Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI-10)

Online Learning of Uneven Terrain for Humanoid Bipedal Walking

Seung-Joon Yi
Biointelligence Laboratory
Seoul National University
Seoul 151-742, Korea

GRASP Laboratory
University of Pennsylvania
Philadelphia, PA 19104, USA
yiseung@seas.upenn.edu

Byoung-Tak Zhang
Biointelligence Laboratory
Seoul National University
Seoul 151-742, Korea
btzhang@cse.snu.ac.kr

Daniel D. Lee
GRASP Laboratory
University of Pennsylvania
Philadelphia, PA 19104, USA
ddlee@seas.upenn.edu

Abstract

We present a novel method to control a biped humanoid robot to walk on unknown inclined terrains, using an online learning algorithm to estimate in real-time the local terrain from proprioceptive and inertial sensors. Compliant controllers for the ankle joints are used to actively probe the surrounding surface, and the measured sensor data are combined to explicitly learn the global inclination and local disturbances of the terrain. These estimates are then used to adaptively modify the robot locomotion and control parameters. Results from both a physically-realistic computer simulation and experiments on a commercially available small humanoid robot show that our method can rapidly adapt to changing surface conditions to ensure stable walking on uneven surfaces.

Introduction

The main advantage of legged locomotion over wheeled locomotion is that legs have the capability of climbing rougher terrain than wheeled or tracked vehicles. Unfortunately, this ideal is often not achieved in reality, especially for the current generation of bipedal humanoid robots. Many walking controller implementations for humanoid robots assume perfectly flat surfaces, and even a slight deviation in the floor can lead to serious instabilities in these controllers. In this manuscript, we propose a novel method to learn and adapt to uneven terrain, and demonstrate how our methods can stabilize walking in humanoid robots.

The problem of legged walking control on uneven surfaces has been studied by a number of researchers (Hashimoto et al. 2006; Hyon 2009; Jenkins, Wrotek, and McGuire 2007; Kalakrishnan et al. 2009; Kim et al. 2007; Kim, Park, and Oh 2007; Mikuriya et al. 2005; Kolter, Rodgers, and Ng 2008; Liu, Chen, and Veloso 2009; Ogino, Toyama, and Asada 2007; Park and Kim 2009; Pongas, Mistry, and Schaal 2007; Ratliff, Bagnell, and Srinivasa 2007; Rebula et al. 2007; Sano et al. 2008; Yamaguchi, Takanishi, and Kato 1994). This problem can be divided into two parts: (a) using sensor information to create a model of the surrounding terrain, and (b) constructing controllers to walk on rough terrains. Much of the previous research in the literature have focused on the second part of this problem, such as building reactive feedback controllers to walk on uneven surfaces without an explicit model of the surface. In these approaches, the terrain model is implicitly represented in parameters of the feedback controller and tuned indirectly. In addition, much of this work has been implemented on proprietary hardware platforms, using special hardware to help mechanically stabilize the robot platform (Sano et al. 2008; Hashimoto et al. 2006). More recent work on quadruped locomotion over uneven terrain assumes that a precise model of the terrain is given, and uses motion planning techniques with perfect state and terrain information (Kalakrishnan et al. 2009; Kolter, Rodgers, and Ng 2008; Pongas, Mistry, and Schaal 2007; Ratliff, Bagnell, and Srinivasa 2007; Rebula et al. 2007). Obviously, this has clear limitations in real world implementations where noise and uncertainty are much greater problems.

In this work, we show how to use existing hardware on bipedal robots to address the sensing part of the problem using online machine learning techniques. By incorporating electronic compliance and foot pressure sensors, the swing foot is used to provide noisy estimates of the local gradient of the contact point, and the computed pose of the foot from joint encoders and the inertial measurement unit is used to rapidly learn an explicit model of the surface the robot is walking upon.

Our proposed framework has a number of key advantages. We show how algorithms for surface measurement, terrain model estimation, and adaptive walking control can be separately analyzed and integrated without having to consider the full complicated dynamics of the robot and environment. Our approach also uses direct sensory information from the foot, thus enabling more rapid learning of the surrounding environmental characteristics compared to other methods that depend upon indirect measurement from the torso alone.

To show the effectiveness of our approach, we first demonstrate our algorithms using a freely available robot simulation package that we have developed. This simulation environment allows us to perform many trials to quantify the performance of our approach. In addition, we have performed experiments using a commercially available small humanoid robot platform without any special hardware modifications. In spite of limitations in the hardware

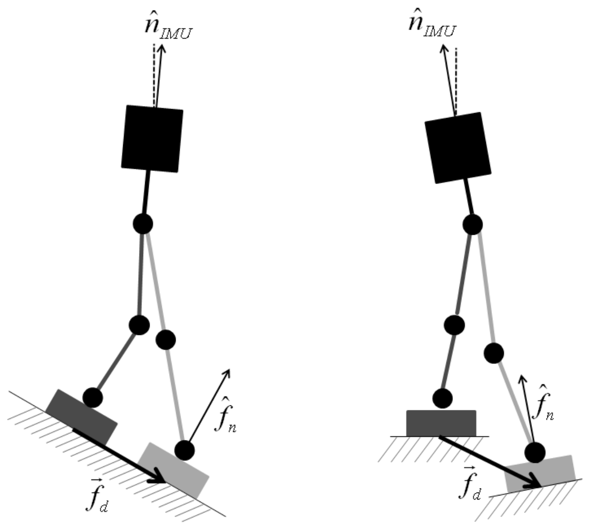

Figure 1: Side and frontal view of surface sensing. Foot position displacement $\vec{f_d}$ and foot normal $\hat{f}_n$ are calculated by forward kinematics using inertial sensors and joint angle readings.

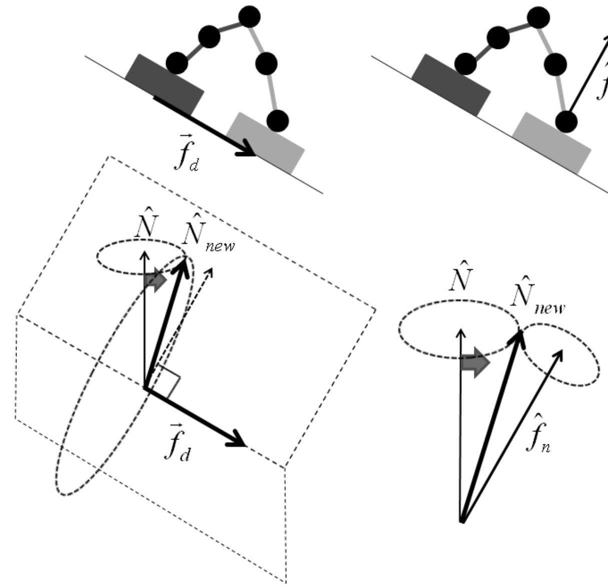

Figure 2: Estimating the surface gradient $\hat{N}$ from foot displacement $\vec{f_d}$ and foot normal $\hat{f}_n$ measurements. For the foot displacement case, we incorporate the constraints that $\hat{N}_{new}$ should lie in the same plane with $\hat{N}$ and $\vec{f_d}$ to uniquely determine $\hat{N}_{new}$.

platform, our results show that our method can rapidly estimate the surface gradient, enabling the robot to successfully walk across unknown surfaces.

This paper is structured as follows. First we describe our method to probe the shape of the walking surface using a compliant swing foot, and show how online learning is used to estimate an explicit surface model in real-time using the measured foot position and angles. Then we demonstrate our walk controller for rough terrain when the surface model is known and estimated from the sensors. We quantify our results from simulation as well as experiments on a humanoid robot, and conclude with a discussion of our findings.

Online Learning of Uneven Terrain

When bipedal animals walk on unknown terrains without visual feedback, they can cautiously sense the local surface using proprioception and force feedback. A foot is placed into an expected landing position until the ground reaction force is felt, and the location and shape of the surface can be estimated from the foot landing position and angle. In situations where the surface is very rough, humans actively probe the surface by sensing a number of positions and the foot trajectory is actively modified according to the learned surface profile. This process is the inspiration for our approach. Instead of relying upon torso-based measurements, we take the foot landing positions and angles as noisy measurements and use them to learn the surrounding surface gradient. We explain this process in more detail here.

Surface sensing using compliant swing leg

In order for the foot to conform to the local surface gradient upon landing, we purposefully give higher compliance to the joints of the swing leg during the latter part of the single support phase. Touchdown is determined using the foot force sensors, and the swing leg is stiffened once the foot has successfully made ground contact. After landing, joint angle values from the encoders are measured, and the pose of the foot is calculated using the forward kinematic model of the robot. An overview of this surface sensing process is shown in Figure 1. The reliability of this calculation depends upon the accuracy of the joint encoders and inertial sensor readings. We compensate for any noise in the estimates of the foot displacement $\vec{f_d}$ and the foot normal $\hat{f}_n$ by using the following online learning algorithm.

Learning the surface model

After the position and angle of the landed foot is determined, we use this information to explicitly estimate the local surface gradient. There are a variety of methods to explicitly model the various terrain surfaces. One possibility is to define the surface as a function of the global 2D position of the robot: $M_{surface} = f(x, y)$.

This requires very good localization capability, as well as building a detailed map of the surrounding environment. Instead we take a simpler approach, using rapid online learning of the local surface gradient via the noisy, built-in sensors. Thus, in this simple approach, the model consists of maintaining an estimate of the local surface normal vector: $M_{surface} = \hat{N}$.

After each footfall, we get two unit vectors as noisy measurements: the foot normal $\hat{f}_n$ and foot displacement $\vec{f_d}$. We will estimate the surface gradient using both of these measurements. In the ideal case, the measured gradient of the landing point will coincide with that of the surface gradient, and the foot displacement vector will be perpendicular to the surface normal:

$$\hat{f}_n = \hat{N} \tag{1}$$

$$\vec{f}_d \cdot \hat{N} = 0 \tag{2}$$

But because of noise in these measurements, we will not use these as hard constraints for the model. Instead, we define the new model N_{new} so that it minimizes an online learning cost function. For the foot normal measurement, we use following cost:

$$C_{fn} = \|N_{new}^{\hat{f}n} - \hat{f}_n\|^2 + \alpha\|N_{new}^{\hat{f}n} - N^{\hat{f}n}\|^2 \tag{3}$$

which results in the following simple update rule:

$$N_{new}^{\hat{f}n} = (1-\alpha')N^{\hat{f}n} + \alpha'\hat{f}_n \tag{4}$$

For the foot displacement measurement, we use following cost:

$$C_{fd} = \|\vec{f}_d \cdot N_{new}^{\hat{f}d}\|^2 + \alpha\|N_{new}^{\hat{f}d} - N^{\hat{f}d}\|^2 \tag{5}$$

Optimizing this cost with minimal norm yields the following update rule:

$$N_{new}^{\hat{f}d} = N^{\hat{f}d} + \alpha''(\vec{f}_d \cdot N^{\hat{f}d})\vec{f}_d \tag{6}$$

Thus, we have two separate ways to estimate the surface gradient from foot measurements, as shown in Figure 2. Both these measurements are inherently quite noisy due to a number of factors. Since the length of the foot is typically longer than its width, and the torso roll angle is constantly varying during dynamic walking, foot roll angle measurements are usually noisier than foot pitch angle measurements. On the other hand, the lateral displacement of the two feet is usually larger than the transverse displacement, so the surface normal estimate from foot displacement has lower roll error than pitch error in many cases. It is thus advantageous to combine two separate filtered estimates, $N^{\hat{f}n}$ and $N^{\hat{f}d}$ using the online learning rules above into a single estimate. In this work, we use a simple weighted filter for this combination, where the weight matrices W_1 and W_2 are determined using cross-validation:

$$\hat{N} = W_1 N^{\hat{f}n} + W_2 N^{\hat{f}d} \tag{7}$$

Because of the simple surface model used, estimation and adaptation of the small number of model parameters is rapidly achieved. We could also incorporate higher order statistics in the model by maintaining a dynamic distribution over measured surface gradients. The walk controller could then be modified accordingly to take into account the measured uncertainty in the surface gradient estimates.

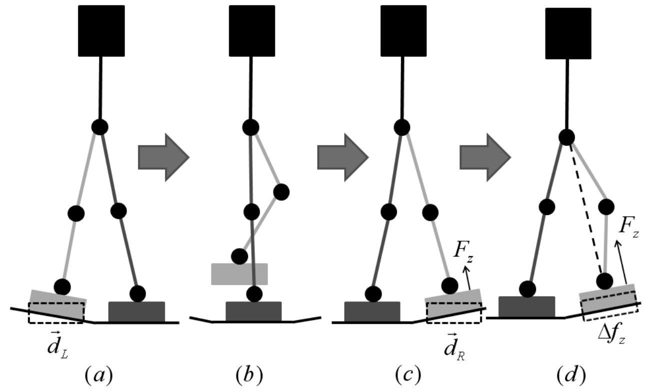

Figure 3: Stepping on uneven terrain. (a) The swing foot starts at ground position with distance variable $\vec{d_L}$. (b) The foot pose follows the computed trajectory during the early phase of the step, and the swing leg becomes compliant during the latter part of the step. (c) Gait timing is delayed until the ground reaction force F_z goes above threshold, and the new disturbance variable $\vec{d_R}$ is measured to keep the torso upright. (d) After landing, the target height of the foot f_z is dynamically adjusted according to the ground reaction force F_z.

Walking Control on Uneven Terrain

In this section we introduce our walking controller that has been extended to handle uneven surfaces. Our basic controller consists of a real-time foot trajectory generator, upper body movement controller using a zero-moment point (ZMP) algorithm, and a fast inverse kinematics solver to obtain desired joint angles. The controller is capable of omnidirectional motion with real-time speed control and a high degree of maneuverability. Figure 3 shows a number of extensions we have made to the controller to enable the robot to walk over uneven surfaces, which are explained in more detail in the following subsections.

Posture control and foot trajectory modification

To ensure that the robot does not tilt on the inclined surface, we use the estimated gradient of the surface to bias the target foot height and angles, and modify the swing foot trajectory accordingly. Furthermore, to allow for local disturbances in surface gradient and height, we modify the walking phase (Hashimoto et al. 2006; Kim, Park, and Oh 2007). Each foot has a local disturbance variable $\vec{d}$ that denotes the difference of the actual foot height and angle from the calculated values using the surface model. The swing foot starts at the previous landed position and angles with nonzero disturbance variable, and it returns to its calculated trajectory as the disturbance variable goes to zero during the early phase of the step. In the latter part of the step, the swing leg becomes compliant to land on the surface with unknown position and gradient. When foot landing is detected by the force sensors, the swing foot stiffens and a new disturbance variable is measured from the pose of the foot.

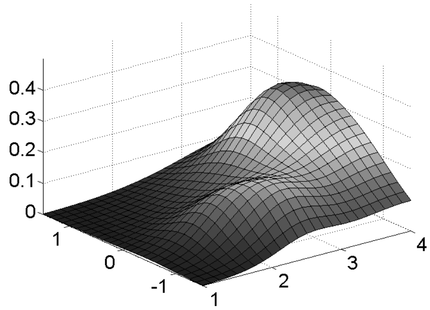

(a)

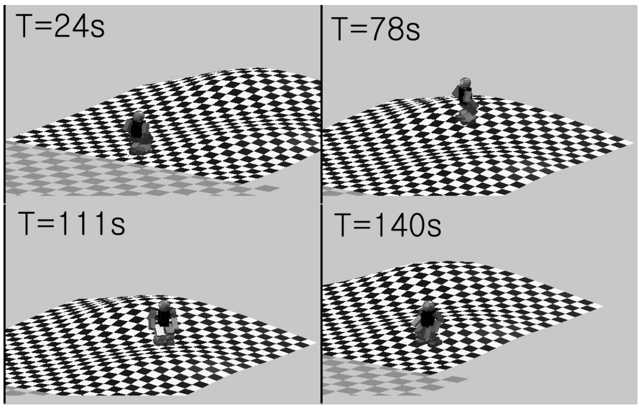

(b)

Figure 4: Simulations of walking on known uneven surfaces. (a) Uneven surface used for the experiment where local inclination exceeds 20 percent. Z axis is scaled to clearly show the shape of the surface. (b) Snapshots of omnidirectional walk trial on the uneven surface.

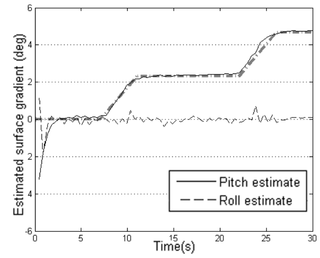

(a)

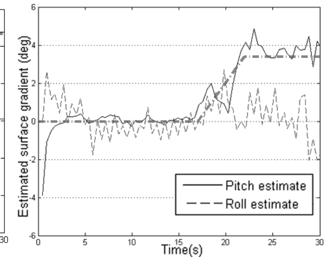

(b)

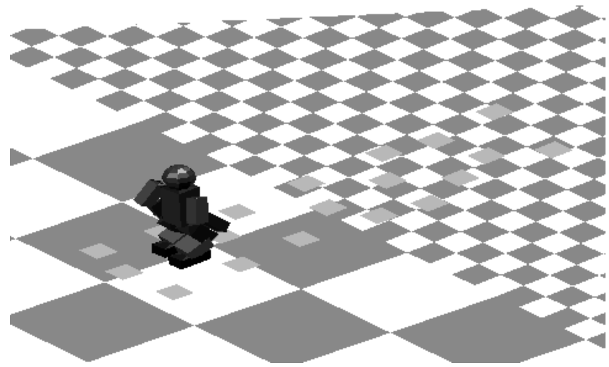

(c)

Figure 5: Simulated walking experiment on inclined surface. (a) Estimated surface pitch and roll angles for a surface with changing inclinations (0%, 4%, 8%). Dash-dotted line shows the real surface gradient. (b) Estimated surface pitch and roll angles for a surface with changing inclinations (0%, 6%) as well as local disturbances. Dash-dotted line shows the real surface gradient. (c) A snapshot of walking on terrain with local disturbances.

Adaptive gait timing

When the robot walks on an uneven surface, there are errors between the expected surface height and actual surface height. This results in landing timing errors, with the foot landing too early or the foot landing too late. If we do not account for these timing errors explicitly, it can seriously impact the stability of walking. In our controller, early landing is handled implicitly in the foot trajectory modification algorithm described above. In the latter part of the walking phase, the swing foot is lowered until it lands on the ground, and stops when the foot lands on the surface. During early landing, the robot stops the swing foot at the moment of landing, effectively preventing the instability induced by continually lowering the swing foot after early landing. Late landing is handled via an extended foot trajectory curve to ensure the swing foot is lowered until touchdown, and the phase of the locomotion gait is prolonged accordingly.

Virtual compliance control

Impact during foot landing can make the robot shake laterally and make walking unstable. This problem is mitigated by using a foot trajectory curve with zero velocity at landing, and fast feedback control to allow the foot to follow the precalculated trajectory very closely. However we have found that due to uncertainty in motors and gear backlash, such a precise control is very difficult to achieve in practice. There will also be footfall disturbances due to rapidly changing terrain and errors in the underlying estimated surface model. One way of handling this is to incorporate special hardware compliance in the robot joints (Sano et al. 2008; Yamaguchi, Takanishi, and Kato 1994). In the commercial platform used in our experiments, such a hardware modification was not possible. Instead, we implemented a software-controlled compliance by using virtual compliance in the walk controller (Kim, Park, and Oh 2007; Kim et al. 2007; Mikuriya et al. 2005) Based upon the ground reaction force F_z measured from the foot, the target height of the foot f_z is dynamically adjusted to mimic a spring and damper system:

$$\Delta f_z(t_n) = k_1 F_z + k_2 \frac{\Delta f_z(t_{n-1}) - \Delta f_z(t_{n-2})}{t_{n-1} - t_{n-2}} \qquad (8)$$

Simulation Results

We first ran trials on a simulation environment incorporating hypothetical terrains. These results are based upon the open-sourced simulation environment we built integrating the Open Dynamics Engine (ODE) with Matlab-based controllers and graphics.

Our first experiment tests the performance of our walking controller on an fully known uneven terrain. We made a test surface modeled with a mixture of Gaussians and let the robot walk over it using accurate information about the robot pose and surface model. In our simulations, the following parameters were used: step size 0.06 m, step height 0.04 m, double support phase ratio 0.5. Figure 4 shows the result of the test. We see that with the known model of the surface, our controller is capable of stable omnidirectional walking on uneven surfaces where the pitch exceeds 20 percent grade. On the other hand, the baseline walk controller, not considering unevenness of the surface, makes the robot fall down at the initial stage of the trial.

We proceeded to test the validity of our surface gradient learning algorithm using a test surface with a number of different inclinations. Figure 5(a) shows the estimated surface gradient during walking. and we see that our method is capable of learning the surface gradient rapidly and reliably in this simulated environment.

Finally we have tested our method on an uneven surface with local disturbances. The test surface is shown in Figure 5(c), which consists of two planes with inclination zero and six percent, with a number of small millimeter thick plates placed randomly on top. Considering that the humanoid model we use has feet width of 70mm, one millimeter of local height difference may induce as much as 3% of local gradient, which is sufficient to make the walking unstable from our experience.

Although the estimated surface gradient shown in Figure 5(b) shows higher variance with the presence of local disturbances, the robot has no difficulty in walking stably on the surface in the presence of both changing surface normals and local disturbances.

Experimental Results

We have also tested our algorithms using a commercially available small humanoid robot, the Nao robot from Aldebaran Robotics[1]. The Nao robot is 58cm tall, weighs 4.3kg, and has 21 joints. It has a number of sensors, including cameras, inertial accelerometer and gyroscopic sensors, ultrasonic range sensors, foot bumper sensors and pressure sensors. It can operate untethered using a built-in lithium polymer battery, and algorithms are run on the embedded Geode-based Linux computer.

As a mass-produced, low-budget robot platform, the Nao has some limitations on maximum motor torques and backlash in its gear trains. Combined with noise in the sensors and limits on control update rates, the platform poses several challenges for successful implementation of control and learning algorithms. Nevertheless, to show that our method is applicable on such a system, we first let the robot walk on a flat surface using a walking controller for flat surfaces, and observed how well the robots learns the flat surface model. We used the same controller as used in our simulation experiments, but with a lower step size of 0.04 m, and a higher double support phase ratio of 0.75.

[1] http://www.aldebaran-robotics.com

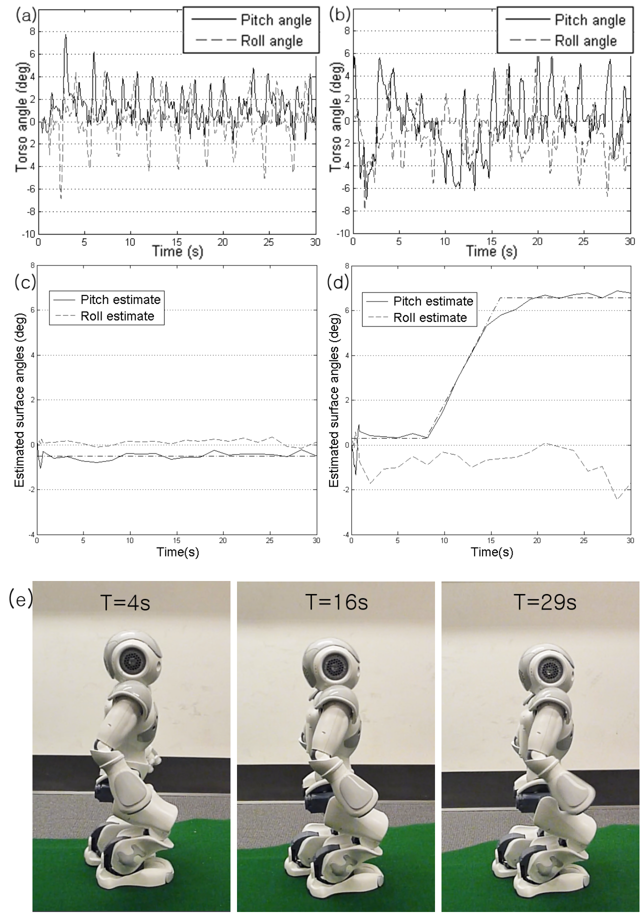

Figure 6: Results of robot walking experiment. (a) Torso angle readings during walking on a flat surface. (b) Torso angle readings during walking on a surface with changing inclination. (c) Estimated surface pitch and roll angles during walking on a flat surface. (d) Estimated surface pitch and roll angles during walking on a surface with changing inclination. (e) Snapshots of Nao robot walking on a surface with changing inclination.

Figure 6(a) shows the torso angle measured by the inertial sensor. We see that the dynamics of the Nao robot are more compliant than in simulation, resulting in higher variance of torso pitch angle. However, as seen in Figure 6(c), the estimated surface gradient remained within one degree of error, which shows that our learning algorithm is accurate enough even with the noisy sensors on the Nao.

After this initial test, we proceeded to test our method on a surface of changing inclination. The testing surface is made of two wooden planks with inclination of roughly zero and ten percent, and we let the robot walk over the transition to see how rapidly our method can estimate the changing surface gradient and make the robot walk stably on both sets of surfaces. Snapshots of this experiment are shown in Figure 6(e). Figure 6(b) shows the torso angle measured by the inertial sensor, and Figure 6(d) shows the estimated surface gradient during walking.

We see that our method can estimate the new surface gradient after a couple of steps, and although the torso angle shows higher variance than the flat surface case, our method can let the robot walk stably over the uneven surface. On the other hands, baseline walk controller, not considering the unevenness of the surface, makes the robot unstable and eventually lets it fall down on the inclined surface.

Conclusions and Future Work

We have proposed a method for having a humanoid robot walk over an unknown, uneven surface. We use compliance in the swing leg to actively probe the surface, and incorporate noisy measurements from built-in sensors in an online learning algorithm to model the surrounding surface. Finally, we use the estimated surface model to enable the humanoid robot to walk stably on the uneven terrain.

We have tested our method in simulation, showing that our method can rapidly estimate and walk on uneven surfaces. Our algorithms are also robust enough to handle local disturbances in simulation. We have also tested our method on the commercial Nao humanoid robot without any special modifications. Our experiments show that our methods can quickly learn surface gradients and enable this robot to walk stably on uneven surfaces, in spite of some limitations of the hardware platform. As the accuracy and speed of the learning is largely dependent on the platform, we think our approach can handle more diverse terrain types with more powerful platforms.

There is much room for improvement within our framework. We can explicitly account for additional uncertainty in the environment by incorporating prior knowledge from vision or ranging sensors in a Bayesian framework. Also, we are currently working to implement these methods on the much larger Hubo humanoid platform.

Acknowledgment

We acknowledge the support of the NSF PIRE program under contract OISE-0730206. This work was also partially supported by IT R&D program of MKE/KEIT (KI002138, MARS), NRF grant of MEST (314-2008-1-D00377, Xtran), and the BK21-IT program of MEST.

References

Hashimoto, K.; Sugahara, Y.; Ohta, A.; Sunazuka, H.; Tanaka, C.; Kawase, M.; Lim, H.; and Takanishi, A. 2006. Realization of stable biped walking on public road with new biped foot system adaptable to uneven terrain. In *BioRob*, 226–231.

Hyon, S.-H. 2009. Compliant terrain adaptation for biped humanoids without measuring ground surface and contact forces. *IEEE Transactions on Robotics* 25(1):171–178.

Jenkins, O.; Wrotek, P.; and McGuire, M. 2007. Dynamic humanoid balance through inertial control. In *IEEE Intl. Conf. on Humanoid Robots*.

Kalakrishnan, M.; Buchli, J.; Pastor, P.; and Schaal, S. 2009. learning locomotion over rough terrain using terrain templates. In *IEEE international conference on intelligent robots and systems*.

Kim, Y.-D.; Lee, B.-J.; Ryu, J.-H.; and Kim, J.-H. 2007. Landing force control for humanoid robot by time-domain passivity approach. *IEEE Transactions on Robotics* 23(6):1294–1301.

Kim, J.-Y.; Park, I.-W.; and Oh, J.-H. 2007. Walking control algorithm of biped humanoid robot on uneven and inclined floor. *J. Intell. Robotics Syst.* 48(4):457–484.

Kolter, J.; Rodgers, M.; and Ng, A. 2008. A control architecture for quadruped locomotion over rough terrain. In *ICRA*, 811–818.

Liu, J.; Chen, X.; and Veloso, M. 2009. Simplified walking: A new way to generate flexible biped patterns. In *CLAWAR*.

Mikuriya, Y.; Sunazuka, H.; Hashimoto, K.; Sugahara, Y.; Kawase, M.; Lim, H. O.; Hosobata, T.; and Takanishi, A. 2005. Walking control method of biped locomotors on inclined plane. In *ICRA*, 1977–1982. IEEE.

Ogino, M.; Toyama, H.; and Asada, M. 2007. Stabilizing biped walking on rough terrain based on the compliance control. In *IROS*, 4047–4052. IEEE.

Park, J. H., and Kim, E. S. 2009. Foot and body control of biped robots to walk on irregularly protruded uneven surfaces. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics* 39(1):289–297.

Pongas, D.; Mistry, M.; and Schaal, S. 2007. A robust quadruped walking gait for traversing rough terrain. In *ICRA*, 1474–1479.

Ratliff, N.; Bagnell, J. A. D.; and Srinivasa, S. 2007. Imitation learning for locomotion and manipulation. In *IEEE-RAS International Conference on Humanoid Robots*.

Rebula, J.; Neuhaus, P.; Bonnlander, B.; Johnson, M.; and Pratt, J. 2007. A controller for the littledog quadruped walking on rough terrain. In *ICRA*, 1467–1473.

Sano, S.; Yamada, M.; Uchiyama, N.; and Takagi, S. 2008. Point-contact type foot with springs and posture control for biped walking on rough terrain. In *AMC*, 480–485.

Yamaguchi, J.; Takanishi, A.; and Kato, I. 1994. Development of a biped walking robot adapting to a horizontally uneven surface. In *IROS*, volume 2, 1156–1163 vol.2.