspond only to a handful of discrete values), extension to include G should pose no problem.

The problem we are trying to solve can be summed up as finding the inverse mapping relationship; that is to identify the environment C given a sensory signature S:

$$C = \phi^{-1}(S) \tag{4}$$

In recent times, Bayesian systems have established themselves as robust methods for dealing with uncertainty in an optimal or near-optimal way. Using the Bayes, this problem can be reformulated in a probabilistic manner as:

$$P(C|S) = \frac{P(S|C)P(C)}{P(S)} \tag{5}$$

with the conditional probability $P(S|C)$ representing a probability density of the sensor signals S given an environment C.

As for many problems, estimating the prior probabilities $P(C)$ is a difficult task. In our case, this prior is highly dependent on the operating conditions of the robot. It should also be noted that this prior is highly non-stationary. In order to sidestep this issue and better concentrate on the environment identification problem, we assume equal priors $P(C)$ for all environments. Letting α be the usual normalizing constant, the classifier is reduced to the maximum-likelihood form:

$$P(C|S) = \alpha P(S|C) \tag{6}$$

A. Conditioning the Sensor Signals on Leg Angle θ_{leg}

A complete and optimal probabilistic analysis of this system would require the set of full conditional probabilities. With the 18 sensors, the size of the full conditional probabilities table would be enormous (on the order of 2^{18}). However, this problem can be greatly simplified by first keeping only one of the leg's signals (I_{leg}) and then conditioning the probabilities on that particular leg's angle θ_{leg}. Using the product rule and having the new sensor space $S' = \dot{\phi}_p, \dot{\phi}_r, \dot{\phi}_y, a_x, a_y, a_z, I_{leg}$ we get

$$P(C|S', \theta_{leg}) = \alpha P(S', \theta_{leg}|C) = \alpha P(S'|\theta_{leg}, C) P(\theta_{leg}|C) \tag{7}$$

As an approximation, we assume the leg is always rotating in any environment C. The probability $P(\theta_{leg}|C)$ is therefore always equal to 1 and Eq. 7 becomes:

$$P(C|S', \theta_{leg}) = \alpha P(S'|\theta_{leg}, C) \tag{8}$$

Finally, assuming conditional independence of the probabilities and selecting the Maximum A Posteriori (MAP) hypothesis C, we get the usual naive Bayes classifier equation:

$$C = \underset{C_i \in C_{env}}{\arg\max} \prod_{s \in S'} P(s|\theta_{leg}, C_i) \tag{9}$$

B. Importance of Using Information Synchronized on Leg Angle

Using the sensor information synchronized on the leg angle θ_{leg} represents on of the key aspects of our approach. This stems from the fact that in the robot steady-state regime, the sensor signals will be periodic, and the period will be tied to the gait period itself. The leg angle is directly computed from
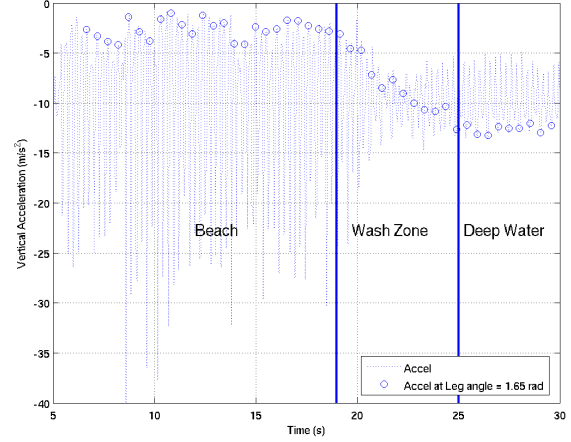

Fig. 2. Vertical acceleration a_z over time. The circles (o) indicate the value of this signal when θ_{leg}=1.65 rad during the leg rotation.

the gait clock for the robot (the robot walks in open-loop, so the leg angle is tied to the gait timing). Changes in the signal should be felt more strongly if one keeps in mind the gait timing. A clear example of this can be seen in Fig. 2. If we only look at the statistics of the acceleration a_z on this graph, we see a reduction of amplitude when the robot starts entering the wash zone, but an increase when it has reached full water depth. Lost in the amplitude envelope is the notion of the phase of that signal, and simply using amplitude, we fail to recognize a phase reversal (a shift of π in the waveform) at time t=23s. Using the synchronized information, this phase reversal is taken into account, improving the accuracy of the information. Therefore, statistical methods based on amplitude would be unable to distinguish between C_{wash} (shallow water) and C_{water}; for both locations, the amplitude of the oscillation is similar. This technique bears some similitude to the sampling of the return map at a particular angle in Tedrake et al.[18], however in our case the sampling phase angles can be multiple and are selected at locations maximizing information.

Another advantage in using the synchronized signal is that the information tends to have a higher signal-to-noise ratio. Since a signal with a lower signal-to-noise ratio intrinsically has less information in it, more data needs to be collected and analyzed, often having to perform some averaging. The immediate impact of this filtering over several gait periods is the introduction of delays in the environment identification process.

V. RESULTS FOR LITTORAL SETTING

In this first experiment, we limited ourselves to a littoral (i.e., surf-zone) environment. The set of possible values for C_{env} was therefore limited to:

$$C_{env} \in \{C_{beach}, C_{wash}, C_{water}\} \tag{10}$$

where C_{beach} represents the sandy beach, C_{wash} the section of the beach covered with shallow water (up to the robot's
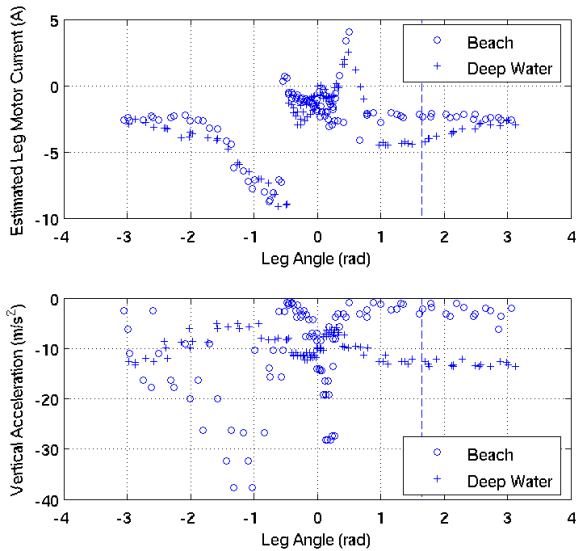
Fig. 3. Motor leg current I_{leg} (above) and vertical acceleration a_z (below) plotted as a function of leg angle θ_{leg} for the beach trials. Only the cases where the robot is on the beach (o) and in deep water (+) are plotted. Each case contains a few leg rotations, in order to show the distribution of the data.

height), and C_{water} the section where the water depth is greater that the robot's height. The sensor space S was reduced to the following signals:

- θ_{leg}, the angle of a particular leg
- I_{leg}, the estimated motor current driving that particular leg
- a_z the vertical acceleration of the robot.

Physically, the current I_{leg} is a good indicator of the torque generated at a given leg.

A. Experimental Results

Data was collected over 15 trial runs on a pebble-covered beach. This particular beach was selected due to its gentle slope and the absence of large waves. This way, the depth of the water could be assumed to be monotonically increasing with the distance away from the shore. Fig. 3 shows the motor current for a leg (top graph) and the vertical acceleration (bottom graph) plotted against the leg angle. The first set of point (circles) corresponds to the beach environment (C_{beach}), and the second set (crosses) corresponds to when the robot is completely submerged by sea water (C_{water}). The intermediate values corresponding to C_{wash} were omitted from this graph to better show the distance between the C_{beach} and the C_{water} signatures. Although the data from this plot comes from a single trial, all other trials on this beach had similar results. This indicates the relative stability of this signature methods.

This graphic suggests that the easiest way to discriminate between these two environments is to look at the location where their probability distributions $P(I_{leg}|\theta_{leg}, C)$ and $P(a_z|\theta_{leg}, C)$ have maximum separation (dashed line in Fig. 3). Using a simple threshold rule, the environment can be

identified. This rule could be applied to either the current I_{leg} or the vertical acceleration a_z, or both.

B. Automatic Gait Switching from Walking to Swimming

Armed with the above information, successful automatic gait switching experiments were performed on a new beach setting. We used the same leg angle ($\theta_{leg} = 1.65 rad$), previously identified as having the largest discrimination. A midpoint threshold value for current was selected ($I_{leg} < -3.0$), ensuring that the depth of water was sufficient to enable the robot to swim. The Fig. 4 shows a time-lapse sequence of frames for one of the switchover experiments. These experiments validated the ease and reliability of identifying the C_{beach} and C_{water} environments. It should be noted that all the above parameters were manually selected. Automatically learning these values is currently outside the scope of this paper, as it would entail the robot being able to measure the true vehicle displacement in order to evaluate the performance of the gait.

VI. IDENTIFYING THREE SIMILAR ENVIRONMENTS: SNOW, ICE AND LINOLEUM

To further evaluate the effectiveness of our technique, we tested it on an additional set of environments with similar properties. This set C_{env} was composed of the following environments:

- indoor lab floor (linoleum),
- rough hard ice,
- unpacked granulated snow (5 cm deep).

This particular set was selected due to commonalities between environments:

- ice and linoleum are both hard and cohesive surfaces,
- ice and snow are both slippery materials.

Since the differences between these environment signatures are less dramatic than in the littoral setting, more information will be required to reliably identify the environments. Consequently, 4 data points were selected (as compared to one for the previous experiment). A crude probability density function was established for each of the manually selected 4 data points. This probability density was related to the density of sample points for each of the three environments, and the classification result was simply the output of the naive Bayes classifier:

$$C = \underset{C_i \in \{C_{snow}, C_{linoleum}, C_{ice}\}}{\text{argmax}} \prod_j P_j(...|\theta_{leg}, C_i) \quad (11)$$

with

$$P_1 = P(I_{leg}|\theta_{leg} = 0.0 rad, C_i) \quad (12)$$

$$P_2 = P(I_{leg}|\theta_{leg} = 0.15 rad, C_i) \quad (13)$$

$$P_3 = P(I_{leg}|\theta_{leg} = 0.45 rad, C_i) \quad (14)$$

$$P_4 = P(\dot{\phi}_y|\theta_{leg} = 1.35 rad, C_i) \quad (15)$$

The angles θ_{leg} were manually selected based on the distance between the distributions. In Fig. 5, the leg motor current is plotted as a function of the leg angle for the

Fig. 4. Frame sequence taken during the automatic gait switch from walking to swimming. The actual switchover happens on the 6^{th} frame, when the robot detects the appropriate depth for swimming.

three environments. The lines indicate the leg angle where the current value I_{leg} is used to compute the probabilities P_1, P_2 and P_3. These angles correspond to certain physical phenomena: P_1 detects the presence of non-cohesive material (early touchdown), whereas P_2 and P_3 are related to the adherence of the leg on the surface. Fig. 6 shows the body yaw angular velocity $\dot{\theta}_y$ at various leg angles. Since large yaw angular velocities are indicators of good traction on the ground, this variable can be used to find the traction characteristics of the ground. The yawing moment comes from the tripod gait used on the robot: two legs are pushing the robot forward on one side compared to one on the other side. This creates a small imbalance in the yaw moment, proportional to the traction generated by the legs. The probability P_4 can be derived from this information. Although Fig. 7 was not used directly in this classification problem, it shows how a given environment signature will be visible for many signals: the snow environment can be differentiated easily from the ice and linoleum. Table I shows the result of the classification. Fig. 8 shows an estimate of the probability densities P_1, P_2, p_3 and P_4.
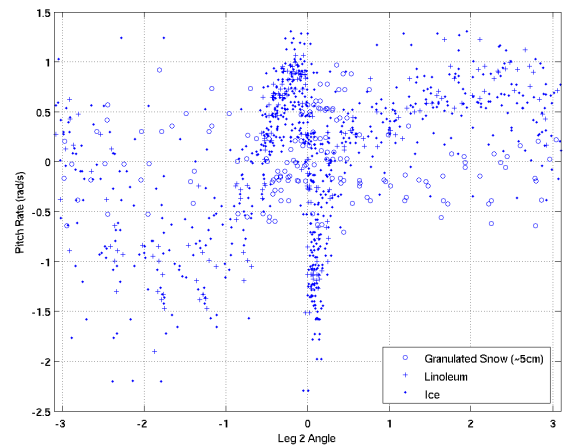

Fig. 7. Pitch angular velocities $\dot{\phi}_p$ plotted as a function of leg angle θ_{leg} for C_{ice}, C_{snow} and $C_{linoleum}$ environments.
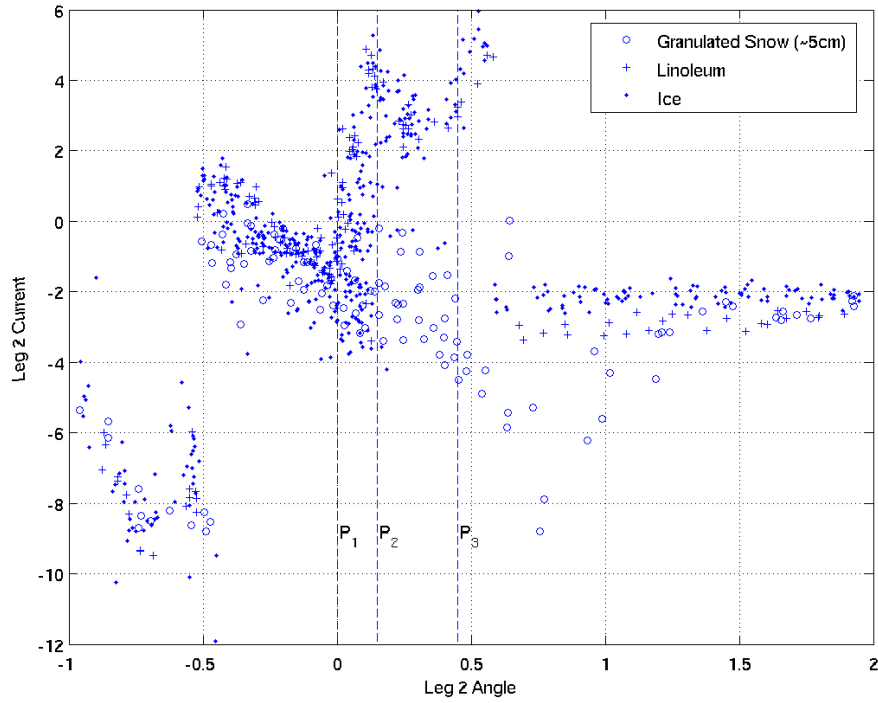
Fig. 5. Motor leg current I_{leg} plotted as a function of leg angle θ_{leg} for C_{ice}, C_{snow} and $C_{linoleum}$ environments. The three vertical dashed lines indicate the angles where the probabilities of Eq. 12, Eq. 13 and 14 are evaluated.
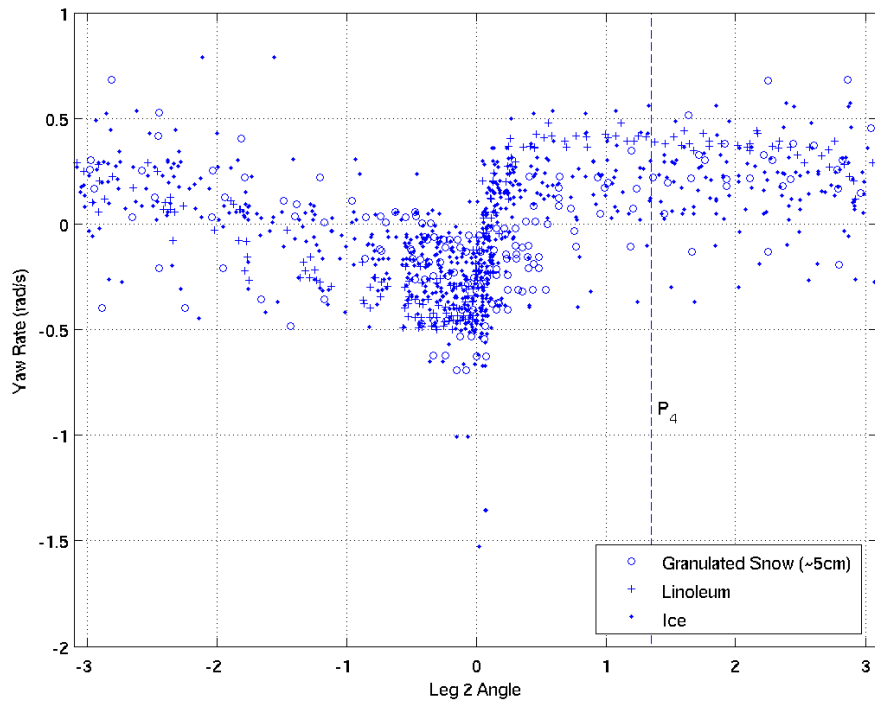

Fig. 6. Yaw angular velocities $\dot{\theta}_y$ plotted as a function of leg angle θ_{leg} for C_{ice}, C_{snow} and $C_{linoleum}$ environments. The vertical dashed line indicates the angle where the probability of Eq. 15 is evaluated.

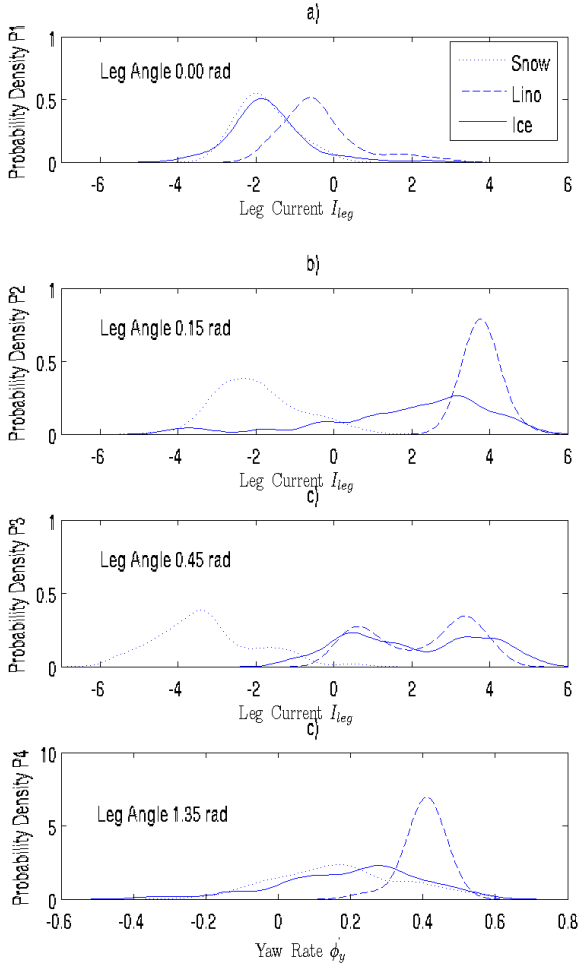| Env.	Samples	C_{snow}	$C_{linoleum}$	C_{ice}	Success Rate
Snow	52	45	1	6	87 %
Linoleum	38	0	31	7	82 %
Ice	81	5	8	68	84 %

Fig. 8. Probability density estimates for a) P1, b) P2, c) P3 and d) P4 via Parzen windows. Gaussian kernels of standard deviation 0.4 for I_{leg} and 0.04 for $\dot{\phi}_y$ were used to estimate the probability densities.

TABLE I

Classification results for the lab floor, snow and ice environments

Env.	Samples	C_{snow}	$C_{linoleum}$	C_{ice}	Success Rate
Snow	53	49	0	4	93 %
Linoleum	37	0	35	2	95 %
Ice	81	1	7	73	90 %

Cross validation was performed by applying the same probability densities to the sensors of a different leg, located on the opposite side of the robot. Given that the robot is symmetric, the inertial cues will be similar, up to a phase shift of π in the signal. The inferior results seen in table II might be interpreted as an indication of over training. However, a more probable explanation is that the power signature for each leg is unique. Indeed, small offsets in the position of large power variations were observed between these two legs. Therefore, the location of the probability distributions P_1, P_2 and P_3 should change from leg to leg, and this would contribute to reducing the success rate by mislabeling data points.

One important note is that the classification was done for the steady-state regime only. Classification of the transient regime that happens when the robot starts moving is a more difficult problem, since it violates the periodicity assumption of the robot and sensor responses. However, this transient regime is rather short lived, corresponding to 2-3 complete leg cycles (1-2s).

VII. Conclusion

In this paper we have presented a rapid and effective way for a dynamic legged robot to identify and thus adapt to a number of different environment classes. This technique relies on the analysis of actuator and inertial sensor information. One of the key ideas is the synchronization of the sensor signal with the leg position (i.e. angle of the leg with respect to the body). This synchronization captures periodic properties of the gait by matching samples acquired at different times but at the same leg posture. This serves to greatly reduce the number of data points needed for identification as well as preserving the phase information of the sensor signals. Gait change experiments in a littoral setting demonstrated the autonomous gait selection capabilities enabled by this technique.

The demonstrated performance represents a lower bound on what could be achieved with this technique. By combining the measurements of more legs for example, performance could be improved without adding delay to the identification. If more accurate identification is needed, averaging across gait periods could be employed, therefore reducing the impact of noise but at the price of a larger latency in the detection. Further improvement could be achieved by increasing the number of modalities of sensing and using a Bayes net to fuse the information. This should allow for more discrimination power, thus enabling more reliable and extensive identification. Initial

work on integrating video information has been completed, and integration of audio information is planned. Also, a more complete and comprehensive classification method based on a nearest neighbors classifier is in progress. Another difficulty is the relatively high bandwidth of the acceleration signal. Given the sampling rate of 50 Hz for our acceleration unit, it makes it hard to capture the rapidly changing acceleration signal at the moment of impact. Work is being done to improve the sampling rate of the inertial unit. Also, we hope to be able to tackle the issue of discovering new environments as they appear. This could be possibly be achieved using standard clustering techniques that would indicate the appearance of a new cluster in the data set. Another extension would be to apply dimensionality reduction techniques such as Principal Component Analysis (PCA), covering the cases of continuously varying environments.

REFERENCES

[1] G. Dudek, M. Jenkin, C. Prahacs, A. Hogue, J. Sattar, P. Giguere, A. German, H. Liu, S. Saunderson, A. Ripsman, S. Simhon, L. A. Torres-Mendez, E. Milios, P. Zhang, I. Rekleitis *A Visually Guided Swimming Robot*, Proceedings of the 2005 IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 1749-1754, 2005

[2] U. Saranli, M. Buehler, D.E Koditschek *RHex: A simple and highly mobile hexapod robot*,   The International Journal of Robotics Research, vol. 20, no.7, pp.616-631, 2001.

[3] G. Dudek, P. Giguere, J. Sattar *Sensor Based Behavior Control for an Autonomous Underwater Vehicle*, Proceedings of the International Symposium on Experimental Robotics (ISER), Rio de Janeiro, Brazil, July, 2006.

[4] D. Wettergreen, H. Thomas, and C. Thorpe *Planning Strategies for the Ambler Walking Robot*, Proceedings of the 1990 IEEE International Conference on Systems Engineering, August, 1990, pp. 198 - 203.

[5] E. Krotkov, R. Hoffman, *Terrain Mapping for a Walking Planetary Rover*, IEEE Transactions on Robotics and Automation, Vol. 10., No. 6, December 1994.

[6] K Suzuki, H Imai, Y Kawamura, Y Sankai, *Gait control of human and humanoid on irregular terrain considering interaction with environment*, Robot and Human Interactive Communication, 2003. Proceedings. RO-MAN 2003. The 12th IEEE International Workshop on, 31 Oct.-2 Nov. 2003 pp. 277-284.

[7] D. Sadhukhan, C. Moore, *Online Terrain Estimation Using Internal Sensors*, in Proceedings of Florida Conference on Recent Advances in Robotics, Boca Raton, FL, 2003.

[8] C. Brooks, K. Iagnemma, S. Dubowsky, *Vibration-based Terrain Analysis for Mobile Robots*, Proceedings of the 2005 IEEE International Conference on Robotics and Automation, Barcelona, 2005.

[9] R. Smith, P. Cheeseman. *On the Estimation and Representations of Spatial Uncertainty*, International Journal of Robotics Research, vol.5, no. 4, Winter 1987, pp 56-68.

[10] J. Weingarten, G. Lopes, M. Buehler, R. Groff, D. Koditschek *Automated gait adaptation for legged robots*, Proceedings of the 2004 IEEE International Conference on Robotics and Automation (ICRA 2004), Volume 3, pp. 2153-2158.

[11] S. Skaff, A. Rizzi, H. Choset, and P. Lin, *A Context-Based State Estimation Technique for Hybrid Systems*, Proceedings of the 2005 IEEE International Conference on Robotics and Automation, IEEE, April, 2005, pp. 3935-3940.

[12] S. Lenser, M. Veloso, *Classification of Robotic Sensor Streams Using Non-Parametric Statistics*, Proceedings of the 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems, IEEE, October, 2004, pp. 2719-2724 vol.3.

[13] S. Lenser and M. Veloso. *Automatic detection and response to environmental change*. In Proceedings of the International Conference of Robotics and Automation, May 2003.

[14] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*,   Morgan Kaufmann, San Mateo, CA, 1988.

[15] D. Kortenkamp et. al. *Mobile Robot Exploration and Navigation of Indoor Spaces Using Sonar and Vision*, Proceedings of CIRFFSS'94, Houston, TX, 21-24 Mar, 1994, pp. 509-519.

[16] J. Carlson, R. Murphy, S. Christopher, J. Casper, *Conflict Metric as a Measure of Sensing Quality*, Proceedings of the 2005 IEEE International Conference on Robotics and Automation, IEEE, April, 2005

[17] K. Lang, *Newsweeder: Learning to filter netnews*, Proceedings of the 12^{th} International Conference on Machine Learning, San Francisco, CA, 1995, pp. 331-339.

[18] R. Tedrake, T. W. Zhang, H. S. Seung, *Stochastic Policy Gradient Reinforcement Learning on a Simple 3D Biped*, Proceedings of the IEEE International Conference on Intelligent Robots and Systems (IROS), Senda, Japan, September 2004, pp. 2849-2854.