Proprioceptive localization for a quadrupedal robot on known terrain

Sachin Chitta, Paul Vernaza, Roman Geykhman and Daniel D. Lee
GRASP Laboratory
University of Pennsylvania
Philadelphia, PA 19104
[sachinc,vernaza,geykhman,ddlee]@grasp.upenn.edu

***Abstract*— We present a novel method for the localization of a legged robot on known terrain using only proprioceptive sensors such as joint encoders and an inertial measurement unit. In contrast to other proprioceptive pose estimation techniques, this method allows for global localization (i.e., localization with large initial uncertainty) without the use of exteroceptive sensors. This is made possible by establishing a measurement model based on the feasibility of putative poses on known terrain given observed joint angles and attitude measurements. Results are shown that demonstrate that the method performs better than dead-reckoning, and is also able to perform global localization from large initial uncertainty.**

I. INTRODUCTION

The problem of localization using visual features has been widely studied in recent years. A typical approach is to use a known map of the environment and update pose online using a combination of odometry and observation of known features in the scene. In another approach, the process of localization and choosing the right set of features is carried out simultaneously giving rise to the approach known as SLAM. In all this work, localization relies heavily on the use of an *exteroceptive* sensor, such as a camera, GPS unit or laser scanner. Proprioceptive sensors, such as joint encoders and angular rate gyros provide odometry information while accelerometers can be used (at rest) to determine the direction of gravity. Very little work has addressed the problem of localization in known environments using *only* proprioceptive sensing.

In this work, we will address the problem of localization of a quadruped robot using only proprioceptive sensing. A good analogy to motivate and explain this concept is that of human motion in a dark but familiar room. Initially, humans have no sense of where they are in the room but taking a couple of steps and encountering, for example, a familiar feature like a step or rise instantly localizes us in the room. There is still an element of uncertainty in the estimate of position in the room, for example, it may be difficult to localize along the length of the step. However, even the limited pose information gleaned from this data is often enough to make meaningful decisions. This technique, that we will henceforth refer to as *proprioceptive localization*, can also prove very useful in situations where the primary external sensors of a robot fail.

We are specifically motivated by the application of this technique to legged robots. Legged robots can sense rich ground features and, if appropriate sensors are available, also actively probe the environment. Although the set of features is not as rich or unique as visual features, the technique provides sufficient information to localize the robot even with reasonable uncertainity in its initial position. Internal sensors (gyroscopes and accelerometers) provide the necessary local information to determine the *local pose* of the robot, *i.e.* its roll, pitch and height above the ground. Internal joint sensors provide joint angles for the legs. Fig. 1 briefly illustrates how this data might be used in a localization strategy.

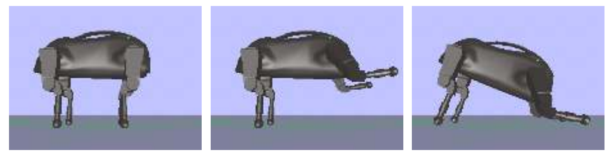

(a) Feasible configuration (b) Infeasible configuration (c) Feasible configuration

Fig. 1. Illustration of how proprioception might be used to distinguish valid poses from invalid ones. Fig. 1(a) illustrates a body pose that is consistent with proprioception, given the assumption of static stability; the zero pitch of the body is consistent with all legs being extended on a flat surface. Fig. 1(b) illustrates a body pose that is inconsistent with proprioception and the given terrain; since the front legs have nothing to rest on in that configuration, we would expect the robot to be pitched forward were it in a statically stable pose, as in Fig. 1(c).

The method we present in this paper could be used to provide robustness to failure of a primary exteroceptive sensor like a camera or GPS system. While the method may not provide enough information to localize the robot very accurately in such a situation, it may still provide enough information to move the robot to a safer location.

This paper is organized as follows. In Section II, we look at related work in other areas of research. In Section III, we present the robot and associated hardware used for this effort. In Section IV, we present the gait controller used to make this robot walk over different types of terrain. In Section V, we present details of the particle filter used for localization. In Section VI, we present details of the experimental procedure and results from the experiments. We conclude with a discussion of the results in Section VII.

II. RELATED WORK

Existing research in localization for legged robots might be divided into approaches that are based primarily on exteroception and those that are primarily based on proprioception. Methods in the former category have primarily been

Fig. 2. The LittleDog robot (with retroreflective markers) on a terrain board.

explored in visually structured environments with vision-based localization approaches. Often times these approaches do not take particular advantage of the structure provided by legged locomotion (such as in [1]). Hoffman et. al. describe a legged robot localization system that uses both vision and proprioceptive data [2]. In this case, proprioception is used to improve the motion model such as to more accurately update uncertainty predictions from odometry. However, proprioceptive data by itself is insufficient to localize the robot in this scheme. The critical difference in our work is that our method allows for global localization from proprioception; i.e., localization uncertainty can decrease in time given the observations.

Lin et. al. demonstrate a form of proprioceptive pose estimation for a hexapedal robot in [3]. Their robot is equipped with a leg pose sensor and an inertial measurement unit. A Kalman filter is used to fuse these measurements, and different process models are considered for different phases of the gait. Again, in contrast to our method, Lin's model does not allow for global localization based only on proprioceptive data and terrain information.

Another closely related line of research relates to localization for robotic assembly, where the goal is to physically assemble parts with fine precision in the presence of uncertainty. Chhatpar ([4]) describes a method for localization in such a scenario using particle filtering. Similar to our method, Chhatpar's method relies on computing likelihoods of contact configurations in order to localize an object, given a predefined map of possible contact configurations. Our method differs in that we consider how a similar technique applies to legged locomotion, and how the localization problem is aided by the particular characteristics of legged locomotion.

III. Experimental Hardware

In this work, we use a quadruped robot called LittleDog (Figure 2) manufactured by Boston Dynamics Inc. The robot has four legs with three joints in each leg. The robot can be powered by onboard batteries or by an external power system. Communication is through a wireless 802.11a connection with a host computer. Onboard sensing includes an accelerometer and gyroscopes. Each foot also has a single axis force sensor at the bottom.

Our experimental setup includes a VICON motion capture system. The system consists of 6 high speed cameras operating at approximately 100 Hz. A set of reflective markers on the body of the robot allow the system to track the position and orientation of the robot. Terrain boards measuring 60 cm by 60 cm are used for testing the robot. The terrain boards are accurately scanned, providing an elevation map of each board. The boards are also registered with respect to a local coordinate system marked out by reflective markers. The reflective markers register the terrain board accurately in the global coordinate system defined by the motion capture system.

IV. Gait Control

In this section, we present details of the gait for the robot. It is important to mention that this controller is executed with feedback provided by the motion capture system. The motion capture system provides information about the complete pose of the robot itself while internal sensors provide information about the joint angles of the robot. The data from the motion capture system was used only for control and not in the estimation process. The experimental procedure for collecting and analysing data from the trials will be described in more detail in Section VI. It should be noted that the focus of this work was on validating the localization procedure and not on the gait control procedure. Hence, we will present the gait control technique only in brief.

The controller implements a statically stable gait with no more than one foot off the ground at any point of time. Figure 3 shows the phasing of the legs for walking. There are two separate parts to the motion of the robot, the motion of the body itself and the motion of individual legs through the air. Our implementation requires the body to stay stationary while a leg is swinging. The gait can be easily adapted to achieve different types of walking (crawl, trot) using just a few parameters. Splines are used to specify smoother motion profiles for the robot body and leg motion. Zero velocity boundary conditions are used to specify smooth touchdowns.

The controller follows the following algorithm:

1) Choose next foot to lift based on phasing sequence (FL,HR,FR,HL).
2) Choose foothold for next foot by choosing a nominal foot position (based on current velocity). Check quality of foothold based on local flatness of area around point chosen, quality of next support triangle and kinematic feasiblity.
3) Execute foot motion - motion executed is spline based with zero velocity conditions at beginning and end of rise, fall and swing.
4) Execute body motion - body is moved to centroid of next triangle of support. Checks are performed on the kinematic feasibility of the motion and only a kinematically feasible motion is executed while ensuring that a sufficent stability margin is achieved for foot pickup.
5) Return to Step 1.

The duty cycle for the motion is thus 0.875 since each leg spends only $1/8$ of each whole cycle off the ground. The legs are phased 0.25 apart. The net motion of the robot in a run is shown in Figure 4. The markers in the figure represent the

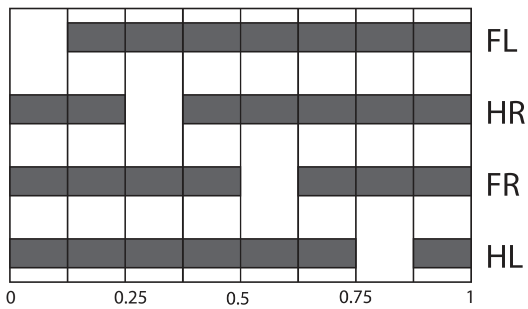

Fig. 3. Phasing diagram for gait. The shaded areas represent parts of the gait cycle when the feet are on the ground. The feet are labeled as FL (front left), FR (front right), HL (hind left) and HR (hind right).

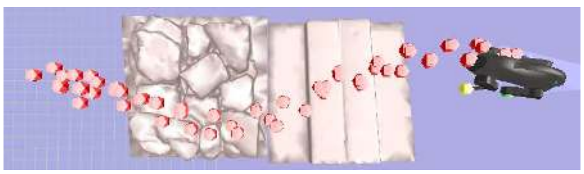

Fig. 4. A trial run for the robot over rough terrain.

position of the center of mass of the robot at the end of each footfall. Also visible is the lateral oscillation of the robot's body, effected in order to achieve a good margin of stability.

V. Particle Filtering

Particle filtering [5] has emerged as the estimation method of choice for many difficult problems in robotics due to its flexibility and ease of use. We employ a particle filter in this work as our localization method for these reasons and others that will be made apparent in this section. We will briefly review particle filtering methods here; for a more detailed description of these methods, the reader is advised to consult one of many instructive references on the subject, such as [6] or [7].

A. Preliminaries

Our ultimate goal is to obtain an estimate of the six-degree-of-freedom pose of the robot body with respect to a global coordinate frame, given priopprioceptive sensor measurements and prior knowledge of the dynamics of the legged robot. This problem can be expressed in a Bayesian setting via the following probabilistic equation:

$$p(x_t|z_{0:t}) = \alpha p(z_t|x_t) \int dx_{t-1} p(x_t|x_{t-1}) p(x_{t-1}|z_{0:t-1}) \tag{1}$$

This expresses that we wish to obtain a probability distribution over the pose at the current time (x_t) given all previous measurements $z_{0:t}$. This distribution is a function of the measurement likelihood function $p(z_t|x_t)$, the system dynamics $p(x_t|x_{t-1})$, and the prior pose distribution $p(x_{t-1}|z_{0:t-1})$. The normalization factor α can be computed from the constraint that $\int dx_t p(x_t|z_{0:t}) = 1$.

The general difficulty in computing Eq. 1 follows from the intractability of the integration and the computation of the normalization factor. Unless strong restrictions are placed on the system (i.e., linearity), approximations must generally be used. Particle filtering approximates the posterior distribution by a weighted sum of point mass distributions (known as particles). The following is an informal description of how the method works. Substituting the point mass approximation into Eq. 1 yields the following, where the superscript i refers to the ith particle, and the w_i are weights associated with the particles.

$$\begin{aligned} p(x_t|z_{0:t}) &= \ldots \int dx_{t-1} \sum_i p(x_t|x_{t-1}) w^i_{t-1} \delta(x_{t-1} - x^i_{t-1}) \\ &= \sum_i w^i_{t-1} p(z_t|x_t) p(x_t|x^i_{t-1}) \end{aligned} \tag{2}$$

Here the properties of the Dirac delta function $\delta(\cdot)$ have been used to perform the integration. The resulting mixture distribution in Eq. 2 can be sampled to yield another point mass distribution for the posterior. The most common method simply samples x_t from the dynamics distribution, $p(x_t|x^i_{t-1})$, since a closed-form distribution for the right hand side of Eq. 2 might not exist. This yields

$$p(x_t|z_{0:t}) = \sum_i w^i_{t-1} p(z_t|x^i_t) \delta(x_t - x^i_t) \tag{3}$$

where x^i_t is a sample from $p(x_t|x^i_{t-1})$. This sampled posterior distribution can then be used as the prior at the next time instant, yielding an efficient recursive inference procedure.

B. Particle filtering for proprioceptive localization

One drawback of this sort of particle filter is the curse of dimensionality; the volume of the state space grows exponentially with the dimension, making it generally very difficult to apply particle filters with state spaces of more than a few dimensions. It is therefore crucial to choose a minimal representation of the state if a particle filter is to be applied.

We exploit several features of legged locomotion over rough terrain in order to find this minimal parameterization of the state. Again, we are ultimately interested in the 6-DOF pose of the robot body. The first assumption we make is that the robot is using a statically stable gait, such as that previously described. This is a reasonable assumption to make for a robot with unknown pose attempting to traverse rough terrain, since this represents the "safest" class of gaits. This assumption allows us to neglect dynamic effects for which we might have to maintain linear and angular velocities and accelerations as part of the state. Although unwanted dynamic effects are still possible regardless, we expect them to be small enough to characterize as noise in an otherwise static gait.

The previous assumption leaves us with a six-dimensional state consisting of the 6-DOF body pose. However, it can be reduced to just three given the fact that not all poses are realizable at all positions given the assumption of static stability, and given known terrain and leg poses. For example, any pose with more than three legs off the ground is unrealizable, as is any pose that requires legs to penetrate the terrain. This inefficiency is resolved by choosing a three-dimensional state vector consisting of two-dimensional translation parallel to the ground (p_t) and a one-dimensional yaw angle θ_t. We reasonably assume height, pitch, and roll can be estimated via other means; the precise rationalization for this assumption is made more clear when the measurement model is described later in this section. The complete state vector is therefore given by $x_t = [p_t^T \; \theta_t]^T$.

Our dynamics model is specified by the desired, commanded movement that induces the motion of the robot's center of mass; we refer to this as "odometry" in analogy to the case of wheeled robots. Due to the nature of the static gait used, odometry updates are only specified at certain time instants. These instants correspond to the periods of quadruple support, where all feet are stationary and the center of mass is moving to the new support triangle centroid; at all other instants of time, the center of mass is stationary. Assuming that the robot wishes to reach a point $p_t + \Delta p$ with yaw angle $\theta_t + \Delta\theta$ (equivalently, $x_t + \Delta x$), the dynamics model for the periods of motion is simply

$$p(x_{t+1}|x_t) = \mathcal{N}(x_{t+1}; x_t + \Delta x, \Sigma_{od}) \tag{4}$$

The notation $\mathcal{N}(x; \mu, \Sigma)$ indicates a Gaussian distribution in x with mean μ and covariance Σ. This indicates that in order to sample from the motion model, we sample a Gaussian distribution with mean equal to the expected destination and covariance Σ_{od} set according to the amount of uncertainty in the movement.

C. Formulation of measurement likelihood function

The observations in our model consist of measurements from "proprioceptors:" in our case, accelerometers, angular rate gyros, and joint encoders. Assuming that gravity is the dominant force on the robot body, the roll and pitch of the robot can be recovered fairly accurately by fusing and filtering the accelerometer and gyro readings with a method such as that presented in [8]. We can therefore transform the inertial observations into observations of roll and pitch, which will be much more useful in formulating a measurement likelihood function.

The measurement likelihood function is then based on the feasibility of a particular pose given the terrain and filtered state, and observations of roll, pitch, and joint angles of the robot. The more infeasible the conjunction of all these things, the less likely the observations are given the state, and vice-versa. We therefore need to specify a distribution over the feasibility of overconstrained configurations specified by all these variables. Figure 5 illustrates how this is accomplished.

Measurements are taken when three feet are thought to be on the ground and one in the air (*swing phase*), according to the known phase of the gait. Applying the discussed constraints on translational coordinates, body Euler angles, and joint angles yields only one free degree of freedom—the height of the robot in the global workspace. Thus, the robot can be considered to be moving vertically up and down along a prismatic actuator that acts only in that direction as shown in Figure 5(a). Since the feet in contact with the ground are known at any point of time, we move the robot down along this actuator until one of the feet in contact touch the ground. This is illustrated in Figure 5(b) where the hind left, hind right and front right feet are supposed to be in contact with the ground.

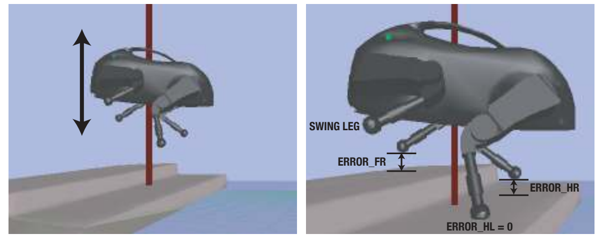

(a) (b)

Fig. 5. An illustration of the measurement model for the system. The robot is "skewered" on an axis through the position of the particle, with yaw angle fixed from the particle as well. Roll and pitch are recovered from inertial readings and are fixed as well. Feasibility of the pose is then expressed in terms of the sum of distances of stance feet from the ground after grounding one of the stance feet.

However, for the given pose, moving the robot down the prismatic actuator grounds the hind left foot first. At this point err_{FR}, err_{HL} and err_{HR} define the errors in the positions of the front right, hind left and hind right feet respectively with $err_{HL} = 0$. These errors can be used as a measure of the feasibility of this pose. If the pose were fully feasible, all these errors would be zero.

Now, the likelihood function is defined for the particular case in Figure 5(b) as a zero-mean Gaussian on the error with an appropriate covariance Σ_z that captures the uncertainty due to inertial and encoder error:

$$z_t = \sqrt{err_{FR}^2 + err_{HR}^2} \tag{5}$$

$$p(z_t|x_t) = \mathcal{N}(z_t; 0, \Sigma_z) \tag{6}$$

Given this measurement model, the particle filtering algorithm is briefly summarized in Algorithm 1.

VI. Experimental Results

In this section, we will present experimental results for the technique described in this paper. We will first present the experimental procedure used to carry out the trials.

As noted earlier, the trials are carried out with the motion capture system being used to provide feedback for the position and orientation of the body. Thus, the controller has full knowledge of the terrain and the position of the body. All the data from a trial is logged. The system has the ability to *playback* trials using the logged data. All our localization experiments are carried out offline using the *proprioceptive* data from the log. It should be noted that the localization could just as easily be performed online without the use of motion capture data; however, this was not possible here due

for i = *1 to* N = *number of particles* **do**
$x_0 \sim p(x_0)$ *// draw poses from initial distribution*
end
while z_t = *new accelerometer,* Δx_t = *new odometry* **do**
for i = *1 to* N **do**
// randomized deviation from nominal motion
$\tilde{x}_t^i \sim \mathcal{N}(0, \Sigma_{od})$
// apply motion model
$x_t^i \leftarrow x_{t-1}^i + \Delta x_t + \tilde{x}_t^i$
// update particle weights with measurement
$w_t^i \leftarrow w_{t-1}^i p(z_t|x_t^i)$
// $p(z_t|x_t)$ *represents the measurement likelihood function*
// resample if necessary
end
end

Algorithm 1: Particle filtering for proprioceptive localization

to LittleDog's lack of adequate touchdown sensors. Without these sensors, the controller must rely on motion capture data to detect touchdown.

The trial starts with the motion capture system turned on and the robot placed at the starting point on the terrain. The trial is then run and the robot proceeds to walk from the starting position to the goal position. Data from all the sensors on the robot, the motion capture system and data corresponding to the odometry is continuously logged as the robot completes the task. The localization algorithm is then run offline using the logs.

The measurement model requires knowing positions when three feet are on the ground and also knowing which foot is off the ground. Since the cycle time of the gait is fixed, this is easy to determine (under the assumption of a statically stable gait). There may be cases where the dynamics of the motion result in the use of spurious data from the logs where a different foot is off the ground. However, such data could be looked at as a noisy measurement which should get filtered out by subsequent observations.

Figure 6 shows snapshots from one offline estimation trial run. The initial position and orientation of the particles is chosen randomly around the starting position of the robot. All the particles have equal weights in the beginning.

Figure 7 plots the actual, filtered and odometry based state of the robot during the trial. The filtered state was calculated as a weighted mean of the state of all the particles. The odometry estimate was calculated separately using only odometry information. The actual position of the robot is obtained from the motion capture system. In Figure 7, x and y represent the position of the robot in a global coordinate system while the yaw represents the yaw of the body of the robot in the global coordinate frame.

Table I shows estimation error statistics from multiple trials with varying terrain configurations. Data from these trials demonstrates a marked improvement in localization performance over dead reckoning, with the improvement averaging 45% over the course of a trial.

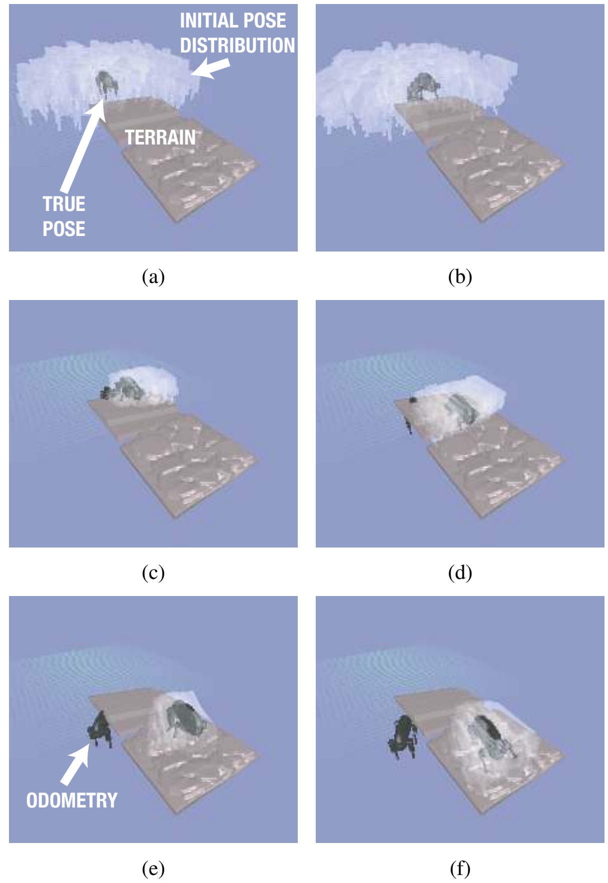

Fig. 6. Visualization of experiment with localization from offline data. Fig. 6(a) shows the initial, highly uncertain pose distribution. Each particle is represented by a semi-opaque "stick figure" robot. The true pose is represented by a solid black robot, as is the odometry estimate. Upon approaching the step (Fig. 6(b)), particles that have already passed the step are eliminated. Note the lateral ambiguity present as the robot approaches the second terrain board (Fig. 6(d)). The ambiguity begins to be resolved as the robot steps over the more informative terrain. The final distribution estimate is consistent with the true final pose, whereas the estimate from odometry is significantly off (Fig. 6(e)).

A. Discussion

The results demonstrate that the use of this technique for localization is feasible. The technique works best when the terrain is uneven since this helps in resolving the ambiguity in the pose of the robot. Flat terrain is featureless and our technique is incapable of localizing the robot when walking over such terrain. As seen in Figure 7(a), when a step is detected the x position of the robot converges to the x value. However, there is still an ambiguity in the y position of the robot and the yaw of the robot (reflected by the spread of the particles along the length of the step in Figure 6(c)). As the robot moves over rougher terrain, its pose estimate gets better since there are now richer features available to get observations from. Our experimental results also suggest that coarser terrain features, such as simple steps, result in better estimation performance than finer features such as small crevices.

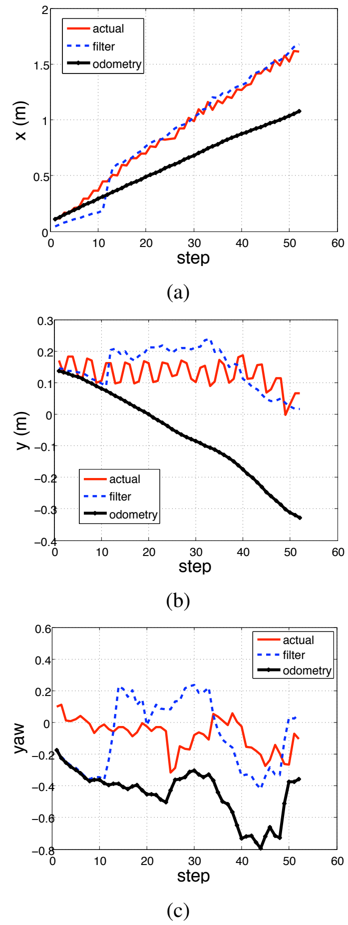

Fig. 7. Plots of actual, filtered and odometry based position (x, y) and yaw of the robot. Note exaggerated scale for plot of y-position.

Trial	Mean distance error (filter)	Mean distance error (odometry)	Improvement (percent)
1	0.099	0.177	43
2	0.108	0.259	58
3	0.157	0.214	26
4	0.115	0.242	52
5	0.123	0.223	44
6	0.131	0.264	51
7	0.131	0.209	37
8	0.177	0.360	51
Average	0.122	0.227	45

TABLE I

ESTIMATION PERFORMANCE OVER SEVERAL TRIALS VARYING TERRAIN

It is possible that the pose of the robot could converge to a different trajectory where the terrain features are spatially and temporally similar to the actual trajectory of the robot. If the terrain features are sufficiently rich, as is the case with the terrain boards used for our trials, the possibility of this happening will be lower. However, on flat terrain, the uncertainty in the estimate will grow with time since there are no features available to correct the estimate and the estimate will be no better that one computed using only odometry information.

VII. Conclusions and Future Work

We have studied the novel problem of global localization for a quadrupedal robot using only proprioceptive sensors, assuming known terrain. We have demonstrated a solution to this problem that is informed by the particular characteristics of legged locomotion over rough terrain. Specifically, this is accomplished using particle filtering with a minimal state representation and a novel measurement model that combines proprioception with terrain information. Our results show that the method is able to perform global localization. A significant improvement was also observed relative to dead-reckoning.

The proprioceptive localization problem contains many interesting issues that we have not yet been able to fully explore. One such avenue for further research is the active localization problem [9], which was also investigated in Chhatpar's work [4]. In active localization, the robot chooses to perform the actions that are expected to minimize localization uncertainty. In this setting, the resulting behavior would be much like the earlier example of the person in a dark room; the robot would "grope around" in order to orient itself. This might greatly aid the performance of any navigational tasks in situations with high initial uncertainty.

Another future avenue of research is the SLAM problem. In a realistic setting, it is expected that the map would not be given a-priori. Existing SLAM algorithms might be used to simulaneously build the map and localize the robot. Additionally, the problem contains interesting structure that could be used to help the process. For example, if there is some notion as to the general expected shape of the terrain, this could be used as a prior to aid the mapping process. We have yet to fully investigate these interesting aspects.

References

[1] S. Lenser and M. Veloso, "Sensor resetting localization for poorly modelled mobile robots," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2000.

[2] J. Hoffmann, M. Spranger, D. Gohring, and M. Jungel, "Exploiting the unexpected: negative evidence modeling and proprioception modeling for improved Markov localization," in *9th International Workshop on RoboCup 2005 (Robot World Cup Soccer Games and Conference)*, 2006.

[3] P.-C. Lin, H. Komsuoglu, and D. E. Koditschek, "Sensor data fusion for body state estimation in a hexapod robot with dynamical gaits," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2005.

[4] S. R. Chhatpar and M. S. Branicky, "Particle filtering for localization in robotic assemblies with position uncertainty," in *Proceedings of the IEEE International Conference on Intelligent Robots and Systems*, 2005.

[5] A. Doucet, "On sequential Monte Carlo sampling methods for Bayesian filtering," *Statistics and Computing*, 2000.

[6] S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for on-line non-linear/non-Gaussian Bayesian tracking," *IEEE Transactions of Signal Processing*, vol. 50, pp. 174–188, 2002.

[7] D. Fox, S. Thrun, W. Burgard, and F. Dellaert, "Particle filters for mobile robot localization," in *Sequential Monte Carlo Methods in Practice*. Springer-Verlag, 2001.

[8] P. Vernaza and D. D. Lee, "Rao-Blackwellized particle filtering for 6-DOF estimation of attitude and position via GPS and inertial sensors," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2006.

[9] D. Fox, W. Burgard, and S. Thrun, "Active markov localization for mobile robots," *Robotics and Autonomous Systems*, 1998.