A Model of Visually Triggered Gait Adaptation

M. Anthony Lewis and Lucia S. Simo

Iguana Robotics, Inc.
Mahomet, IL 61853

tlewis@iguana-robotics.com
http://www.iguana-robotics.com/

Abstract-*Walking machines can walk over obstacles without touching them only if they can anticipate contact and make suitable gait modifications. Existing visually guided machines use computationally intense approaches that require construction of a geometrically correct model of both the environment and the robot. We present a model, inspired by research in vertebrates, in which the stride length modification that takes place before stepping over the obstacle is learned based on experience. The key hypothesis introduced here is the use of temporal gating of the visual signal encoding the distance to an obstacle. This hypothesis enables the formulation of the problem as a direct mapping of perception to action. In addition, the use of temporal gating also facilitates learning by simplifying the credit assignment problem. Our approach does not require that a geometric representation of the environment be created and updated based on new observations. Our simulation results indicate that the desired mapping can be learned quickly. The resulting gait modulation is smooth and coordinated with the phase of the central pattern generator controlling the robot. Our model qualitatively reproduces human data where the uncertainty in footsteps decreases with approach to an object.*

1 Introduction

Locomotion and perception have been treated as separate problems in the field of robotics. Under this paradigm, one first solves the 'vision problem' of recovering the three-dimensional geometry of the scene. This information is then passed to a planning system that has access to an explicit model of the robot. A good trajectory is found for each individual leg to move over the obstacle. This solution is computationally intense and, as demonstrated for the Ambler walking machine (Krotkov and Hoffman, 1994; Krotkov and Simmons, 1996), too slow for real-time control using moderate power CPUs. Furthermore, this approach does not exploit the fact that the walking machine will be presented with a similar situation again and again.

The approach considered here is to eliminate the intermediate explicit model and consider creating a direct coupling of perception to action, with the mapping being adaptive and based on experience. For this approach we use a temporal gating hypothesis by which sensory data (distance to the object) is temporally gated to modify the output of the locomotor controller.

Recently, a number of studies have pointed out the necessity of gating mechanisms to control the flow of sensory signals in the brain of vertebrates (Prochazka, 1989; Chapman, 1994; Apps, 1999). In particular, temporal gating during a visual discrimination task prevents extraneous signals occurring around the time of the critical visual event to affect performance (Seidemann et al., 1998).

Continuous visual input is not necessary for accurate stepping. Not all visual samples have the same potential for control limb movements. Samples taken when the foot to be controlled is in stance are by far more effective in modulating gait. It has been suggested that during stepping visual information is used during the stance phase in a feedforward manner to plan and initiate changes in the swing limb trajectory (Holland and Marple-Horvat, 1996; Patla et al., 1996).

Finally, behavioral studies in humans have shown that the regulation of the step depends on the distance to the obstacle. Data from athletes in the long jump have demonstrated that just prior to lift-off the athlete modulates his/her stride length over the last three steps (Lee et al., 1982). Also, the standard deviation of the footsteps decreases over the last three steps.

Taken together, this may indicate that gait is modulated at discrete intervals. This modulation may be a highly stereotyped program that depends on a brief sampling of the visual environment to instantiate it (c.f. Patla et al., 1991). This hypothesis is intriguing because it implies that after a brief sample it is not necessary to store an internal representation of the world that needs to be shifted and updated during movement.

This shifting and updating is problematic for both neural and traditional robotics models.

2 Elegant Stepping Model

The adaptation problem that we will address can be described abstractly as follows. We wish to make associations between a distance to the obstacle and a change in stride length. We wish to adjust this mapping adaptively and based on experience.

We choose the occurrence of a paw extension and paw placement reflex as training signals. If a reflex is triggered while the leg is extending, then the paw had almost cleared the obstacle. In this case we adjust previous associations between distance and stride length to make longer strides in the future. If a paw placement reflex is triggered when the leg is flexing, we adjust the previous associations between distance and stride length to make shorter strides in the future.

One key difficulty in learning is how to propagate the error back in time in a biologically plausible way. Note that visual information flows into the animal's eyes continuously. However, we note that changes in the step cycle are most effective during narrow time windows. Therefore, we hypothesize that sensory information from visual areas (e.g. distance) is gated periodically and in synchrony with the step cycle. This is our temporal gating hypothesis. This information is then held, decaying exponentially, and is used to modulate the gait over the following step cycle. Thus, as the robot approaches an obstacle, it makes at most three discrete decisions prior to going over the obstacle. These decisions occur at the three footsteps prior to going over the obstacle. This discretization simplifies the credit assignment problem.

The model has four main parts, referring to Fig. 1:

(1) *Range Encoder*— encodes distance to the obstacle using nonoverlapping cells. No spatial ordering of units is assumed. These elements are gating into short-term memory.

(2) *Locomotory Generator*— the central pattern generator (CPG) is modeled as a ring oscillator (Lewis, 1996) that drives two output functions. One drives the muscle of the leg and the other indicates the beginning of each step cycle and is used for the sensory gating. In addition, a "lift reflex" increases the amplitude of the CPG output and is hardwired.

(3) *Mechanical System*— this is the model of environment/leg interaction. We simulate the muscle as a low-pass filter. This muscle drives the flexion of one degree of freedom leg. Each obstacle is simulated as being a rectangle.

(4) *Learning System*— the activity of the units in the range encoder are one-to-one gated into short-term memory cells (STM) in synchrony with the step cycle. The gate used to accomplish this is a shunting inhibition signal originating in the CPG. An adaptive premotor module receives a weighted signal from the STM, and controls the stride length by modulating the burst length (parameter of the CPG controlling the flexion of the leg). The STM activates synapses in the adaptive module. Traces in these synapses maintain a brief memory of having being activated. If a reflex is triggered, then a heuristic is used to modify the weights of the adaptive module. If a paw placement reflex has occurred, then all synapses contributing to this decision should be incrementally decreased. If a paw extension reflex occurs, they should be increased.

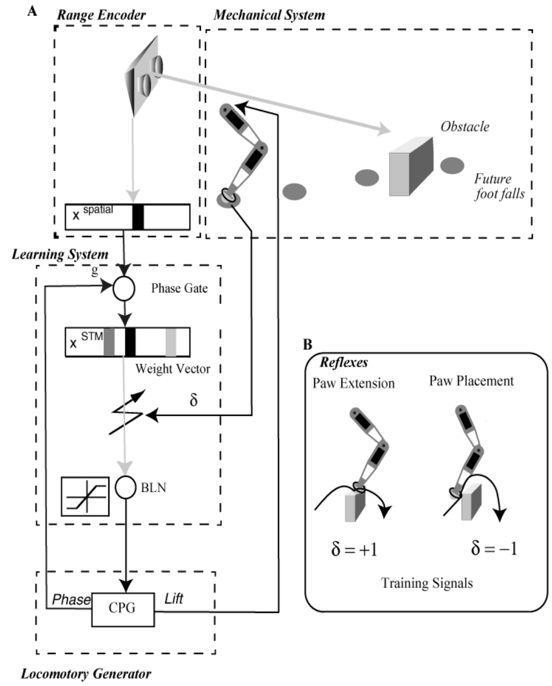

Figure 1. Model of Elegant stepping. (A)Visual input is periodically gated into short term memory by a phase signal ascending from the CPG. Short term memory elements are weighted and the resulting output is used to shorten or lengthen the stride length. (B) Certain reflexes signal error conditions. This supervisory signal is propogated back through time, in a biologically plausible way, and adjusts the short term memory weights.

3 Simulation Experiments

Figure 2 shows a typical foot trajectories before and after learning. As can be seen, the adaptive gait allows the foot to be in a position to clear the obstacle. If stride length adjustments are not made, it may be nearly impossible for the leg to clear the obstacle.

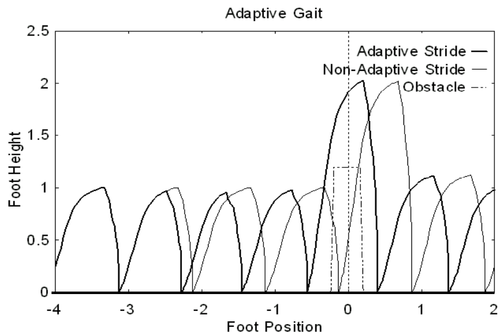

Figure 2. Typical gait trajectories. Example of gait trajectory before and after learning.

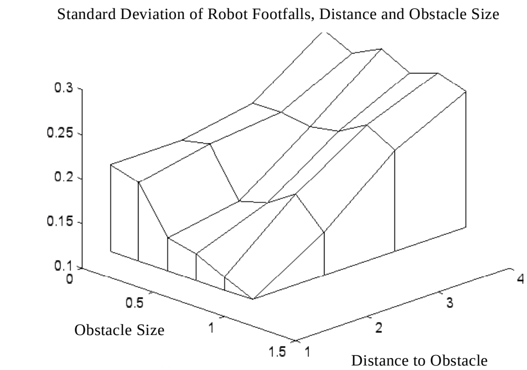

Figure 3. Standard deviation for varying object sizes and distances to object. As the robot approaches the obstacle, its foot fall variance decreases. Variance also decreases with increasing object size.

Notice that the stride length is adjusted three times before the animal reaches the obstacle.

The learning takes place quickly. The algorithm performs well after about 20 training cycles. After about 100 trials, no more mistakes are made in the gait. Learning is smooth. As the robotic leg moves toward the obstacle, the burst length (parameter of the CPG controlling the flexion of the leg) is gradually altered. After passing the obstacle, the burst length gradually relaxes to its former value. Thus, the gait is altered smoothly.

Interestingly, the variance in footsteps decreases as the robotic leg approaches the obstacle (Fig 3). Just as in long jump athletes the standard deviation of the footsteps decreases just before the final footstep. Thus, the robot found a ‘sweet’ spot to land on just before going over the obstacle. Furthermore, the variance in footsteps also decreases with increasing object size. The ‘sweet’ spot is small if the object is large.

The weight distribution after learning is periodic (Fig. 4). The perceptual space is divided into periodic regions.

4 Discussion

In our model perception and action are tightly coupled. The mapping is adaptive and based on experience. The goal of the adaptation is to use distance measurements to smoothly modulate a CPG controlling gait. A key element in our model is the use of a temporal gating hypothesis which simplifies the learning problem.

Our approach does not require that a geometric representation of the environment be created and updated. This is in strong contrast to current practice in machine vision and robotics of surface reconstruction as a prerequisite to planning.

4.1 Separation of Obstacle Clearance and Stride Length adaptation.

The model presented here separates the task of stepping over the obstacle into two components: stride length adjustment and foot elevation going over the obstacle. The focus of the adaptation in the model is the stride length adjustment. It can be argued that if the stride length is adjusted in anticipation of the obstacle, the task of stepping over the obstacle will be easier. Thus, there is some interaction between the two components. If stride length is poor, then the final step may fail.

Future work should entail strategies for learning the sensory motor transformation for the last step. That is, how does the animal step over the obstacle, while, presumably, optimizing other criteria such as stability, comfort and perhaps energy usage.

Currently, information about the height of the object only impacts stride length. Training occurs for a given set of weights for a single object height only. In the future, object size should be used to give the weights a certain context.

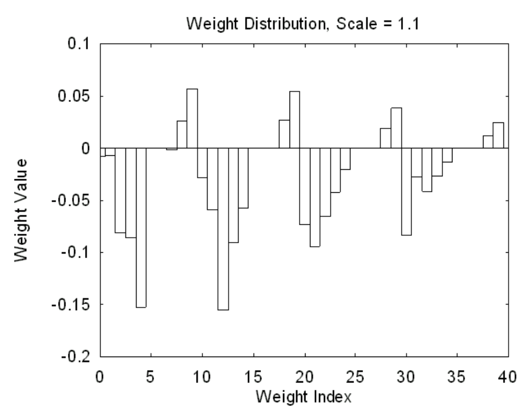

Figure 4. Typical weight distribution after learning. Notice that the weights, or adjustment commands are periodic. Here the weights are ordered according to their correspondence to distances from the obstacle. This ordering is done for the sake of presentation. The algorithm does not assume any particular structure of the sensory space.

4.2 Temporal Gating

Recent studies in cats suggest that during stepping over obstacles premotor signals from the motor cortex may be gated onto the spinal CPG network in synchrony with the step cycle (Drew et al., 1996).

In our model sensory signals (distance) are gated in synchrony with the step cycle. This is our temporal gating hypothesis. Sensory gating has been shown for signals exiting the middle temporal visual area during a visual discrimination task (Seidemann et al., 1998). In addition, movement-related gating of sensory input to the cerebellum via climbing fibers has also been suggested (Apps, 1999).

A recent article by Taga (1996) addresses the problem of adjusting the parameters of a CPG so that a biped figure is able to walk over an object. In that paper Taga proposes a method for synchronizing corrective input to muscles in synchrony with the step cycle. He was inspired by work of Drew and others (see Drew et al., 1996 for a review).

The Taga work is a complement to the work presented here. While our model address the acquisition of a visuomotor mapping, and proposes *sensory gating.* The Taga work supposes that there is a kind of *motor gating of* commands to the muscles. These are compatible interpretations. In practice our adjustment signal might need to be broken up into discrete time intervals to control individual muscles. While our CPG system is rudimentary, Taga's CPG is more complex. It is likely that Taga's biomechanical model could be substituted for the rudimentary biomechanical model presented here. The results should be similar even with a more complex model.

Secondly, the Taga model is concerned with the details of stepping over the object, the system described here considers changes in stride length before the animal or robot reaches the obstacle. We design a system that assumes such programs exists. We are concerned with providing input parameters to such a motor program.

Finally, the model presented by Taga is not concerned with learning.

4.3 Learning Visuomotor Behavior

Asada et al. (1996) describe a system that uses reinforcement learning to automatically generate associations between perceptual stimuli and action. In general the reinforcement learning problem is more difficult than the problem examined here. Using some knowledge of the problem, we were able to deduce that a reflex signal would be an ideal training input. This signal gives the algorithm feedback as to *what* it should do when an error occurs. Thus the learning algorithm used here is a supervised learning problem. The formulation of the this problem as a supervised learning problem undoubtedly accounts for the quick learning observed.

Acknowledgements

The authors acknowledge support of Grant No. N00014-99-0984 from the Office of Naval Research.

5 REFERENCES

Apps, R. (1999) Movement-related gating of climbing fibre input to cerebellar cortical zones. *Progress in Neurobiology*, **57**, 537-562.

Asada, M., Noda, S., Tawaratsumida, S. and Hosada K. (1996) Purposive Behavior Acquisition for a Real Robot by Vision-Based Reinforcement Learning, *Machine Learning*, **23**, 279-303.

Chapman C.E. (1994) Active versus passive touch: factors influencing the transmission of somatosensory signals to primary somatosensory cortex. *Canadian Journal of Physiology and Pharmacology*, **72**, 558-570.

Drew T., Jian W., Kably B., and Lavoie S. (1996) Role of the motor cortex in the control of visually triggered gait modifications. *Canadian Journal of Physiology and Pharmacology*, **74**, 426-442.

Grillner S. and Wallén P. (1985) Central pattern generators for locomotion, with special reference to vertebrates. *Annual Review of Neuroscience*, **8**, 233-61.

Hollands, M.A. and Marple-Horvat, D.E. (1996) Visually guided stepping under conditions of step cycle related denial of visual information. *Experimental Brain Research*, **109**, 343-356.

Krotkov, E. and Simmons, R. (1996) Perception, planning and control for autonomous walking with the ambler planetary rover. *IEEE Transactions on Robotics and Automation*, **15**, 155-180.

Krotkov, E. and Hoffman, R. (1994) Terrain Mapping for a walking planetary rover. *IEEE Transactions on Robotics and Automation*, **10**, 728-736.

Lee, D.N., Lishman, J.R. and Thomson, J.A. (1982) Regulation of gait in long jumping. *Journal of Experimental Psychology: Human Perception and Performance*, **8**, 448-459.

Lewis, M. A. Self Organization of Locomotory Controllers in Animals and Robots, Ph.D. Dissertation, Electrical Engineering Department, University of Southern California, Los Angeles, 1996.

Patla, A.E., Prentice, S.D., Robinson, C. and Neufeld, J. (1991) Visual control of locomotion: strategies for changing direction and for going over obstacles. *Journal of Experimental Psychology: Human Perception and Performance*, **17**, 603-634.

Patla, A.E., Adkin, A., Martin, C., Holden, R. and Prentice, S.D. (1996) Characteristics of voluntary visual sampling of the environment for safe locomotion over different terrains. *Experimental Brain Research*, **112**, 513-522.

Prochazka A. (1989) Sensorimotor gain control: a basic strategy of motor systems? *Progress in Neurobiology*, **33**, 281-307.

Seidemann E, Zohary E, and Newsome WT (1998) Temporal gating of neural signals during performance of a visual discrimination task. *Nature*, **394**, 72-75.